

SCRT
30.1

Using the Sub-Capacity Reporting Tool



Note

Before you use this information and the product it supports, read the information in [“Notices” on page 531](#).

This edition applies to version 30, release 1, modification 0 of the Sub-Capacity Reporting Tool (SCRT) and to all subsequent releases and modifications until otherwise indicated in new editions. This release of SCRT reports on the execution-based sub-capacity products that run on the following operating system products and all subsequent releases and modifications of these products until otherwise indicated in new editions:

- Version 3 Release 1 of IBM® z/OS® (5655-ZOS) and later
- Version 2 Release 1 of IBM z/OS (5650-ZOS) and later
- Version 1 Release 1 of IBM z/OS (5694-A01) and later
- Version 1 Release 3 of IBM z/OS.e (5655-G52) and later
- Version 1 Release 1.0 of IBM z/TPF Enterprise Edition (5748-T15) and later
- Version 6 Release 3.0 of 21st Century Software VSEⁿ with restrictions
- Version 6 Release 1.0 of IBM z/VSE® (5686-VS6) and later
- Version 5 Release 1.0 of IBM z/VSE (5609-ZV5) and later
- Version 4 Release 1.0 of IBM z/VSE (5609-ZV4) and later

This edition replaces SC23-6845-21.

Last updated: 2025-09-20

© **Copyright International Business Machines Corporation 2000, 2025.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Tables.....	xi
About this document.....	xv
Who should read this document?.....	xv
Skill requirements for SCRT users.....	xv
Where to find more information.....	xvi
Summary of changes.....	xix
Summary of changes for SC23-6845-22, SCRT 30.1.....	xix
Part 1. SCRT concepts and planning.....	1
Chapter 1. About the Sub-Capacity Reporting Tool.....	3
What is SCRT?.....	3
What is the IBM Z software pricing framework?.....	4
Monthly license charge (MLC) pricing.....	5
IBM Z International Product License Agreement (IPLA) Software.....	6
Select Application License Charges.....	6
IBM Z New Application License Charges (zNALC).....	7
What types of sub-capacity pricing does SCRT support?.....	7
How does sub-capacity pricing work?.....	8
Overview of the SCRT process.....	8
Using SCRT for planning purposes.....	10
Chapter 2. Preparing to use the Sub-Capacity Reporting Tool.....	11
Understanding the relationship between LPARs and sub-capacity pricing.....	11
Preparing your hardware and software.....	11
Using z/OS SMF records with SCRT.....	12
Using SCRT89 records with SCRT.....	12
How SCRT uses SMF and SCRT89 records.....	13
Collecting z/OS SMF records.....	13
Collecting SCRT89 records.....	13
Extracting z/OS SMF records you have collected.....	13
Extracting SCRT89 records you have collected.....	14
Retaining SMF and SCRT89 records you have collected.....	14
SMF and SCRT89 records from different time zones.....	14
Operating systems that run with the time set to the future or the past.....	15
How SCRT determines the license capacity.....	15
What is the 4-hour rolling average utilization?.....	16
How the IBM zAAP and IBM zIIP affect sub-capacity pricing.....	16
The SCRT reporting period.....	16
SCRT data collection reporting.....	17
Unusual situations that affect utilization.....	17
Part 2. Using the Sub-Capacity Reporting Tool.....	19
Chapter 3. Using SCRT on z/OS systems.....	21
Requirements for SCRT.....	21
Obtaining and installing SCRT on a z/OS system.....	21
Running SCRT on a z/OS system.....	24

Customizing the SCRT JCL.....	25
SCRT messages.....	67
Chapter 4. Using SCRT on Windows and Linux systems.....	69
About SCRT for Windows and Linux.....	69
Requirements for SCRT for Windows and Linux.....	69
Obtaining and installing SCRT for Windows and Linux.....	70
Installing SCRT in console mode.....	71
Transferring SMF or SCRT89 data from a host system.....	71
Using the SCRT graphical user interface.....	73
Running the SCRT graphical user interface.....	74
Using the SCRT command line interface.....	90
Running the SCRT command line interface.....	91
Sample SCRT command line invocation.....	97
Using the SCRT CLI sample scripts.....	97
Part 3. Understanding and interpreting sub-capacity reports.....	101
Chapter 5. Sub-capacity reporting for z/OS native systems.....	103
Identities reported for z/OS native systems.....	103
Advanced image naming.....	104
Sample base z/OS single-version configuration.....	104
Sample sub-capacity report for single-version z/OS native systems.....	105
Billing-related section.....	105
Customer verification section	118
Sub-capacity reporting for multiple versions of z/OS.....	127
Chapter 6. Sub-capacity reporting for z/OS systems that are z/VM guests.....	133
Summary of MSUs reported for sub-capacity eligible products running on z/OS guest systems under z/VM.....	133
Capacity examples for z/OS guest systems.....	134
Example 1: Single LPAR with all guests at z/OS V1.6 or lower.....	134
Example 2: Single LPAR with all guests at z/OS V1.7 or higher.....	134
Example 3: Single LPAR with mixed z/OS guest systems.....	135
Example 4: Multiple LPARs on the same CPC; one with mixed z/OS guests and one with only z/OS V1.7 guests.....	136
Example 5: Multiple LPARs on the same CPC; one with mixed z/OS guests and one with a native z/OS V1.7 system.....	137
Example 6: Multiple LPARs on the same CPC; one with z/OS guests at V1.7 and one with a native z/OS system at any level.....	137
Identities reported for z/OS guest systems.....	138
Sample report showing any z/OS V1.6 or earlier guest systems.....	139
Interpreting the Detail Data sections for V1.6 or earlier guests.....	141
Sample sub-capacity report showing only z/OS V1.7 (or later) guest systems.....	141
Billing-related fields in the sub-capacity report for guest systems.....	142
Customer verification section	142
Chapter 7. Sub-capacity reporting for System z New Application License Charges (zNALC).....	145
Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads...	145
Sample sub-capacity report for z/OS V1 and z/OS V2 running traditional and zNALC workloads..	149
Sample product summary information section for a CPC only running zNALC workloads.....	153
Sub-capacity reporting for z/OS zNALC guest systems.....	153
Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in different LPARs.....	154
Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in the same LPAR.....	157

Sample sub-capacity report for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads.....	161
Chapter 8. Sub-capacity reporting for Getting Started Sub-Capacity Pricing (GSSP).....	165
Considerations for Getting Started Sub-Capacity Pricing.....	165
Restrictions for z/OS 1.6 guest systems.....	165
Restrictions for partially supported products.....	165
Restrictions for power save mode.....	166
Considerations for WebSphere Transformation Extender V8 (5655-R95).....	166
Getting Started Sub-Capacity Pricing algorithm.....	168
Sub-capacity report externals for Getting Started Sub-Capacity Pricing.....	168
Getting Started Sub-Capacity Pricing footnote.....	169
Getting Started Sub-Capacity Pricing report sections.....	169
Suppress_Report_Section control statement.....	170
Getting Started Sub-Capacity Pricing trace records.....	171
Sample configuration for Getting Started Sub-Capacity Pricing.....	173
Sample control statement input to SCRT for Getting Started Sub-Capacity Pricing.....	174
Sample sub-capacity report for Getting Started Sub-Capacity Pricing.....	174
Billing-related section.....	175
Customer verification section.....	177
Chapter 9. Sub-capacity reporting for Integrated Workload Pricing (IWP).....	183
Qualifications for Integrated Workload Pricing.....	185
SCRT externals for Integrated Workload Pricing.....	186
Integrated_Workload_Pricing control statement.....	186
Addition to the Special Conditions section.....	186
Integrated Workload Pricing footnote.....	187
Integrated Workload Pricing trace records.....	187
Sample configuration for Integrated Workload Pricing.....	188
Sample control statement input to SCRT for Integrated Workload Pricing.....	189
Sample sub-capacity report for Integrated Workload Pricing.....	189
Billing-related section.....	190
Customer verification section.....	192
Chapter 10. Sub-capacity reporting for Mobile Workload Pricing (MWP).....	199
Overview of Mobile Workload Pricing.....	199
Qualifications for Mobile Workload Pricing.....	201
Data collection requirements for Mobile Workload Pricing.....	202
MWP monthly reporting process.....	202
SCRT externals for Mobile Workload Pricing.....	203
SCRT externals for INPUTCSV-based Mobile Workload Pricing.....	203
SCRT externals for WLM classification-based Mobile Workload Pricing.....	205
Reporting sample for INPUTCSV-based Mobile Workload Pricing.....	207
Sample configuration for INPUTCSV-based Mobile Workload Pricing.....	207
Sample control statement input to SCRT for INPUTCSV-based Mobile Workload Pricing.....	208
Sample sub-capacity report for INPUTCSV-based Mobile Workload Pricing.....	209
Reporting sample for WLM classification-based Mobile Workload Pricing.....	217
Sample configuration for WLM classification-based Mobile Workload Pricing.....	217
Sample control statement input to SCRT for WLM classification-based Mobile Workload Pricing.....	218
Sample sub-capacity report for WLM classification-based Mobile Workload Pricing.....	219
Chapter 11. Sub-capacity reporting for IBM Z Collocated Application Pricing (zCAP).....	229
Qualifications for IBM Z Collocated Application Pricing.....	232
Data collection requirements for IBM Z Collocated Application Pricing.....	233
zCAP monthly reporting process.....	233
SCRT externals for IBM Z Collocated Application Pricing.....	233
INPUTCSV DD statement for zCAP.....	234

IBM Z Collocated Application Pricing footnote.....	234
Generate_Detailed_Data control statement.....	234
Mobile/zCAP Pricing Detailed Data section.....	235
Format of the INPUTCSV input data for IBM Z Collocated Application Pricing.....	235
Sample configuration for IBM Z Collocated Application Pricing.....	237
Sample control statement input to SCRT for IBM Z Collocated Application Pricing.....	238
Sample sub-capacity report for IBM Z Collocated Application Pricing.....	239
Billing-related section.....	239
Customer verification section.....	242
Chapter 12. Sub-capacity reporting for IBM Z Workload Pricing for Cloud (zWPC).....	251
Overview of IBM Z Workload Pricing for Cloud.....	251
Qualifications for IBM Z Workload Pricing for Cloud.....	253
Data collection requirements for IBM Z Workload Pricing for Cloud.....	254
zWPC monthly reporting process.....	254
SCRT externals for IBM Z Workload Pricing for Cloud.....	255
SCRT externals for INPUTCSV-based IBM Z Workload Pricing for Cloud.....	255
SCRT externals for WLM classification-based IBM Z Workload Pricing for Cloud.....	258
Reporting sample for INPUTCSV-based IBM Z Workload Pricing for Cloud.....	260
Sample configuration for INPUTCSV-based IBM Z Workload Pricing for Cloud.....	260
Sample control statement input to SCRT for INPUTCSV-based IBM Z Workload Pricing for Cloud.....	260
Sample sub-capacity report for INPUTCSV-based IBM Z Workload Pricing for Cloud.....	261
Reporting sample for WLM classification-based IBM Z Workload Pricing for Cloud.....	269
Sample configuration for WLM classification-based IBM Z Workload Pricing for Cloud.....	269
Sample control statement input to SCRT for WLM classification-based IBM Z Workload Pricing for Cloud.....	270
Sample sub-capacity report for WLM classification-based IBM Z Workload Pricing for Cloud.....	271
Chapter 13. Sub-capacity reporting for z/TPF systems.....	281
Requirements for sub-capacity pricing for z/TPF.....	281
Identities reported for z/TPF systems.....	282
Sub-capacity reporting without SCRT89 records.....	282
Sub-capacity reporting using SCRT89 records.....	283
Running z/TPF and non-z/TPF guests in the same LPAR.....	283
Multiple sub-capacity operating systems.....	284
Steps for z/TPF data collection and processing.....	284
Installing and using the SCRT89 record utility.....	287
Summary of sub-capacity reporting for sub-capacity eligible products running under z/TPF.....	289
Sub-capacity report for a CPC running z/TPF systems and non-z/TPF systems.....	290
Interpreting the detail data section of the sub-capacity report for z/TPF.....	293
Sub-capacity report for a CPC running z/OS and z/TPF.....	294
Interpreting the detail data sections of the sub-capacity report for z/OS and z/TPF.....	297
Chapter 14. Reporting for the z/TPF Transformation Engines Offering.....	299
Qualifications for the z/TPF Transformation Engines Offering.....	299
SCRT restriction on the z/TPF Transformation Engines Offering.....	300
SCRT externals for the z/TPF Transformation Engines Offering.....	300
z/TPF_Transformation_Engines_Offering control statement.....	301
Addition to the Customer Information section.....	301
Addition to the Special Conditions section.....	301
Changes to the Product Summary Information section.....	302
Changes to the Product Max Contributors section.....	302
Changes to the Product Grid Snapshot section.....	303
Chapter 15. Sub-capacity reporting for z/VSE systems.....	305
Introduction to sub-capacity pricing for z/VSE systems.....	305
Requirements for sub-capacity pricing for z/VSE systems.....	306

Reporting of IBM z/VSE products on 21 st Century Software VSE ⁿ	307
Identities reported for z/VSE systems.....	308
Sub-capacity reporting for multiple z/VSE versions.....	308
Overview of the CMT.....	309
Configuring the CMT files.....	310
DLBL and EXTENT statements.....	310
System GETVIS 24-bit requirements.....	310
Overview of skeletons for CMT.....	310
Steps for preparing z/VSE to use CMT.....	313
Sub-capacity reporting using SCRT89 records.....	314
Processing SCRT89 data on a z/OS system.....	314
Overview of processing SCRT89 data on a z/OS system.....	314
Preparing to process SCRT89 data on a z/OS system.....	316
Steps for processing SCRT89 data on a z/OS system.....	317
Processing SCRT89 data on a Windows or Linux system.....	320
Customizing the SCRT control statements.....	320
Customizing the PARMs control statements.....	320
Customizing the NO89 control statements.....	320
Customizing the SPECIAL control statements.....	321
Backup considerations.....	322
Sample sub-capacity reports for z/VSE.....	322
Sample sub-capacity report for z/VSE V4 only.....	323
Sample sub-capacity report for z/VSE V4 and V5.....	327
Sample sub-capacity report for z/VSE V6, z/VSE V5, and z/VSE V4.....	330
Sample sub-capacity report for z/OS and z/VSE V4.....	333
Sub-capacity reporting for z/VSE systems in a country multiplex report.....	336
Sample Country Multiplex Pricing report for z/OS and z/VSE systems.....	336
 Chapter 16. Sub-capacity reporting for Country Multiplex Pricing (CMP).....	 343
Qualifications for Country Multiplex Pricing.....	343
SCRT externals for Country Multiplex Pricing.....	344
Country_Multiplex_Pricing control statement.....	344
Generate_Detailed_Data control statement.....	345
Sample configuration for Country Multiplex Pricing.....	345
Sample control statement input to SCRT for Country Multiplex Pricing.....	346
Sample sub-capacity report for Country Multiplex Pricing.....	347
Multiplex summary section.....	347
Machine Summary section for CPC M1C1.....	351
Machine summary section for CPC M2C1.....	355
Customer verification section.....	357
Detailed Interval Data section.....	362
 Chapter 17. Sub-capacity reporting and Multi-Version Measurement (MVM).....	 367
Overview of Multi-Version Measurement.....	367
Prerequisites for Multi-Version Measurement.....	367
SCRT externals for Multi-Version Measurement.....	367
Sub-capacity report differences for Multi-Version Measurement.....	368
Sample sub-capacity report for Multi-Version Measurement.....	368
 Chapter 18. Sub-capacity reporting for products that report a non-MSU metric.....	 371
Non-MSU metric types.....	371
SCRT externals for reporting of non-MSU metrics.....	371
Sample sub-capacity report for non-MSU-based products.....	372
 Chapter 19. Reporting for Tailored Fit Pricing for IBM Z.....	 375
Requirements for Tailored Fit Pricing for IBM Z.....	375
Format of a Tailored Fit Pricing solution ID.....	376
Interactions with traditional sub-capacity offerings.....	377

SCRT externals for Tailored Fit Pricing for IBM Z.....	378
Tailored Fit Pricing for IBM Z and ISV reporting.....	381
Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z.....	381
Sample configuration for a collocated solution for Tailored Fit Pricing.....	381
Sample control statement input for a collocated solution for Tailored Fit Pricing.....	381
Sample sub-capacity report for a collocated solution for Tailored Fit Pricing.....	382
Sample multiplex report for a standalone LPAR solution for Tailored Fit Pricing for IBM Z.....	392
Sample configuration for a standalone solution for Tailored Fit Pricing.....	392
Sample control statement input for a standalone solution for Tailored Fit Pricing.....	393
Sample multiplex report for a standalone solution for Tailored Fit Pricing.....	393
 Chapter 20. Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution.....	401
Requirements for Tailored Fit Pricing for IBM Z - Software Consumption Solution.....	401
SCRT externals for Software Consumption reporting.....	401
Sample configuration for Software Consumption reporting.....	402
Sample control statement input for Software Consumption reporting.....	402
Sample report for Software Consumption reporting.....	403
 Chapter 21. Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution.....	405
Requirements for Hardware Consumption reporting.....	405
SCRT externals for Hardware Consumption reporting.....	405
Sample report for Hardware Consumption reporting.....	405
 Chapter 22. Workload Classification Supplement.....	409
SCRT externals for the Workload Classification Supplement.....	409
The Workload Classification Supplement report.....	410
 Chapter 23. Reporting for independent software vendor programs.....	413
Overview of reporting for ISV programs.....	413
SCRT externals for ISV program reporting.....	413
Considerations for character set encoding.....	414
Restrictions for ISV program reporting.....	414
Sample SCRT JCL for ISV program reporting.....	415
Sample SCRT product report for ISV programs.....	415
 Part 4. Viewing, editing, and submitting a sub-capacity report.....	417
 Chapter 24. Viewing and editing the sub-capacity report.....	419
Considerations for viewing and editing the sub-capacity report.....	419
Spreadsheet method for viewing and editing the sub-capacity report.....	420
Using placeholder fields in the sub-capacity report.....	420
 Chapter 25. Preparing to submit the sub-capacity report to IBM.....	423
Sub-capacity report checklist.....	423
How IBM handles the sub-capacity report.....	425
 Chapter 26. Submitting sub-capacity reports to IBM.....	427
Getting access to the License Management Support (LMS) application.....	427
LMS web interface.....	427
Viewing reports before submitting them.....	429
Confirming and submitting sub-capacity reports.....	434
Viewing previously submitted reports.....	437
Subscribing to LMS informational email.....	437
 Chapter 27. Troubleshooting SCRT errors and problems.....	439
SCRT fixes.....	439
Obtaining technical support from IBM.....	439

SCRT limitations.....	439
LPARs with z/OS V1R6 and earlier systems running as z/VM guests.....	440
SMF type 89 record-related limitations.....	440
Products with features that can run in different LPARs.....	440
Multi-Version Measurement.....	440
Single Version charging.....	440
Limitations that have been removed.....	440
Support for up to 128 native and guest systems.....	440
LPARs with z/OS V1R7 and higher systems running as z/VM guests.....	441
Dedicated CPs and Wait Completion=Yes.....	441
Machine replacements and machine software model changes.....	441
Machine software model changes.....	441
Machine replacements.....	442
z/OS issues message IWM063I.....	442
Common SCRT problem scenarios.....	442
Product MSUs exceed defined capacity value set for an LPAR.....	442
Data from the same CPC reported on more than one sub-capacity report.....	443
Changing values in the Tool MSUs field of the sub-capacity report or deleting SMF records....	443
Providing alternate MSU values.....	443
Testing Capacity BackUp, disaster recovery, or OOCOD.....	443
Unexpected product reported on a sub-capacity report.....	445
SCRT generates sub-capacity reports instead of a multiplex report.....	446
SCRT JAR file not found.....	446
Out-of-memory conditions.....	447
Using a Java version earlier than the minimum required.....	447
Environments in which you can and cannot generate a full report.....	447
SCRT return codes.....	448
Appendix A. SCRT messages.....	449
SCRT messages for z/OS systems.....	449
SCRT messages for internal SCRT restrictions.....	494
SCRT89 Record Utility messages.....	495
Appendix B. Advanced topics in sub-capacity pricing and reporting.....	497
Capping individual LPARs.....	497
Group capacity.....	498
How SCRT uses group capacity.....	499
Group capacity and Tailored Fit Pricing for IBM Z.....	499
Examples of capping enforcement by SCRT.....	500
Capping enforcement on a traditional z/OS configuration.....	500
Capping enforcement on a z/OS traditional and z/OS zNALC configuration.....	503
Capping enforcement on a z/OS V1 and z/OS V2 configuration.....	505
Static power save mode.....	506
SCRT support for power save mode.....	507
z/OS support for power save mode.....	508
z/TPF support for power save mode.....	511
z/VSE support for power save mode.....	513
Sending TSAD data.....	514
Appendix C. SCRT country code list.....	517
Appendix D. Changes and fixes in prior SCRT versions.....	521
Summary of changes in prior SCRT versions.....	521
Summary of changes for SC23-6845-21, SCRT 29.1.0.....	521
Summary of changes for SC23-6845-19, SCRT Version 28.2.0.....	522
Summary of changes for SC23-6845-18, SCRT Version 28.1.0.....	523
Summary of changes for SC23-6845-17, SCRT Version 27.2.0.....	525

Summary of changes for SC23-6845-16, SCRT Version 27.1.0.....	526
Appendix E. Accessibility.....	529
Notices.....	531
Policy for unsupported hardware.....	532
Minimum supported hardware.....	532
Trademarks.....	532
Index.....	535

Tables

1. Roles, responsibilities, and minimum skill requirements for SCRT users.....	xv
2. Additional and revised product support in SCRT 30.1.....	xxi
3. Example of how SCRT calculates the license capacity for different product and LPAR combinations....	15
4. Simple example of a 4-hour rolling average utilization.....	16
5. Examples of SCRT Customer_Number values.....	35
6. SCRT SPECIAL control statements.....	40
7. SCRT command line options.....	91
8. Sample base z/OS single-version hardware and software configuration.....	105
9. Sample base z/OS V1 and z/OS V2 hardware and software configuration.....	128
10. Calculating MSUs for z/OS products in configurations with z/VM guest systems.....	133
11. Example of a single LPAR with all guests at z/OS V1.6 or lower.....	134
12. How MSUs are calculated for example 1: Single LPAR with all guests z/OS V1.6 or lower.....	134
13. Example of a single LPAR with all guests at z/OS V1.7 or higher.....	135
14. How MSUs are calculated for example 2: Single LPAR with all guests z/OS V1.7 or higher.....	135
15. Example of a single LPAR with mixed guest systems.....	135
16. How MSUs are calculated for example 3: Single LPAR with mixed guest systems.....	136
17. Example of two LPARs with mixed guest z/OS systems.....	136
18. How MSUs are calculated for example 4: Two LPARs with mixed guest z/OS systems.....	136
19. Example of multiple LPARs, one with mixed guests and one with a native z/OS system at any level.	137
20. How MSUs are calculated for example 5: Multiple LPARs; one with mixed guests and one with a native z/OS system at any level.....	137
21. Example of multiple LPARs, one with z/OS V1.7 guests and one with a native z/OS system at any level.....	138
22. How MSUs are calculated for example 6: Multiple LPARs; one with z/OS V1.7 guests and one with a native z/OS system at any level.....	138

23. Sample configuration with z/OS V1.6 and V1.7 systems as guests of z/VM.....	140
24. Sample configuration with z/OS V1.7 (or later) systems as guests z/VM.....	142
25. Sample configuration for native z/OS systems running traditional and zNALC workloads.....	146
26. Sample configuration for native z/OS V1 and z/OS V2 systems running traditional and zNALC workloads.....	150
27. Sample configuration for a CPC where native z/OS systems only ran zNALC workloads.....	153
28. Sample configuration for z/OS V1.6 guest systems running traditional and zNALC workloads in different LPARs.....	154
29. Sample configuration for z/OS V1.6 guest systems running traditional and zNALC workloads in the same LPAR.....	158
30. Sample configuration for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads in the same CPC.....	162
31. Sample configuration with Getting Started Sub-Capacity Pricing eligible products.....	174
32. IWP adjusted programs.....	183
33. IWP defining programs.....	184
34. Sample configuration for Integrated Workload Pricing.....	189
35. MWP defining programs.....	199
36. Sample configuration for Mobile Workload Pricing.....	207
37. Sample configuration for Mobile Workload Pricing.....	218
38. zCAP defining programs.....	229
39. Sample configuration for IBM Z Collocated Application Pricing.....	238
40. zWPC defining programs.....	251
41. Sample configuration for IBM Z Workload Pricing for Cloud.....	260
42. Sample configuration for IBM Z Workload Pricing for Cloud.....	270
43. Summary of z/TPF system identities reported by SCRT.....	282
44. How product MSUs are calculated for z/TPF-based programs using SCRT89 records.....	289
45. Sample configuration with z/TPF native and guest systems and non-z/TPF systems.....	291
46. Sample configuration with z/OS and z/TPF systems.....	294

47. Recommended APAR levels for z/VSE CMT.....	306
48. Summary of z/VSE system identities reported by SCRT.....	308
49. Sample configuration with z/VSE V4 systems.....	323
50. Sample configuration with z/VSE V4 and V5 systems.....	327
51. Sample configuration with z/VSE V6, V5, and V4 systems.....	330
52. Sample configuration with z/OS and z/VSE systems.....	333
53. Sample configuration for Country Multiplex Pricing.....	346
54. Non-MSU metric types and footnotes.....	372
55. Sub-capacity report fields in which LMS accepts DBCS data double-byte character set (DBCS) values in report fields acceptable by LMS.....	421
56. Countries, regions, and code pages for which LMS accepts DBCS values.....	427
57. Examples of conditions where CPCs have been replaced.....	442
58. SCRT reporting behavior in various hardware and operating system environments.....	447
59. Sample z/OS traditional configuration with individual LPAR capping and group capping.....	500
60. MSUs for hour 51 in the sample z/OS traditional configuration.....	501
61. MSUs for hour 80 in the sample z/OS traditional configuration.....	502
62. Sample z/OS traditional and z/OS zNALC configuration with individual LPAR and group capping.....	503
63. MSUs for hour 80 in the sample z/OS traditional and z/OS zNALC configuration.....	504
64. Sample z/OS traditional configuration with z/OS V1 and z/OS V2.....	505
65. MSUs for the sample z/OS z/OS V1 and z/OS V2 configuration.....	505
66. Explanations for power save mode values listed in Detail LPAR Data Section for z/OS systems.....	510
67. Explanations for power save mode values listed in Detail LPAR Data Section for z/TPF systems.....	512
68. Explanations for power save mode values listed in Detail LPAR Data Section for z/VSE systems.....	514
69. SCRT country code list.....	517
70. Additional and revised product support in SCRT 29.1.0.....	522
71. Additional and revised product support in SCRT V28.2.0.....	523

72. Additional and revised product support in SCRT V28.1.0.....	524
73. Additional and revised product support in SCRT V27.2.0.....	526
74. Additional and revised product support in SCRT V27.1.0.....	527

About this document

This information describes how to use the Sub-Capacity Reporting Tool (SCRT), which supports IBM sub-capacity software pricing metrics, including Workload License Charges (WLC), Advanced Workload License Charges (AWLC), Entry Workload License Charges (EWLC), Advanced Entry Workload License Charges (AEWLC), System z® New Application License Charges (zNALC), and Midrange Workload License Charge (MWLC). SCRT also supports enhancements to the sub-capacity pricing metrics for System z International Product License Agreement (IPLA) products, IBM Getting Started Sub-Capacity Pricing for z/OS IPLA software (GSSP), and Integrated Workload Pricing (IWP).

SCRT also optionally reports on the hardware-related z/TPF Transformation Engines Offering.

SCRT also optionally reports on IBM Mobile Workload Pricing (MWP), IBM Z® Collocated Application Pricing (zCAP), IBM Country Multiplex Pricing (CMP), and Tailored Fit Pricing for IBM Z offerings.

Note that new releases of the Sub-Capacity Reporting Tool immediately apply to *all* eligible products on all releases of z/OS—not just to the most current release of each.

Some SCRT functions are only available for certain releases of z/OS. For example, reporting for z/OS systems that are z/VM® guests is more detailed with z/OS V1R7 and later systems. Where there are specific operating system level requirements for certain functions, those requirements are noted in the text.

For z/TPF, new releases of SCRT apply to z/TPF release 1.1 and above.

For z/VSE, new releases of SCRT apply to z/VSE 4.1.0 and above.

Unless specifically noted, all references to z/OS also apply to z/TPF and z/VSE.

Who should read this document?

This document is intended to be read, understood, and used by anyone who will use SCRT, including those who will perform the roles and responsibilities shown in [Table 1 on page xv](#). For each role, the table indicates the minimum set of skills required to carry out those responsibilities.

Skill requirements for SCRT users

[Table 1 on page xv](#) lists the roles, responsibilities, and minimum skill requirements for SCRT users. This list is not intended to be all-inclusive; rather, it highlights the most critical responsibilities and expected skill levels to help ensure the successful use of SCRT and sub-capacity reporting.

Table 1. Roles, responsibilities, and minimum skill requirements for SCRT users

Role / responsibility	Skill requirements
Systems programmers responsible for customizing their systems to collect SMF and/or SCRT89 accounting records	<ul style="list-style-type: none">For z/OS systems:<ul style="list-style-type: none">Familiarity with the MVS system management facilities (SMF) component and ability to configure SMF to collect the required accounting recordsFamiliarity with your organization's operational procedures for saving and processing SMF dataFamiliarity with setting defined capacity and LPAR group capacityFor z/TPF systems using SCRT89 records, familiarity with the ZSUBC command, familiarity with transferring data from z/TPF systems to a z/OS system (such as by using FTP), and, optionally, familiarity with the SCRT89 record utilityFor z/VSE systems, familiarity with configuring and using the Capacity Measurement Tool (CMT) to produce and record SCRT89 accounting records

Table 1. Roles, responsibilities, and minimum skill requirements for SCRT users (continued)

Role / responsibility	Skill requirements
Those responsible for obtaining and installing SCRT	<ul style="list-style-type: none"> • Ability to access the World Wide Web and download the latest SCRT deliverable to a workstation • Ability to extract the contents of the downloaded deliverable onto your workstation and then upload them to the target z/OS or z/VSE host system on which you intend to run SCRT • Ability to allocate the necessary data sets or files on the target host system into which to upload the SCRT deliverable • Ability to access and use z/OS UNIX System Services, including familiarity with Java™ on z/OS, to install SCRT
Those responsible for generating sub-capacity reports	<ul style="list-style-type: none"> • Knowledge of your organization's hardware configuration, including CPC machine types, model numbers, serial numbers, LPAR names, and SYSIDs • Knowledge of the software products that your organization uses, including product IDs and the LPARs in which they run • Knowledge of any unusual operational conditions which might have occurred during the SCRT reporting period, such as testing Capacity BackUp (CBU) or disaster recovery, or running a software product on a CPC for which it is not licensed because the CPC for which it is licensed experienced a hardware outage • Ability to edit, customize, and submit the JCL to run SCRT • Knowledge of your organization's batch job processing guidelines and JCL standards • Ability to perform basic problem determination by: <ul style="list-style-type: none"> – Checking for any non-zero return codes and error messages in the SCRT error report and SYSPRINT output – Referring to diagnostic information in this publication to look up return codes and error message explanations, and performing the indicated corrective action – When necessary, gathering required diagnostic information and reporting the problem to IBM by opening a service request, as described in “Obtaining technical support from IBM” on page 439. • Access to the z/OS UNIX System Services directory where SCRT is installed. • Ability to run Java programs. • For data from z/VSE systems: <ul style="list-style-type: none"> – Ability to use CMTREPO to extract the recorded SCRT89 data – Ability to transfer VSAM data between z/VSE systems, between z/VSE systems and a z/OS system (such as by using FTP) – Ability to transfer files containing SCRT89 records from z/VSE systems to a z/OS system, Windows workstation, or a Linux® system on which to run SCRT • For any IBM Z server, familiarity with submitting Transmit System Availability Data (TSAD) using either the Remote Support Facility (RSF) or the Transmit Service Data task on the Support Element (SE)
Software asset managers or others responsible for overseeing software licensing and billing	<ul style="list-style-type: none"> • Knowledge of your organization's hardware, software, and workload configurations • Knowledge of your organization's inventory of software products • Understanding of your organization's contractual obligations for software licensing, sub-capacity reporting, and billing
Those responsible for submitting sub-capacity reports to IBM	<ul style="list-style-type: none"> • Ability to review the sub-capacity report (.csv file) and to supply optional and required information, such as customer MSU values and customer comments, as needed, using a spreadsheet application or the LMS web application • Ability to use the LMS web application to submit the sub-capacity report to IBM

Where to find more information

You might find it helpful to consult the following publications in conjunction with *Using the Sub-Capacity Reporting Tool*:

- For sub-capacity data collected from one or more z/OS systems:
 - *z/OS Planning for Sub-Capacity Pricing*, SA23-2301
 - *z/OS MVS System Management Facilities (SMF)*, SA38-0667

- *z/OS MVS Initialization and Tuning Reference*, SA23-1380
- *z/OS UNIX System Services Command Reference*, SA23-2280
- *JZOS Installation Guide*, SA38-0696
- For sub-capacity data collected from one or more z/TPF systems:
 - *z/TPF Operations*, SH31-0162
- For sub-capacity data collected from one or more z/VSE systems:
 - *z/VSE Operations*, SC33-8309

The following websites contain more information about SCRT, z/OS, and IBM Z software pricing, including Workload License Charges and sub-capacity pricing for certain IBM Z IPLA software products.

- [Pricing for IBM Z \(www.ibm.com/z/pricing\)](http://www.ibm.com/z/pricing)
- [IBM Terms \(www.ibm.com/support/customer/csol/terms/\)](http://www.ibm.com/support/customer/csol/terms/)
- [IBM z/OS home page \(www.ibm.com/products/zos\)](http://www.ibm.com/products/zos)
- [IBM z/Transaction Processing Facility \(www.ibm.com/products/z-transaction-processing-facility\)](http://www.ibm.com/products/z-transaction-processing-facility)

Summary of changes

This topic summarizes the changes made to this document in the current version of the Sub-Capacity Reporting Tool (SCRT). For changes made in prior versions of SCRT, see [Appendix D, “Changes and fixes in prior SCRT versions,”](#) on page 521.

IBM requires that you use the latest version of SCRT. For the latest available version of SCRT, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](#).

Summary of changes for SC23-6845-22, SCRT 30.1

This document contains information previously presented in SC23-6845-21 and includes terminology, maintenance, and editorial changes.

New and changed information

Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Documentation changes

The following list highlights the technical changes made to this documentation in this release of SCRT. For the complete list of functional and technical changes in this release of SCRT, including changes in any interim modification levels since the last release, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](#).

- References to "Java" throughout Chapter 3, “Using SCRT on z/OS systems,” on page 21 are updated to include IBM Semeru Runtime® Certified Edition for z/OS.
- The sample SCRT JCL has been updated in [“Sample JCL to run the 64-bit SCRT on a z/OS system”](#) on page 25.
- The following SPECIAL DD control statements are added in [“SPECIAL DD statement”](#) on page 40:

CONTAINER_INTERVAL_CONSUMPTION

DETAIL_INTERVAL_RATE_DATA

INTERVAL_RATE_DATA

- PDSE data sets are supported for the data set specified on the OUTPUT DD statement. See [“OUTPUT DD statement”](#) on page 64.
- The topic, “Windows and Linux software prerequisites for SCRT” on page 69, is updated.
- The topic, “Sample base z/OS V1 configuration,” is renamed to [“Sample base z/OS single-version configuration”](#) on page 104.
- The topic, “Sample sub-capacity report for z/OS V1 native systems,” is renamed to [“Sample sub-capacity report for single-version z/OS native systems”](#) on page 105 and the introduction is updated.
- The descriptions of the virtual products are updated in [“Product Summary Information section”](#) on page 110.
- The zIIP Usage Data section (I8) is added in [“zIIP Usage Data section”](#) on page 117.
- The description of the Detail LPAR Usage Data section is updated in [“Detail LPAR Usage Data section”](#) on page 120.
- The topic, “Sub-Capacity reporting for both z/OS V1 and z/OS V2 systems,” is renamed to [“Sub-capacity reporting for multiple versions of z/OS”](#) on page 127 and is updated to include version 3.x of z/OS.
- Information about the **CONTAINER_INTERVAL_CONSUMPTION** control statement and the Container Interval Consumption report section (I7) are added in [Chapter 19, “Reporting for Tailored Fit Pricing for IBM Z,”](#) on page 375.

- The **Var Model** and **Var Capacity** columns are added to the Detail Interval Data section (V9) in [“Sample report for Hardware Consumption reporting”](#) on page 405.
- The **Var Model** and **Var Capacity** columns are changed to **Var Model or Rep Model** and **Var Capacity or Rep Capacity**, respectively, in the Detail Interval Data section (V9) in [“Sample report for Hardware Consumption reporting”](#) on page 405.
- Information about Workload Classification Supplement is added in Chapter 22, “Workload Classification Supplement,” on page 409. Additionally, the WCSUPP DD statement is added in [“WCSUPP DD statement”](#) on page 65, and the **-wcsupp** option is added in [“Specifying command line options”](#) on page 91.
- [“Spreadsheet method for viewing and editing the sub-capacity report”](#) on page 420 is updated to recommend using a plain-text editor to edit placeholder fields.
- The telephone number to contact LMS customer support is updated in [“LMS web interface”](#) on page 427.
- Support has been added for new z/OS sub-capacity products, as described in [“Product support changes”](#) on page xx.

SCRT message changes

- The following new messages have been added in [Appendix A, “SCRT messages,”](#) on page 449:
[SCRTTOOL416](#)
- The following existing messages have been updated:
[SCRTTOOL055](#)
[SCRTTOOL134](#)

Product support changes

SCRT 30.1 adds sub-capacity support for the products listed in Table 2 on page xxi. The NO89 products are marked with an asterisk (*), and updates to the NO89 DD of the SCRT JCL for the z/OS deliverable are required to self-report usage or non-usage of these products.

Table 2. Additional and revised product support in SCRT 30.1

Operating system	Product ID	Product name
z/OS	5755-AB1	IBM Application Delivery Foundation for z/OS 5
	5755-AB4	IBM Application Performance Analyzer for z/OS 16
	5755-AB6	IBM Debug for z/OS 17
	5655-CA1	IBM Decision Runtime for z/OS 9
	5755-AB5	IBM Developer for z/OS Enterprise Edition 17
	5755-AB2	IBM Fault Analyzer for z/OS 16
	5755-AB3	IBM File Manager for z/OS 16
	5698-HG8	IBM HourGlass for z/OS 8
	5655-KAF	IBM Open Enterprise SDK for Apache Kafka 1
	5655-NOJ	IBM Open Enterprise SDK for Node.js 18
	5655-NOS	IBM Open Enterprise SDK for Node.js 20
	5655-NDE	IBM Open Enterprise SDK for Node.js 22
	5655-ABG	IBM Security zSecure Alert 3
	5655-ABD	IBM Security zSecure CICS® Toolkit 3
	5655-RA1	IBM Sterling™ Transformation Extender for z/OS 10
	5655-RA2	IBM Sterling Transformation Extender for z/OS 11
	5698-WMC	IBM Watson® Machine Learning for z/OS Core Edition
	5655-ZCX	IBM zCX Foundation for Red Hat® OpenShift®
	5698-BR3*	IBM Z Backup Resiliency 2*
	5698-ZDH	IBM Z Digital Integration Hub
	5698-ZWH	IBM Z Distribution for Zowe™ 3

Software fixes in SCRT

For information about APARs and other software fixes in this release of SCRT, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](https://ibm.biz/SCRTCurrentLevels).

Deleted information

Preventive service planning (PSP) buckets for many IBM products, including z/OS 2.5 and 3.1, are no longer updated. For more information, see [PSP bucket information for IBM Z products \(www.ibm.com/support/pages/node/7127792\)](https://www.ibm.com/support/pages/node/7127792).

Part 1. SCRT concepts and planning

This part describes some general concepts and planning information that you must understand to properly prepare for and use the Sub-Capacity Reporting Tool (SCRT). These concepts apply to all platforms on which SCRT runs.

This part contains the following chapters:

- [Chapter 1, “About the Sub-Capacity Reporting Tool,” on page 3](#)
- [Chapter 2, “Preparing to use the Sub-Capacity Reporting Tool,” on page 11](#)

Chapter 1. About the Sub-Capacity Reporting Tool

This chapter describes the Sub-Capacity Reporting Tool (SCRT). It also provides an overview of software pricing for IBM Z environments, especially the following sub-capacity pricing metrics and offerings:

- Workload License Charges (WLC)
- Advanced Workload License Charges (AWLC)
 - Integrated Workload Pricing (IWP)
- Entry Workload License Charges (EWLC)
- Advanced Entry Workload License Charges (AEWLC)
 - Integrated Workload Pricing (IWP)
- System z New Application License Charges (zNALC)
- Midrange Workload License Charges (MWLC)
- Sub-capacity System z International Product License Agreement (IPLA) pricing
 - Getting Started Sub-Capacity Pricing (GSSP)
- Mobile Workload Pricing (MWP)
- IBM Z Collocated Application Pricing (zCAP)
- IBM Z Workload Pricing for Cloud (zWPC)
- Tailored Fit Pricing for IBM Z
- Tailored Fit Pricing for IBM Z - Software Consumption Solution

In addition, SCRT supports IBM Country Multiplex Pricing (CMP) for Country Multiplex License Charges (CMLC) and Multiplex zNALC (MzNALC). For details, see [Chapter 16, “Sub-capacity reporting for Country Multiplex Pricing \(CMP\),”](#) on page 343.

SCRT also supports the z/TPF Transformation Engines Offering, which is different than software sub-capacity pricing. For a details about the z/TPF Transformation Engines Offering, see [Chapter 14, “Reporting for the z/TPF Transformation Engines Offering,”](#) on page 299.

What is SCRT?

The Sub-Capacity Reporting Tool (SCRT) is a no-charge IBM tool that reports the required license capacity for sub-capacity eligible products that run on z/OS, z/TPF, z/VSE, or 21st Century Software VSEⁿ systems. SCRT analyzes a month's worth of data for IBM Z environments and produces a *sub-capacity report*.¹ The report indicates the required license capacity, in MSUs, of each sub-capacity eligible product.

z/TPF, z/VSE, and limited 21st Century Software VSEⁿ support for sub-capacity differs from z/OS support, for example, by using SCRT89 records instead of SMF records. Unless specifically noted, the same concepts described for z/OS also apply to z/TPF and z/VSE. The concepts that apply to z/VSE in general mostly also apply to 21st Century Software VSEⁿ with the notable exception that SCRT only reports on IBM z/VSE products running on 21st Century Software VSEⁿ systems. SCRT does not report on the 21st Century Software VSEⁿ operating system nor on any 21st Century Software VSEⁿ products. Neither does SCRT reflect any 21st Century Software VSEⁿ operating system nor any 21st Century Software VSEⁿ products in the optional (A11) entries used for Multi-Version Measurement (MVM).

¹ You might also see the terms *SCRT report* and *multiplex report*. The terms *SCRT report* and *sub-capacity report* are generally synonymous, although *SCRT report* is a broader term that can refer collectively to any of the various report types that SCRT produces. This documentation uses the specific terms *sub-capacity report* and *multiplex report*, as appropriate.

For details about z/TPF support and its unique considerations, see [Chapter 13, “Sub-capacity reporting for z/TPF systems,”](#) on page 281 and [Chapter 14, “Reporting for the z/TPF Transformation Engines Offering,”](#) on page 299.

For details about z/VSE and limited 21st Century Software VSEⁿ support and its unique considerations, see [Chapter 15, “Sub-capacity reporting for z/VSE systems,”](#) on page 305.

Sub-capacity products are charged based on the 4-hour rolling average (4HRA) utilization of the LPARs in which the sub-capacity products run during the reporting period. For each hour in the reporting period, SCRT determines the required license capacity by examining:

- The 4-hour rolling average utilization, by LPAR
- The eligible products that were active in each LPAR

SCRT then cross references LPAR utilization and product execution by LPAR to determine the maximum concurrent LPAR 4-hour rolling average utilization—the highest combined utilization of LPARs where each product executes during the reporting period.

SCRT processes the following input data:

- From z/OS environments: System Management Facilities (SMF) data, specifically SMF type 70, subtype 1 (CPU Activity) records and SMF type 89, subtypes 1 and 2 (Product Use) records
- From z/TPF and z/VSE environments: SCRT89 records

Rule: The IBM sub-capacity contract requires you to save all SMF type 70 records, SMF type 89 records, and SCRT89 records for at least six months in order to be able to resolve questions and problems.

SCRT processes the input data that you specify and produces report output. To comply with contractual terms and conditions, the input data must contain one reporting period of SMF type 70 and type 89 records for all the z/OS images and SCRT89 records for all z/TPF and z/VSE images on a z/Architecture[®] CPC. (For more information about the reporting period and data collection, see [“The SCRT reporting period”](#) on page 16.) The report output can be:

- On a z/OS system:
 - A partitioned data set (PDS) with one member for each CPC found in the input data stream. Each member is one sub-capacity report. The report is a text file in comma-separated value (.csv) format, which can most easily be read by a spreadsheet program. For more information, see [“Collecting z/OS SMF records”](#) on page 13 and [“Spreadsheet method for viewing and editing the sub-capacity report”](#) on page 420.
 - A sequential data set that contains a single text file in comma-separated value (.csv) format. This file contains all the reports for the CPCs in the input stream to SCRT and can be used to make uploading reports to the LMS web application easier (see [Chapter 26, “Submitting sub-capacity reports to IBM,”](#) on page 427). For details on how to set this up, see the description of the OUTPUT DD statement in [“OUTPUT DD statement”](#) on page 64.
- On a Windows or Linux system:
 - A text file in comma-separated value (.csv) format.

SCRT also does the following:

- Provides a simple, non-intrusive means for you to generate sub-capacity reports that show required license capacity for all eligible sub-capacity products that execute on a CPC
- Allows IBM to receive, view and store sub-capacity reports received from customers
- Allows IBM to generate customer bills according to the sub-capacity reports received from customers

What is the IBM Z software pricing framework?

IBM Z software pricing options fall into the following categories.

Monthly license charge (MLC) pricing

Monthly license charges (MLC) apply to many IBM Z software products, including z/OS, Db2®, CICS, IMS, MQSeries®, z/TPF, z/VSE and others. Pricing and terms and conditions for MLC products are based on the pricing metric you select. Pricing metrics generally fall into two categories: *full capacity based* pricing metrics and *sub-capacity capable* pricing metrics.

- **Full capacity based pricing metrics**

Under a full capacity-based pricing metric, software charges for all products are determined by the capacity of the CPCs in which the products run. Parallel Sysplex® License Charges (PSLC) and zSeries Entry License Charges (zELC) are examples of full capacity based metrics.

- **Sub-capacity capable pricing metrics**

Under a sub-capacity pricing metric, software charges for certain products are based on the utilization capacity of the LPARs in which those products run. The following are examples of sub-capacity capable pricing metrics:

- Workload License Charges (WLC)
- Advanced Workload License Charges (AWLC), including Integrated Workload Pricing (IWP), when appropriate
- Entry Workload License Charges (EWLC)
- Advanced Entry Workload License Charges (AEWLC), including Integrated Workload Pricing (IWP), when appropriate
- Midrange Workload License Charges (MWLC)
- System z New Application License Charges (zNALC)

WLC, AWLC, EWLC, AEWLC, MWLC, and zNALC can be implemented in either sub-capacity or full capacity mode. Sub-capacity pricing is available only on CPCs for which a sub-capacity report has been submitted and all other terms and conditions are met. If no sub-capacity report is submitted in a given month for a CPC for which sub-capacity pricing has been established, any sub-capacity-eligible products on that CPC are charged based on the full capacity of the CPC for that month.

When you elect WLC for a particular operating system on a CPC, all of that operating system family's MLC products are licensed under WLC on that CPC. Sub-capacity eligible MLC products are called *Variable WLC (VWLC)* products. Non-sub-capacity eligible MLC products are called *Flat WLC (FWLC)* products and their pricing is a fixed monthly charge unrelated to the capacity of the CPC on which they run.

When you elect AWLC for a particular operating system on a CPC, all of that operating system family's sub-capacity eligible MLC products are licensed under AWLC on that CPC. Non-sub-capacity eligible MLC products are considered FWLC products and their pricing is a fixed monthly charge unrelated to the capacity of the CPC on which they run.

When you elect EWLC for a particular operating system on a CPC, all of that operating system family's MLC products that are sub-capacity eligible are licensed under EWLC on that CPC. The MLC products for that operating system on that CPC that are not sub-capacity eligible are charged a capacity-based price using the zELC pricing metric (for z800 CPCs) or the Tiered EWLC pricing metric (for z10 BC, z9® BC, and z890 CPCs).

When you elect AEWLC for a particular operating system on a CPC, all of that operating system family's MLC products that are sub-capacity eligible are licensed under AEWLC on that CPC. The MLC products for that operating system on that CPC that are not sub-capacity eligible are charged a capacity-based price using the Tiered EWLC pricing metric.

When you elect MWLC for a z/VSE V4 or later operating system on a CPC, those z/VSE MLC products that are sub-capacity eligible are licensed under MWLC on that CPC. Note that sub-capacity MWLC pricing is only available when all of the VSE systems running on the CPC are at a V4 or higher level. On a System z9® BC or System z10® BC CPC, all other z/VSE programs will be priced according to Tiered EWLC. On an IBM z13®, zEnterprise® EC12, zEnterprise 196, System z10 EC, or System z9 EC CPC, all other z/VSE

programs will be priced according to either Flat Workload License Charges, Graduated Monthly License Charges, or Extended License Charges.

When you cancel the last license for an IBM operating system family on a CPC, any remaining IBM MLC products for that operating system family on that CPC are no longer sub-capacity eligible and change to a full-capacity pricing metric. Specifically, when you cancel your last IBM z/VSE license for a CPC, your remaining IBM z/VSE MLC products are no longer sub-capacity eligible.

For a complete description of all IBM Z software pricing metrics, see [Pricing for IBM Z \(www.ibm.com/z/pricing\)](http://www.ibm.com/z/pricing).

IBM Z International Product License Agreement (IPLA) Software

Some software products that run on the IBM Z platform have an up-front license fee and an optional annual maintenance charge. These products are licensed under the *International Program License Agreement (IPLA)*. IPLA products include some data management tools, CICS tools, IMS tools, application development tools, WebSphere® products, and Tivoli® products.

Many IPLA products that run on z/OS can be licensed at a sub-capacity level. Sub-capacity charging for these products is available to customers who are already using sub-capacity pricing on that CPC for their MLC products using one of the pricing metrics described earlier.

This information helps you decide whether or not to use sub-capacity MLC pricing. Once you have decided to use sub-capacity MLC pricing, it is always advantageous to use sub-capacity IPLA pricing for those products that offer it.

IPLA software products that are sub-capacity eligible are licensed according to either *execution-based terms*, *reference-based terms*, or *z/OS-based terms*.

Execution-based licensing terms apply to z/Architecture IPLA products whose value is based on the processing power of the LPAR in which they run. These products are priced based on the utilization capacity of the LPARs in which they run. Execution-based IPLA products will appear on the sub-capacity report if those products execute on a CPC where you run the SCRT.

IBM also offers a pricing enhancement for a subset of the execution-based IPLA products, called *IBM Getting Started Sub-Capacity Pricing for z/OS IPLA software*, which offers entry pricing for customers with low usage on select IBM WebSphere programs. For more information, see [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),”](#) on page 165.

Reference-based licensing terms apply to z/Architecture IPLA products whose value is based on the value of another product, called the *parent product*. For example, a reference-based Db2 tool would be priced based on the license capacity of Db2. In this case, Db2 is the parent product of the Db2 tool. Reference-based IPLA products will not appear on the sub-capacity report because the parent products already appear on the report.

z/OS-based licensing terms apply to IPLA products that run under z/OS and provide value to the particular machine where the product is used. These products are priced based on the license capacity of z/OS on the CPC where the product is used. z/OS-based IPLA products will not appear on the sub-capacity report because z/OS already appears on the report.

For a list of sub-capacity-eligible IPLA products and how they are licensed, see [IPLA Sub-Capacity Overview \(www.ibm.com/downloads/cas/AOJR8PEA\)](http://www.ibm.com/downloads/cas/AOJR8PEA).

Select Application License Charges

Select Application License Charges (SALC) are available solely for MQSeries and only on sub-capacity CPCs with one of the sub-capacity MLC pricing metrics. SCRT reports MSU values for these products, but the reported values are not used to calculate the license charges. For more information about SALC and the products it applies to, see [S/390 Usage Pricing and Select Application License Charges \(SALC\) \(www.ibm.com/downloads/cas/PY98L9LR\)](http://www.ibm.com/downloads/cas/PY98L9LR).

IBM Z New Application License Charges (zNALC)

IBM Z New Application License Charges (zNALC) is a monthly license charge price metric available on IBM z/Architecture servers running z/OS and offers a reduced price for z/OS on LPARs where qualified applications are running. zNALC applies to the z/OS base feature and z/OS priced features. zNALC is available for z/OS on LPARs dedicated to qualified applications, among other requirements. Similar to Workload License Charges and Entry Workload License Charges, customers may implement zNALC in either full-capacity or sub-capacity mode.

One of the actions that customers must perform in order to receive zNALC pricing is to identify the z/OS systems which qualify for zNALC pricing. zNALC systems can be identified by either:

- IPLing a native z/OS system or a z/OS guest system running under z/VM in an LPAR named ZNALxxxx where xxxx is any combination of letters or numbers, **or**
- IPLing a native z/OS system or a z/OS guest system running under z/VM using a LICENSE value of ZNALC. The LPAR can have any name when you specify LICENSE=ZNALC at IPL time. The LICENSE=ZNALC parameter is supported by the IEASYSxx member of SYS1.PARMLIB on z/OS 1.6 to z/OS 1.8 systems with APAR OA20314 applied.

More information on zNALC is available in [Chapter 7, “Sub-capacity reporting for System z New Application License Charges \(zNALC\),”](#) on page 145 and on the web at [z Systems New Application License Charges \(zNALC\) \(www.ibm.com/downloads/cas/WGVAN8A2\)](http://www.ibm.com/downloads/cas/WGVAN8A2).

What types of sub-capacity pricing does SCRT support?

SCRT supports the various sub-capacity price metrics and sub-capacity offerings for eligible products based upon the operating systems and CPCs to which they apply.

- In z/OS environments, SCRT supports the WLC, AWLC, EWLC, AEWLC, and zNALC price metrics, as well as sub-capacity IPLA pricing (including GSSP). SCRT also supports the IWP, MWP, zCAP, zWPC, MVM, and Tailored Fit Pricing for IBM Z price offerings.
- In a z/VSE environment, SCRT supports MWLC and the AEWLC price metrics and the MVM price offering.
- In a z/TPF environment, SCRT supports the WLC, AWLC, EWLC, and AEWLC price metrics.
- In a z/OS or z/TPF Country Multiplex Pricing environment, SCRT supports the Country Multiplex License Charges (CMLC) and Multiplex zNALC (MzNALC) price metrics.

Sub-capacity pricing for z/OS products applies only to CPCs where z/OS is the only MVS-based operating system in use. For instance, there can be no LPARs running OS/390 or earlier operating systems.

Sub-capacity pricing for z/VSE products applies only to CPCs where z/VSE V4 or higher is the only VSE-based operating system in use.

Sub-capacity pricing for z/TPF products applies only to CPCs where z/TPF is the only TPF-based operating system in use.

Sub-capacity eligibility is on an operating system family basis, not a CPC basis. Whenever multiple operating system families (such as z/OS, z/VSE, and z/TPF) run on the same CPC, any of the eligible operating system families can be placed under sub-capacity pricing while other operating system families on that same CPC might not meet eligibility requirements and are not eligible for sub-capacity pricing.

Example: z/TPF might meet the eligibility requirements for sub-capacity pricing and receive sub-capacity pricing, even though a z/OS system running on the same CPC is not eligible for sub-capacity pricing because OS/390® is still licensed to that CPC. Note that all of an operating system family's products that are licensed to the same CPC must be priced using the same method.

Products that use sub-capacity IPLA pricing (including those that are eligible for GSSP) must be running under z/OS using sub-capacity pricing.

To be eligible for sub-capacity pricing, you must use LPAR technology to define the appropriate execution environment, and you must use SCRT to submit monthly sub-capacity reports to IBM.

For more information about the different sub-capacity eligible price metrics and price offerings and the CPCs upon which they are eligible to run, see [IBM Z Software Pricing Reference Guide \(www.ibm.com/downloads/documents/us-en/107a02e95d48f807\)](http://www.ibm.com/downloads/documents/us-en/107a02e95d48f807).

How does sub-capacity pricing work?

Traditional sub-capacity pricing² is based on the concept of the highest observed *4-hour rolling average utilization* of the LPAR or LPARs where a sub-capacity product runs.

Suppose a CPC has a specific sub-capacity eligible product running concurrently in three LPARs. The simultaneous combined utilization of these three LPARs is determined for each hour in the reporting period and the highest observed combined utilization is used as the basis for the traditional sub-capacity pricing of the product.

A subset of products are eligible for sub-capacity pricing offerings other than the traditional sub-capacity pricing described earlier:

- For Getting Started Sub-Capacity Pricing (GSSP) programs, the basis for sub-capacity pricing is the GSSP algorithm described in [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#).
- For IWP adjusted programs, the basis for sub-capacity pricing is the IWP algorithm described in [Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing \(IWP\),” on page 183](#).
- For MWP defining programs, the basis for sub-capacity pricing is the MWP algorithm described in [Chapter 10, “Sub-capacity reporting for Mobile Workload Pricing \(MWP\),” on page 199](#).
- For zCAP defining programs, the basis for sub-capacity pricing is the zCAP algorithm described in [Chapter 11, “Sub-capacity reporting for IBM Z Collocated Application Pricing \(zCAP\),” on page 229](#).
- For CMP defining programs, the basis for sub-capacity pricing is the CMP algorithm described in [Chapter 16, “Sub-capacity reporting for Country Multiplex Pricing \(CMP\),” on page 343](#).
- For certain sub-capacity-eligible IPLA programs, the basis for sub-capacity pricing is the MSU value reported for the parent program or for the z/OS operating system. For an example of how LPAR utilization capacity is calculated, see [“How SCRT determines the license capacity” on page 15](#).

Overview of the SCRT process

This is a high-level overview of the SCRT process for data collected on z/OS, z/TPF, and z/VSE systems, regardless of whether you intend to run SCRT on a z/OS system or on a Windows or Linux system.

To enable sub-capacity pricing for a qualifying CPC, you must perform the following general steps:

1. Apply all required software service to support sub-capacity reporting.

Prior to collecting any SMF or SCRT89 records for SCRT to use in generating sub-capacity reports, apply all required service for z/OS, z/TPF, and z/VSE systems, including maintenance in the SMP/E **IBM.Function.PricingInfrastructure** fix category (FIXCAT). Apply service updates *before the start of the reporting period* for which you will be collecting data for sub-capacity reporting. Failure to apply all of the required service may prevent SCRT from creating a sub-capacity report that will be acceptable to IBM.

2. For the SCRT deliverable that runs on a z/OS system, customize the JCL to meet your installation's requirements, and customize the SCRT control statements to supply the required customer-related information, software configuration information, and any other optional processing information, as necessary. For details, see [“Running SCRT on a z/OS system” on page 24](#).
3. Collect one reporting period's worth of System Management Facility (SMF) data for z/OS systems and collect SCRT89 records for z/TPF and z/VSE systems running on your CPCs that use sub-capacity pricing. A reporting period is from midnight (the very beginning) of the second day of the month up to midnight (the very beginning) of the second day of the next month.

² Some pricing contracts use the term *standard sub-capacity pricing*, which is synonymous with the term *traditional sub-capacity pricing* used in this documentation.

4. Move all of the SMF and SCRT89 data that you collected for each CPC to the system on which you intend to run SCRT. This can be a z/OS system, a Linux system, or a Windows system.
5. Run the SCRT program to generate your sub-capacity reports.
6. Review the return code from the SCRT program.

A return code of 0 means that SCRT successfully generated a sub-capacity report for every CPC that provided SMF or SCRT89 records for processing by that execution of SCRT. Return codes greater than 0 indicate that SCRT encountered one or more problems generating one or more sub-capacity reports.

For the SCRT deliverable that runs on a z/OS system, look up any SCRTTOOLxxx messages that are listed by the SYSPRINT DD in the job log from the SCRT step in [Appendix A, “SCRT messages,” on page 449](#) for an explanation of the problem. Follow the actions indicated for each message to address each problem before attempting to rerun the SCRT job.

If the required actions do not correct the problem, open a service request, as described in [“Obtaining technical support from IBM” on page 439](#), with a description of the problem, an attachment containing the job log for the SCRT job (if you ran SCRT on a z/OS system), and attachments containing each of the sub-capacity reports that may have been created.

After these problems have been solved, proceed to the next step in this process.

7. Review the sub-capacity report. You can choose any of the following methods to review the sub-capacity report:

- If you generated your reports on a z/OS system:
 - Use one of the tools available on the z/OS system on which the report was generated (such as the ISPF browse or edit functions, under TSO/E).
 - Move the report to your workstation as a .csv (comma-separated values) text file and use a spreadsheet application on your workstation.
 - If you have signed up for and have access to the LMS web application, move the report to your workstation as a .csv text file and use the LMS web application.
- If you generated your reports on a Windows or Linux system:
 - Use a spreadsheet application on your workstation to review the .csv text file.
 - If you have signed up for and have access to the LMS web application, use the LMS web application.

Based on the results of your review, you might need to rerun SCRT using different SCRT control statement values to account for any unexpected changes to your hardware or software configuration, or to handle special conditions which might require you to use optional SCRT control statements, such as the **Exclude** control statement.

While reviewing the sub-capacity report, add comments to the report, if necessary, before submitting it to IBM. For more information, see [Chapter 24, “Viewing and editing the sub-capacity report,” on page 419](#) and [Chapter 25, “Preparing to submit the sub-capacity report to IBM,” on page 423](#).

8. Contact your IBM representative or IBM Business Partner representative to obtain all of the IBM contracts and addenda that are required for any sub-capacity pricing metrics and pricing options. Verify that all of the required contracts and addenda have been signed prior to submitting any sub-capacity reports. Failure to have the required contracts or addenda signed or otherwise agreed to may cause IBM to reject the affected sub-capacity report.
9. Submit the sub-capacity report to IBM. If you have signed the appropriate Customer Agreement and the report is being used for billing, you can use the LMS web process to submit the sub-capacity report to IBM. Remember that the LMS application only accepts as input the comma-separated value format with a file extension of .csv. For details, see [Chapter 26, “Submitting sub-capacity reports to IBM,” on page 427](#).
10. The CPC must be configured to send weekly Transmit System Availability Data (TSAD, also known as Call Home). All of the sub-capacity pricing metrics supported by SCRT require CPCs to submit TSAD.

This configuration is usually done when the CPC is installed. For more information, see [“Sending TSAD data”](#) on page 514.

11. Retain the SMF and SCRT89 data. The IBM sub-capacity contract requires you to save all SMF type 70 records, SMF type 89 records, and SCRT89 records for at least six months in order to be able to resolve questions and problems.

For example, if you begin to collect data for SCRT in the month of January, you must submit the sub-capacity report by 9 February. Your submitted data will be reflected in the software billing for March.

Using SCRT for planning purposes

For z/OS, z/TPF, and z/VSE environments that meet the SCRT requirements for support, you can use SCRT to generate a report for your own planning purposes at any time. A sub-capacity report is not used by IBM for billing unless you submit it officially and meet the sub-capacity eligibility requirements. The information in the report can be useful if your configuration meets the prerequisites for using SCRT and you are considering implementing sub-capacity pricing, but have not yet met all the sub-capacity pricing terms and conditions.

Chapter 2. Preparing to use the Sub-Capacity Reporting Tool

Before you install and run the Sub-Capacity Reporting Tool (SCRT), you will need to learn about and understand the following topics:

- Where to get more information on the relationship between LPAR sizes and software costs
- How to prepare your hardware and software
- How to collect SMF and SCRT89 records
- How to extract and file the SMF and SCRT89 records
- How SCRT determines the reporting period

Understanding the relationship between LPARs and sub-capacity pricing

With sub-capacity pricing you are charged for certain software products based on the utilization of the LPARs in which these products run. This means that sub-capacity pricing can only be applied to CPCs that have LPARs configured. To configure LPARs on an IBM Z CPC, use the Processor Resource/System Manager (PR/SM) facility.

For more information on PR/SM, see the appropriate edition of *Processor Resource/System Manager Planning Guide* for your server's product line, available from the [Resource Link home page](http://www.ibm.com/servers/resourcelink) (www.ibm.com/servers/resourcelink).

For more information on sub-capacity pricing and LPAR utilization for z/OS systems, see *z/OS Planning for Sub-Capacity Pricing* (https://www.ibm.com/docs/en/SSLTBW_3.1.0/pdf/e0zi100_v3r1.pdf).

Preparing your hardware and software

Before collecting any SMF or SCRT89 records for SCRT, verify that all of the required hardware and software service has been applied, that the appropriate hardware has been IMLed (if required), and that the appropriate operating systems have been IPLed (if required). Failure to apply and activate all of the required service may prevent SCRT from generating a sub-capacity report that will be acceptable to IBM for sub-capacity pricing. For details, see step 1 in “Overview of the SCRT process” on page 8.

The prerequisites for sub-capacity pricing (and SCRT) are:

- z/Architecture hardware platforms only:
 - IBM z17 (z17) CPCs
 - IBM z16 (z16) CPCs
 - IBM z15 (z15) CPCs
 - IBM z14 (z14) CPCs
 - IBM z13[®] (z13) and z13s[®] CPCs
 - IBM zEnterprise EC12 (zEC12) and zEnterprise BC12 (zBC12) CPCs
 - IBM zEnterprise 196 (z196) and zEnterprise 114 (z114) CPCs
 - IBM System z10 platform: System z10 Enterprise Class (z10 EC) and System z10 Business Class (z10 BC) CPCs
 - IBM System z9 platform: System z9 Enterprise Class (z9 EC) and System z9 Business Class (z9 BC) CPCs

- IBM eServer zSeries hardware: zSeries 990 (z990), zSeries 900 (z900), zSeries 890 (z890), and zSeries 800 (z800) CPCs
- At least one LPAR must be configured on the CPC.
- When z/OS is licensed for sub-capacity pricing:
 - All instances of z/OS running on the CPC must be running in z/Architecture (64-bit) mode.
 - There can be *no* OS/390 or MVS operating systems licensed or executing on the CPC. (Any older systems that were running on that CPC must be discontinued before you start collecting SMF records for SCRT reporting.)
 - You must configure the systems to collect SMF type 70 subtype 1 records and SMF type 89 subtypes 1 and 2 records.
- When z/TPF is licensed for sub-capacity pricing:
 - You must configure the system to collect SCRT89 records.
- When z/VSE is licensed for sub-capacity pricing:
 - All instances of z/VSE running on the CPC must be at the Version 4 level or higher.
 - You must configure the system to collect SCRT89 records.
 - IBM Z CPCs only
- SCRT must be installed and customized (IBM strongly recommends testing SCRT prior to creating an official report for submission to IBM)
- You must submit monthly sub-capacity reports to IBM for each CPC that uses sub-capacity pricing.
- You must submit monthly TSAD data to IBM for each CPC that uses sub-capacity pricing.
- You must save all SMF type 70 records, SMF type 89 records, and SCRT89 records for at least six months in order to be able to resolve questions and problems.

Using z/OS SMF records with SCRT

Before running SCRT, you must collect SMF type 70 subtype 1 records and type 89 subtypes 1 and 2 records for z/OS for one reporting period. IBM billing procedures define a *reporting period* as the start of the second day of a month through the end of the first day of the next month (for example, from the start of the day on 2 January to the end of the day on 1 February). The SMF interval must not be greater than one hour.

The input for SCRT is one or more data sets or files containing SMF type 70 subtype 1 and SMF type 89 subtypes 1 and 2 records, together with the SCRT89 records, for all the LPARs on the specified CPCs. (You can include data for multiple CPCs in an input data set.) See *z/OS MVS System Management Facilities (SMF)* for a description of the fields in these records. SCRT ignores SMF type 70 and 89 records that are outside the reporting period being analyzed.

To produce SMF type 70 subtype 1 records, a systems management product is required. For example, you can use Resource Measurement Facility (RMF) or an equivalent product. No additional product beyond the base z/OS operating system is required to produce SMF 89 records.

Using SCRT89 records with SCRT

Before running SCRT, you might also need to collect SCRT89 records for z/TPF systems and/or z/VSE systems for one reporting period. IBM billing procedures define a *reporting period* as the start of the second day of a month through the end of the first day of the next month. For example, from the start of the day on 2 January to the end of the day on 1 February.

SCRT89 intervals are set by the operating systems creating those records and cannot be adjusted by you. The input for SCRT is one or more sequential data sets containing SMF type 70 subtype 1 and SMF type 89 subtypes 1 and 2 records, together with the SCRT89 records, for all the LPARs on the specified CPCs. (You can include data for multiple CPCs in an input data set.) SCRT ignores SCRT89 records that are outside the reporting period being analyzed. To generate SCRT89 records on a z/TPF system see

[“Sub-capacity reporting using SCRT89 records” on page 283](#). To generate SCRT89 records on a z/VSE systems, see [“Overview of the CMT” on page 309](#).

How SCRT uses SMF and SCRT89 records

In each hour, for SMF records from z/OS systems, SCRT compares the 4-hour rolling average utilization (the average of the SMF70LAC values in that hour) to the defined capacity, if set; otherwise to the maximum capacity (the average of the SMF70WLA values in that hour) and chooses the lesser of the two. That chosen value becomes the utilization value for the z/OS system for that hour. See [“SCRT limitations” on page 439](#) for information about cases in which SCRT is not able to precisely calculate utilization. Note that under the appropriate conditions, SCRT may choose a lower 4-hour rolling average value based on group capping. See [“Group capacity” on page 498](#) for information about group capacity. SCRT simply uses the 4-hour rolling average value provided on SCRT89 records because the z/TPF and z/VSE systems that generate those records do not support defined capacity or group capping.

SCRT calculates MSU values for all products, but you must use the NO89 DD statement to provide SCRT with the information it needs for products that do not produce SMF type 89 records. Use the NO89 parameter of the SCRT JCL to list the LPARs in which such products are active. See [“NO89 DD statement” on page 61](#) for an explanation of the NO89 parameter.

Collecting z/OS SMF records

If the z/OS systems in your installation are not already collecting SMF type 70 subtype 1 and SMF type 89 records, you must update the **SYS** and **SUBSYS** parameters of the SMFPRMxx PARMLIB member to specify that these records should be collected. For example, the following SMFPRMxx parameters enable SMF to collect type 70 subtype 1 and type 89 records:

```
SYS(TYPE(70,89))
SUBSYS(STC,TYPE(70,89))
```

Tip: Because SCRT assigns each SMF type 70 and type 89 record to a specific hour of the month, ensure that the SMF record intervals are synchronized to the start of the hour for all z/OS systems in the report.

Since SMF type 70 subtype 1 and type 89 records represent a very small percentage of the total volume of SMF records produced and saved in a typical installation, you may want to modify your collection procedures to put a copy of only these two record types in a separate data set just for the use of SCRT. Having only the records needed for SCRT processing in a dedicated data set causes the SCRT job to run very quickly.

SCRT expects its input data set to contain only one month of SMF data. If you use year-to-date files or files with more than one month of data, refer to the discussion of the **Report_Period** parameter in [“PARMS DD statement” on page 34](#).

For more information on collecting SMF records, see chapter 4 of *z/OS MVS System Management Facilities (SMF)*. For syntax information about the SMFPRMxx member, see *z/OS MVS Initialization and Tuning Reference*.

SCRT discards duplicate SMF records so, if you accidentally collect the same records more than once, you can still use what you collected as input to SCRT.

Collecting SCRT89 records

To generate SCRT89 records on a z/TPF system, see [“Sub-capacity reporting using SCRT89 records” on page 283](#). To generate SCRT89 records on a z/VSE system, see [“Overview of the CMT” on page 309](#).

Extracting z/OS SMF records you have collected

Before you can run SCRT, you must extract the SMF type 70 subtype 1 and SMF type 89 records that you have collected. If your installation already has a procedure in place for dumping SMF data sets to tape (or to another set of DASD), you can continue to follow your normal procedure for collecting SMF data. If not, you will have to modify your existing procedures to extract the SMF type 70 subtype 1 and SMF type 89

records. The JCL that runs SCRT needs to point to the place where the extracted SMF type 70 and type 89 records reside. For more information, see chapters 2 and 3 of *z/OS MVS System Management Facilities (SMF)*.

Extracting SCRT89 records you have collected

Before you can run SCRT, you must extract the SCRT89 records that you have collected. See [“Sub-capacity reporting using SCRT89 records” on page 283](#) for a description of the process for extracting the SCRT89 records generated by a z/TPF system. See [“Sub-capacity reporting using SCRT89 records” on page 314](#) for a description of the process for extracting the SCRT89 records generated by a z/VSE system.

Retaining SMF and SCRT89 records you have collected

The IBM sub-capacity contract requires you to save all SMF type 70 records, SMF type 89 records, and SCRT89 records for at least six months in order to be able to resolve questions and problems.

SMF and SCRT89 records from different time zones

The SMF and SCRT89 records contain various date and time stamps that indicate the measurement interval represented by those records, as well as the timing of other events. Those date/time stamps (date, hour, second, hundredth of a second) are all saved on the SMF and SCRT89 records in terms of the local time (time zone) set for the operating system. Operating system time zones are usually set to match the time zone of the majority of the end users who are using the applications running on those operating systems. Examples of time zones include Eastern Standard Time (New York City, USA) and Central European Time (Boeblingen, Germany).

Whenever all of the operating systems running on the same CPC are set to the same local time zone, SCRT reports all of the dates and times on the sub-capacity report in terms of that local time zone because the date/time stamps from all of the operating systems on that CPC are synchronized to represent the same time. For instance, the local time stamp of 1:00 AM on January 1, 2011 represents the same time for all of the operating systems on the CPC. However, when one or more of the operating systems uses a different time zone (for instance, because the applications used by its end users are in different time zones) the same local time stamp of 1:00 AM represents different hours in different time zones.

Since SCRT adds each product's sub-capacity values together on a per-hour basis across all of the operating systems in which the product runs on a given CPC, SCRT must synchronize the dates/times for all of those operating systems so that the synchronized dates/times represent the same dates/times in terms of Coordinated Universal Time (UTC). That way, the hour represented by 1:00 AM on January 1, 2011 in UTC time represents the same date/time everywhere, regardless of the local date/time stamp.

SCRT converts the local dates/times to UTC dates/times whenever a difference between any of the operating systems' local time zones lasts longer than two hours on the same CPC during the same SCRT reporting period. SCRT tolerates a time zone difference which lasts up to two hours in order to avoid converting all of the dates/times to UTC dates/times when you change to or from Daylight Saving Time, provided the time offsets of all of the operating system instances can be changed within two hours. SCRT uses the offset from what the operating system instance believes is a hardware clock set to the actual UTC time to determine whether there are time zone differences between different operating system instances. SCRT identifies such converted dates/times by appending a suffix of UTC after those dates/times on the sub-capacity report. When this condition occurs, the sub-capacity report shows a suffix of UTC after all UTC time values.

This documentation refers to the LPAR Sysplex Timer offset value and the z/VM virtual clock (VTOD) value as *hypervisor offsets* because each of those values is controlled by a hypervisor, either the PR/SM hypervisor or the z/VM hypervisor. SCRT does not count the additional date/time offsets imposed by a Sysplex Timer offset or a VTOD value when determining whether the local time zones for any of the operating system instances are different and should be converted to UTC values.

Operating systems that run with the time set to the future or the past

Suppose you need to test certain applications by running them as if they were executing in the future or in the past. When such testing is required, you might adjust the LPAR or operating system date to run as if it were the end of the month, the end of the quarter, or the end of the year; or you might adjust them to run before a certain event occurred. Perhaps you need to run several such tests with different adjusted times.

Starting with the original z/TPF and z/VSE support for SCRT and with z/OS V1R6 or later, z/OS, z/TPF, and z/VSE systems using future or past dates will automatically be correctly handled without any additional changes to the current SCRT process, provided the following conditions are met:

- The CPC clock must be set to the actual UTC date and time. This has always been a requirement for sub-capacity pricing.
- One or more of the following methods are used to adjust the local date/time reported by the z/OS, z/TPF, or z/VSE operating system to a future or past date/time:
 - Setting a Sysplex Timer offset value on the image profile for the LPAR
 - Choosing a future or past date when replying to the CLOCK prompt during the IPL of the operating system
 - Using the MVS **SET** command to reset the date to a future or past date after IPLing z/OS
 - Setting z/VM's virtual clock (VTOD value) to a future or past date for a z/OS, z/TPF, or z/VSE guest system. Note that the Sysplex Timer offset cannot be used in conjunction with this method.

This support can handle cases where any combination of the supported methods for setting a different date are used simultaneously in different LPARs or different z/OS systems on the same CPC (except that the VTOD reset and a Sysplex Timer offset are not supported together). This support also handles cases where the same dates are used repeatedly during the reporting period.

Note: If the local system time results in a UTC offset greater than 24 hours in either direction, SCRT converts the local dates and times to UTC dates and times.

z/OS V1R6 systems must have the following APARs applied in order for SCRT to properly handle future or past dates/times:

- RMF APAR OA11375 must have been applied before the SMF type 70 records were collected.
- SMF APAR OA11036 must have been applied before the SMF type 89 records were collected.

If your reports show problems that you believe are related to the use of a past or future time even though you have met all of the criteria for handling these conditions, open a service request, as described in “Obtaining technical support from IBM” on page 439. The support team will determine whether SCRT can handle your particular situation and will send you special instructions for handling it.

How SCRT determines the license capacity

Table 3 on page 15 shows an example of how SCRT calculates the license capacity for different product and LPAR combinations.

Table 3. Example of how SCRT calculates the license capacity for different product and LPAR combinations

Hour	1		2		3		...	719		720	
LPAR	A	B	A	B	A	B	...	A	B	A	B
4-hour rolling average, in MSUs	70	30	75	30	50	50	...	100	30	100	10
z/OS	yes	yes	yes	yes	yes	yes	...	yes	yes	yes	yes
Db2	yes	no	yes	no	yes	no	...	no	no	no	no

The peak interval is the highest utilization determined from the sum of the utilization for all LPARs in which a particular product ran in a given hour. It is *not* the sum of highest utilization for individual LPARs in which a particular product ran during different hours.

In this example, the peak interval for z/OS is in hour 719. The z/OS utilization value for the month is the sum of the z/OS utilization in both LPARs during hour 719, or 130 MSUs.

The peak interval for Db2 is in hour 2. Since Db2 only runs in LPAR A, the Db2 utilization value for the month is the value for LPAR A in hour 2, or 75 MSUs.

What is the 4-hour rolling average utilization?

The 4-hour rolling average utilization measures sustained use of CPC capacity and helps compensate for high and low values of short duration. [Table 4 on page 16](#) shows approximately how SCRT records the 4-hour rolling average for a 12-hour period. (SCRT uses a slightly more complex calculation than the simple average shown here.)

Table 4. Simple example of a 4-hour rolling average utilization

Hour:	1	2	3	4	5	6	7	8	9	10	11	12
Utilization this hour:	38	35	36	40	42	45	50	47	46	46	44	40
4-hour rolling average utilization:	–	–	–	37	38	41	44	46	47	47	46	44

You can see that the highest and lowest values do not appear in the average, since they are uncommon.

How the IBM zAAP and IBM zIIP affect sub-capacity pricing

On IBM zEnterprise CPCs that support them, specialty engines (processor units) called IBM zEnterprise Application Assist Processors (zAAP) support eligible application workloads, such as Java and XML.

On IBM Z CPCs that support them, specialty engines (processor units) called IBM z Integrated Information Processors (zIIP) support specific data and transaction processing workloads for business intelligence (BI), ERP, and CRM, and select network encryption workloads.

LPARs can consist of a combination of general purpose CPs and zAAPs or zIIPs.

- When LPARs are set up with both general purpose CPs and zAAPs, some or all of any zAAP-eligible workload that runs in that LPAR runs on the zAAP processors (depending on the capacity of the zAAPs).
- When LPARs are set up with both general purpose CPs and zIIPs, some or all of specific zIIP-eligible workloads that run in that LPAR run on the zIIP processors (depending on the capacity of the zIIPs).

In either case, the 4-hour rolling average measured by SCRT is based only on the processing of the general purpose CPs in the LPAR. Adding zAAPs to an LPAR that runs zAAP-eligible workloads can decrease the 4-hour rolling average utilization for all products on that LPAR. Adding zIIPs to an LPAR that runs specific zIIP-eligible workloads can decrease the 4-hour rolling average utilization for all products on that LPAR.

Contact your IBM representative for help in determining how many zAAP or zIIP processors you could apply to support workloads in your LPARs.

The SCRT reporting period

An SCRT *reporting period* runs from 00:00:00 on the second day of a month to 24:00 on the first day of the next month (that is, from the very start of the second day of a month to the very end of the first day of the next month). For instance, the March reporting period runs from 00:00 on 2 March to 24:00 on 1 April.

SCRT determines the reporting period by choosing the reporting month with the highest percentage of data collected (by SMF) and provided in the input to SCRT. In the case where multiple reporting periods have the same percentage of data collected, SCRT chooses the latest reporting period. SCRT discards data from any days that are outside of the chosen reporting period.

You can use the PARMS **Report_Period** control statement to force SCRT to use a different reporting period. For more information about this control statement, see [“PARMS parameters that apply to all CPCs” on page 34](#).

For a particular hour to be included in the calculations, there must be both an SMF type 70 subtype 1 record and an SMF type 89 record for that hour. The hour is counted as part of the percentage of data collected as long as there are SMF type 70 and SMF type 89 records for one or more LPARs on the CPC. That is, if a CPC has two LPARs but only one has valid SMF data during a particular interval, that interval is still counted.

SCRT data collection reporting

SCRT reports information about data collection incidents whenever it detects that expected data for the reporting period is missing (also known as *data gaps* or *data collection incidents*).

SCRT reports this information in the following sections:

Missing LPAR Data Details (H4)

Section H4 reports two types of message records: one for LPARs that appear to have been active but for which no SMF or SCRT89 records were found in the input data, and another for mismatches between SMF type 70 and type 89 records. SCRT reports one message record in section H4 for each contiguous time interval during which an LPAR had neither SMF type 70 nor type 89 records (or SCRT89 records), or had either an SMF type 70 record or a type 89 record but not both.

Missing CPC Data Details (H6)

SCRT reports one message record in section H6 for each contiguous time interval when no data from any LPAR on a CPC was found in the input data.

You must provide a justification for each incident reported in sections H4 and H6. SCRT automatically inserts a justification for data gaps with a one-hour duration, but this does not change, reduce, or eliminate your contractual obligation to provide all of the required data for sub-capacity reporting.

The following examples are valid reasons for an LPAR to report no data:

- The LPAR was used as a coupling facility.
- The LPAR only ran z/VM or Linux for z Systems®.
- The LPAR was excluded for a valid reason, such as a non-production CBU test or Backup (DR) test.

The following examples are not valid reasons for an LPAR to report no data and are unacceptable:

- The LPAR is used as a "sandbox" for systems programmers.
- The LPAR is used for maintenance testing.
- The LPAR is a test LPAR that does not run RMF.

IBM considers all of the preceding invalid reasons to be billable activities, along with any other similar activities.

If SCRT detects no data gaps during the reporting period, sections H4 and H6 report that all expected data was collected.

Unusual situations that affect utilization

Customers have the primary responsibility for preventing uncontrolled loops, operator errors, or unwanted utilization spikes. However, IBM understands that, occasionally, situations that could not be prevented (especially situations related to disaster recovery) might cause exceptional utilization values. In these situations, IBM does not normally expect customers to pay for the increased utilization associated with the unusual situation.

Use your best judgement to determine if an unusual situation has occurred. IBM does not publish a list of unusual situations because, by their nature, they will be unpredictable.

For products that do not have the (optional) and (conditional) comment placeholders in the Product Summary Information section of the sub-capacity report, if you know the CPC, LPAR names,

product IDs, and date and time ranges for when your unusual situation occurred, you might be able to use one or more SCRT **Exclude** control statements to exclude the data that is in error and have SCRT list the appropriate MSU values on the sub-capacity report. For details about the **Exclude** control statement, see [“EXCLUDE control statement” on page 44](#).

In cases that involve NO89 products, you cannot use the **Exclude** control statement to adjust the product MSUs because the **Exclude** control statement does not support NO89 products. If you know the correct MSU value (whole numbers only, with a maximum of four digits), you can indicate an unusual situation by replacing the (optional) comment placeholder in the **Customer MSUs** field of the sub-capacity report with the correct number of MSUs. You can use such tools as the Sub-Capacity Planning Tool, MXG Software reports, or RMF to determine the appropriate value for **Customer MSUs**. Describe the reason for using the **Customer MSUs** field in the **Customer Comments** field. See the description of these fields in [“Product Summary Information section” on page 110](#).

IBM normally trusts your judgment when you indicate that an unusual situation has occurred. However, if you repeatedly override the **Tool MSUs** value (for example, if you always have unusual situations occur at the end of a quarter) or if large discrepancies are found during IBM audits (IBM Call Home TSAD data is compared with data from sub-capacity reports), IBM will contact you at the time the pattern is noticed.

For MSU consumption-based offerings, including Tailored Fit Pricing for IBM Z - Software Consumption Solution and Software Consumption License Charges, it is not normally appropriate to exclude all of the MSU consumption from an LPAR for a qualified event that causes unusual utilization, such as looping products. Instead, you can use the **EXCLUDE MSU_CONSUMPTION** control statement to exclude only the MSU consumption associated with a qualified event. The **EXCLUDE MSU_CONSUMPTION** control statement allows you to exclude consumed MSU from a specific solution based upon the excess CPU time experienced during an unusual situation. You should use your best judgment and available tooling to estimate the CPU time associated with the incident. The **EXCLUDE MSU_CONSUMPTION** control statement should only be used in cases where the impact of the unusual activity can be quantified (for instance, unusual CPU consumption caused by an event such as a looping IBM program). In cases where the event occurs over a prolonged period of time or has an impact that is difficult to quantify, work with your IBM representative to determine the best approach.

Guideline: Only customers who are running the z/TPF operating system should use LPAR names that begin with the characters TPFT or TPFP. However, if you have LPARs whose names begin with those characters but are not running z/TPF, and you are unable to change those LPAR names, SCRT might incorrectly report z/TPF usage. In this case, replace the (optional) comment placeholder in the **Customer Comments** row at the very beginning of the sub-capacity report with an explanation stating that the LPAR names beginning with TPFT or TPFP are not being used to run z/TPF. Note that you cannot use **Exclude** control statements to exclude the z/TPF product because **Exclude** control statements do not support excluding a product ID for an operating system. For more information about SCRT and z/TPF, see Chapter 13, [“Sub-capacity reporting for z/TPF systems,” on page 281](#).

Rule: If an unusual situation occurs, *do not* delete any SMF records from the SCRT input data set and *do not* change any values in the **Tool MSUs** field of the sub-capacity report.

For more information about unusual situations, see [“Providing alternate MSU values” on page 443](#).

Part 2. Using the Sub-Capacity Reporting Tool

This part describes how to obtain, install, set up, and run SCRT on each platform on which it runs.

This part contains the following chapters:

- [Chapter 3, “Using SCRT on z/OS systems,” on page 21](#), which describes how to generate sub-capacity reports by running SCRT on a z/OS system
- [Chapter 4, “Using SCRT on Windows and Linux systems,” on page 69](#), which describes how to generate sub-capacity reports by running SCRT on a Microsoft Windows system or on a Linux system

Chapter 3. Using SCRT on z/OS systems

This chapter describes how to download, install, and use the Sub-Capacity Reporting Tool (SCRT) to generate sub-capacity reports by running SCRT on a z/OS system.

This chapter also describes the level of double-byte character set (DBCS) support provided by SCRT and by different parts of the sub-capacity reporting process.

Requirements for SCRT

In addition to the prerequisites discussed in Chapter 2, “Preparing to use the Sub-Capacity Reporting Tool,” on page 11, prerequisites for running SCRT on z/OS systems include:

- z/OS software prerequisites
- TSO user ID requirements

z/OS software prerequisites for SCRT

SCRT requires the following software on the z/OS system on which you will run SCRT:

- One of the following Java environments:
 - IBM SDK, Java Technology Edition, Version 8 or later with all current maintenance
For more information and installation instructions, see [IBM SDK, Java Technology Edition \(www.ibm.com/docs/en/sdk-java-technology\)](http://www.ibm.com/docs/en/sdk-java-technology).
 - IBM Semeru Runtime Certified Edition for z/OS, Version 11 or later with all current maintenance
For more information and installation instructions, see [IBM Semeru Runtime Certified Edition for z/OS \(https://www.ibm.com/products/semeru-runtime-certified-zos\)](https://www.ibm.com/products/semeru-runtime-certified-zos).
- IBM JZOS Batch Launcher and Toolkit (included with IBM SDK, Java Technology Edition, or IBM Semeru Runtime Certified Edition for z/OS)

For installation instructions, see the JZOS Batch Launcher and Toolkit documentation for your Java environment:

- IBM SDK, Java Technology Edition: JZOS Batch Launcher and Toolkit (<https://www.ibm.com/docs/en/sdk-java-technology/8?topic=components-jzos-batch-launcher-toolkit>)
- IBM Semeru Runtime Certified Edition for z/OS: JZOS Batch Launcher and Toolkit (<https://www.ibm.com/docs/en/semeru-runtime-ce-z/11.0.0?topic=guide-jzos-batch-launcher-toolkit>) Use the **Change version** drop-down list to select the version of the Semeru Runtime that you are using.

TSO user ID requirements for SCRT

SCRT is to be installed under z/OS UNIX System Services (z/OS UNIX) and requires a user ID that has a defined OMVS segment.

Obtaining and installing SCRT on a z/OS system

Follow these instructions to obtain and install SCRT on your z/OS system.

Obtaining SCRT for supported z/OS releases

Starting with z/OS V2R3, SCRT is included with z/OS. Updates are provided via PTF for all releases of z/OS that are in-support as of the SCRT update. If your version or release of z/OS is out-of-support, you must follow the instructions in [“Obtaining SCRT for out-of-support z/OS releases” on page 22.](#)

About this task

SCRT is already installed in the `/usr/lpp/scrt/IBM` directory.

SCRT sample jobs (JCL) exist in the following members:

- `SYS1.SAMPLIB(ISCJCL80)` — Invokes the 32-bit Java runtime
- `SYS1.SAMPLIB(ISCJCL86)` — Invokes the 64-bit Java runtime

At z/OS 3.1 and later, additional sample jobs (JCL) exist in the following members:

- `SYS1.SAMPLIB(ISCJCL16)` — Invokes IBM Semeru Runtime Certified Edition for z/OS, Version 11
- `SYS1.SAMPLIB(ISCJCL17)` — Invokes IBM Semeru Runtime Certified Edition for z/OS, Version 17
- `SYS1.SAMPLIB(ISCJCL21)` — Invokes IBM Semeru Runtime Certified Edition for z/OS, Version 21

Procedure

No further action is needed to obtain or install SCRT.

What to do next

Continue with [“Running SCRT on a z/OS system” on page 24.](#)

Obtaining SCRT for out-of-support z/OS releases

Perform this task to download the SCRT deliverable.

Before you begin

Ensure that you understand the requirements, as described in [“Requirements for SCRT” on page 21.](#)

About this task

The SCRT deliverable includes five sample JCL parts in the `/jcl` directory. These include a 64-bit and a 31-bit sample for each version of Java, as shown in the following list:

- `SCRT31.jcl` — Invokes the Java 8 31-bit JZOS batch launcher and JDK
- `SCRT64.jcl` — Invokes the Java 8 64-bit JZOS batch launcher and JDK
- `SCRT16.jcl` — Invokes IBM Semeru Runtime Certified Edition for z/OS, Version 11
- `SCRT17.jcl` — Invokes IBM Semeru Runtime Certified Edition for z/OS, Version 17
- `SCRT21.jcl` — Invokes IBM Semeru Runtime Certified Edition for z/OS, Version 21

Note: The support for SCRT control statements in these JCL files is identical. The difference provided by these JCL files is in the underlying environment in which SCRT runs. This is a necessary step because the resource usage for SCRT can vary widely depending on the number of CPCs and LPARs that exist in a given reporting environment.

Procedure

Download the SCRT deliverable to your workstation from [IBM MRS Tool - System z Software Licensing Tool Downloads \(ibm.biz/SCRTzOSDownload\)](#).

What to do next

Continue with [“Installing SCRT for out-of-support z/OS releases”](#) on page 23.

Installing SCRT for out-of-support z/OS releases

After downloading the `scrttool.tar.gz` file from IBM, perform this task to install SCRT.

Before you begin

If you have an existing release of SCRT V30.1.0 or later installed, you can, in most cases, continue to use your existing SCRT JCL. You can use the following procedure to install the deliverable; however, instead of copying over the new SCRT JCL, you can simply modify your existing copy so that the **APP_HOME** environment variable points to the new installation location.

Procedure

1. On your z/OS system, ensure that OMVS (z/OS UNIX) is active and the file system where you want to install SCRT is available.
2. Upload (using FTP or other utility) the `scrttool.tar.gz` file in binary format to a z/OS UNIX directory.

The following steps demonstrate one method of accomplishing this:

- a) Open a command session on your workstation.
- b) Change to the directory that contains the `scrttool.tar.gz` file.
- c) Begin an FTP session with your z/OS host system:

```
ftp hostname
```

- d) When prompted, enter your user ID and password on the host system.
- e) Change the FTP transfer mode to binary:

```
bin
```

- f) Move the `scrttool.tar.gz` file to a directory in the z/OS UNIX file system on the z/OS host system:

```
put scrttool.tar.gz /u/myuserid/mydir
```

Replace `/u/myuserid/mydir` with the desired z/OS UNIX path.

- g) End the FTP session:

```
quit
```

3. Log in to z/OS UNIX on your host system.
You can do this in one of the following ways:
 - Log on using a secure shell (SSH) client program.
 - Enter the OMVS command from the ISPF command line.
4. Navigate to the directory containing the `scrttool.tar.gz` file, as determined by you earlier in step “2.f” on page 23.
5. Use the **pax** or **tar** command to extract the file.
The following example uses the **pax** command:

```
pax -rvf scrttool.tar.gz
```

This expands the SCRT deliverable and its directory structure into the directory where the `scrttool.tar.gz` file resides.

In the directory that was just created, locate the /jcl subdirectory. The /jcl directory contains the SCRT JCL samples that are used to launch SCRT. Only one of these is necessary. For information about how to determine which JCL sample to use, see [“Running SCRT on a z/OS system”](#) on page 24.

6. Choose the appropriate sample JCL file, such as SCRT64 . JCL, and copy it to an MVS data set.

An easy way to do this is to create a new sequential data set or PDS member, as in the following example.

- a) Create a data set on the z/OS host.

Use ISPF 3.2 to allocate a new 80-byte, fixed data set called *userid* . T00L . JCL. Specify the following attributes to allocate the data set:

Attribute	Value
Record format	FB (fixed-block)
Record length	80
Directory blocks	0
Block size	0 (system-determined block size)

- b) Follow these steps to copy the JCL to the data set.

- i) Edit the data that you allocated in the previous step.

- ii) From the editor command line, enter:

```
copy '/u/myuserid/mydir/scrtv301/jcl/SCRT64.JCL'
```

Replace */u/myuserid/mydir* with the location into which you expanded the *scrttool.tar.gz* file. If you chose to use a different JCL sample, replace SCRT64 . JCL with the appropriate file name.

7. Modify the JCL.

Perform the following steps to modify the JCL in addition to any modifications that are normally required by SCRT:

- a) Modify the SCRT EXEC PROC name to point to a valid JZOS procedure.

The sample provided with SCRT might not need to be changed; for more information, see [“SCRT EXEC statement”](#) on page 33.

- b) Make the following changes to the STDENV DD statement:

- i) Modify the **APP_HOME** line to point to the location of the *scrtv301* folder. For example:

```
APP_HOME=/u/myuserid/mydir/scrtv301
```

- ii) Modify the **JAVA_HOME** line to point to the Java 8 or later SDK.

- c) Modify the remaining SCRT DD statements according to the instructions in [“Customizing the SCRT JCL”](#) on page 25.

Results

You can now use the JCL to submit the job and produce sub-capacity report output.

What to do next

Continue with [“Running SCRT on a z/OS system”](#) on page 24.

Running SCRT on a z/OS system

SCRT is a Java application. Like other platforms, Java applications on z/OS require a Java virtual machine (JVM) in which to execute. The JVM provides the environment under which SCRT will execute.

Note: All references to Java apply equally to IBM Semeru runtimes.

There are various ways to start a JVM on z/OS. The approach that SCRT uses is to start SCRT and, thereby, the JVM in which it will run, by using the JZOS batch launcher, which is shipped with the IBM Java runtime environment. This approach facilitates keeping SCRT as a batch application with JCL that is similar to the classic SCRT.

The JZOS batch launcher starts the JVM using the options specified on the STDENV DD statement in the JCL. The SCRTJCL job contains the recommended set of options for starting the JVM based on both the requirements of the SCRT application and the results of IBM testing. Details about these options are described in [“Customizing the SCRT JCL”](#) on page 25.

Guidelines:

- Choose an appropriate JCL sample based on the Java level that is installed on the CPC on which you intend to generate sub-capacity reports.
- In most cases, choose a 64-bit JCL sample (for instance, SCRT64.jcl), which are preconfigured to support the majority of environments.
- When running with the IBM Semeru Runtime Certified Edition for z/OS, Version 21 or later, include -Dfile.encoding=COMPAT in the **IBM_JAVA_OPTIONS** environment variable that is set in the STDENV DD. This ensures that SCRT reads and writes using the platform native encoding (usually an EBCDIC variant) instead of UTF-8 and operates in a manner compatible with earlier Java and Semeru versions.
- The 31-bit JCL samples are only appropriate when you intend to run SCRT using input data that contains a small number of CPCs and LPARs.

In all cases, the JVM arguments in the JCL are based on IBM internal testing of SCRT.

Customizing the SCRT JCL

[“Sample JCL to run the 64-bit SCRT on a z/OS system”](#) on page 25 shows sample JCL that is similar to what you receive when you configure SCRT using the 64-bit JCL sample, SCRT64.jcl, that ships with SCRT.

Note: The sample JCL for the 31-bit version of SCRT (SCRT31.jcl) and the IBM Semeru versions (SCRT16.jcl, SCRT17.jcl, and SCRT21.jcl) are not shown, but have the same format as the SCRT64.jcl sample.

The JCL comments describe the most recent changes to SCRT and provide hints for customizing the JCL to meet the rules established for running batch jobs on the z/OS system on which you intend to run SCRT.

Sample JCL to run the 64-bit SCRT on a z/OS system

The following example shows the contents of the SCRT64.jcl file that ships with SCRT:

```
//SCRTJCL JOB (????,????),MSGCLASS=0,NOTIFY=HLQ,REGION=0M,TIME=20
//*****
//* Licensed Materials - Property of IBM
//* SubCapacity Reporting Tool V30 (SCRT)
//* Copyright IBM Corp. 2001, 2025
//*
//* SCRT VERSION 30.1.7 - INTENDED TO RUN ON z/OS
//*
//* For recent changes to the Sub-Capacity Reporting Tool (SCRT), see
//* the SCRT Current Levels document and/or the SCRT User's Guide.
//*
//* Requirements and Installation:
//*   - SCRT Java requirements:
//*     - IBM SDK for z/OS, Java Technology Edition,
//*       Version 8 or later with all current maintenance.
//*     or
//*     - IBM Semeru Runtime Certified Edition for z/OS,
//*       Version 11 or later with all current maintenance.
//*
//*   - This JCL invokes the Java 8.0 SDK.
//*   - This JCL invokes the 64-bit JVM and so usage of this
//*     JCL requires that the 64-bit Java 8.0 SDK is installed.
```

```

/**
/** Modifications necessary to run SCRT in the JAVA VM:
/**
/** - Tailor the proc and job for your installation:
/** - Modify the Job card per your installation's requirements
/** - Modify the EXEC PROC= statement to point to the PDS
/** containing the installation version of the JZOS PROC.
/** - Edit JAVA_HOME to point to the location of the Java JDK.
/** - Edit APP_HOME to point to the SCRT installation directory
/** containing the SCRT JAR (scrt.jar) file.
/**
/** - Do a global change on hlq to your desired high level qualifier.
/**
/** - Point the SMF DD at your SMF dataset containing type 70 and 89
/** records for the month to be reported. Note that for reporting
/** purposes a month begins at 00:00 on the second, and ends at
/** 24:00 on the first of the following month.
/**
/** - Update the PARMS and N089 DD statements below.
/**
//DELETE EXEC PGM=IEFBR14
//DD1 DD DSN=HLQ.SCRTT00L.CSV,DISP=(MOD,DELETE),
// SPACE=(TRK,(1)),UNIT=SYSDA
/** FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS
/** UNCOMMENT THE DD IF YOU CREATE A CSV FILE TO SHOW
/** TAILORED FIT PRICING DETAILED DATA
/** //CPSTRACE DD DISP=(MOD,DELETE),DSN=HLQ.CPSTRACE.CSV,
/** // SPACE=(TRK,(1)),UNIT=SYSDA
/** FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS
/** UNCOMMENT THE DD IF YOU CREATE A CSV FILE TO SHOW THE REASONS
/** THAT GETTING STARTED PRICING MSUS WERE NOT USED IN THE REPORT
/** //GSTRACE DD DISP=(MOD,DELETE),DSN=HLQ.GSTRACE.CSV,
/** // SPACE=(TRK,(1)),UNIT=SYSDA
/** FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS
/** UNCOMMENT THE DD IF YOU CREATE A CSV FILE TO SHOW THE REASONS
/** THAT INTEGRATED WORKLOAD PRICING WASN'T USED IN THE REPORT
/** //IWPTRACE DD DISP=(MOD,DELETE),DSN=HLQ.IWPTRACE.CSV,
/** // SPACE=(TRK,(1)),UNIT=SYSDA
/** FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS
/** UNCOMMENT THE DD IF YOU GENERATE A WORKLOAD CLASSIFICATION
/** SUPPLEMENT REPORT
/** //WCSUPP DD DISP=(MOD,DELETE),DSN=HLQ.WCSUPP.CSV,
/** // SPACE=(TRK,(1)),UNIT=SYSDA
/**
/**/*****
//SCRT EXEC PROC=JVMPC86,
// JAVACLS='com.ibm.scrt.SCRTe'
//STDENV DD *
# This is a shell script which configures
# any environment variables for the Java JVM.
# Variables must be exported to be seen by the launcher.

. /etc/profile
export JAVA_HOME=/usr/lpp/java/J8.0_64

export PATH=/bin:${JAVA_HOME}/bin

LIBPATH=/lib:/usr/lib:${JAVA_HOME}/bin
LIBPATH=${LIBPATH}:${JAVA_HOME}/lib/s390x
LIBPATH=${LIBPATH}:${JAVA_HOME}/lib/s390x/j9vm
LIBPATH=${LIBPATH}:${JAVA_HOME}/bin/classic
export LIBPATH

CLASSPATH=${JAVA_HOME}/lib:${JAVA_HOME}/lib/ext

APP_HOME=/u/userid/scrtv301
for i in ${APP_HOME}/*.jar; do
    CLASSPATH=${CLASSPATH}:${i}
done
export CLASSPATH

export JZOS_ENABLE_MVS_COMMANDS=false
# SCRT default 64bit JVM options
IJO="-Xms128m -Xmx2048m"
IJO=${IJO} -Xquickstart
IJO=${IJO} -Xcompressedrefs
export IBM_JAVA_OPTIONS=${IJO}
/**
/** FOR SEQUENTIAL OUTPUT (ALL CPCS IN ONE DS)
/** CHANGE SPACE PARAMETER TO SPACE=(TRK,(15,15))
//OUTPUT DD DISP=(,CATLG),DSN=HLQ.SCRTT00L.CSV,
// SPACE=(TRK,(15,15,15)),UNIT=SYSDA

```



```

//SMF      DD      DISP=SHR,DSN=HLQ.SMF.DATA
//PARMS    DD      *
*
* THERE ARE 5 REQUIRED PARM KEYWORDS:
*   THEY MUST BE FILLED IN OR THE PROGRAM WILL NOT PROCEED.
*
* ENTER YOUR CUSTOMER_NAME
Customer_Name=
*
* ENTER YOUR IBM COUNTRY CODE FOLLOWED BY YOUR SITE ID
CUSTOMER_NUMBER=
*
* IF COUNTRY_MULTIPLEX_PRICING OR ENTERPRISE_TAILORED_FIT_PRICING
* WAS SPECIFIED AND YOU MUST ASSIGN MULTIPLE SITE IDS
* WITHIN THE ENVIRONMENT, UNCOMMENT THE FOLLOWING LINE TO ASSIGN
* ADDITIONAL SITE IDS. NOTE THE BASE CUSTOMER_NUMBER
* PARAMETER WILL SERVE AS THE DEFAULT VALUE.
* ASSIGN CUSTOMER_NUMBER= ,CPC=
*
* ENTER THE NAME OF THE PERSON SUBMITTING THE REPORT
Contact_Name=
*
* ENTER THE EMAIL ADDRESS OF PERSON SUBMITTING THE REPORT
Contact_EMail=
*
* ENTER THE PHONE NUMBER OF THE PERSON SUBMITTING THE REPORT
Contact_Phone=
*
* OPTIONALLY UNCOMMENT THIS LINE AND ENTER A PO NUMBER
* PO_Number=
*
* OPTIONALLY UNCOMMENT ONE OF THESE LINES TO SELECT A REPORT PERIOD
* Report_Period=YYYY/MM
* Report_Period=Last_Month
*
* OPTIONALLY UNCOMMENT THESE LINES TO PROVIDE AN LPAR COMMENT
* LPAR_Comment CPC= ,LPAR= ,
* Comment="....."
*
//* FOLLOWING DD IS FOR OPTIONAL USE BY z/OS, z/TPF CUSTOMERS
//* UNCOMMENT DESIRED CONTROL STMTS AND PROVIDE MISSING VALUES
//SPECIAL DD *
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT CONTROL STMT TO REQUEST THAT SCRT INCLUDE ONLY THE
* IMAGE OR IMAGES SPECIFIED. THE COMMENT KEYWORD IS OPTIONAL.
* THE COMMENT CAN SPAN MULTIPLE LINES BY ENCLOSING THE COMMENT
* IN DOUBLE QUOTATION MARKS.
* Include CPC= ,Image_Id= ,
* Comment=
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT CONTROL STMT TO REQUEST THAT SCRT EXCLUDE AN IMAGE
* AND/OR PRODUCT FOR A SPECIFIED TIME INTERVAL. THE COMMENT
* KEYWORD IS OPTIONAL. THE COMMENT CAN SPAN MULTIPLE LINES BY
* ENCLOSING THE COMMENT IN DOUBLE QUOTATION MARKS.
* Exclude CPC= ,Image_Id= ,Product_Id= ,
* Start= ,Resume= ,
* Comment=
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT CONTROL STMT TO REQUEST THAT SCRT EXCLUDE MSU
* CONSUMPTION FOR A SOLUTION ID. THE COMMENT KEYWORD IS
* OPTIONAL. THE COMMENT CAN SPAN MULTIPLE LINES BY ENCLOSING
* THE COMMENT IN DOUBLE QUOTATION MARKS.
* Exclude MSU_CONSUMPTION,CPC= ,HOUR= ,CPUTIME= ,
* ID= ,
* Comment=
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE FOR z/OS
* UNCOMMENT CONTROL STMT TO REQUEST INTEGRATED WORKLOAD PRICING
* Integrated_Workload_Pricing
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE FOR z/TPF
* UNCOMMENT STMT TO REQUEST Z/TPF TRANSFORMATION ENGINES OFFERING
* z/TPF_Transformation_Engines_Offering
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST COUNTRY MULTIPLEX PRICING
* Country_Multiplex_Pricing
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST A TAILORED FIT PRICING ENTERPRISE

```

```

* SOLUTION REPORT
* Enterprise_Tailored_Fit_Pricing
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO ASSIGN Z WORKLOAD PRICING FOR CLOUD
* WLM CATEGORY ASSIGNMENT
* Assign zWPC=CATEGORYA
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST DETAILED INTERVAL DATA
* Generate_Detailed_Data
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST SCRT IGNORE WLM MOBILE DATA
* IGNORE_WLM_MWP_Data
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST WLM MOBILE DIAGNOSTIC MESSAGES
* DIAGMSG=WLMMOB
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST WLM CLOUD DIAGNOSTIC MESSAGES
* DIAGMSG=WLMZWPC
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST MLC PRODUCT FAMILY LINE WHEN
* ONLY ONE FAMILY PRODUCT IN USE.
* MLC_ALL
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST IPLA PRODUCT FAMILY LINE WHEN
* ONLY ONE FAMILY PRODUCT IN USE.
* IPLA_ALL
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST INTERVAL RATE DATA REPORTING
* INTERVAL_RATE_DATA
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST DETAIL INTERVAL RATE DATA REPORTING
* DETAIL_INTERVAL_RATE_DATA
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST CONTAINER INTERVAL
* CONSUMPTION REPORTING
* CONTAINER_INTERVAL_CONSUMPTION
*
* FOLLOWING SPECIAL CONTROL STMT IS REQUIRED FOR
* TAILORED FIT PRICING SOLUTIONS. FOR
* EVERY CONTAINER SOLUTION ID IN THE DATA, YOU MUST ASSIGN
* THE ASSOCIATED LMS PROVIDED NAME AND YOU MAY OPTIONALLY
* SPECIFY THE N089 PRODUCTS THAT EXECUTE WITHIN THE CONTAINER.
* Update Container,
*   Id= ,
*   Set_Name="...",
*   N089_Product_Id=
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO DEDICATE AN LPAR TO A TAILORED FIT PRICING
* SOLUTION.
* Container CPC= ,Image_Id= ,
*   Id=
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REPLACE A SOLUTION ID ASSOCIATED WITH A
* CONTAINER TRG WITH A NEW OR CORRECTED SOLUTION ID.
* Update Solution_Id,
*   Old= ,
*   New=
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO REQUEST THAT SCRT IGNORE A CONTAINER
* PRICING SOLUTION FOR A SPECIFIED TIME INTERVAL
* Ignore Container,Start= ,Resume= ,
*   Id=
*
* FOLLOWING SPECIAL CONTROL STMT IS FOR OPTIONAL USE
* UNCOMMENT STMT TO SPECIFY SPECIAL SCRT FEATURES.
* Feature ...
*
//* FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS
//* UNCOMMENT THE DD TO GET A CSV FILE SHOWING CONTAINER

```

```

//* PRICING DETAILED DATA
//* //CPSTRACE DD      DISP=(,CATLG),DSN=HLQ.CPSTRACE.CSV,
//* //              SPACE=(TRK,(15,15)),UNIT=SYSDA
//*
//* FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS
//* UNCOMMENT THE DD TO GET A CSV FILE SHOWING REASONS THAT
//* GETTING STARTED PRICING MSUS WERE NOT USED IN THE REPORT
//* //GSTRACE DD      DISP=(,CATLG),DSN=HLQ.GSTRACE.CSV,
//* //              SPACE=(TRK,(15,15)),UNIT=SYSDA
//*
//* FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS
//* UNCOMMENT THE DD TO GET A CSV FILE SHOWING REASONS THAT
//* INTEGRATED WORKLOAD PRICING WASN'T USED IN THE REPORT
//* //IWPTRACE DD     DISP=(,CATLG),DSN=HLQ.IWPTRACE.CSV,
//* //              SPACE=(TRK,(15,15)),UNIT=SYSDA
//*
//* FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS
//* UNCOMMENT THE DD TO GENERATE A WORKLOAD CLASSIFICATION
//* SUPPLEMENT REPORT
//* //WCSUPP DD       DISP=(,CATLG),DSN=HLQ.WCSUPP.CSV,
//* //              SPACE=(TRK,(15,15)),UNIT=SYSDA
//*
//* FOLLOWING DD IS FOR OPTIONAL USE BY z/OS CUSTOMERS.
//* UNCOMMENT THE DD TO SPECIFY CSV FILE(S) PROVIDING INPUT FOR
//* EITHER MOBILE WORKLOAD PRICING OR IBM Z SYSTEMS COLLOCATED
//* APPLICATION PRICING.
//* //INPUTCSV DD *
//*
//N089      DD *          NON-89 PRODUCT SECTION
*
* FOR EACH PRODUCT LISTED, YOU MUST ENTER ONE OF THESE 3 OPTIONS:
*
* 1. ENTER *ALL IF THE PRODUCT RUNS IN ALL LPARS IN THE INPUT SMF
*    DATA FILE.
*    E.G. 5655-043=*ALL
* 2. ENTER *NONE IF THE PRODUCT RUNS IN NONE OF THE LPARS IN THE
*    INPUT SMF DATA FILE.
* 3. ENTER A LIST OF LPAR NAMES SEPARATED BY COMMAS OF THE LPARS IN
*    WHICH THE PRODUCT RUNS. (YOU MAY CONTINUE ONTO THE NEXT LINE BY
*    ENDING THE PRIOR LINE WITH A COMMA.)
*
* NOTE: IF YOU HAVE THE SAME LPAR NAME ON DIFFERENT CPCS, AND YOU
*       NEED TO DISTINGUISH BETWEEN THEM, THE FOLLOWING SYNTAX WILL
*       ALLOW IT.  FOR EXAMPLE: YOU HAVE 57XX-XXX RUNNING IN LPARA
*       ON CPC 2064-12345, BUT NOT IN LPARA ON 2064-98765. SPECIFY:
*       57XX-XXX=2064-12345:LPARA
*
* ----- Z/OS N089S -----
*
* DELETE THE FOLLOWING STMTS IF Z/OS IS *NOT* RUNNING ON THE CPC
*
*
* NETVIEW PERF MONITOR V2
5655-043=(required)
* TIVOLI NETVIEW FOR Z/OS V6
5697-NV6=(required)
* TIVOLI NETVIEW FOR Z/OS V5
5697-ENV=(required)
* TIVOLI NETVIEW FOR OS/390
5697-B82=(required)
* TIVOLI WORKLOAD SCHEDULER FOR Z/OS V9
5698-T08=(required)
* TIVOLI WORKLOAD SCHEDULER FOR Z/OS V8
5697-WSZ=(required)
* OPC V2
5697-OPC=(required)
* SYSTEM AUTOMATION FOR OS/390
5645-005=(required)
* SYSTEM AUTOMATION OS/390 V2
5645-006=(required)
* LOTUS DOMINO FOR S/390
5655-B86=(required)
* LOTUS ACTIVE INSIGHT FOR Z/OS V6
5655-U15=(required)
* COBOL FOR OS/390 & VM V2
5648-A25=(required)
* VA PL/I FOR OS/390 V2
5655-B22=(required)
* IBM ENTERPRISE PL/I FOR Z/OS AND OS/390 V3
5655-H31=(required)
* IBM ENTERPRISE PL/I FOR Z/OS V4
5655-W67=(required)

```

```

* IBM ENTERPRISE COBOL FOR Z/OS AND OS/390 V3
5655-G53=(required)
* IBM ENTERPRISE COBOL FOR Z/OS V4
5655-S71=(required)
* QMF MVS VERSION 3
5706-254=(required)
* FAULT ANALYZER FOR OS/390 V1
5697-F19=(required)
* FAULT ANALYZER FOR Z/OS AND OS/390 V2
5655-G74=(required)
* FAULT ANALYZER FOR Z/OS AND OS/390 V3
5655-J47=(required)
* FAULT ANALYZER FOR Z/OS V4
5655-L25=(required)
* FAULT ANALYZER FOR Z/OS V5
5655-M20=(required)
* FAULT ANALYZER FOR Z/OS V6
5655-P16=(required)
* FAULT ANALYZER FOR Z/OS V7
5655-R46=(required)
* FILE MANAGER FOR OS/390 V1
5697-F20=(required)
* FILE MANAGER FOR Z/OS AND OS/390 V2
5655-G75=(required)
* FILE MANAGER FOR Z/OS AND OS/390 V3
5655-J48=(required)
* FILE MANAGER FOR Z/OS V4
5655-L26=(required)
* FILE MANAGER FOR Z/OS V5
5655-M21=(required)
* FILE MANAGER FOR Z/OS V6
5655-P17=(required)
* FILE MANAGER FOR Z/OS V7
5655-R47=(required)
* FILE EXPORT FOR Z/OS V1
5697-I12=(required)
* APPLICATION MONITOR FOR Z/OS AND OS/390 V1
5697-H63=(required)
* APPLICATION MONITOR FOR Z/OS V2
5655-L22=(required)
* DEBUG TOOL UTILITIES AND ADV FUNCTIONS FOR Z/OS & OS/390 V3
5655-J18=(required)
* DEBUG TOOL UTILITIES AND ADV FUNCTIONS FOR Z/OS V4
5655-L23=(required)
* DEBUG TOOL UTILITIES AND ADV FUNCTIONS FOR Z/OS V5
5655-M19=(required)
* DEBUG TOOL UTILITIES AND ADV FUNCTIONS FOR Z/OS V6
5655-P15=(required)
* DEBUG TOOL UTILITIES AND ADV FUNCTIONS FOR Z/OS V7
5655-R45=(required)
* DEBUG TOOL UTILITIES AND ADV FUNCTIONS FOR Z/OS V8
5655-S16=(required)
* CLOUD 9 FOR SCLM FOR OS/390 V1
5697-G57=(required)
* CLOUD 9 FOR SCLM FOR Z/OS V2
5655-G93=(required)
* BREEZE FOR SCLM OS/390 V1
5697-G58=(required)
* IBM TIVOLI MONITORING FOR NETWORK PERFORMANCE V1
5698-A10=(required)
* IBM TIVOLI SYSTEM AUTOMATION FOR OS/390 V2
5698-A14=(required)
* IBM TIVOLI SYSTEM AUTOMATION FOR OS/390 V3
5698-SA3=(required)
* IBM TIVOLI SYSTEM AUTOMATION FOR Z/OS V4
5698-SA4=(required)
* IBM TIVOLI WORKLOAD SCHEDULER FOR HOST EDITION V8
5698-A17=(required)
* IBM TIVOLI MONITORING FOR NETWORK PERFORMANCE V2
5698-FNP=(required)
* WEBSPPHERE TRANSFORMATION EXTENDER FOR Z/OS V8
5655-R95=(required)
* IBM WEBSPPHERE PORTAL FOR Z/OS AND OS/390 V4
5655-K12=(required)
* IBM WEBSPPHERE HOST PUBLISHER FOR V4 FOR ZSERIES
5655-I03=(required)
* IBM WEBSPPHERE MQ INTEGRATOR BROKER FOR Z/OS V2
5655-I58=(required)
* IBM WEBSPPHERE MQ INTEGRATOR FOR Z/OS V2
5655-G97=(required)
* WEBSPPHERE STUDIO APPLICATION MONITOR V1
5655-J67=(required)

```

```

* WEBSPPHERE STUDIO ASSET ANALYZER V4
5655-M22=(required)
* WEBSPPHERE STUDIO ASSET ANALYZER V3
5655-L21=(required)
* WEBSPPHERE STUDIO ASSET ANALYZER V2
5655-I49=(required)
* WEBSPPHERE STUDIO WORKLOAD SIMULATOR V1
5655-I14=(required)
* WEBSPPHERE STUDIO APPLICATION MONITOR FOR Z/OS V2
5655-L42=(required)
* WEBSPPHERE STUDIO APPLICATION MONITOR FOR Z/OS V3
5697-J18=(required)
* IMS DATA COLLECTOR FOR WEBSPPHERE STUDIO APPL MONITOR V3
5697-J01=(required)
* CICS DATA COLLECTOR FOR WEBSPPHERE STUDIO APPL MONITOR V2
5655-L50=(required)
* CICS DATA COLLECTOR FOR WEBSPPHERE STUDIO APPL MONITOR V3
5697-J17=(required)
* WEBSPPHERE DASHBOARD FRAMEWORK FOR Z/OS V6
5655-U13=(required)
* APPLICATION WORKLOAD MODELER FOR Z/OS V1
5655-J62=(required)
* ENHANCED ACCESS CONTROL FOR SCLM V1
5697-H59=(required)
* WORKLOAD SIMULATOR FOR Z/OS AND OS/390 V1
5655-I39=(required)
* MERGE TOOL FOR Z/OS AND OS/390 V1
5697-H70=(required)
* MIGRATION UTILITY FOR Z/OS V1
5655-I18=(required)
* MIGRATION UTILITY FOR Z/OS V2
5697-I89=(required)
* SCLM SUITE ADMIN WORKBENCH FOR Z/OS V1
5655-L03=(required)
* SCLM DEVELOPER TOOLKIT V2
5655-R37=(required)
* SCLM DEVELOPER TOOLKIT V1
5655-M99=(required)
* APPLICATION PERF ANALYZER FOR Z/OS V1
5697-N37=(required)
* APPLICATION PERF ANALYZER FOR Z/OS V7
5697-N53=(required)
* APPLICATION PERF ANALYZER FOR Z/OS V8
5697-N63=(required)
* TIVOLI COMPOSITE APPLICATION MANAGER FOR CICS V6
5698-A69=(required)
* TIVOLI COMPOSITE APPLICATION MANAGER FOR IMS V6
5698-A70=(required)
* TIVOLI COMPOSITE APPLICATION MANAGER FOR APPL DIAG V7
5698-B48=(required)
* TIVOLI COMPOSITE APPLICATION MANAGER FOR WEBSPPHERE V6
5698-A71=(required)
* SCLM ADMIN TOOLKIT V2
5697-J19=(required)
* IBM MIGRATION UTILITY FOR Z/OS V3
5697-N44=(required)
* TIVOLI OMEGAMON XE FOR MESSAGING V7
5698-B23=(required)
* TIVOLI OMEGAMON XE FOR MESSAGING V6
5698-A87=(required)
* AUTOMATED TAPE ALLOC MANAGER V1
5697-H62=(required)
* DATA STUDIO PUREQUERY RUNTIME FOR Z/OS V1
5655-U93=(required)
* OPTIM PUREQUERY RUNTIME FOR Z/OS V3
5655-W92=(required)
* OPTIM PUREQUERY RUNTIME FOR Z/OS V2
5655-V80=(required)
* WEBSPPHERE ILOG BUSINESS RULES FOR Z/OS V7
5655-ILG=(required)
* DATA STUDIO PUREQUERY RUNTIME FOR Z/OS V2
5655-V57=(required)
* WEBSPPHERE ILOG JRULES FOR SYSTEM Z V7
5655-V82=(required)
* WEBSPPHERE ILOG RULE TEAM SERVER FOR SYSTEM Z V7
5655-V84=(required)
* WEBSPPHERE ILOG RULES COBOL FOR SYSTEM Z V7
5655-V87=(required)
* WEBSPPHERE ILOG DECISION VALIDATION SVCS FOR SYSTEM Z V7
5655-V89=(required)
* IBM CPLEX Optimizer for z/OS V12
5655-Y21=(required)

```

```

* MIGRATION UTILITY FOR Z/OS V4
5655-MGU=(required)
* IBM AUTOMATION CONTROL FOR Z/OS V1
5698-LSA=(required)
* IBM Tivoli Netview Monitoring for GDPS V6
5698-BMP=(required)
* IBM Standards Processing Engine for Healthcare Payer V8
5655-Q19=(required)
* IBM Standards Processing Engine for Supply Chain EDI V8
5655-Q22=(required)
* IBM Transformation Extender Advanced for z/OS V9
5655-Q35=(required)
* IBM Z Service Automation Suite V1
5698-SA1=(required)
* IBM Workload Scheduler distributed Agent for z/OS V9
5698-AAR=(required)
* IBM Tivoli Workload Scheduler dist. Agent for z/OS V8
5698-AA3=(required)
* IBM Migration Utility for z/OS V5
5698-MG5=(required)
* IBM Z JCL Expert
5698-020=(required)
* IBM Z Workload Scheduler 10
5698-T09=(required)
* (NEW 30.1.4) IBM Z Backup Resiliency 2
5698-BR3=(required)
*
* ----- Z/VSE N089S -----
*
*   DELETE THE FOLLOWING STMTS IF Z/VSE IS *NOT* RUNNING ON THE CPC
*
* DB2 SERVER FOR VSE & VM
5697-F42=(required)
* DITTO/ESA FOR VSE
5648-099=(required)
* DL/1 DOS/VS
5746-XX1=(required)
* HIGH LEVEL ASSEMBLER VSE & VM
5696-234=(required)
* IBM C FOR VSE/ESA V1
5686-A01=(required)
* IBM COBOL VSE/ESA
5686-068=(required)
* IBM DFSORT/VSE V3
5746-SM3=(required)
* IBM PL/1 FOR VSE/ESA
5686-069=(required)
* MQSERIES FOR VSE/ESA V2
5686-A06=(required)
* WEBSPHERE MQ FOR VSE/ESA V3
5655-U97=(required)
* TCP/IP FOR VSE
5686-A04=(required)
* TCP/IP FOR Z/VSE V2
5686-CS1=(required)
* IBM RATIONAL COBOL RUNTIME FOR Z/VSE V7
5648-F66=(required)
* IPv6/VSE V1
5686-BS1=(required)
*
//*
```

The following topics describe the JCL statements shown in [“Sample JCL to run the 64-bit SCRT on a z/OS system”](#) on page 25.

JOB statement

You must change the sample JOB statement in [“Sample JCL to run the 64-bit SCRT on a z/OS system”](#) on page 25 to meet the rules established for running batch jobs on the z/OS system on which you intend to run SCRT.

The sample JOB statement sets the **TIME** parameter to a value of ten minutes. However, you must customize this value to match the specific needs of your SCRT job, as this value depends on several factors:

- The CPC model on which the job runs

- The RMF and SMF recording intervals
- The number of LPARs and guest z/OS systems providing data to SCRT
- The number of products being reported

Because there is some additional overhead associated with the new SCRT execution environment, do not base the value of the **TIME** parameter solely on your past experience with older versions of SCRT.

Note that some data centers restrict the use of the **TIME** parameter or restrict the maximum value that you can specify. If your JCL is rejected when you use the **TIME** parameter, check your data center's rules for using the TIME parameter.

The value specified for the **REGION** parameter on the JOB statement is 0M, which allows SCRT to access all of the storage available to it when it runs. This value allows SCRT to automatically obtain any increased storage needs that may be required as SCRT supports new function in future releases. You must verify that your data center's JCL restrictions permit a value of 0M for the **REGION** parameter before making this change to the JCL for your SCRT job.

The value that you specify for the **REGION** parameter must consider the following factors:

- Maximum heap size
- Stack storage requirements
- Language Environment® (LE) and Java code
- LE internal control blocks

Here again, do not base the value of the **REGION** parameter solely on your past experience with older versions of SCRT.

DELETE EXEC statement

The DELETE job step in [“Sample JCL to run the 64-bit SCRT on a z/OS system”](#) on page 25 deletes the .csv report file created by the previous execution of SCRT whenever you rerun SCRT one or more times for the same month.

The DELETE job step also contains an optional (commented out) DD statement to allocate and delete the GSTRACE file which is optionally created by the SCRT job step. If your SCRT job step uses the GSTRACE DD statement to create a GSTRACE file, uncomment the GSTRACE DD statements in the preceding DELETE job step so that the GSTRACE file can be deleted whenever you rerun SCRT for the same month.

You should assign unique file names for the SCRT report file that is created for each month and for the GSTRACE file, if that file is also created each month. If you need to rerun SCRT for a particular month due to a problem with the SCRT control statements or a problem with the data that is being processed, this step deletes the prior version of the SCRT report file and GSTRACE file for that month so that the next step can allocate new SCRT report and GSTRACE files.

SCRT EXEC statement

You must modify the EXEC statement to point to an IBM JZOS procedure. A sample procedure (JVMPRCxx) is distributed with each IBM Java version in the mvstools/samples/jc1 directory.

If it has not already been done, copy the sample procedure to an appropriate PROCLIB. You can generally use the sample procedure as is, as the SCRT JCL contains any necessary customization. The only exception is if the JZOS batch launcher (that is, the JVMLDMxx module) is not in your LNKLIST. In that case, uncomment the STEPLIB statement and modify it to point to the correct PDSE.

You must use the procedure that corresponds to the Java or Semeru version that you are using.

For details about Java environment configuration and the JZOS batch launcher, see the JZOS Batch Launcher and Toolkit documentation, as described in [“z/OS software prerequisites for SCRT”](#) on page 21.

STDENV DD statement

The STDENV file is a required input to the JZOS batch launcher. It contains a shell script that is used to set the environment variables to configure the Java runtime environment.

The sample SCRT JCL contains a set of suggested variables and requires minimal customization. Specifically:

- Change the **APP_HOME** variable to point to the location of the SCRT . jar file in the z/OS UNIX file system.
- Verify that the location of the **JAVA_HOME** variable points to a valid location for the Java SDK on your system. The location specified in the SCRT JCL is a common install location.

For details about Java environment configuration and the JZOS batch launcher, see *JZOS Batch Launcher and Toolkit: Installation and User's Guide*.

Note: Disable sequence numbers in the SCRT JCL. The presence of sequence numbers can cause errors in the interpretation of the STDENV DD statement.

PARMS DD statement

The parameters supported by the PARMS DD statement fall into two categories:

1. PARMS parameters whose values apply to *all* of the CPCs processed by that execution of SCRT
2. PARMS parameters that must specify the CPC to which they apply, even when only a single CPC is processed by that execution of SCRT.

Because the syntax differs for each parameter category, each category is described separately.

PARMS parameters that apply to all CPCs

The parameters in this PARMS category conform to the following rules:

1. Each parameter value is applied to all of the CPCs processed by that execution of SCRT. This means that those same values will appear on every sub-capacity report, one per CPC, that is generated by that execution of SCRT. For instance, whenever the optional **Report_Period** parameter is used, it will control the reporting period for all of the CPCs processed by that execution of SCRT.
2. Each parameter consists of a keyword, immediately followed by an equal sign, immediately followed by the desired value. The value may consist of a string of one or more blank-delimited words, as appropriate. For example:

```
keyword=value string
```

3. Each parameter keyword must start in column one of the control statement.
4. Each parameter keyword and its value must fit on a single 71-character control statement.
5. Parameter specifications that exceed 71 characters will be truncated.
6. Parameter keywords can be in upper, lower, or mixed case.
7. Control statements beginning with an asterisk (*) are considered comments and are ignored by SCRT.

The PARMS DD statement must contain the following five parameters that apply to all CPCs:

Customer_Name
Customer_Number
Contact_Name
Contact_EMail
Contact_Phone

If you do not specify these parameters, SCRT issues a message and exits with return code 16.

You can specify the following PARMS keywords and their values in any order:

Customer_Name

(Required) The name of your company.

Example:

```
Customer_Name=Big National Bank
```

Customer_Number

(Required) The country or region code followed by the IBM customer number that applies to the CPC(s) on which this execution of SCRT is reporting.

The SCRT **Customer_Number** parameter supports six- and ten-digit IBM customer numbers. Older IBM customer numbers are six digits long, depending on the country or region. New IBM customer numbers are 10 digits long and IBM is phasing them in by region over time. (IBM will notify you by email with your new customer number and instructions for when and where to use the new customer number.)

- If you have a six-digit IBM customer number, provide your three-digit country or region code, immediately followed by your six-digit customer number, and only include leading zeros when needed to form a six-digit number (`Customer_Number=cccnnnnnn`). All country and region codes *except* for 897 support six-digit customer numbers.
- If you have a 10-digit IBM customer number, provide your two-character ISO country or region code, followed by a dash (or minus sign), followed by your ten-digit customer number (`Customer_Number=cc-xnnnnnnnnn`). SCRT accepts lower case characters and automatically converts them to upper case characters when reporting the value on the sub-capacity report.

You must use the same customer number for SCRT as that which you used to register for access to the License Management Support (LMS) application (to submit your sub-capacity reports to IBM). Note that LMS or other sources might refer to your customer number as your *site ID*.

SCRT verifies that the customer number meets the previously described format. If the format is not valid, SCRT issues an error message and terminates without generating any sub-capacity reports.

Examples: Table 5 on page 35 shows some examples of valid SCRT **Customer_Number** values.

Table 5. Examples of SCRT **Customer_Number** values

Number of digits in customer number	Country	Country code	Example of IBM customer number	Value to enter for SCRT Customer_Number parameter
6	France	706	004321	706004321
10	Germany	DE	S123456789	DE-S123456789

See Appendix C, “SCRT country code list,” on page 517 for a list of country codes. The two-digit ISO country codes will be added to each affected country or region as IBM phases in the new 10-digit customer numbers. If you do not know your customer number, contact your IBM representative.

Contact_Name

(Required) The name of the person who IBM should contact for billing-related issues.

The SCRT control statement for this parameter only supports single-byte character set data; however, after the sub-capacity report has been created, the contents of this field can be changed to use double-byte character set data for a select set of countries and regions. For details, see [“Spreadsheet method for viewing and editing the sub-capacity report”](#) on page 420 and [“LMS web interface”](#) on page 427.

Example:

```
Contact_Name=John Q Programmer
```

Contact_Email

(Required) The email address of the person who IBM should contact for billing-related issues.

Example:

```
Contact_EMail=JQProgrammer@bigbank.com
```

Contact_Phone

(Required) The telephone number of the person who IBM should contact for billing-related issues.

Example:

```
Contact_Phone=+44 99 7777 8888
```

PO_Number

(Optional) The purchase order number assigned by your company's procurement department.

You can omit the optional **PO_Number** parameter if you do not wish to specify it. The sub-capacity report will then show (optional) for that parameter value to indicate that you can also either provide a purchase order number by using the LMS application or by using a spreadsheet application, such as Lotus 1-2-3 or Microsoft Excel, as described in [“Spreadsheet method for viewing and editing the sub-capacity report”](#) on page 420.

Your company might order software through a purchase order, which represents a commitment from your procurement department that a software order is authorized and will be paid. If your company requires a purchase order, you must submit your sub-capacity reports with the purchase order number assigned by your procurement department. If you submitted software orders through a purchase order before you began using sub-capacity pricing, you must enter a purchase order number on your sub-capacity reports. Contact your procurement department for more information.

Example:

```
PO_Number=XYZ1234
```

Report_Period

(Optional) A specific reporting period to appear in the sub-capacity report, other than the reporting period that SCRT would select by default. The value for the reporting period can be:

yyyy/mm

A specific year and month, in the format *yyyy/mm*. Specify a date in this format to override the year and month that SCRT automatically sets.

Example: To specify that the month of July 2006 appear in the report, specify:

```
Report_Period=2006/07
```

Last_Month

Specifies that SCRT should use the last reporting period which ended before the day on which SCRT is run. For instance, if you run SCRT on any day from 2 May to 1 June, then specifying a value of Last_Month causes SCRT to create a report for the reporting period *2 April 2006 to 1 May 2006*. However, if you run SCRT on 1 May, then a value of Last_Month causes SCRT to create a report for the reporting period *2 March 2006 to 1 April 2006*.

The specified reporting period applies to *all* machines being processed by a given execution of SCRT. When SCRT sets the reporting period based on the **Report_Period** parameter, SCRT lists the value specified by the **Report_Period** parameter, followed by a space and (customer override) to indicate that you chose the reporting period. Both sets of information are listed on the **Reporting Period** row of the sub-capacity report.

When you omit the **Report_Period** parameter, SCRT calculates the report period separately for each CPC encountered during that execution of SCRT. This means that SCRT could generate a sub-capacity report with one reporting period for the first CPC and a sub-capacity report with a different reporting period for a second CPC and so on. When you omit the **Report_Period** parameter, SCRT uses the following process to calculate the reporting period for each CPC:

1. Reads the input file and calculates a data collection percentage for each reporting period for each CPC in the file
2. Selects the reporting period with the highest data collection percentage for each CPC
3. When multiple reporting periods for a CPC have an identical highest data collection percentage, then SCRT chooses the reporting period with the highest collection percentage that has the latest date for that CPC.

PARMS parameters that apply to specific CPCs

The parameters in this PARMS category are in the form of *control statements* and conform to the following rules:

1. Each parameter control statement is applied to a specific CPC, as specified by the appropriate subparameter. The subparameter that specifies the CPC value is required even when SCRT only processes data for a single CPC.
2. Each parameter control statement consists of a parameter keyword and one or more subparameter keywords and values. The general syntax is:

```
Parameter Subparameter=value,Subparameter=value,...
```

3. Each parameter keyword must start in column one of the control statement.
4. Subparameters can be continued onto additional statements (lines) subject to the following conditions:
 - a. A subparameter keyword and its value must be contained on the same statement line.
 - b. Whenever the second and following subparameters are continued onto a following line, the prior line must end with a comma after the last subparameter's value.

For example:

```
Parameter subparm1=aaaa,  
subparm2=bbbb
```

5. Parameter and subparameter specifications that exceed 71 characters on the same line will be truncated and might result in a syntax error.
6. Parameter and subparameter keywords can be in upper, lower, or mixed case.
7. Control statements beginning with an asterisk (*) are considered comments and are ignored by SCRT.

The following PARMS parameters apply to specific CPCs:

ASSIGN Customer_Number
LPAR_Comment

ASSIGN Customer_Number=nnnnnn, CPC=tttt-sssss

(Optional) Use the **ASSIGN Customer_Number** parameter control statement to override the default **Customer_Number** parameter assignment for a specific CPC. You must use this parameter control statement in cases where SCRT must process within a single execution multiple CPCs that are associated with different customer numbers.

Note: When Country Multiplex Pricing is specified, SCRT verifies that all assigned customer numbers have the same country code. If SCRT detects that this condition has not been met, it issues message SCRTTOOL221 and terminates processing.

ASSIGN

(Required) Identifies the beginning of an **ASSIGN** control statement. The **ASSIGN** keyword must start in column one.

Customer_Number

(Required) The customer number that is to be assigned to the specified CPC. The specification of the **Customer_Number** here is the same as for “[Customer_Number](#)” on page 35.

CPC

(Required) The CPC machine type (*tttt*) and serial number (*sssss*) to which this customer number applies, in the format *tttt-sssss* (for instance, 2094-12345). This subparameter does not support special values (such as *ALL) nor a list of values. If SCRT does not find the specified CPC on any of the SMF or SCRT89 records being processed by this execution of SCRT, it issues an error message and terminates processing with return code 16.

LPAR_Comment

(Optional) Use the **LPAR_Comment** parameter control statement to explain why no SMF or SCRT89 records were provided for an LPAR.

You must provide comments in the Missing LPAR Data Details (H4) section of the sub-capacity report for every LPAR for which no SMF or SCRT89 data was provided to SCRT. Use the **LPAR_Comment** control statement for those explanations that apply to the specified LPAR for the current reporting period and all future reporting periods. (For example, when an LPAR is only used by Linux guests.) Using the **LPAR_Comment** control statement eliminates the need to update the Missing LPAR Data Details section of your sub-capacity report to add the same comment for the same LPAR each time you run SCRT.

Note: If SCRT processes an **LPAR_Comment** control statement for an LPAR that provided data at any point during the reporting period, SCRT issues the SCRTTOOL158 message and terminates processing without generating a sub-capacity report.

Whenever you use an **Exclude** control statement to exclude *all* of the data (both within and outside of the reporting period) from an LPAR, SCRT places the comment text that you supply on the **LPAR_Comment** statement in both the Detail Data Collection section and the **Customer Comments** column in the Special Conditions section for the **Exclude** control statement that excludes the data for that LPAR, but only if that **Exclude** control statement excludes the data from that one specific LPAR and not from any other LPARs or guest systems.

Following the instructions provided with the Solution Edition offering, use a separate **Exclude** control statement to exclude each LPAR being used for a Solution Edition offering, and set the **Comment** parameter on the **LPAR_Comment** control statement for the LPAR to indicate the title of the Solution Edition offering being used on that LPAR. When a Solution Edition offering is being used on multiple LPARs, use a separate **LPAR_Comment** control statement for each LPAR. An example of the contents of a Solution Edition comment would be:

```
Solution Edition for SAP Applications
```

Guideline: Use the **LPAR_Comment** control statement to indicate the title of the Solution Edition offering being used in the LPAR that is reporting a value of zero in its **Report Period % Data** column for these reasons:

1. The **Exclude** control statement will exclude all of the data for that Solution Edition LPAR, which will result in SCRT reporting a value of zero in the **Report Period % Data** column, *and*
2. The same LPAR comment will be used every month while the Solution Edition offering is being used on that LPAR, *and*
3. The same LPAR comment will also be listed in the Special Conditions section as the required customer comment for the corresponding **Exclude** control statement, as long as that **Exclude** statement only excluded that specific LPAR.

Furthermore, if you use the LMS web interface to add customer comments to sub-capacity reports and you are in a country or region where LMS supports DBCS data in comment fields, use the **LPAR_Comment** control statement to provide this comment because it must be entered using a single-byte character set in US English, which may be difficult to enter when using LMS in DBCS mode.

The general syntax for the **LPAR_Comment** control statement is:

```
LPAR_Comment CPC=tttt-sssss,LPAR=nnnnnnnn,Comment="ccc...ccc"
```

or, for comments that must span multiple lines:

```
LPAR_Comment CPC=tttt-sssss,LPAR=nnnnnnnn,Comment="ccc...ccc",
"ccc...ccc",
.
"ccc...ccc"
```

where the keywords and values are:

LPAR_Comment

(Required) Identifies the beginning of an **LPAR_Comment** control statement. The **LPAR_Comment** keyword must start in column one. Each **LPAR_Comment** control statement supports a single LPAR. Whenever comments are to be provided for multiple LPARs, each LPAR must be specified on a separate **LPAR_Comment** control statement.

CPC

(Required) The CPC machine type (*tttt*) and serial number (*sssss*) to which this **LPAR_Comment** control statement applies, in the format *tttt-sssss* (for instance, 2094-12345). This subparameter does not support special values (such as *ALL) nor a list of values. If SCRT does not find the specified CPC on any of the SMF or SCRT89 records being processed by that execution of SCRT, it issues an error message and terminates processing with return code 16.

LPAR

(Required) The LPARs to which this **LPAR_Comment** control statement applies. You can specify the value in the following ways:

LPAR=nnnnnnnn

A single LPAR name.

LPAR=(nnnnnnnn,nnnnnnnn,...)

A list of LPAR names. Enclose the list in parentheses and separate each name with a comma.

LPAR=*ALL

All LPARs on the specified CPC.

If SCRT does not find the specified LPAR and CPC combination on any of the SMF or SCRT89 records being processed by that execution of SCRT, it issues an error message and terminates processing with return code 16.

Comment

(Required) The comment text ("*ccc...ccc*") to display on the sub-capacity report for the specified LPAR. This must be the last subparameter on the **LPAR_Comment** control statement, as it is the only parameter whose value can span multiple lines. The comment text must be enclosed within double quotation marks and can contain blanks but should not contain commas nor any special characters that cannot be correctly translated from EBCDIC to ASCII encoding. The comment text is limited to a maximum of 255 characters and can be continued on multiple lines, as long as each continuation line is preceded by a line that ends with a comma. The last line of the comment text should contain no comma.

The SCRT control statement for this parameter only supports SBCS data; however, after the sub-capacity report has been created, the contents of this field can be changed to use DBCS data for a select set of countries and regions. For details, see [“Spreadsheet method for viewing and editing the sub-capacity report” on page 420](#) and [“LMS web interface” on page 427](#).

Note that characters beyond column 71 on each statement are ignored and might result in a syntax error.

Tip: IBM recommends placing any **LPAR_Comment** parameter control statements at the end of the PARMS DD statements because the **LPAR_Comment** control statement contains multiple subparameters that can be continued on multiple lines.

Example: The following **LPAR_Comment** control statements are equivalent.

SPECIAL DD statement

```
LPAR_Comment CPC=2096-12345,LPAR=LPARVM1,  
Comment="This LPAR is only used for Linux"  
  
LPAR_Comment  
CPC=2096-12345,  
LPAR=LPARVM1,Comment="This LPAR",  
"is only used",  
"for Linux"
```

SPECIAL DD statement

You can use the optional SPECIAL DD statement to specify one or more of the SCRT control statements listed in [Table 6 on page 40](#).

Table 6. SCRT SPECIAL control statements

Control statement name	Usage
ASSIGN ZWPC	Specify the WLM reporting attribute to use for IBM Z Workload Pricing for Cloud.
CONTAINER	Associate an LPAR with a Tailored Fit Pricing solution ID.
CONTAINER_INTERVAL_CONSUMPTION	Request the generation of sub-hour, per-container MSU consumption data.
COUNTRY_MULTIPLEX_PRICING	Request the use of Country Multiplex Pricing.
DETAIL_INTERVAL_RATE_DATA	Request the generation of detailed interval rate data.
DIAGMSG	Direct SCRT to provide additional diagnostic messages about WLM classification-based adjustments for mobile or public cloud workloads.
ENTERPRISE_TAILORED_FIT_PRICING	Request that SCRT generate an Enterprise Tailored Fit report.
EXCLUDE EXCLUDE MSU_CONSUMPTION	Handle unusual reporting situations where you need to exclude SMF or SCRT89 records from SCRT processing.
GENERATE_DETAILED_DATA	Request the generation of certain sub-capacity report sections that provide additional details about the applied pricing options and specific MSU reductions.
IGNORE CONTAINER	Ignore a Tailored Fit Pricing solution for a specified time interval.
IGNORE_WLM_MWP_DATA	Request that the Mobile Workload Pricing algorithm not be applied to mobile workload data from WLM.
INCLUDE	Select specific images to be included in the report and have all other images fully excluded.
INTEGRATED_WORKLOAD_PRICING	Request the use of the Integrated Workload Pricing option.
INTERVAL_RATE_DATA	Request the generation of interval rate data.
IPLA_ALL	Direct SCRT to report an IPLA family line for MVM eligible products in the Product Summary Information (E5) section even when only one product in the product family is in use.
ISV_SHOW_CONTAINERS	Requests container processing for independent software vendor (ISV) product reports. By default, SCRT does not process containers in ISV mode.
MLC_ALL	Direct SCRT to report an MLC family line for MVM eligible products in the Product Summary Information (E5) section even when only one product in the product family is in use.
REPORT_TIME	Direct SCRT to report dates and times from z/OS systems on the sub-capacity report using a special form of local time or using UTC time.
SUPPRESS_REPORT_SECTION	Suppress certain report conditions related to Getting Started Sub-Capacity Pricing.
UPDATE CONTAINER	Associate an LMS-provided solution name with a Tailored Fit Pricing solution ID and, optionally, designate the NO89 products that are running in the container.

Table 6. SCRT SPECIAL control statements (continued)

Control statement name	Usage
UPDATE SOLUTION_ID	Replace a solution ID associated with a container tenant resource group with a new or corrected solution ID.
z/TPF_TRANSFORMATION_ENGINES_OFFERING	Request that SCRT report on the z/TPF Transformation Engines Offering.

All of the keywords on the control statements in the SPECIAL DD can be in upper, lower, or mixed case. Each control statement must start in column one with a keyword indicating the control statement name. Each control statement can be up to 71 bytes long. Control statements that are longer than 71 characters will result in syntax errors. If any errors are found on SPECIAL control statements, SCRT issues a message and terminates processing with return code 16 without generating reports for any CPCs.

Any statement beginning with an asterisk (*) in column one is treated as a comment and is ignored by SCRT.

Details about each SPECIAL control statement follow.

ASSIGN ZWPC control statement

Use the **ASSIGN ZWPC** control statement to direct SCRT to use the IBM Z Workload Pricing for Cloud (zWPC) offering and to use SMF type 70 records as a source of eligible public cloud transaction CPU time.

The syntax for the **ASSIGN ZWPC** control statement is:

```
ASSIGN ZWPC=classification
```

In order to request WLM-based zWPC, you must use the **ASSIGN ZWPC** control statement and specify the appropriate WLM reporting attribute (*classification*) that you defined in the WLM administrative application. You can specify one of the following values for *classification*:

CATEGORYA

Specifies that you are using the CATEGORYA classification in WLM to identify eligible public cloud workloads.

Example: ASSIGN ZWPC=CATEGORYA

CATEGORYB

Specifies that you are using the CATEGORYB classification in WLM to identify eligible public cloud workloads.

Example: Assign ZWPC=CATEGORYB

For more information, see [Chapter 12, “Sub-capacity reporting for IBM Z Workload Pricing for Cloud \(zWPC\),” on page 251.](#)

CONTAINER control statement

In cases where an LPAR is dedicated to a qualified Tailored Fit Pricing solution, use the **CONTAINER** control statement (or the appropriate SCRT GUI panel) to associate the LPAR with a container solution ID.

Whenever a valid **CONTAINER** control statement is encountered, SCRT removes the specified LPARs from the traditional sub-capacity environment and aggregation, and associates them with the Tailored Fit Pricing solution indicated by the provided solution ID.

IBM authorizes the use of the **CONTAINER** control statement only in cases where you have signed or otherwise agreed to the terms and conditions of a Tailored Fit Pricing solution. Any LPARs that are part of a Tailored Fit Pricing solution must be dedicated to the solution or workload indicated as part of the agreement. The solution ID that you specify on the **CONTAINER** statement must match the solution ID provided on the IBM License Management Support (LMS) website for the qualified workload.

Each **CONTAINER** control statement starts with the **CONTAINER** keyword. All subsequent parameters are delimited by commas. At least one parameter must follow the **CONTAINER** keyword. Remaining parameters may be specified on subsequent continuation lines. Each line that is being continued must end with a comma.

The **CONTAINER** control statements are processed in the order in which they are specified. Later control statements override earlier control statements. For instance, you can assign all images to a single solution ID using the **IMAGE_ID=*ALL** format, then override the solution ID assignment for a subset of images by specifying additional **CONTAINER** control statements.

The general syntax for the **CONTAINER** control statement is as follows:

```
CONTAINER CPC=tttt-sssss,IMAGE_ID=(image1,image2,...,imageN),ID=solutionID
```

The keywords and values are:

CONTAINER

(Required) Identifies the beginning of the control statement.

CPC=tttt-sssss

CPC=*ALL

(Required) Specifies the CPC machine type (*tttt*) and serial number (*sssss*) to which this control statement applies, in the format *tttt-sssss* (for instance, 2094-12345). This parameter does not support a list of values. If SCRT does not find the specified value in any of the SMF records being processed, it issues an error message and terminates processing with return code 16.

You can specify **CPC=*ALL** to assign the specified solution ID to multiple images in a single **CONTAINER** control statement. The solution ID is assigned to all images on all CPCs that match the **IMAGE_ID** specification, or to all images if the **IMAGE_ID** keyword is not specified.

There are conditions under which the CPUs on a CPC might be assigned a serial number that differs from the serial number assigned to the CPC. For native z/OS systems, this can occur due to a hardware CPU serial number assignment by special request to IBM. For guest systems, you can assign a unique CPU serial number to the virtual machine on which those guest systems run. Whenever the CPU serial numbers differ from the CPC serial number, you must use the appropriate serial number for the **CONTAINER** control statement. Otherwise, SCRT either will not include the intended data or will issue an error message indicating that the CPC serial number is not found in the SMF data being processed.

ID=solutionID

(Required) Specifies the IBM-provided 64-character solution ID (*solutionID*) provided to you through the IBM LMS website.

Only a single solution ID can be specified on the **CONTAINER** command. This should be the solution ID associated with the Tailored Fit Pricing solution running on the named LPARs.

IMAGE_ID=image_id

IMAGE_ID=(image_id1,...,image_idn)

IMAGE_ID=*ALL

(Required) Specifies the operating system image ID, in one of the following formats:

image_id

A single image ID.

(image_id1,...,image_idn)

A list of image IDs. The list can be split across multiple lines of the control statement.

***ALL**

Assigns the specified solution ID to multiple images in a single **CONTAINER** control statement. The solution ID is assigned to all images on the specified CPC or, if **CPC=*ALL** is specified, to all images on all CPCs.

- For native z/OS systems, use the LPAR name as the image ID.
- For guest z/OS systems, use the **SID** value from the SMFPRMxx member as the image ID.

If SCRT does not find the specified value or list of values in any of the SMF records being processed, it issues an error message and terminates processing for this CPC with return code 16.

CONTAINER_INTERVAL_CONSUMPTION control statement

Use the **CONTAINER_INTERVAL_CONSUMPTION** control statement to see MSU consumption on a per-container basis with 5-minute granularity. This control statement is optional and should be used only at the direction and authorization of IBM. An SMF interval length of 5 minutes is required for all systems that are part of the container.

The per-container MSU consumption data appears in the Container Interval Consumption section (I7) of the SCRT report for Tailored Fit Pricing. For more information, see [Chapter 19, “Reporting for Tailored Fit Pricing for IBM Z,”](#) on page 375.

COUNTRY_MULTIPLEX_PRICING control statement

Use the **COUNTRY_MULTIPLEX_PRICING** control statement to request the use of the Country Multiplex Pricing (CMP) option.

In order to request CMP, you must use the **COUNTRY_MULTIPLEX_PRICING** control statement to specifically direct SCRT to use the CMP algorithm when generating sub-capacity reports. The default SCRT processing does not use the CMP algorithm.

Note: SPECIAL DD control statements must immediately follow the SPECIAL DD line. JCL does not pass control statements to SCRT if they follow JCL comment lines (lines that start with `/` or `*`), and, without the **COUNTRY_MULTIPLEX_PRICING** control statement in effect, SCRT will generate standard sub-capacity reports instead of multiplex reports.

Note: IBM only authorizes the use of the **COUNTRY_MULTIPLEX_PRICING** control statement by customers who have signed or otherwise agreed to a CMP contract. IBM will reject all sub-capacity reports with CMP whenever a CMP contract has not been signed or otherwise agreed to. For more information about CMP and its requirements, see [Chapter 16, “Sub-capacity reporting for Country Multiplex Pricing \(CMP\),”](#) on page 343.

DETAIL_INTERVAL_RATE_DATA control statement

Use the **DETAIL_INTERVAL_RATE_DATA** control statement to request the generation of detailed interval rate data. The detailed interval rate data appears in the Detailed Interval Data section (V9) of the SCRT report for Hardware Consumption. Section V9 appears per CPC and provides finer granularity for interval rate data by showing a (typically) 15-minute view of the CPC MSU rate. For more information, see [Chapter 21, “Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution,”](#) on page 405.

DIAGMSG control statement

(Optional) Use the **DIAGMSG** control statement to direct SCRT to provide additional diagnostic messaging about the application of WLM classification-based adjustments for mobile or public cloud workloads. The additional diagnostic messages can help you determine why an expected adjustment does not appear on the sub-capacity report.

You can specify one of the following values for *value*:

WLMMOB

Requests additional diagnostic messaging for WLM classification-based mobile adjustments. For details, see [“DIAGMSG=WLMMOB control statement”](#) on page 206.

Example: DIAGMSG=WLMMOB

WLMZWPC

Requests additional diagnostic messaging for WLM classification-based zWPC adjustments. For details, see [“DIAGMSG=WLMZWPC control statement”](#) on page 259.

Example: DIAGMSG=WLMZWPC

ENTERPRISE_TAILORED_FIT_PRICING control statement

Use the **ENTERPRISE_TAILORED_FIT_PRICING** control statement to generate an Enterprise Tailored Fit report. Enterprise Tailored Fit reports are for use by customers with a Tailored Fit Pricing Enterprise Solution agreement or contract. For more information, see [Chapter 20, “Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution,”](#) on page 401.

EXCLUDE control statement

The **EXCLUDE** special control statement is for use under very rare conditions that require you to direct SCRT to exclude the processing of certain data that is reported on SMF and SCRT89 records. Several special conditions, such as products running in loops and testing Capacity BackUp, previously required you to use the IFASMFDP SMF dump utility to exclude certain records from SCRT processing.

Depending on conditions, the SCRT support team might also instruct you to use **EXCLUDE** control statements to resolve specific problems with the SMF and SCRT89 records that are being processed by SCRT. For a detailed explanation about how to use **EXCLUDE** control statements to handle the cases for products running in loops and Capacity BackUp testing, see [“Common SCRT problem scenarios”](#) on page 442.

You must use **EXCLUDE** control statements, instead of the SMF dump utility, to exclude SMF and SCRT89 data from within the SCRT reporting period. Note that the SMF dump utility is still required during the normal process of moving SMF records from the active SMF data sets to sequential files for processing by post processor programs, such as RMF and SCRT.

You can use **EXCLUDE** control statements to direct SCRT to exclude records in any combination of the following ways:

- Exclude all of the SMF and SCRT89 records from a specified CPC between a set of starting and ending dates and times within that reporting period.
- Exclude only an hour’s worth of the data for a specific product running on a specific operating system image (LPAR) running on a specific CPC.

Depending on the exact circumstances, several different **EXCLUDE** control statements might be needed when multiple time ranges are involved or when different combinations of products and operating system images are involved.

Whenever valid **EXCLUDE** control statements are encountered, SCRT lists each statement in the optional Special Conditions section of the sub-capacity report for the CPC specified on that **EXCLUDE** statement. The Special Conditions section also reports on the net effect of each of the **EXCLUDE** statements in terms of the actual start and end time that was excluded and the number of hours from SMF or SCRT89 records that were excluded.

IBM requires you to provide an explanation in the appropriate comments column for each **EXCLUDE** statement that appears in the Special Conditions section of the sub-capacity report. You can use the **LPAR_Comment** control statement to provide the required comment for the **EXCLUDE** control statement when all of the requirements are met. See the description of the **LPAR_Comment** control statement in [“PARMS parameters that apply to specific CPCs”](#) on page 37 for the conditions that must be satisfied in order for the **LPAR_Comment** statement to provide this support.

The LMS application prompts you for the required explanations, if the comments are not provided prior to LMS processing.

Per the instructions provided as a part of a Solution Edition offering, all of the data from each Solution Edition LPAR must be excluded from the normal sub-capacity report which you are submitting to IBM. Use the **EXCLUDE** control statement to exclude that data. The value of the **RESUME** parameter on the **EXCLUDE** statement can be set to the future date on which that Solution Edition offering ends. When that is done, the same **EXCLUDE** control statement can be used, unchanged for that LPAR, for subsequent months' sub-capacity report.

If you use separate **EXCLUDE** statements to exclude all the SMF data from each Solution Edition LPAR, SCRT copies the comment provided on the corresponding **LPAR_Comment** control statements to the **Customer Comment** fields for those **EXCLUDE** statements in the Special Conditions section. (See the description of **LPAR_Comment** in “PARMS parameters that apply to specific CPCs” on page 37 for more information.) Note that when you provide the comment by manually updating the .csv file or by entering the comment through the LMS web application, the comment must indicate the title of the Solution Edition offering being used on the LPAR, such as:

Solution Edition for SAP Applications

Note: One or more **EXCLUDE** control statements might result in the exclusion of all data for all LPARs on a CPC. In this case, SCRT generates a report for the fully excluded CPC that reflects no usage during the reporting period. For multiplex reports, the fully excluded CPC is included in the report, but reflects no usage.

Each **EXCLUDE** control statement starts with the **EXCLUDE** keyword in column 1. All of the remaining **EXCLUDE** parameters are delimited by commas. At least one parameter must be specified on the control statement after the **EXCLUDE** keyword. The remaining parameters must be specified on subsequent continuation lines of the control statement. Each line that is being continued must end in a comma. The **IMAGE_ID** and **PRODUCT_ID** parameters can be specified as lists whose values can be continued onto subsequent lines.

The general syntax for the **EXCLUDE** control statement is:

```
EXCLUDE CPC=tttt-sssss,IMAGE_ID=image_id,PRODUCT_ID=xxxx-xxx,
      Start=yyyy/mm/dd/hh,Resume=yyyy/mm/dd/hh,
      Comment=comment-string
```

The keywords and values are:

EXCLUDE

(Required) Identifies the beginning of an **EXCLUDE** control statement. The **EXCLUDE** keyword must start in column one.

CPC

(Required) The CPC machine type (tttt) and CPC serial number (sssss) to which this **EXCLUDE** control statement applies, in the format tttt-sssss (for instance, 2094-12345). This parameter does not support the special value *ALL nor a list of values. If SCRT does not find the specified value on any of the SMF or SCRT89 records being processed by this execution of SCRT, it issues an error message and terminates processing with return code 16.

Note that there are conditions under which the CPUs on a CPC might be assigned a serial number that is different than the serial number assigned to the CPC. For native z/OS, z/TPF and z/VSE systems, this can occur due to a hardware CPU serial number assignment by special request to IBM. For z/OS, z/TPF and z/VSE guest systems, you can assign a unique CPU serial number to the virtual machine on which those guest systems run. Whenever the CPU serial numbers are different than the serial number for the CPC, you must use the appropriate serial number for the **EXCLUDE** control statement. Otherwise, SCRT will either not exclude the intended data or will issue an error message indicating that the CPC serial number is not in the data being processed by that execution of SCRT.

- To exclude all of the systems running on the CPC for a specified date range, set the **CPC** value to the last five digits of the value reported by SCRT in the **Machine Serial Number** field at the beginning of the sub-capacity report.
- To exclude specific z/OS systems, set the **CPC** value for the z/OS systems to be excluded to the last five digits of the value reported by SCRT in the **Machine Serial Number** field at the beginning of the

sub-capacity report, whenever that serial number is not followed by another value in parentheses. If the value in the **Machine Serial Number** field is followed by another value in parentheses, set the **CPC** value to 0 plus the last four digits of the value in the parentheses. However, if SCRT does not accept that CPC value for a z/OS guest system, you must use the last five digits of serial number of the virtual machine in which the z/OS guest system was running. The virtual machine serial number is reported as the CPU serial number on the RMF reports for that z/OS system.

- For all z/TPF and z/VSE systems (native systems and guest systems), always set the **CPC** value to the last five digits of the value reported by SCRT in the **Machine Serial Number** field at the beginning of the sub-capacity report.

IMAGE_ID

(Required) The operating system image ID, in any of the following formats:

***ALL**

Indicates that the statement applies to all operating system images in the SMF or SCRT89 records being processed by this execution of SCRT for all images on the specified CPC. Note that using the *ALL value for this parameter will generate a data collection incident in the Missing CPC Data Details (H6) section of the sub-capacity report.

image_id**(image_id1,...,image_idn)**

A single image ID or a list of image IDs, as follows:

- For the native z/OS, z/TPF, and z/VSE systems, use the LPAR name.
- For z/OS V1.7 (and later) guest systems, use the SID value from the SMFPRMxx member.
- For z/OS V1.6 (and earlier) guest systems and for later z/OS V1.7 systems being reported under z/OS V1.6 guest rules, use the value assigned to the LPAR name by the matching VMGUESTS DD control statement.
- For z/TPF and z/VSE guest systems, use the VM user ID assigned to the virtual machine in which the guest system runs.

The image ID values within the (*image_id1*, . . . , *image_idn*) list format can be split across multiple lines of the control statement.

If SCRT does not find the specified value or list of values on any of the SMF or SCRT89 records being processed by this execution of SCRT, it issues an error message and terminates processing for this CPC with return code 16.

PRODUCT_ID

(Optional) A single product ID in the format *xxxx-xxx* or a list of product IDs in the format (*xxxx-xxx*, . . . , *xxxx-xxx*) that are to be excluded from SCRT processing for the specified **CPC** and **IMAGE_ID** values between the specified **START** and **RESUME** date/time values. Any product IDs that are not specified as values on the **PRODUCT_ID** keyword are not excluded from processing.

When the product ID is for an IWP defining program, SCRT excludes the processing associated with reporting that product's sub-capacity value for the hours being excluded. However, SCRT continues to use the busy times associated with that product to determine the adjustment to the IWP adjusted programs. See Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing (IWP),” on page 183 for more information.

If the **EXCLUDE** control statement is to apply to *all* products (including NO89 products and operating system products, such as z/OS, z/TPF, z/TPFDF and z/VSE), omit the optional **PRODUCT_ID** keyword from the **EXCLUDE** control statement. For more information, see [Example 1 — Excluding records for all products by measurement interval](#).

Each *xxxx-xxx* value represents the product ID for a product supported by SCRT, subject to the following restrictions:

- Exclude processing does not accept the product IDs for the operating system products (that is, for z/OS, z/TPF, z/TPFDF and z/VSE themselves) and does not accept the product IDs for NO89 products.

- If SCRT finds any of these restricted product IDs or does not find a specified product ID value on any of the SMF or SCRT89 records being processed in the current execution of SCRT, it issues an error message and terminates processing for this CPC with return code 16.

The product ID values within the $(xxx-xxx, \dots, xxx-xxx)$ list format can be split across multiple lines of the control statement.

Note: When ISV reporting is active, the **EXCLUDE** statement supports the specification of ISV products.

START

(Required) The start date and time, in local SMF time or UTC time, for the measurement interval for the SMF and SCRT89 records that are to be excluded from SCRT processing.

There are two formats for this value:

- The format when specifying the date/time values in local time is:

```
yyyy/mm/dd/hh
```

- The format when specifying the date/time values in UTC time is:

```
yyyy/mm/dd/hh/UTC
```

where:

yyyy

The year. Valid values are from 2000 to 2025, inclusive. (For instance, 2007.)

mm

The month. Valid values are from 01 to 12, inclusive. (For instance, 08 for August.)

dd

The day of the month. Valid values are from 01 to 31, inclusive, depending on the month. (For instance, 01 for the first day of the month.)

hh

The hour. Valid values are from 00 to 23, inclusive. (For instance, 00 starts the day at midnight.)

UTC

Indicates that the date/time value is in UTC time.

The following rules summarize how to determine whether you must use the local format or UTC format to specify the date/time values for the **START** and **RESUME** parameters on **EXCLUDE** control statements:

1. UTC date/time formats are acceptable for all date/time conditions *except* when using the **REPORT_TIME** control statement with the **Special_Local_Time** option. Note that UTC date/time formats are required for a specific set of date/time conditions.
2. Local date/time formats are acceptable for all date/time conditions *except*:
 - Whenever the data for the excluded date/time range contains a change in time zones (change in the UTC offset), **or**
 - Whenever the data for the excluded date/time range contains multiple image IDs and the image IDs are in different time zones (different UTC offsets), **or**
 - Whenever you use the **REPORT_TIME** control statement with the **UTC_Time** option

When SCRT requires the use of UTC time:

The only way for SCRT to address time zone ambiguities is by requiring that you specify the date/time values for the **START** and **RESUME** parameters as UTC date/time values, using the UTC format $(yyyy/mm/dd/hh/UTC)$. The following conditions require the use of UTC times:

- Whenever the excluded date/time range includes one or more hours where a time zone change has occurred. For instance, excluding the hours 1:00 to 2:00 for the day when the transition between Standard Time and Daylight Savings Time occurs requires that the **START** and **RESUME**

date/time values be specified as UTC date/time values in UTC format (*yyyy/mm/dd/hh/UTC*) because the time zone change typically occurs at 2:00. However, exclusions for date/time ranges that start and end before 1:00 on that day can use local date/time values in local date/time format (*yyyy/mm/dd/hh*). Likewise, exclusions for date/time ranges that start and end after 3:00 on that day can also use local date/time values in local date/time format.

- Whenever the same **EXCLUDE** control statement applies to multiple image IDs and the image IDs do not have the same time zones (UTC offset values) for the date/time range being excluded. The multiple image IDs can be due to a list of image IDs or the use of the *ALL value. You have two options to handle this condition:
 1. You can split the single **EXCLUDE** control statement into separate **EXCLUDE** control statement for each image ID. When you do this, the date/time values for the **START** and **RESUME** parameters on each of the **EXCLUDE** control statements can remain as local date/time values in local date/time format, unless there is time zone change (change in UTC offset) for the date/time range being excluded for that image ID. In that case, SCRT requires that you use UTC date/time values in UTC date/time format.
 2. You can convert the local date/time values for the **START** and Resume parameters on the single **EXCLUDE** control statement into their corresponding UTC date/time values and specify them in the UTC date/time format on the control statement.

Tip: In order to specify the date/time values as UTC times, you must convert the local date/time values to the actual UTC date/time values by correctly applying the UTC offsets in the SMF records to the local times in those same records, as follows:

- Locations to the *east* of Greenwich, England, must *subtract* the absolute value of the UTC offset from the local time to obtain the UTC time.
- Locations to the *west* of Greenwich must *add* the absolute value of the UTC offset to the local time to obtain the UTC time.

Note that the date can change when converting a local time to a UTC time. You might also need to contact your operations manager or system programmer to determine which UTC offsets were used on which dates for each image ID being excluded.

Requirement: SCRT requires that the date/time values for the **START** and **RESUME** parameters on the same **EXCLUDE** control statement be in the same format—either both in local format (*yyyy/mm/dd/hh*) or both in UTC format (*yyyy/mm/dd/hh/UTC*). This requirement applies on an individual **EXCLUDE** control statement basis. That means that SCRT will accept one **EXCLUDE** control statement using local date/time formats and a different **EXCLUDE** control statement using UTC date/time formats in the same execution of SCRT (as long as there are no date/time ambiguities for using the local date/time format in the first place).

RESUME

(*Required*) The date and time, in local SMF time or UTC time, for the measurement interval for which SCRT is to resume processing SMF and SCRT89 records. There are two formats for this value.

The format when specifying the date/time values in local time is:

```
yyyy/mm/dd/hh
```

The format when specifying the date/time values in UTC time is:

```
yyyy/mm/dd/hh/UTC
```

where:

yyyy

The year. Valid values are from 2000 to 2025, inclusive. (For instance, 2007.)

mm

The month. Valid values are from 01 to 12, inclusive. (For instance, 08 for August.)

dd

The day of the month. Valid values are from 01 to 31, inclusive, depending on the month. (For instance, 01 for the first day.)

hh

The hour. Valid values are from 00 to 23, inclusive. (For instance, 00 starts the day at midnight.)

UTC

Indicates that the date/time value is in UTC time.

The date/time value that you specify for the **RESUME** parameter must be in the *same* local format or UTC format as the date/time value for the matching **START** parameter. Otherwise, SCRT issues an error message and terminates processing for this CPC with return code 16. See the description of the **START** parameter for a discussion about local and UTC times and formats.

COMMENT

(*Optional*) A comment that justifies the use of the **EXCLUDE** control statement. The specified value will be inserted into the Customer Comments column of the corresponding Exclude Control Statement message in section D5 of the sub-capacity report. Comments may be up to 255 characters in length and should not contain commas or quotation marks. The comment may span multiple lines of the SPECIAL DD by enclosing the text in quotation marks.

Syntax examples:

1. The following **EXCLUDE** control statement demonstrates how the values for the **IMAGE_ID** and **PRODUCT_ID** keyword parameters can be split over multiple statement lines:

```
EXCLUDE CPC=tttt-sssss,
START=yyyy/mm/dd/hh,
RESUME=yyyy/mm/dd/hh,
IMAGE_ID=(image1,
image2),
PRODUCT_ID=(xxxx-xxx,
xxxx-xxx)
```

This preceding above is equivalent to the following statement:

```
EXCLUDE CPC=tttt-sssss,START=yyyy/mm/dd/hh,RESUME=yyyy/mm/dd/hh,
IMAGE_ID=(image1,image2),PRODUCT_ID=(xxxx-xxx,xxxx-xxx)
```

2. The following **EXCLUDE** control statement demonstrates how the comment string can be split over multiple statement lines:

```
EXCLUDE CPC=tttt-sssss,
START=yyyy/mm/dd/hh,RESUME=yyyy/mm/dd/hh,
IMAGE_ID=image1,
COMMENT="This sample comment
spans multiple lines"
```

The following examples help explain the use of **EXCLUDE** control statements to exclude SMF records. Additional examples appear in [“Common SCRT problem scenarios”](#) on page 442.

Example 1 — Excluding records for all products by measurement interval: In this example, the customer needs to remove damaged SMF records that were generated for the measurement intervals on two different days. The SMF records were generated by a native z/OS system running in LPAR BANKS on a 2097 with serial number 12345. One damaged record was generated for the measurement period for 15 October 2024 at 7:00 PM; the other damaged records were generated for the measurement period for 20 October 2024 between 1:00 PM and 3:00 PM.

Due to the different dates and times involved, two different **EXCLUDE** control statements are needed, as follows:

```
* exclude bad SMF records for 15 October 2024 for the 7 PM hour
EXCLUDE CPC=2097-12345,IMAGE_ID=BANKS,START=2024/10/15/19,
RESUME=2024/10/15/20
*
*exclude bad SMF records for 20 October 2024 for the 1 PM and 2 PM hours
```

```
EXCLUDE CPC=2097-12345,IMAGE_ID=BANKS,START=2024/10/20/13,
RESUME=2024/10/20/15
```

The first **EXCLUDE** control statement removes the SMF records generated on 15 October 2024 at 7:00 PM (19:00) and the second **EXCLUDE** control statement removes the SMF records for 20 October 2024 between 1:00 PM and 2:59 PM. Note that both **EXCLUDE** statements had to be split across multiple lines. In both cases, the first line of the control statement ends with a comma and the parameter that continues on the next line begins in column 3 (the indentation is purely for readability).

The **START** and **RESUME** values for the control statement for 15 October exclude only hour 19 from the BANKS LPAR, as the **RESUME** hour specified on that statement (20) marks the hour when SCRT resumes processing the SMF records. The **START** and **RESUME** values for the control statement for 20 October exclude the hours 13 and 14 (for 1:00 PM and 2:00 PM), with SCRT processing resuming with hour 15 (3:00 PM).

Because all of the products from the BANKS LPAR are excluded, SCRT subtracts the three excluded hours when it calculates the hours to be reported for the **Report Period % Data** value for the BANKS LPAR in the Detail Data Collection section of the sub-capacity report. If BANKS was the only LPAR on that CPC that supplied SMF records for those three hours, then SCRT would also subtract those three hours from the hours used to calculate the **% Data Collected** value in the Tool Information section of the report.

Note that this customer must use the appropriate comment column in the Special Conditions section to explain the reason for using each **EXCLUDE** control statement.

Example 2 — Excluding records for selected products by measurement interval: In this example, the customer needs to remove the statistics for two products that, at the beginning of the month, ran on a CPC while those products were not licensed for that CPC. This occurred because the customer ran both products as part of a workload that had to be temporarily moved to that CPC because the regular CPC for which these products were licensed was down, due to a hardware problem; however, at the time, those two products were not licensed for the CPC on which the customer was temporarily running them. The customer eventually decided, later in the month, to license both products for this CPC and so the product usage on that CPC later in the month should rightly appear on the sub-capacity report.

The following **EXCLUDE** control statement removes the two products from the SMF records generated for the measurement period from midnight (00:00) at the very beginning of 5 October 2024 up to, but not including, midnight (00:00) at the very beginning of 6 October 2024:

```
* exclude MQSeries V6 for 24 hours for 5 October
* exclude Lotus Domino V7 for 24 hours for 5 October
EXCLUDE CPC=2097-23456,IMAGE_ID=PRODEAST,
PRODUCT_ID=(5655-L82,5655-N14),
START=2024/10/05/00,RESUME=2024/10/06/00
```

Note that the SMF records that were generated for the measurement period starting at midnight at the very beginning of 6 October are included in SCRT processing. Note, too, that the **EXCLUDE** control statement had to be split across multiple lines and that the **PRODUCT_ID** parameter contains a list of products.

Because the SMF records are excluded at the product level, SCRT does not subtract any of the excluded hours when it calculates the hours to be reported for the **Report Period % Data** value for the PRODEAST LPAR in the Detail Data Collection section of the sub-capacity report. SCRT also does not subtract those hours from the hours used to calculate the **% Data Collected** value in the Tool Information section of the report.

Note that this customer must use the appropriate comment column in the Special Conditions section to explain the reason for using the **EXCLUDE** control statement.

Example 3 — Excluding records using a mix of local and UTC times: This example is based on the same conditions as in Example 1, with one exception: There was a time zone change (UTC offset change) at 2:00 PM on 20 October. If the customer needs to exclude SMF records for 15 October at 7:00 PM, the customer can continue to use local date/time values on the **EXCLUDE** control statement because there were no time zone changes for the time range to be excluded.

However, if the customer needs to exclude the SMF data collected between 1:00 PM and 3:00 PM local time (not counting 3:00 PM) on 20 October, then the customer must use UTC date/time values because there was a time zone change during the excluded interval. Since this customer was running with a UTC offset of -4 at 1:00 PM and a UTC offset of -5 at 3:00 PM, the customer must add the UTC offsets to the corresponding local timestamps on the SMF records for those hours to calculate their corresponding **START** and **RESUME** UTC times. The local time of 13:00 (1:00 PM) time translates into 17:00 UTC time (13 + 4). The local time of 15:00 (3:00 PM) translates to 20:00 UTC time (15 + 5). Since none of the translated times crosses a day boundary, the UTC dates are the same as the local dates.

The correct **EXCLUDE** control statements for the set of conditions in this example are:

```
* exclude bad SMF records for 15 October 2024 for the 7 PM hour
EXCLUDE CPC=2097-12345,IMAGE_ID=BANKS,START=2024/10/15/19,
RESUME=2024/10/15/20
*
*exclude bad SMF records for 20 October 2024 for the 1 PM and 2 PM hours
EXCLUDE CPC=2097-12345,IMAGE_ID=BANKS,START=2024/10/20/17/UTC,
RESUME=2024/10/20/20/UTC
```

This example demonstrates that SCRT accepts a mixture of **EXCLUDE** control statements, some that use local date/time formats and others that use UTC date/time formats.

EXCLUDE MSU_CONSUMPTION control statement

The **EXCLUDE MSU_CONSUMPTION** control statement is for use under very rare conditions that require you to direct SCRT to deduct some MSU consumption from a Tailored Fit Pricing container that is based on a consumption metric. This control statement can be used to adjust for unusual incidents that significantly impact MSU consumption. For further guidance, see [“Unusual situations that affect utilization”](#) on page 17.

An **EXCLUDE MSU_CONSUMPTION** control statement specifies the impacted container solution ID, a CPC, an incident hour, and the number of CPU seconds attributed to the unusual incident. SCRT converts the CPU time to MSUs based on the machine capacity at the hour specified and deducts the resulting value from the MSU consumed by the container. You are responsible for correctly accounting CPU time related to the unusual incident.

Depending on the exact circumstances, several different **EXCLUDE MSU_CONSUMPTION** control statements might be needed if the event being excluded occurred in multiple hours and during those hours the impacted CPC capacity model was changed, or the incident occurred on multiple CPCs.

The Special Conditions section of the report shows the values specified on each **EXCLUDE MSU_CONSUMPTION** control statement and the number of MSUs that SCRT calculated for each statement.

Report sections that show MSU consumption values will reflect the reduced value.

IBM requires that you provide an explanation in the appropriate comments column for each **EXCLUDE MSU_CONSUMPTION** control statement that appears in the Special Conditions section of the report.

The general syntax for the **EXCLUDE MSU_CONSUMPTION** control statement is:

```
EXCLUDE MSU_CONSUMPTION,CPC=tttt-sssss,HOUR=yyyy/mm/dd/hh,CPUTIME=seconds,ID=solutionId,
COMMENT=comment-string
```

The keywords and values are:

EXCLUDE MSU_CONSUMPTION

(Required) Identifies the beginning of an **EXCLUDE MSU_CONSUMPTION** control statement.

CPC

(Required) The CPC machine type (tttt) and CPC serial number (sssss) to which this control statement applies. Note that the same conditions apply to this parameter as for the **CPC** parameter in [“EXCLUDE control statement”](#) on page 44.

HOUR

(Required) The hour in which the incident started. SCRT uses the machine capacity from this hour to convert CPU time into MSU consumption. If an incident spans multiple hours and the CPC changed capacity during those hours, you must specify multiple **EXCLUDE MSU_CONSUMPTION** control statements.

There are two formats for this value:

- The format when specifying the date/time values in local time is:

```
yyyy/mm/dd/hh
```

- The format when specifying the date/time values in UTC time is:

```
yyyy/mm/dd/hh/UTC
```

where:

yyyy

The year. Valid values are from 2000 to 2025, inclusive. (For instance, 2019.)

mm

The month. Valid values are from 01 to 12, inclusive. (For instance, 08 for August.)

dd

The day of the month. Valid values are from 01 to 31, inclusive, depending on the month. (For instance, 01 for the first day of the month.)

hh

The hour. Valid values are from 00 to 23, inclusive. (For instance, 00 starts the day at midnight.)

UTC

Indicates that the date/time value is in UTC time.

The following rules summarize how to determine whether you must use the local format or UTC format to specify the date/time values for the **HOUR** parameter on **EXCLUDE MSU_CONSUMPTION** control statements:

1. UTC date/time formats are acceptable for all date/time conditions *except* when using the **REPORT_TIME** control statement with the **Special_Local_Time** option. Note that UTC date/time formats are required for a specific set of date/time conditions.
2. Local date/time formats are acceptable for all date/time conditions *except*:
 - Whenever the data for the excluded date/time range contains a change in time zones (change in the UTC offset), **or**
 - Whenever the data for the excluded date/time range contains multiple image IDs and the image IDs are in different time zones (different UTC offsets), **or**
 - Whenever you use the **REPORT_TIME** control statement with the **UTC_Time** option

CPUTIME

(Required) The CPU time, in seconds, attributed to the incident.

ID

(Required) The ID of the Tailored Fit Pricing solution from which the MSUs should be deducted.

COMMENT

(Optional) A comment that justifies the use of the **EXCLUDE** control statement. The specified value will be inserted into the Customer Comments column of the corresponding Exclude Control Statement message in section D5 of the sub-capacity report. Comments may be up to 255 characters in length and should not contain commas or quotation marks. The comment may span multiple lines of the SPECIAL DD by enclosing the text in quotation marks. (An example of this syntax appears in [“EXCLUDE control statement”](#) on page 44.)

GENERATE_DETAILED_DATA control statement

Use the **GENERATE_DETAILED_DATA** control statement to request the generation of certain sub-capacity report sections that provide additional details about the applied pricing options and specific MSU reductions. The following additional report sections can be generated:

- Mobile/zCAP Pricing Detailed Data section (V5), for sub-capacity reports that use the MWP or zCAP algorithms
- Sub-hour Data Collection Details section (V6), for sub-capacity reports for Tailored Fit Pricing solutions that use an MSU consumption metric.
- Detailed Interval Data section (W3), for sub-capacity reports (multiplex reports) that use the CMP algorithm

These new report sections provide additional details about the application of pricing metrics. These details are provided on a per-product basis for each hour in the reporting period.

Because these sections of the sub-capacity report can generate a large amount of output, they are suppressed by default; you must use the **GENERATE_DETAILED_DATA** control statement to generate them.

Typically, you should not submit a sub-capacity report containing these optional sections to IBM.

For more information about section V5, see [“Mobile/zCAP Pricing Detailed Data section for MWP” on page 214](#) and [“Mobile/zCAP Pricing Detailed Data section for zCAP” on page 246](#).

For more information about section W3, see [“Detailed Interval Data section” on page 362](#).

IGNORE CONTAINER control statement

Use the optional **IGNORE CONTAINER** control statement to request that SCRT ignore a Tailored Fit Pricing solution for a specified time interval.

The **IGNORE CONTAINER** control statement is only intended for use under rare circumstances and as directed by IBM.

For the interval specified on an **IGNORE CONTAINER** control statement, SCRT does not apply the Tailored Fit Pricing algorithm and considers any Tailored Fit Pricing tenant resource groups (TRGs) and dedicated LPARs that are associated with the specified solution ID to be part of the standard sub-capacity environment. Specifically:

- Tailored Fit Pricing LPARs are aggregated into standard sub-capacity product reporting.
- The 4-hour rolling average contribution and product data for the Tailored Fit Pricing TRG are not removed from the product values for the LPAR.
- Container peak calculations ignore the specified time range.

Each **IGNORE CONTAINER** control statement starts with the **IGNORE CONTAINER** keywords. All subsequent parameters are delimited by commas. At least one parameter must follow the **IGNORE CONTAINER** keywords. Parameters may be specified on subsequent continuation lines. Each line that is being continued must end with a comma.

The general syntax for the **IGNORE CONTAINER** control statement is as follows:

```
IGNORE CONTAINER,
ID=solutionID,
START=yyyy/mm/dd/hh[/UTC],
RESUME=yyyy/mm/dd/hh[/UTC]
```

The keywords and values are:

IGNORE CONTAINER

(Required) Identifies the beginning of the control statement.

ID=solutionID

(Required) Specifies the 64-character solution ID (*solutionID*) of the Tailored Fit Pricing solution that is to be ignored for the specified time interval, or *ALL. When specifying ID=*ALL, the **START** and **RESUME** keywords are not required; if specified, they are ignored.

START=yyyy/mm/dd/hh[/UTC]

(Required) Specifies the starting date and time (hour), in local SMF time or UTC time, of the measurement interval for the container that is to be ignored by the Tailored Fit Pricing algorithm. Specify the **START** value in one of the following formats:

yyyy/mm/dd/hh

Use this format to specify the date and time values in local time.

yyyy/mm/dd/hh/UTC

Use this format to specify the date and time values in UTC time.

The values are:

yyyy

The year. Valid values are 2000 - 2025.

mm

The month. Valid values are 01 - 12.

dd

The day of the month. Valid values are 01 - 31, depending on the month.

hh

The hour. Valid values are 00 - 23. (For instance, 00 starts the day at midnight.)

UTC

(Optional) Indicates that the date and time values are specified in UTC time; otherwise, SCRT assumes local time.

RESUME=yyyy/mm/dd/hh[/UTC]

(Required) Specifies the date and time (hour), in local SMF time or UTC time, of the measurement interval for the container when the Tailored Fit Pricing algorithm is to resume processing for the container. The formats and values are the same as for the **START** parameter.

IGNORE_WLM_MWP_DATA control statement

Use the **IGNORE_WLM_MWP_Data** control statement to request that the Mobile Workload Pricing algorithm not be applied to mobile workload data from WLM. For details, see [“IGNORE_WLM_MWP_DATA control statement” on page 206](#).

INCLUDE control statement

Use the **INCLUDE** control statement to specify a subset of images to be reported by SCRT. Images that are not explicitly included are not reported.

The **INCLUDE** control statement does not support specifying a date and time range; images are included for the entirety of the reporting period. You can specify both **INCLUDE** and **EXCLUDE** control statements. The **INCLUDE** control statements are processed before any **EXCLUDE** control statements.

Whenever valid **INCLUDE** control statements are encountered, SCRT lists each statement in the Special Conditions section of the report. The Special Conditions section identifies the images that were found in the SMF data but were not explicitly included. IBM requires that you provide an explanation in the appropriate comments column to provide the justification for including only a subset of images in the report.

The general syntax for the **INCLUDE** control statement is:

```
INCLUDE CPC=tttt-sssss,IMAGE_ID=image_id,
COMMENT=comment-string
```

The keywords and values are:

INCLUDE

(Required) Identifies the beginning of an **INCLUDE** control statement.

CPC=tttt-sssss

(Required) The CPC machine type (*tttt*) and CPC serial number (*sssss*) to which this **INCLUDE** control statement applies, in the format *tttt-sssss* (for instance, 2094-12345). This parameter does not support the special value *ALL nor a list of values. If SCRT does not find the specified value on any of the SMF or SCRT89 records being processed by this execution of SCRT, it issues an error message and terminates processing with return code 16.

IMAGE_ID=image_id

IMAGE_ID=(image_id1,...,image_idn)

(Required) The single image ID or list of image IDs to be included.

COMMENT

(Optional) A comment that justifies the use of the **EXCLUDE** control statement. The specified value will be inserted into the Customer Comments column of the corresponding Exclude Control Statement message in section D5 of the sub-capacity report. Comments may be up to 255 characters in length and should not contain commas or quotation marks. The comment may span multiple lines of the SPECIAL DD by enclosing the text in quotation marks. (An example of this syntax appears in [“EXCLUDE control statement”](#) on page 44.)

INTEGRATED_WORKLOAD_PRICING control statement

Use the **INTEGRATED_WORKLOAD_PRICING** control statement to specifically request SCRT to use the IWP algorithm when generating sub-capacity reports. The default SCRT processing is not to use the IWP algorithm. However, even when the **INTEGRATED_WORKLOAD_PRICING** control statement is used, SCRT only applies the IWP algorithm to the CPCs, hours, and LPARs that meet all of the IWP requirements. There is no guarantee that resulting sub-capacity reports will show any benefit due to IWP.

Note that IBM only authorizes use of the **INTEGRATED_WORKLOAD_PRICING** control statement by customers who have signed or otherwise agreed to the IWP Addendum (Z125-8702 / INTC-8702). IBM rejects all sub-capacity reports that show an IWP benefit whenever the IWP addendum has not been signed or otherwise agreed to. See [Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing \(IWP\),”](#) on page 183 for more information about IWP and its requirements.

There are no parameters on the **INTEGRATED_WORKLOAD_PRICING** control statement. The general syntax for the **INTEGRATED_WORKLOAD_PRICING** control statement is:

```
INTEGRATED_WORKLOAD_PRICING
```

where:

INTEGRATED_WORKLOAD_PRICING

(Required) Identifies the **INTEGRATED_WORKLOAD_PRICING** control statement. The **INTEGRATED_WORKLOAD_PRICING** control statement must start in column one.

Whenever the **INTEGRATED_WORKLOAD_PRICING** control statement is used, it applies to all of the sub-capacity reports generated by that execution of SCRT.

Example: The following **INTEGRATED_WORKLOAD_PRICING** control statement requests that SCRT use the IWP option when generating sub-capacity reports for all CPCs processed by that execution of SCRT.

```
* Request the IWP option
INTEGRATED_WORKLOAD_PRICING
```

INTERVAL_RATE_DATA control statement

Use the **INTERVAL_RATE_DATA** control statement to request the generation of interval rate data. Interval rate data appears in the CPC MSU Rate Above Permanent Capacity Rating section (I5) of the SCRT report for Hardware Consumption. For more information, see [Chapter 21, “Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution,” on page 405.](#)

IPLA_ALL control statement

Use the **IPLA_ALL** control statement to request that SCRT report an IPLA family line for MVM eligible products in the Product Summary Information section (E5), even when only one product in the product family is in use.

ISV_SHOW_CONTAINERS control statement

Use the **ISV_SHOW_CONTAINERS** control statement to request container processing for independent software vendor (ISV) product reports. By default, SCRT does not process containers in ISV mode.

There are no parameters on the **ISV_SHOW_CONTAINERS** control statement. The general syntax for the **ISV_SHOW_CONTAINERS** control statement is:

```
ISV_SHOW_CONTAINERS
```

where:

ISV_SHOW_CONTAINERS

(Required) Identifies the **ISV_SHOW_CONTAINERS** control statement. The **INTEGRATED_WORKLOAD_PRICING** control statement must start in column one.

When you use the **ISV_SHOW_CONTAINERS** control statement and ISV reporting is active (by having specified the ISVLIB DD statement), SCRT uses the Tailored Fit Pricing for IBM Z algorithm to report on any configured Tailored Fit Pricing solutions.

MLC_ALL control statement

Requests SCRT to report an MLC family line for MVM eligible products in the Product Summary Information section (E5), even when only one product in the product family is in use.

REPORT_TIME control statement

Use the **REPORT_TIME** control statement to direct SCRT to report dates and times from z/OS systems on the sub-capacity report using either a special form of local time or UTC time.

By default, SCRT reports dates and times on the sub-capacity report in terms of local time whenever the operating system's local time offsets to UTC time are the same for all of the operating system images running on the same CPC during the reporting period. If any of the operating system offsets differ from each other for two or more hours, SCRT reports all of the dates and times in terms of UTC dates and times. For more information about how SCRT handles local time versus UTC time, see [“SMF and SCRT89 records from different time zones” on page 14.](#)

Whenever SCRT processes a valid **REPORT_TIME** control statement, it lists that control statement in the Special Conditions section of every sub-capacity report generated by that execution of SCRT.

The general syntax for the **REPORT_TIME** control statement is:

```
REPORT_TIME=option
```

The keywords and options are:

REPORT_TIME

(Required) Identifies the beginning of the **REPORT_TIME** control statement. The **REPORT_TIME** keyword must start in column one.

You may only specify one **REPORT_TIME** control statement in any given execution of SCRT, and that control statement applies to all of the SMF records for all of the CPCs being processed by that execution, including the data before and after the reporting period.

option

(Required) Must be one of the following keyword options:

Special_Local_Time

Directs SCRT to use the actual date and time values listed on the SMF records to report local time. Without the **REPORT_TIME** control statement, SCRT applies the hardware timer offset values to the date and time values on the SMF records in order to calculate the local dates and times that appear on the sub-capacity report. The list of restrictions for using this value are described later. For SMF records with hardware timer offset values of zero, there is no difference in the dates and times reported by SCRT with or without using the **REPORT_TIME** control statement with the **Special_Local_Time** option.

This support is provided for customers who use hardware timer offsets instead of the z/OS offset (the **TIMEZONE** parameter in the **CLOCKxx** member of **SYS1.PARMLIB**) to set the z/OS local offset to UTC time. The use of this parameter forces the dates and times listed on the sub-capacity report to match the local dates and times reported by other tools, such as RMF.

Restrictions for the Special_Local_Time option:

- The special local time reporting only applies to the data from z/OS systems. Use of the **Special_Local_Time** option with z/TPF or z/VSE data is unpredictable.
- You cannot use the **Special_Local_Time** option to force SCRT to report in terms of special local time if previous executions of SCRT using the same data caused SCRT to report in terms of UTC time, because the data being processed violates one of the following restrictions.

In addition, *all* of the following conditions must be met. If any condition is not met, SCRT issues an error message and terminates processing with return code 16 and without generating any sub-capacity reports.

- The hardware time offset to UTC time, as determined by the Sysplex Timer Offset, LPAR Timer Offset, and the z/VM virtual clock (VTOD), must be the same offset value for all data processed by that execution of SCRT, including data before and after the reporting period.
- Hardware timer offsets greater than 16 hours or less than -16 hours are not supported.
- The z/OS local time offset to UTC time must be the same for all data processed by the same execution of SCRT, including data before and after the reporting period, except for a single two-hour grace period when local offsets change when moving to or from daylight saving time.
- z/OS timer offsets greater than 16 hours or less than -16 hours are not supported.
- UTC times are not supported for the **START** and **RESUME** parameters on the **EXCLUDE** control statement. If UTC times are used on any **EXCLUDE** control statements when the **Special_Local_Time** option is in effect, SCRT terminates processing without generating any sub-capacity reports.

UTC_Time

Directs SCRT to report all z/OS-based dates and times on the sub-capacity report in terms of UTC dates and times. SCRT calculates UTC dates and times by applying both the hardware timer offsets and the z/OS local offsets to the date and time values listed on the SMF records. The restrictions for using the **Special_Local_Time** option do not apply to the **UTC_Time** option. There are no restrictions on the hardware or operating system timer offsets in the data processed by SCRT when using the **UTC_Time** option. This option can only be used for z/OS data.

Restrictions for the UTC_Time option:

SPECIAL DD statement

- The **UTC_Time** option only applies to the data from z/OS systems. Use of the **UTC_Time** option with z/TPF or z/VSE data is unpredictable.
- Local time values (time values without the UTC suffix) are not supported for the **START** and **RESUME** parameters on the **EXCLUDE** control statement whenever a **REPORT_TIME** control statement is used with the **UTC_Time** option. When the **UTC_Time** option is used under these conditions, SCRT issues an error message and terminates without generating any sub-capacity reports.

Examples: The following examples demonstrate the only two supported variations of the **REPORT_TIME** control statement. In both examples, the **REPORT_TIME** control statement is preceded by one or more comments, which start with an asterisk (*) in column 1.

• Example 1:

```
* Force SCRT to report using dates and times directly from SMF
*
REPORT_TIME=Special_Local_Time
```

• Example 2:

```
* Force SCRT to report using UTC dates and times
*
REPORT_TIME=UTC_Time
```

SUPPRESS_REPORT_SECTION control statement

Use the **SUPPRESS_REPORT_SECTION** special control statement to suppress the reporting of the three sub-capacity report sections:

- Getting Started Sub-capacity Pricing Assigned - Hours
- Getting Started Sub-capacity Pricing Equal to Traditional Sub-capacity - Hours
- Getting Started Sub-capacity Pricing Not Assigned - Hours

These new report sections track the hours related to Getting Started Sub-Capacity Pricing products and are described in [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#). SCRT automatically generates these report sections whenever it processes data for one or more of the Getting Started Sub-Capacity Pricing products for that subcapacity report. Customers who do not meet all of the requirements for Getting Started Sub-Capacity Pricing might want to suppress the generation of these report sections and can do so by using this control statement.

Currently, the **SUPPRESS_REPORT_SECTION** control statement only supports a parameter value that requests the suppression of the three Getting Started Sub-Capacity Pricing report sections.

The general syntax for the **SUPPRESS_REPORT_SECTION** control statement is:

```
SUPPRESS_REPORT_SECTION=section_identifier
```

where:

SUPPRESS_REPORT_SECTION

(Required) Identifies the beginning of a **SUPPRESS_REPORT_SECTION** control statement. The **SUPPRESS_REPORT_SECTION** keyword must start in column one.

section_identifier

(Required) Indicates the sub-capacity report section or sections that are to be suppressed. Currently, the only supported value for *section_identifier* is:

GSSP

Suppresses the three sub-capacity report sections related to Getting Started Sub-Capacity Pricing.

Whenever the **SUPPRESS_REPORT_SECTION** control statement is used, it applies to all of the sub-capacity reports generated by that execution of SCRT.

Example: The following **SUPPRESS_REPORT_SECTION** control statement prevents the automatic generation of the three Getting Started Sub-Capacity Pricing report sections for all of the sub-capacity reports generated by that execution of SCRT:

```
* prevent generation of GSSP reports, since GSSP usage too high
SUPPRESS_REPORT_SECTION=GSSP
```

UPDATE CONTAINER control statement

Use the **UPDATE CONTAINER** control statement to associate the LMS-provided solution name with a Tailored Fit Pricing solution based on the associated solution ID and to designate NO89 products that run in the container. The **UPDATE CONTAINER** statement is required for Tailored Fit Pricing solutions.

Each Tailored Fit Pricing solution, whether found in SMF data or specified using the **CONTAINER** control statement, must have a solution name associated with it via an **UPDATE CONTAINER** statement. The name specified on the **UPDATE CONTAINER** statement must match the name specified for the container (as identified by the unique solution ID) in the LMS web application.

Each **UPDATE CONTAINER** control statement starts with the **UPDATE CONTAINER** keywords. All subsequent parameters are delimited by commas. At least one parameter must follow the **UPDATE CONTAINER** keywords. Parameters may be specified on subsequent continuation lines. Each line that is being continued must end with a comma.

The general syntax for the **UPDATE CONTAINER** control statement is:

```
UPDATE CONTAINER,
ID=solutionID,
SET_NAME="solutionName",
NO89_PRODUCT_ID={productId|(productID1,...,productIDn)}
```

The keywords and values are:

UPDATE CONTAINER

(Required) Identifies the beginning of the control statement.

ID=solutionID

(Required) Specifies the 64-character IBM-provided solution ID (*solutionID*) for the qualified Tailored Fit Pricing solution.

SET_NAME="solutionName"

(Required) Specifies the IBM-provided solution name (*solutionName*) for the qualified Tailored Fit Pricing solution. The value specified here must correspond to the solution ID specified in LMS. The text must be enclosed within double quotation marks and can contain blanks but should not contain commas nor any special characters that cannot be correctly translated from EBCDIC to ASCII encoding.

NO89_PRODUCT_ID=pid

NO89_PRODUCT_ID=(pid1,...,pidn)

(Optional) Specifies a NO89 product ID (PID) or a list of NO89 product IDs that are to be assigned to the associated container. To specify a list of product IDs, separate the product IDs with commas and enclose the list in parentheses.

UPDATE SOLUTION_ID control statement

Use the optional **UPDATE SOLUTION_ID** control statement to replace a solution ID associated with a container TRG with a new or corrected solution ID. The replacement solution ID must correspond to the qualified workload that ran in the container.

The **UPDATE SOLUTION_ID** control statement is only intended for use under rare circumstances, such as when a TRG is defined with a solution ID that is incorrect or not applicable.

INPUTCSV DD statement

Each **UPDATE SOLUTION_ID** control statement starts with the **UPDATE SOLUTION_ID** keywords. All subsequent parameters are delimited by commas. At least one parameter must follow the **UPDATE SOLUTION_ID** keywords. Parameters may be specified on subsequent continuation lines. Each line that is being continued must end with a comma.

The general syntax for the **UPDATE SOLUTION_ID** control statement is as follows:

```
UPDATE SOLUTION_ID,  
OLD=oldSolutionID,  
NEW=newSolutionID
```

The keywords and values are:

UPDATE SOLUTION_ID

(*Required*) Identifies the beginning of the control statement.

OLD=*oldSolutionID*

(*Required*) Specifies the 64-character solution ID (*oldSolutionID*) that is to be replaced in the input SMF data. This solution ID is defined to WLM TRGs and, due to a rare circumstance, must be replaced with a new value.

NEW=*newSolutionID*

(*Required*) Specifies the 64-character solution ID (*newSolutionID*) that is to replace the solution ID specified by the **OLD** parameter. This is the solution ID that represents the qualified Tailored Fit Pricing solution, as identified in the LMS web application.

z/TPF_TRANSFORMATION_ENGINES_OFFERING control statement

Use the **z/TPF_TRANSFORMATION_ENGINES_OFFERING** special control statement to request that SCRT report on the z/TPF Transformation Engines Offering. The default is not to report on this offering. However, even when the **z/TPF_TRANSFORMATION_ENGINES_OFFERING** control statement is used, SCRT only reports on this offering when it finds z/TPF data in the reporting period.

You must meet all of the requirements for the z/TPF Transformation Engines Offering before IBM will apply the MSUs listed for that offering on the sub-capacity report.

There are no parameters on the **z/TPF_TRANSFORMATION_ENGINES_OFFERING** control statement. The general syntax for the control statement is:

```
z/TPF_TRANSFORMATION_ENGINES_OFFERING
```

where:

z/TPF_TRANSFORMATION_ENGINES_OFFERING

(*Required*) Identifies the **z/TPF_TRANSFORMATION_ENGINES_OFFERING** control statement. The **z/TPF_TRANSFORMATION_ENGINES_OFFERING** keyword must start in column one.

Whenever the **z/TPF_TRANSFORMATION_ENGINES_OFFERING** control statement is used, it applies to all of the sub-capacity reports generated by that execution of SCRT.

Example: The following **z/TPF_TRANSFORMATION_ENGINES_OFFERING** control statement requests that SCRT report on the z/TPF Transformation Engines Offering when generating sub-capacity reports for all CPCs processed by that execution of SCRT:

```
* Request reporting for the z/TPF Transformation Engines Offering  
z/TPF_TRANSFORMATION_ENGINES_OFFERING
```

INPUTCSV DD statement

The INPUTCSV DD statement identifies the IBM Mobile Workload Pricing, IBM Z Workload Pricing for Cloud, or IBM Z Collocated Application Pricing input data set or data sets. These data sets contain data in plain-text, comma-separated value format. The specified data sets can either be sequential data sets or a

member of a partitioned data set (PDS). SCRT also accepts concatenations of data sets containing MWP or zCAP CSV file data as input on the INPUTCSV DD statement.

Guidelines:

- Sequence numbers must be disabled (or removed) for input CSV data sets.
- Partitioned data sets must include a member name.
- A separate CSV file must be provided for each product family (such as CICS, Db2, IMS, MQ, or WebSphere).
- PDSE data sets are not supported.
- The INPUTCSV data does not support a continuation character. If you require additional columns, use a larger LRECL value.
- If you specify multiple INPUTCSV data sets using DD concatenation, they must all have the same LRECL.

The specific format of the input data depends on the pricing option that the data represents. For more information, see [“Format of the INPUTCSV input data for Mobile Workload Pricing” on page 204](#) and [“Format of the INPUTCSV input data for IBM Z Collocated Application Pricing” on page 235](#).

ISVLIB DD statement

The ISVLIB DD statement identifies the input data provided by an independent software vendor (ISV) that describes the supported ISV programs for the purpose of generating an ISV report. Obtain this data set from the vendor on whose products you wish to generate an ISV report, and supply it as input to SCRT via the ISVLIB DD statement.

NO89 DD statement

Use this section to enter information about any listed products you are using that do not produce SMF type 89 records or SCRT89 records. See the JCL in [“Sample JCL to run the 64-bit SCRT on a z/OS system” on page 25](#).

SCRT supports the z/OS, z/TPF and z/VSE operating system families of products. Whenever both z/OS and another operating system that is eligible for sub-capacity pricing, like z/TPF or z/VSE, are running on the same CPC, the data from all of the eligible operating systems must be input into the same execution of SCRT so that SCRT can create a single report containing all of the sub-capacity products running on that CPC. In order to meet this requirement, the SCRT deliverable contains a set of NO89 control statements for z/OS products and a set of NO89 control statements for z/VSE products. z/TPF products do not use NO89 statements.

Both the z/OS and z/VSE sets of NO89 control statements are included as part of the SCRT deliverable, preceded by comments indicating the operating system to which each belongs. Only use both sets of NO89 control statements when the input to SCRT contains SMF/SCRT89 records from both z/OS and z/VSE systems. Otherwise, delete the NO89 control statements that do not apply to your configuration.

Special considerations: The following types of products require special attention when determining how to report the product’s use on the NO89 control statement:

- WebSphere Transformation Extender for z/OS V8 (5655-R95)

This product supports multiple product features, some of which require that the NO89 control statements indicate where they are running. Other features do not require that the NO89 control statements indicate where they are running because they are being tracked on SMF type 89 records. Since some service levels of the product require the support of the NO89 control statement, SCRT always requires the use of a NO89 control statement for this product, even when the appropriate NO89 value is *NONE. You must be careful to *determine which features are being used, where they are being used, and at what service levels* in order to determine the correct NO89 values to use. For instance, there are cases where the only features being used on a CPC do not require the NO89 control statement to track the product because, at the service level being used, those features are being tracked on SMF type 89 records. In such cases, you must set the value for the required

NO89 control statement for 5655-R95 to *NONE. For a description of all of the cases that need to be considered when setting the NO89 value for this complex product, see [“Considerations for WebSphere Transformation Extender V8 \(5655-R95\)”](#) on page 166.

- Products that are also delivered as part of a bundle or service pack

Whenever the standalone version of a NO89 product was ordered and is being used, there are no special considerations to determine how to specify the NO89 value(s) for that product beyond the general NO89 instructions described in the following sections. However, when the same function (product) is delivered as part of a bundle or a service pack, and that function is deployed under the terms and conditions for the bundle or service pack, then in most cases the NO89 value for the standalone function (product) should be set to *NONE because the function available in the standalone product is to be reported under the product ID for the bundle or service pack under which it is deployed. You must be careful to *follow all of the installation instructions for bundles and service packs*, as the installation instructions should explain how to use the NO89 control statements, if any, for that bundle or service pack.

- If you specify the **ENTERPRISE_TAILORED_FIT_PRICING** control statement, z/OS NO89 products are no longer required in the NO89 DD statement. If z/OS NO89 products are not specified in the NO89 DD, you must use the **NO89_PRODUCT_ID** keyword of the **UPDATE CONTAINER** control statement to specify active products. If your environment also contains z/VSE, you must still specify the z/VSE NO89 products in the NO89 DD.

Guidelines: You must observe the following requirements for the NO89 control statements:

1. Except for the deletion of an entire set of NO89 control statements that do not apply to your configuration, do *not* delete any individual product IDs from the list of IBM-supplied NO89 control statements in the SCRT deliverable that you are using. If you do, you will receive error message SCRTTOOL212 and no sub-capacity report will be generated. If you want to indicate that you did not use a particular product ID in any LPAR during the reporting period, use the *NONE keyword value.
2. Do *not* add any additional product IDs into the list of IBM-supplied NO89 control statements in the SCRT deliverable that you are using. If you do, you will receive error message SCRTTOOL050 and no sub-capacity report will be generated. It is not possible to add NO89 control statements to force SCRT to report on a product which is not already identified by IBM as being a NO89 product in the SCRT deliverable that you are using. Note that if a newer level of SCRT is available from the SCRT web page, that newer level may support the product IDs that you are trying to add to the list of NO89 control statements. If so, obtain, customize, and use that new level of SCRT. Other reasons for not adding additional product IDs into the list of NO89 control statements include:
 - The product that you are trying to add might be instrumented to be tracked on SMF type 89 records and does not require a NO89 control statement for SCRT to support it.
 - The product that you are trying to add might be licensed as a reference-based or z/OS-based product, and such products are not reported by SCRT because their sub-capacity value is the same as the sub-capacity value of their parent product or of z/OS, respectively.

See [Sub-Capacity Eligible and MVM Eligible MLC Programs \(www.ibm.com/downloads/documents/us-en/10a99803f2afd789\)](http://www.ibm.com/downloads/documents/us-en/10a99803f2afd789) for the lists of products that are and are not tracked on SMF type 89 records.

Every NO89 statement can be a maximum of 71 characters long and must start in column one. Whenever a single NO89 statement must be continued on multiple lines, the line being continued must end with a comma and the next line must start in column two or later.

The general syntax for NO89 statements is:

```
product_number=tttt-sssss:image_id,tttt-sssss:image_id,...,
tttt-sssss:image_id,...
```

When the assignment applies to all of the native and guest systems being processed by this execution of SCRT, including native and guest systems on multiple CPCs, the NO89 statement syntax is:

```
product_number=*ALL
```

When the assignment applies to none of the native and guest systems being processed by this execution of SCRT, including native and guest systems on multiple CPCs, the NO89 statement syntax is:

```
product_number=*NONE
```

The NO89 statement parameters are:

product_number

(Required) Identifies the product number for which the assignment is being made. The product number must be in the format *xxxx-xxx*. For the NO89 DD statement in the SCRT JCL in “[Sample JCL to run the 64-bit SCRT on a z/OS system](#)” on page 25, all of the product numbers are provided as incomplete NO89 control statements. Each of these sample NO89 control statements is preceded by a comment statement that contains the name of the product.

tttt-sssss

(Optional) Identifies the CPC machine type (*tttt*) and serial number (*sssss*) that applies to the *image_id* value that immediately precedes the colon (:). This CPC value does not apply to any other *image_id* values on the NO89 control statement unless that CPC value also immediately precedes each of those *image_id* values. This optional *tttt-sssss* value is only required when multiple CPCs are being processed by that execution of SCRT and the *image_id* that follows only applies to the CPC being specified. When this optional *tttt-sssss* value is missing, the NO89 assignment is applied to all of the LPARs named *image_id* that are processed by that execution of SCRT.

image_id

*(Required whenever the keyword values *ALL or *NONE are not used instead)* Identifies the native and guest systems in which the specified NO89 product ran at any time during the reporting period for this sub-capacity report.

Use the following criteria to determine the value to specify for *image_id*:

- For the native z/OS, z/TPF, and z/VSE systems, use the LPAR name.
- For z/OS V1.7 (and later) guest systems, use the SID value from the SMFPRMxx member.
- For z/OS V1.6 (and earlier) guest systems and for later z/OS V1.7 systems being reported under z/OS 1.6 guest rules, use the value assigned to the LPAR name by the corresponding VMGUESTS DD control statement.
- For z/TPF and z/VSE guest systems, use the VM user ID assigned to the virtual machine in which the guest system runs.

***ALL**

*(Optional and mutually exclusive with image_id and *NONE)* Indicates that the specified NO89 product ran on all of the native and guest systems being processed by this execution of SCRT, including all of the CPCs being processed.

***NONE**

*(Optional and mutually exclusive with image_id and *ALL)* Indicates that the specified NO89 product did not run on any of the native and guest systems being processed by this execution of SCRT, including all of the CPCs being processed.

The following instructions apply when the SMF data being processed by SCRT contain records for a z/OS system. You must enter a valid value for each product in the list of z/OS NO89 control statements.

- If you do not have a product installed or do not use that product on any of the CPCs being processed by that execution of SCRT, enter *NONE next to the product number. Do *not* delete products from the list of z/OS NO89 control statements.
- If a product in the list runs in all native and guest systems on all of the CPCs being processed by that execution of SCRT, enter *ALL.
- If a product runs in one or more native and guest systems, enter a list of the *image_id* values for the native and guest systems in which the product runs. The examples that follow show how to handle the case where the same LPAR name is used on multiple CPCs that are being processed by the same execution of SCRT.

OUTPUT DD statement

If the data being processed by this execution of SCRT only contains SMF records for z/OS systems, then delete all of the NO89 control statements for z/VSE products from the set of NO89 statements used for that implementation of SCRT.

When the SMF/SCRT89 records being processed by SCRT include data for z/TPF systems, see [Chapter 13, “Sub-capacity reporting for z/TPF systems,”](#) on page 281.

When the SMF/SCRT89 records being processed by SCRT include data for z/VSE systems, see [Chapter 15, “Sub-capacity reporting for z/VSE systems,”](#) on page 305.

If you use duplicate LPAR names on different CPCs and process those CPCs in the same execution of SCRT, you might have to qualify the LPAR name with the machine type and serial number. Otherwise, the NO89 product will be assigned to all of the CPCs where those LPAR names exist.

Note: When ISV reporting is active, the NO89 DD statement must only include NO89 products for that ISV.

Example 1: To enter the LPARs named LPARA and LPARB for the NO89 product 5655-B86 where:

- Only one CPC is being processed by this execution of SCRT, **or**
- Multiple CPCs are being processed by this execution of SCRT and those LPAR names only exist on one of the CPCs, **or**
- Multiple CPCs are being processed by this execution of SCRT and the NO89 assignment applies to the specified LPAR names regardless of the CPC,

specify the following NO89 control statement:

```
5655-B86=LPARA, LPARB
```

Example 2: To enter the LPAR named LPARA on CPC 2064-12345 for the NO89 product 5655-B86 when you have two LPARs named LPARA on different CPCs, one of which does not have 5655-B86 running on it, specify:

```
5655-B86=2064-12345:LPARA, LPARB
```

The LPAR name that is not qualified with a machine type and serial (LPARB) applies to all CPCs.

Example 3: If both of the LPARs named LPARA from the previous example are running product 5655-B86, you need only specify:

```
5655-B86=LPARA, LPARB
```

Example 4: If product 5655-B86 runs in all LPARs on CPC 2064-12345 and you have multiple CPCs in your input, you can specify:

```
5655-B86=2064-12345:*ALL
```

Example 5: If you have duplicate LPAR names on different CPCs and some of the LPARs do not have the NO89 product running on them, you must specify the machine type and serial number for all the LPARs that do, even if they are on the same CPC. To enter the LPARs named LPARA and LPARB on CPC 2064-12345 for the NO89 product 5655-B86, when you have LPARs named LPARA and LPARB on other CPCs and those other LPARs do not have 5655-B86 running on them, specify:

```
5655-B86=2064-12345:LPARA, 2064-12345:LPARB
```

OUTPUT DD statement

The OUTPUT DD statement defines the output file for the sub-capacity reports. SCRT internally sets the **LRECL** value to 4096 tool and requires no changes to the JCL for the SCRT job. This value supports the reporting of up to 128 native and guest systems in the Product Max Contributors section of the sub-capacity report. This value cannot be overridden by the JCL.

You can choose to have SCRT place the sub-capacity reports for all CPCs processed in a single execution of SCRT into a single, sequential .csv file or into multiple members in a .csv file that is a partitioned data set (PDS).

SCRT checks the data set organization attribute of the OUTPUT DD statement to determine the file format, as follows:

- If the data set is a PDS, SCRT creates a member for each CPC encountered in the SMF input. Each member has the name *#xxyyyyyy*, where *xx* is the plant (location) where the CPC was manufactured, and *yyyyyy* is the CPC serial number.

Note that if you upgrade a CPC during the reporting period and the second CPC has a different machine type but the same serial number, a different report is created for each CPC. The first member name will start with the *#* character and the next will start with a different character.

When the **Country_Multiplex_Pricing** control statement has been specified in the SPECIAL DD, the default PDS member name has a different format than for a traditional sub-capacity report. The member name format is *#mmmyyyy*, where *mmm* is the 3-character month of the reporting period, and *yyyy* is the 4-digit year of the reporting period.

If you wish, you can specify your own member name on the **DSN** value for a PDS.

- If the data set is a sequential data set, all reports created from the input data stream will be placed, one after another, into this data set. Having all reports in a single data set makes uploading reports to the LMS web application easier. If you do not use LMS, submit separate reports for each CPC.

Note that if you upgrade a CPC during the reporting period and the second CPC has a different machine type but the same serial number, a different report is created for each CPC. There will be one report for the first CPC (the first type-serial combination) followed by a report for the second CPC (the second type-serial combination).

A sub-capacity report must be in a separate .csv file for a CPC when you want to submit the report to IBM as an estimate for a price quote or when you submit reports to IBM via email.

The member name assigned to each CPC successfully processed by SCRT is listed in SCRT message SCRTTOOL005, which is written to the SYSPRINT data set when the output is a PDS.

All reports are in comma-separated value (.csv) format and must be validated by the software asset manager and submitted to IBM.

WCSUPP DD statement

The optional WCSUPP DD statement defines the file that is to contain the Workload Classification Supplement. This file can be a sequential file or a member of a PDS or PDSE. If you do not specify a member name for an existing PDS or PDSE, SCRT generates a member name. The member name format is *@mmmyyyy* where *mmm* is the 3-character month of the reporting period, and *yyyy* is the 4-digit year of the reporting period.

For more information, see [Chapter 22, “Workload Classification Supplement,” on page 409](#).

SMF DD statement

Identifies the SMF input data set or data sets. These data sets contain the SMF 70 subtype 1 and SMF 89 records for all z/OS systems and the SCRT89 records for all z/TPF and z/VSE systems running in all LPARs defined on one or more CPCs.

You can use concatenations of SMF/SCRT89 data sets as input to the SMF DD statement. When concatenating data sets, the data set with the largest block size must be the first data set processed by the SMF DD statement. Otherwise, the SCRT job step will terminate with a S002 ABEND and z/OS error message IEC036I 002-04.

The data sets defined by the SMF DD statement must point to the SMF/SCRT89 data for the entire reporting period for all LPARs. See the recommendation for also providing SCRT89 records for one day prior and one day after the reporting period whenever SCRT reports both z/OS and z/TPF and reports in

UTC times, as described in [“Steps for z/TPF data collection and processing”](#) on page 284. SMF/SCRT89 data does not need to be sorted before running the SCRT.

If the SMF data sets contain record types other than 70 and 89, SCRT will ignore those SMF records. However, SCRT may terminate with an ABEND with several different system error messages if the SMF or SCRT89 data sets contain non-SMF or non-SCRT89 records.

Rule: When both z/OS and another operating system that is eligible for sub-capacity pricing, like z/TPF or z/VSE, are running on the same CPC, the data from all of those eligible operating systems must be input into the same execution of SCRT, so that SCRT can create a single sub-capacity report containing all of the sub-capacity products running on that CPC.

CPSTRACE DD statement

The optional CPSTRACE DD statement specifies the data set or file that is to contain detailed, hour-by-hour reporting information for Tailored Fit Pricing for IBM Z.

The record fields in this file are all separated by commas (that is, they are in comma-separated value format) so that the file can be read and processed by a spreadsheet application.

This file can be a sequential file or a member of a PDS or PDSE. If you do not specify a member name for an existing PDS or PDSE, SCRT uses the CPSTRACE member name. (The sample SCRT JCL for the z/OS deliverable uses a sequential data set.) A **RECFM** value of FB and **LRECL** value of 200 are internally set by SCRT and require no changes to the JCL for the SCRT job. The internally set values cannot be overridden by the JCL.

For details about this file, see [“CPSTRACE DD statement”](#) on page 379.

GSTRACE DD statement

The optional GSTRACE DD statement defines the file that is to contain the list of reasons why SCRT did not assign the Getting Started Sub-Capacity Pricing value to an eligible Getting Started Sub-Capacity Pricing product for the specified hour and LPAR. This file also contains the list of reasons why the assigned Getting Started Sub-Capacity Pricing value might be higher than it would have been had the specified actions been taken.

The record fields in this file are all separated by commas (that is, they are in comma-separated value format) so that the file can be read into and processed by a spreadsheet application. This file can be a sequential file or a member of a PDS or PDSE. If you do not specify a member name for an existing PDS or PDSE, SCRT uses the GSTRACE member name. (The sample SCRT JCL for the z/OS deliverable uses a sequential data set.) A **RECFM** value of FB and **LRECL** value of 200 are internally set by SCRT and require no changes to the JCL for the SCRT job. The internally set values cannot be overridden by the JCL.

For more information about this file, see [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),”](#) on page 165.

IWPTRACE DD statement

The optional IWPTRACE DD statement defines the file that is to contain the list of reasons why the IWP defining program was not eligible for use by the IWP algorithm for the specified hour and LPAR.

The record fields in this file are all separated by commas (that is, they are in comma-separated value format) so that the file can be read into and processed by a spreadsheet application. This file can be a sequential file or a member of a PDS or PDSE, as long as you specify the member name as part of the data set name. (The sample SCRT JCL for the z/OS deliverable uses a sequential data set.) A **RECFM** value of FB and **LRECL** value of 200 are set internally by SCRT and require no changes to the JCL for the SCRT job. The internally set values cannot be overridden by the JCL.

For more information about this file, see [Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing \(IWP\),”](#) on page 183.

VMGUESTS DD statement

The optional VMGUESTS DD statement is only used for z/OS guests running under z/VM. VMGUESTS control statements are not required for and must not be used for z/TPF guests or z/VSE guests. You must specify this statement if you are using SCRT to process SMF records by z/OS guests and **any** of those guests are running z/OS V1.6 or earlier releases. If *all* of your z/OS guests on the same CPC are running z/OS V1.7 or later releases, then you do not need this statement and, if it is present, it is ignored.

When the same z/OS system (same SYSID) runs both natively and as a guest at V1.6 or earlier in the same LPAR, the data for that system is shown at the LPAR level.

The syntax for the VMGUESTS DD statement is:

```
//VMGUESTS DD *
*
lparnam1=sys1,sys2,sys3,...
lparnam2=sys4,sys5,sys6,...
```

where:

lparnam1, lparnam2

The one- to eight-character names you select to represent z/VM LPARs. IBM recommends that you use the actual LPAR names, since these names show up as LPAR names in several sections of the sub-capacity report and because SCRT will list more LPARs than are physically present whenever the same LPAR is used by both a native z/OS system and z/OS guest system at different times during the same reporting period. If you choose arbitrary names, they can be confusing to others in your organization who read the report.

You can continue to specify other z/VM LPARS using names of your choosing (again, use of the actual LPAR name is recommended). Each name must be specified on a separate line. A single LPAR name specification can be continued from one line to the next by breaking the line after a comma (after any z/OS system ID instance).

sys1,sys2,sys3,...

The z/OS system IDs specified in the **SID** parameter in the SMFPRMxx member for each z/OS guest in this z/VM LPAR, when at least one of the systems on that CPC is running z/OS V1.6 or earlier.

Restrictions on DD statements

- You cannot specify a mix of DASD and TAPE data sets within a DD statement.

SCRT messages

SCRT messages are described in [Appendix A, “SCRT messages,”](#) on page 449.

SCRT issues error messages to the STDERR DD specified in the JZOS procedure, and issues other non-error or informational messages to the STDOUT DD specified in the JZOS procedure. By default, the JZOS procedure (JVMPCxx) redirects STDERR and STDOUT to SYSOUT.

Chapter 4. Using SCRT on Windows and Linux systems

SCRT for Windows and Linux uses a graphical user interface (GUI) to guide you through the steps to generate a sub-capacity report. A command line interface (CLI) is also available.

This topic describes the requirements for SCRT on these platforms, how to install and set up SCRT, and how to generate sub-capacity reports.

The requirements for the various SCRT program parameters are identical to those you use when you run SCRT on a z/OS system. Similarly, the requirements for the various pricing metrics and offerings supported by SCRT are the same as when you run SCRT on a z/OS system.

About SCRT for Windows and Linux

IBM provides SCRT for Windows and Linux systems to enable you to run Java-based releases of SCRT on non-z/OS platforms. This is particularly important under the following circumstances:

- Users of the SCRT deliverable for z/VSE systems.
- z/OS users who require a Java-based version of SCRT, but cannot meet the requirements to run SCRT on a z/OS system.

Tip: IBM urges customers who are not subject to these circumstances to use a z/OS-based version of SCRT. This is especially important for customers with large environments that require a great deal of customization of the SCRT parameters. This is primarily due to the large amount of SMF data that SCRT must process in these environments and the additional flexibility offered by the z/OS-based tool.

You should understand that SCRT for Windows and Linux still processes the same SMF or SCRT89 data that was generated on z/OS, z/TPF, or z/VSE systems. The ability to run SCRT on a Windows or Linux platform is simply to provide you with flexibility in producing your sub-capacity reports; no native data is accepted from these platforms.

Unless otherwise noted, SCRT for Windows and Linux supports the same offerings as the Java-based SCRT releases that run on z/OS systems.

Requirements for SCRT for Windows and Linux

Ensure that your intended platform on which to run SCRT meets the minimum technical requirements.

Windows and Linux software prerequisites for SCRT

SCRT can run in any modern Windows or Linux environment (including Linux on IBM Z) that supports an appropriate Java runtime environment. Note that the SCRT installer for Windows and Linux might require a later version of the Java runtime than that which SCRT itself would otherwise require.

The current minimum Java runtime requirement is IBM SDK, Java Technology Edition Version 8 or later with all current maintenance or an equivalent Java runtime from another vendor.

Selecting 64-bit, 32-bit, or Java 18+ support

At installation time, you can select to install 32-bit, 64-bit, or Java 18+ support. The default is 64-bit support. If your existing Java runtime environment is version 18 or later, you should select Java 18+ support. You should only select 32-bit support when another environment is not available. The use of a 32-bit environment can impact the ability to generate SCRT reports in very large environments.

Additional disk space requirements for data

A substantial amount of additional disk space might be required to accommodate all of the SMF or SCRT89 data that SCRT needs to generate sub-capacity reports for your environment. This requirement is in addition to the amount of disk space required by the SCRT program.

Obtaining and installing SCRT for Windows and Linux

Perform this task to obtain and install SCRT for Windows and Linux systems.

About this task

Requirement: IBM requires that you use the latest version of SCRT. Consult [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](http://ibm.biz/SCRTCurrentLevels) for the latest available version of SCRT.

Procedure

1. Download the SCRT deliverable for your target environment.

You can find links to the current version of SCRT in the SCRT Current Levels document.

The SCRT deliverable is contained in a compressed (.zip) file. Note that there are multiple SCRT deliverables, depending on the target environment:

- SCRTWin.zip, for Microsoft Windows systems
- SCRTLinux.zip, for Linux or Linux for z Systems systems

These deliverables are in addition to the SCRT deliverable for z/OS systems.


2. Extract the SCRT installer from the .zip file.
3. If necessary, move the file to the system where you intend to run SCRT.
4. Launch the installer according to your target platform.
 - To launch the installer on a Windows system:
 - a. Navigate to the folder that contains the extracted SCRT installer (installSCRT.exe).
 - b. Double click the installSCRT.exe program.
 - To launch the installer on a Linux or Linux for z Systems system:
 - a. Change to the directory where SCRT was extracted.
 - b. At the command prompt, enter:

```
sh ./installSCRT.bin
```

The installation process begins. A wizard guides you through the installation steps; you will need to select the location where you want to install SCRT.

Results

During the installation process, the installer validates that the minimum required level of Java is available on the system. If an appropriate Java environment is not found, the installer displays the following message:



Installing SCRT in console mode

To install SCRT on a supported Linux system that does not have a graphical environment, you must install in console mode.

Procedure

Issue the following command to invoke the installer in console mode:

```
./installSCRT.bin -i console
```

The console mode installer guides you through the rest of the installation process.

If you have problems installing SCRT in console mode, try the following options:

- Disable X11 forwarding in your SSH client.
- Unset your **DISPLAY** environment variable. Under the bash shell, issue the following command before installing SCRT:

```
unset DISPLAY
```

Transferring SMF or SCRT89 data from a host system

To use SCRT for Windows and Linux, you must copy your SMF or SCRT89 data from your host system to the (non-z/OS) target system where you intend to run the tool.

Guidance for transferring z/OS SMF data

SCRT needs SMF type 70 subtype 1 records and type 89 subtype 1 and 2 records to generate sub-capacity reports.

Tip: Remove other SMF record types from the SMF data being copied. There is no problem with including additional record types; however, a large number of unnecessary records will slow SCRT processing.

When you transfer the SMF data from your mainframe (host) system to the target system where you will run SCRT, observe the following protocols:

- Transfer the SMF data in binary mode.
- Preserve the record descriptor word (RDW).

The particular FTP client software or command that you use to transfer the SMF data varies by platform. The following example uses the FTP command on a Windows workstation:

1. Open a command prompt window.
2. Start an FTP session with your host system:

```
ftp yourHostName
```

3. When prompted, enter your host user ID.
4. When prompted, enter your password.
5. Change to binary transfer mode:

```
binary
```

6. Tell the remote (host) system to preserve the RDW:

```
quote Site RDW
```

7. Transfer (download) the data sets containing your SMF records:

```
get hostFileName localFileName
```

Repeat this step to transfer additional data sets, if necessary.

8. End the FTP session:

```
quit
```

Notes:

- The preceding example assumes that you initiate the FTP session from the receiving workstation. If you initiate the FTP session from the z/OS host, replace step 6 with `locsite RDW` (instead of `quote Site RDW`), and replace step 7 with `put hostFileName remoteFileName` where *remoteFileName* is the file name on the receiving system. The *hostFileName* data set must be a cataloged data set on disk. Note that the RDW parameter does not seem to support temporary data sets.
- If you wish to FTP the SMF data from tape, step 6 requires an additional parameter:

```
quote Site RDW readtapeformat=s
```

- All of the SMF records do not need to be in one file. SCRT can combine the data from multiple files to produce a report.

Support for SMF data without RDWs

SCRT provides limited support for processing SMF data that has been transferred from a z/OS system to a Windows or Linux workstation without preserving the record descriptor word (RDW) in the SMF records. This support allows for the use of transfer protocols other than FTP that do not support preserving the RDW.

Tip: Whenever possible, use a transfer protocol, such as FTP, that preserves the RDW, as the presence of the RDW enables more robust parsing of the SMF records.

For SCRT to process SMF records without RDWs, the SMF data must consist of only the following record types:

- SMF type 2
- SMF type 3
- SMF type 70, subtype 1
- SMF type 89, subtypes 1 and 2

SCRT attempts to automatically detect whether the input records include RDWs and processes them accordingly. However, if SCRT incorrectly determines the presence or absence of RDWs for the records in a given data file, you can use the following file naming conventions to explicitly specify the status of the RDWs to SCRT:

- To explicitly specify that the SMF records in a file do not include RDWs, change the file extension from `.SMF` to `.NORDW`.

- To explicitly specify that the SMF records in a file include RDWs, change the file extension from .SMF to .RDW.

Guidance for transferring SCRT89 data

The blocked format record files that contain SCRT89 data originating from z/VSE systems must have file names with a .bin, .scrt89, or .vse extension. Files that contain SCRT89 data originating from z/TPF systems must have file names with a .smf extension.

When you transfer these files from your mainframe (host) system to the target system where you will run SCRT, observe the following protocols:

- Transfer the SCRT89 data in binary mode.
- For transfer from z/VSE systems, transfer the SCRT89 data in block mode (MODE B).

The particular FTP client software or command that you use to transfer the SMF data varies by platform. The following example uses the FTP command on a Windows workstation:

1. Open a command prompt window.
2. Start an FTP session with your host system:

```
ftp yourHostName
```

3. When prompted, enter your host user ID.
4. When prompted, enter your password.
5. Change to binary transfer mode:

```
binary
```

6. For z/VSE systems, change to block transfer mode:

```
quote mode B
```

7. Transfer (download) the data sets containing your SMF records:

```
get hostFileName localFileName
```

Repeat this step to transfer additional data sets, if necessary.

8. End the FTP session:

```
quit
```

For more information about transferring files containing SCRT89 records from z/TPF and z/VSE systems, see:

- "Scenario 2 — Using FTP" in [“Steps for z/TPF data collection and processing”](#) on page 284
- [Figure 163 on page 318](#) in [“Steps for processing SCRT89 data on a z/OS system”](#) on page 317

Using the SCRT graphical user interface

In order to generate sub-capacity reports on Windows or Linux systems, you must:

1. Copy the SMF or SCRT89 data from your host system to the target Windows or Linux system on which you intend to run SCRT, as described in [“Transferring SMF or SCRT89 data from a host system”](#) on page 71.
2. Run the SCRT for Windows and Linux graphical user interface (GUI), as described in [“Running the SCRT graphical user interface”](#) on page 74.

For details about the sub-capacity report output for the various pricing metrics and offerings that SCRT supports, see the appropriate chapter for that offering.

Running the SCRT graphical user interface

SCRT for Windows and Linux uses a common graphical user interface (GUI) across all the platforms on which it runs, including Windows, Linux, and Linux for z Systems.

SCRT also provides a command line interface, described in [“Using the SCRT command line interface” on page 90](#).

The steps to start the SCRT GUI differ based on the platform on which it runs.

Starting the SCRT GUI on a Windows system

1. From the **Start** menu, navigate to the IBM Sub-Capacity Reporting Tool program group.
2. Click **SCRT** to start SCRT for IBM reporting or **SCRT for ISV** to start SCRT for ISV reporting.

Starting the SCRT GUI on a Linux or Linux system

1. Navigate (**cd**) to the directory where you installed SCRT.
2. Start SCRT in one of the following ways:

- To start SCRT for IBM reporting, enter:

```
./SCRT
```

- To start SCRT for ISV reporting, enter:

```
./SCRT\ for\ ISV
```

The SCRT user profile

SCRT maintains a profile to store user and environment information between program executions.

- On Windows systems, the profile is stored in the %APPDATA%\SCRT directory.
- On Linux systems, the profile is stored in the \$XDG_CONFIG_HOME/SCRT directory. If the **XDG_CONFIG_HOME** environment variable is not set, the profile is stored in the \$HOME/.config/SCRT directory.

Prior to SCRT V24.11.0, on Linux systems, the profile was stored in the user's home directory. To use a profile from an earlier SCRT release, copy the profile (`profile.txt`) into the new location prior to starting SCRT.

Notes:

1. Only the values that were used during the previous program execution are saved.
2. The profile is not saved until report generation time.

Navigating the SCRT panels

SCRT uses a series of dialog panels to guide you through the report generation process—a wizard-based approach.

This topic describes the wizard process to help you understand all of the required steps to generate sub-capacity reports or multiplex reports.

Customer parameter panel

The customer parameter panel appears when you start SCRT.

On this panel, you enter information similar to that on the PARMS DD statement in the JCL for the SCRT program that runs on z/OS. For detailed information about the parameter requirements, see [“PARMS DD statement” on page 34](#).

Figure 1 on page 75 shows an example of the customer parameter panel. (When running SCRT for ISV reporting, the customer parameter panel also contains the **ISV Library File** field where you can specify the vendor-supplied library file.)

Figure 1. SCRT for Windows and Linux: Customer parameter panel

Fill in the parameter values as described in [“PARMS DD statement”](#) on page 34. When finished, click **Next**.

After SCRT successfully generates a report using these parameter values, it saves the values in the user profile. When running SCRT for ISV reporting, the parameters are saved in a vendor-specific profile. The next time you start SCRT, the program automatically populates the customer parameter panel with the saved values.

SMF or SCRT89 data selection panel

On the SMF or SCRT89 data selection panel, specify one or more files that contain all of the SMF or SCRT89 data (or both) that SCRT is to process. Your selection must include all of the files necessary for SCRT to generate a sub-capacity report or multiplex report; it is analogous to the SMF DD statement in the JCL for the SCRT program that runs on z/OS.

Figure 2 on page 76 shows an example of the SMF or SCRT89 data selection panel.



Figure 2. SCRT for Windows and Linux: SMF or SCRT89 data selection panel

Important: SCRT expects the SMF or SCRT89 data to retain a specific data format. See [“Transferring SMF or SCRT89 data from a host system” on page 71](#) for instructions about copying data to your target system so as to preserve the data format.

Click **Add** to add data files to the list. If necessary, click **Remove** to remove files. When finished, click **Next**.

When you click **Next**, SCRT performs an initial scan of the data to identify all of the CPCs, LPARs, and operating systems that are in use in the input environment. Large data files might require several minutes to complete before advancing to the next panel.

Notes:

- Data that is transmitted from a z/OS system must have file names with a `.smf` extension.
- Data that is transmitted from a z/VSE system in blocked mode can have file names with a `.bin`, `.scrt89`, or `.vse` extension.

Customer number panel

Use the customer number panel to assign additional customer numbers to the CPCs that SCRT found in the SMF or SCRT89 data. By default, SCRT assigns the customer number that you entered on the customer parameter panel to all of the CPCs.

[Figure 3 on page 77](#) shows an example of the customer number panel.

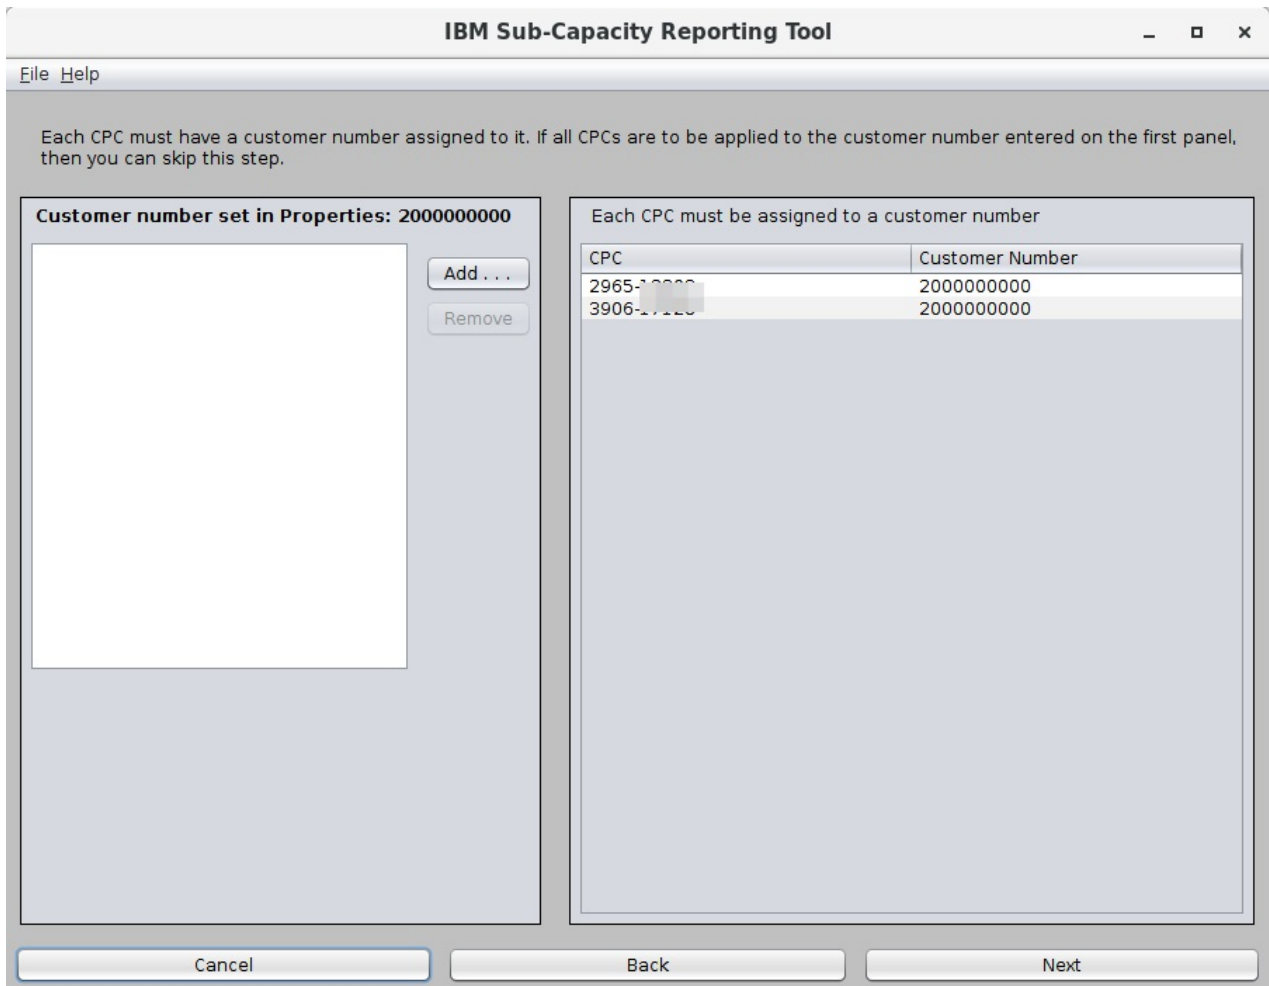


Figure 3. SCRT for Windows and Linux: Customer number panel

To enter additional customer numbers, click **Add**.

After entering alternate customer numbers, you can assign them to the CPCs listed in the **CPC** column. To do so, click inside the **Customer Number** cell for the desired CPC and select a customer number from the drop-down list.

When finished, click **Next**.

Reporting period panel

The reporting period panel allows you to select the reporting period that SCRT is to use to evaluate all of the input data.

Figure 4 on page 78 shows an example of the reporting period panel.

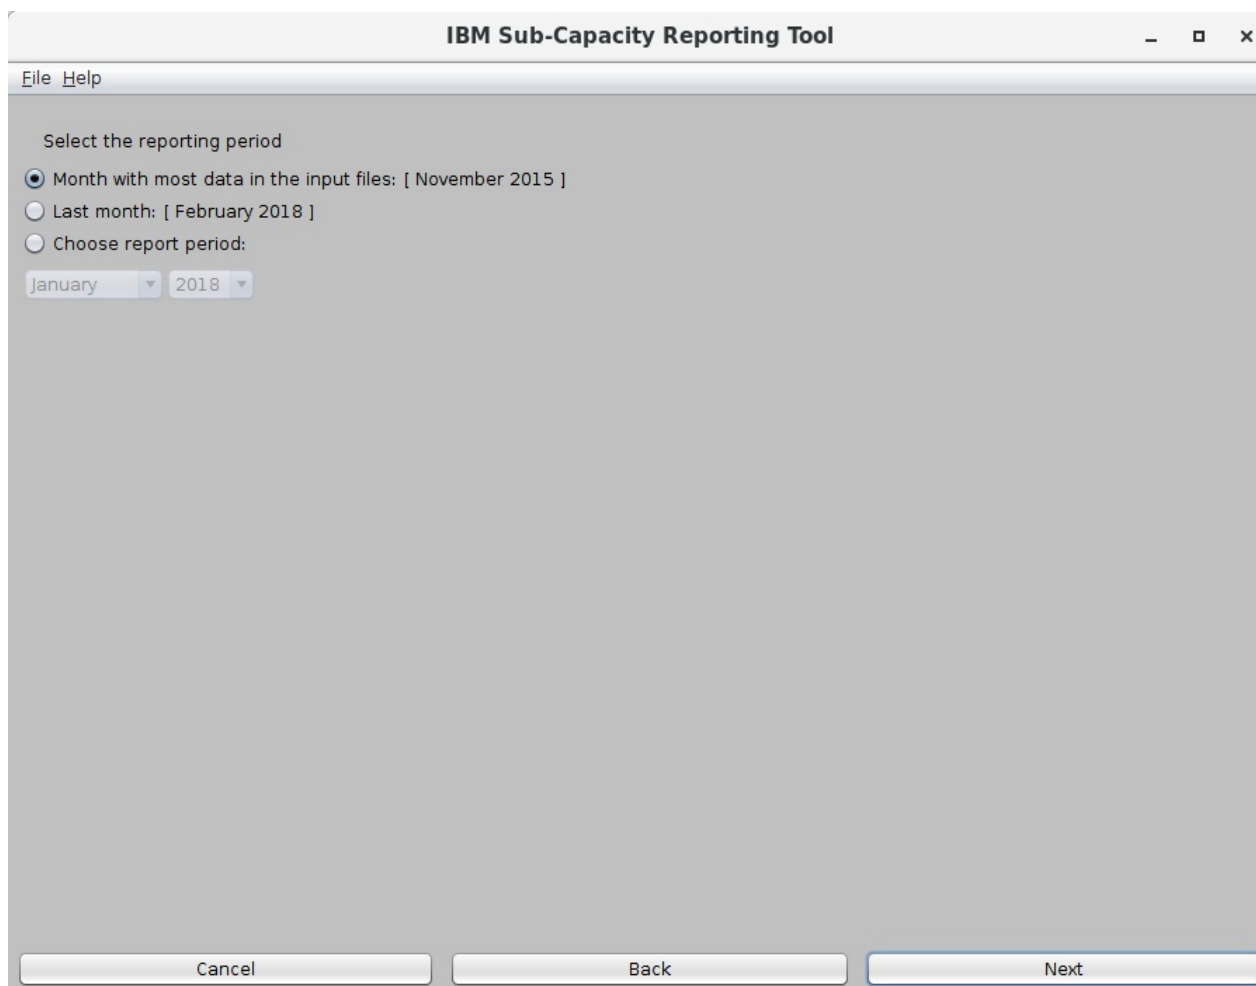


Figure 4. SCRT for Windows and Linux: Reporting period panel

You have three options to select the reporting period:

- **Month with the most data in the input files**

This option is analogous to running SCRT on a z/OS system without specifying a reporting period. SCRT uses the month that is most represented in the input data, as indicated in brackets.

- **Last month**

The option uses the month prior to the current calendar month.

- **Choose report period**

This option allows you to specify any reporting period. Note that if you select a reporting period that is not represented in the input data, SCRT will not generate any reports.

Select the desired option, then click **Next**.

NO89 product selection panel

Use the NO89 product selection panel to select the NO89 products that you use and specify the location where they run. This is a two-step process, as represented by the two tabs on the panel: the **NO89 Product List** tab and the **Product Assignment** tab.

Figure 5 on page 79 shows an example of the NO89 product selection panel with the **NO89 Product List** tab selected.




Figure 5. SCRT for Windows and Linux: NO89 product selection panel — NO89 Product List tab

On the **NO89 Product List** tab, select a NO89 product. SCRT automatically switches to the **Product Assignment** tab so that you can assign that NO89 product to the correct LPARs. Repeat this process for each NO89 product that you use.

Figure 6 on page 80 shows an example of the NO89 product selection panel with the **Product Assignment** tab selected.

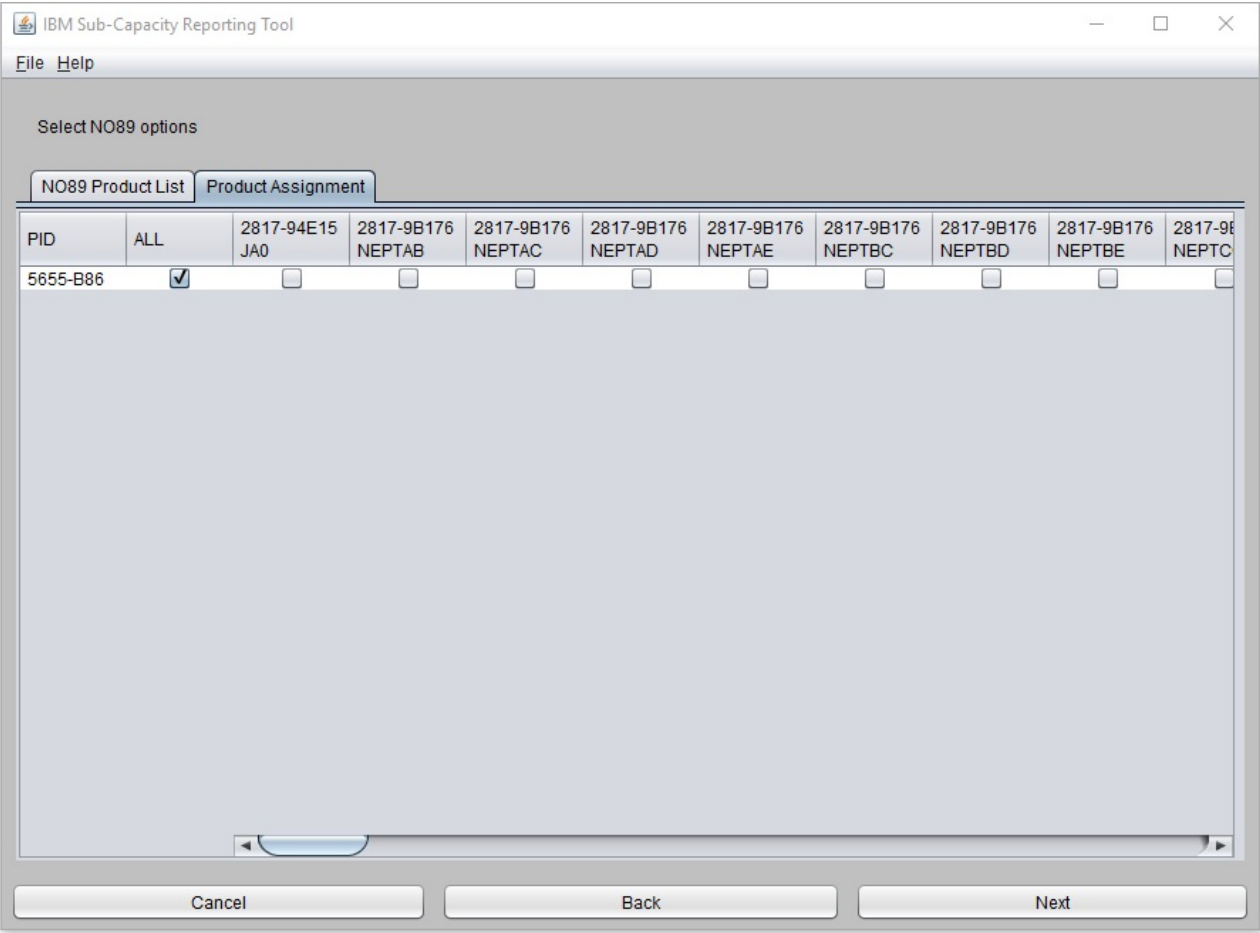


Figure 6. SCRT for Windows and Linux: NO89 product selection panel – Product Assignment tab

The **Product Assignment** tab lists the NO89 products that you selected and allows you to assign the products to the specific LPARs on which they run. Click the check boxes to select or clear the LPARs on which each NO89 product ran during the reporting period. When finished, click **Next**.

Upon generating a report, SCRT saves the selections you made on the NO89 product selection panel and will populate the panel with these selections the next time you run the tool.

For more information about selecting NO89 products, see [“NO89 DD statement” on page 61](#).

Processing options panel

The processing options panel contains two tabs: **Input CSV** and **Processing Options**.

Note: This panel is not applicable and does not appear when using SCRT for ISV reporting.

Use the **Input CSV** tab to specify one or more .csv files to provide input data for Mobile Workload Pricing (MWP), IBM Z Collocated Application Pricing (zCAP), or IBM Z Workload Pricing for Cloud (zWPC).

[Figure 7 on page 81](#) shows an example of the **Input CSV** tab.

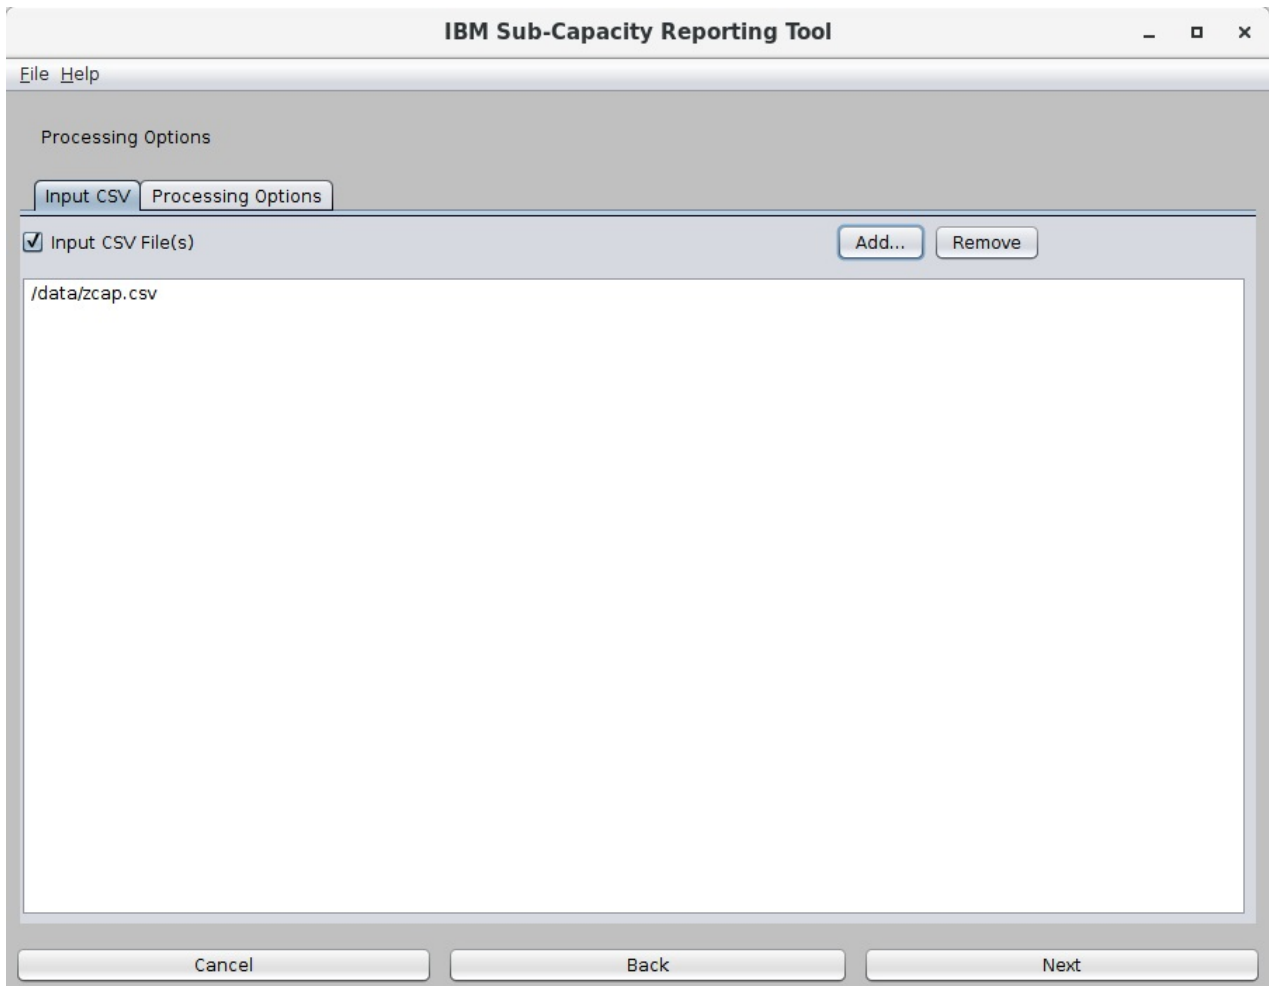


Figure 7. SCRT for Windows and Linux: Input CSV tab

To submit MWP or zCAP pricing data, you must first select the **Input CSV File(s)** check box. Then, click **Add** to locate and specify the necessary input files. Note that the format and rules for using these files are the same as for using the INPUTCSV DD statement in the JCL for the SCRT program that runs on z/OS. For details, see the information about using the INPUTCSV DD in the appropriate chapter for MWP or zCAP.

Use the **Processing Options** tab to select the report type and any pricing offerings that you might require for SCRT processing. [Figure 8 on page 82](#) shows an example of the **Processing Options** tab.

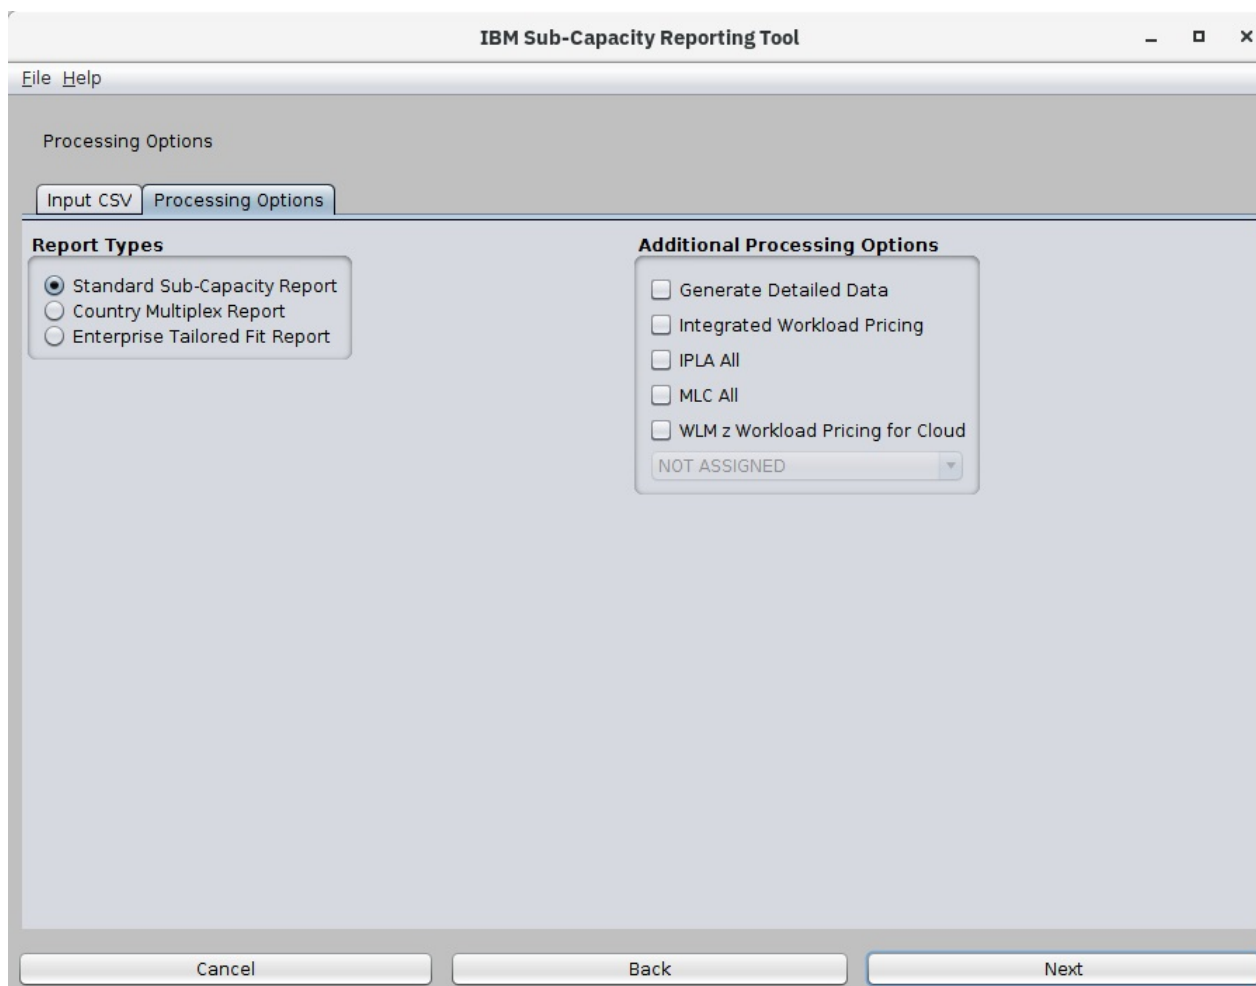


Figure 8. SCRT for Windows and Linux: Processing Options tab

To select the report type, select one of the **Report Types** buttons. You can select a standard sub-capacity report, a Country Multiplex report (for IBM Country Multiplex Pricing), or an Enterprise Tailored Fit report (for Tailored Fit Pricing for IBM Z - Software Consumption Solution). The report types are mutually exclusive; you can select only one.

To select a special pricing option, select the **Additional Processing Options** check box for that option. You can select one or more options, as appropriate. For more information about these offerings, see [“SPECIAL DD statement” on page 40](#) or the appropriate chapter about that offering.

When IBM Z Workload Pricing for Cloud (zWPC) is in use, select the **WLM z Workload Pricing for Cloud** check box and select the appropriate WLM reporting attribute from the drop-down list. This is analogous to specifying the **Assign ZWPC** control statement on the SPECIAL DD statement. See [“Assign ZWPC SPECIAL DD statement” on page 258](#) for more information.

When finished, click **Next**.

Detected solution IDs panel

The detected solution IDs panel lists all of the solution IDs that were assigned to tenant resource groups in the input data and allows you to update solution IDs.

[Figure 9 on page 83](#) shows an example of the detected solution IDs panel.

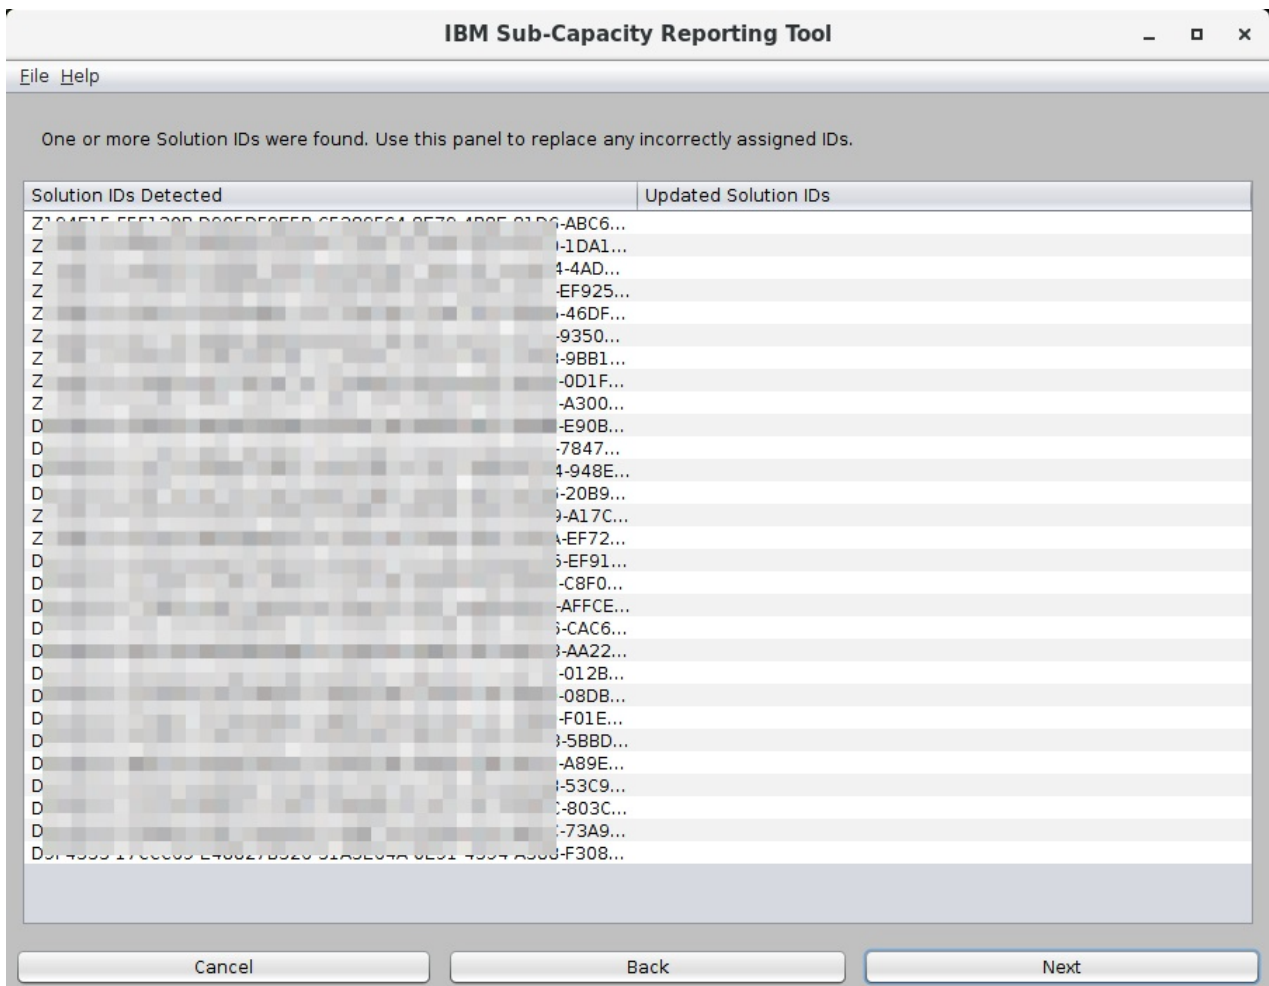


Figure 9. SCRT for Windows and Linux: Detected solution IDs panel

To update a solution ID in the **Solution IDs Detected** column, enter the new solution ID on the corresponding row in the **Updated Solution IDs** column. This is analogous to using the **UPDATE SOLUTION_ID** control statement in the SPECIAL DD in the SCRT JCL or using the **--update-solution-id** option in the SCRT CLI. For more information, see [Chapter 19, “Reporting for Tailored Fit Pricing for IBM Z,”](#) on page 375.

When finished, click **Next**.

Container settings assignment panel

On the container settings assignment panel, you can:

- Assign solution names to the containers that SCRT detected in the input data.
- Assign an entire LPAR (or LPARs) to be part of a container. You can assign an LPAR to an existing container or to a new container.
- Ignore individual containers on a per-hour basis for purposes of SCRT reporting.

[Figure 10 on page 84](#) shows an example of the container settings assignment panel.

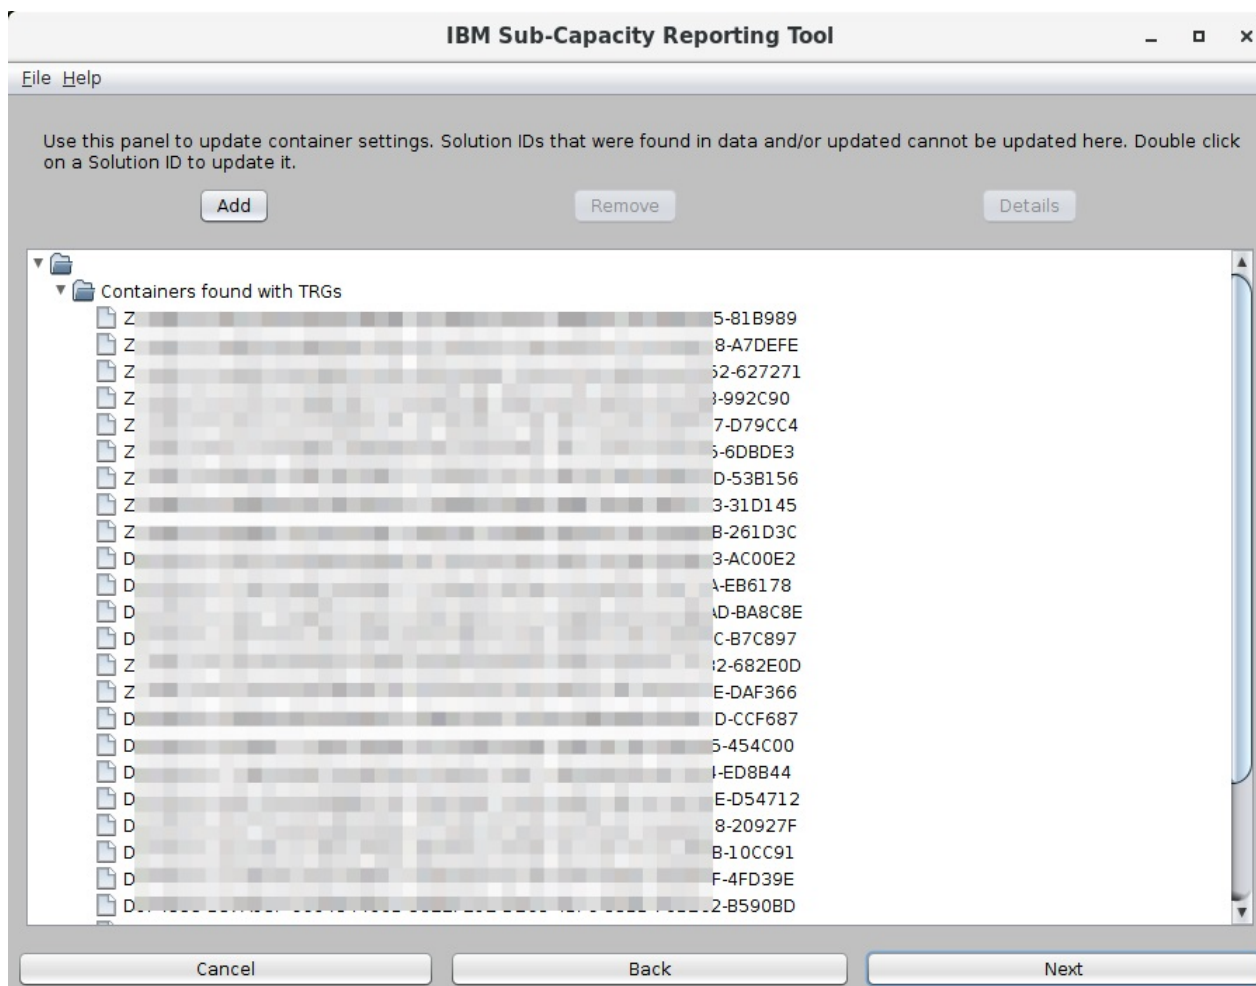


Figure 10. SCRT for Windows and Linux: Container settings assignment panel

- To assign a name to a container, double click a solution ID to launch the Update Container dialog and enter the container name. This is analogous to using the **UPDATE CONTAINER** control statement in the SPECIAL DD in the SCRT JCL or using the **--update-container** option in the SCRT CLI.
- To assign an entire LPAR to an existing container, double click the desired solution ID to launch the Update Container dialog and select the LPAR (or LPARs) from the table. To assign an entire LPAR to a new container, click **Add** and use the Update Container dialog to set the solution ID and select the LPAR (or LPARs) from the table. This is analogous to using the **CONTAINER** control statement in the SPECIAL DD in the SCRT JCL or using the **--container** option in the SCRT CLI.
- To ignore individual containers for SCRT reporting, double click the desired solution ID to launch the Update Container dialog and use the **Ignore Container** tab to add, edit, or remove the time intervals that you want SCRT to ignore. This is analogous to using the **IGNORE CONTAINER** control statement in the SPECIAL DD in the SCRT JCL or using the **--ignore-container** option in the SCRT CLI.
- To assign NO89 products to a container, double-click the desired solution ID to launch the Update Container dialog, and use the **NO89 Products** tab to add products to the container. This is analogous to using the **NO89_PRODUCT_ID** keyword on the **UPDATE CONTAINER** control statement in the SPECIAL DD in the SCRT JCL or using the **--update-container** option in the SCRT CLI.

For more information, see [Chapter 19, “Reporting for Tailored Fit Pricing for IBM Z,”](#) on page 375.

When finished, click **Next**.

Exclude panel

Use the Exclude panel to direct SCRT to exclude the processing of certain data that is reported on SMF or SCRT89 records. This panel is analogous to the **Exclude** control statement on the SPECIAL DD statement in the JCL for the SCRT program that runs on z/OS, and the same rules and restrictions apply here. See “[EXCLUDE control statement](#)” on page 44 for details.

Figure 11 on page 85 shows an example of the Exclude panel.

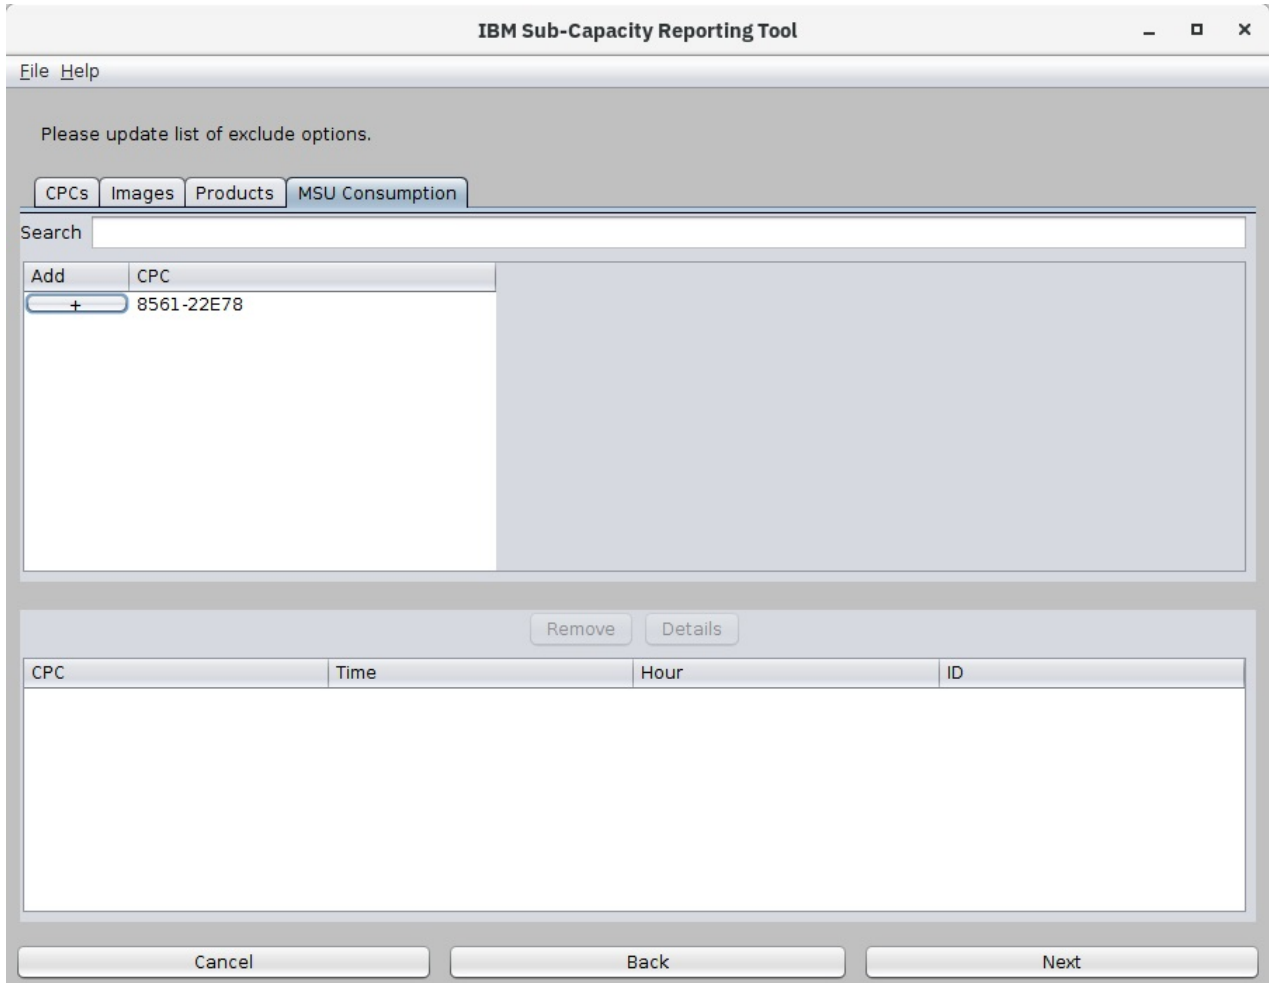


Figure 11. SCRT for Windows and Linux: Exclude panel

The Exclude panel has four tabs: **CPCs**, **Images**, **Products**, and **MSU Consumption**. The first three tabs allow you to begin specifying an Exclude statement at the CPC, image, or product level. The **MSU Consumption** tab allows you to begin specifying an Exclude statement for excluding an amount of consumed MSUs.

To create an Exclude statement that is not based on MSU consumption:

1. Select the appropriate tab. A list of items of the selected type (CPCs, images, or products, depending on which tab you selected) appears, as found in the input data.
2. To exclude an item, click the check box in the **Add** column for that item. This generates an exclude statement, which appears in the list of exclude statements.
3. To further customize the Exclude statement, either double-click it or select it and click **Details**. The **Set Exclude options** window appears, as shown in [Figure 12 on page 86](#).

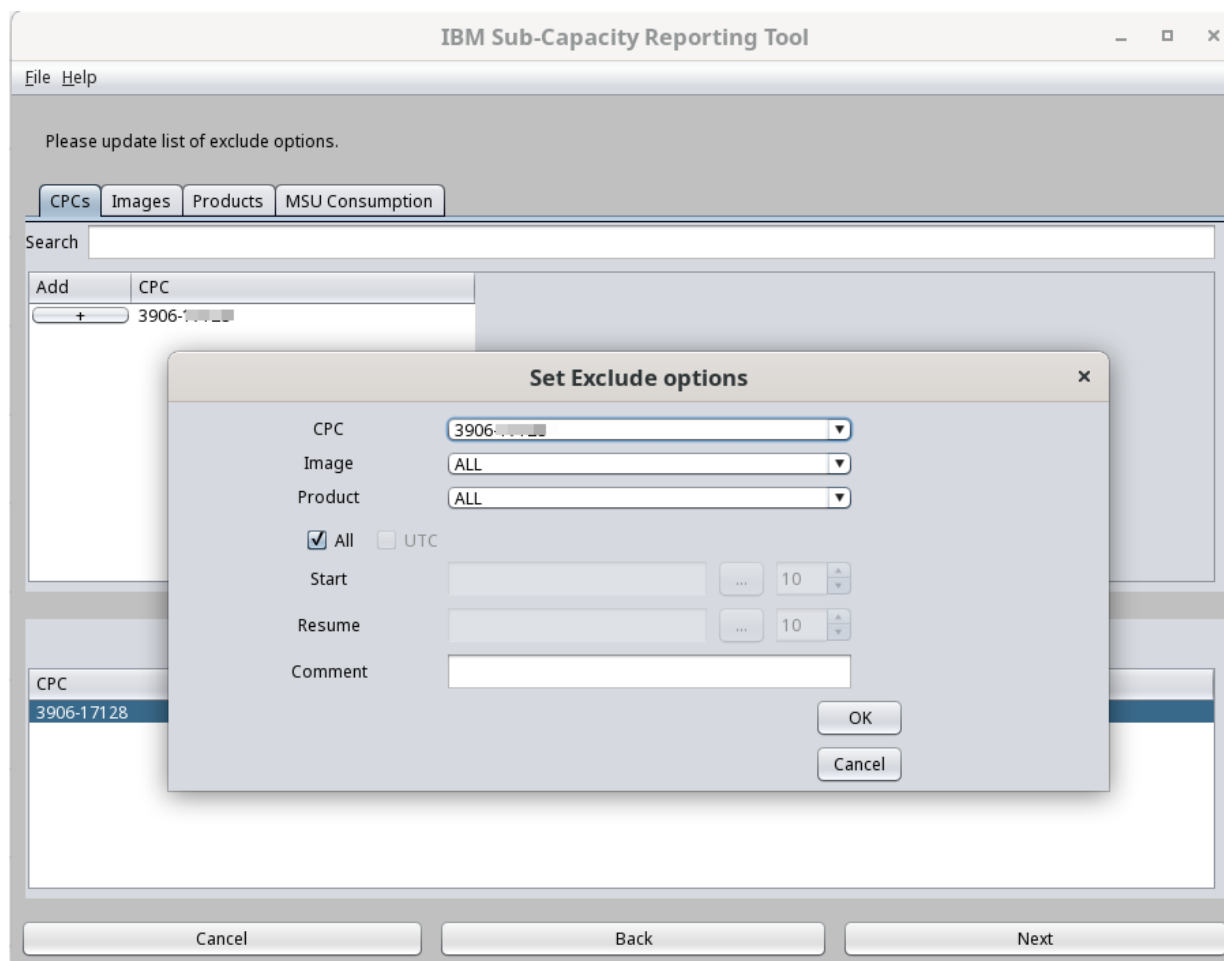


Figure 12. SCRT for Windows and Linux: Exclude panel – Set Exclude options

4. Customize the Exclude statement options by providing additional information, including the **Start** and **Resume** dates and hours.

To specify the entire reporting period, select the **All** check box. To indicate that the start and resume values are in UTC, select the **UTC** check box.

5. Click **OK**.
6. If necessary, repeat steps 1 - 5 to create additional Exclude statements.
7. When finished, click **Next**.

To create an Exclude statement for MSU consumption:

1. Select the **MSU Consumption** tab. A list of CPCs appears, as found in the input data.
2. To exclude MSU consumption from a Tailored Fit Pricing solution on a CPC, click the check box in the **Add** column for that CPC. The Set Exclude Options dialog is displayed, as shown in [Figure 13 on page 87](#).
3. On the Set Exclude Options dialog, select the solution ID to which the exclude should apply, enter the amount of CPU time, in seconds, to exclude, and select a date and hour at which the incident being excluded occurred. Click **OK** to continue.

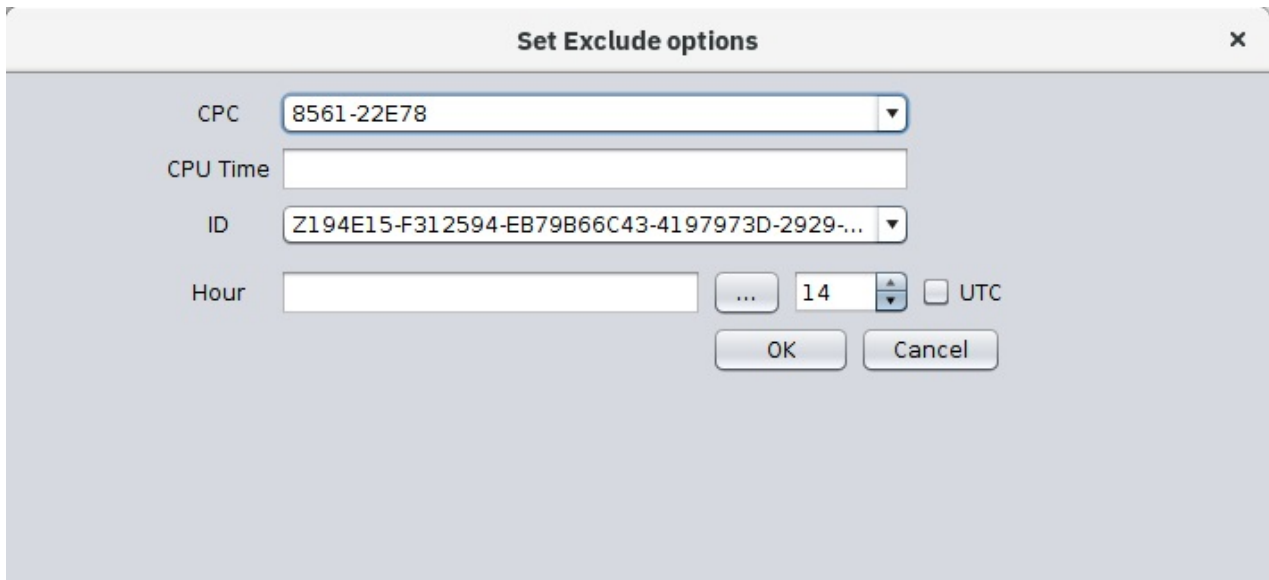


Figure 13. SCRT for Windows and Linux: Exclude panel — Set Exclude options for MSU consumption

Report preview panel

The report preview panel is the last panel before SCRT generates a report. The panel provides a summary of all of the processing options that you requested on the preceding panels.

[Figure 14 on page 88](#) shows an example of the exclude panel.

IBM Sub-Capacity Reporting Tool

File Help

Last step before we launch the calculation and report generation.

Review Input Data Trace Options

====+ Customer Information +====

Profile name:
 Customer name: Customer
 Customer number: 2000000000
 Contact name: Mr. Customer
 Contact email: customer@example.com
 Contact phone: 555-555-5555

====+ SMF File Records Data +====

SMF Filepaths: /data/z2964.smf
 /data/z3906.smf
 OS: ZOS
 Reporting period: Nov 2015
 Processing options: Input CSV /data/zcap.csv

====+ CPC/Customer Numbers +====

CPC	Customer Number
2965	2000000000
3906	2000000000

====+ N089 Products +====

PRODUCT ID	CPC	LPAR

Output file path report.csv Browse

Cancel Back Generate report

Figure 14. SCRT for Windows and Linux: Report preview panel

The report preview panel provides the option to request detailed report output. This is the same function as provided by the **Generate_Detailed_Data** control statement in the SPECIAL DD statement in the JCL for the SCRT program that runs on z/OS. To request a detailed report, select the **Generate detailed report** check box.

In the **Output file name** field, specify the fully-qualified path and file name where SCRT is to write the report output.

Optionally, to enable GSTRACE or CPSTRACE output, go to the **Trace Options** tab, select the **GSTRACE** or **CPSTRACE** checkbox, and specify the path and file name to which SCRT is to write the trace output, as shown in [Figure 15 on page 89](#).

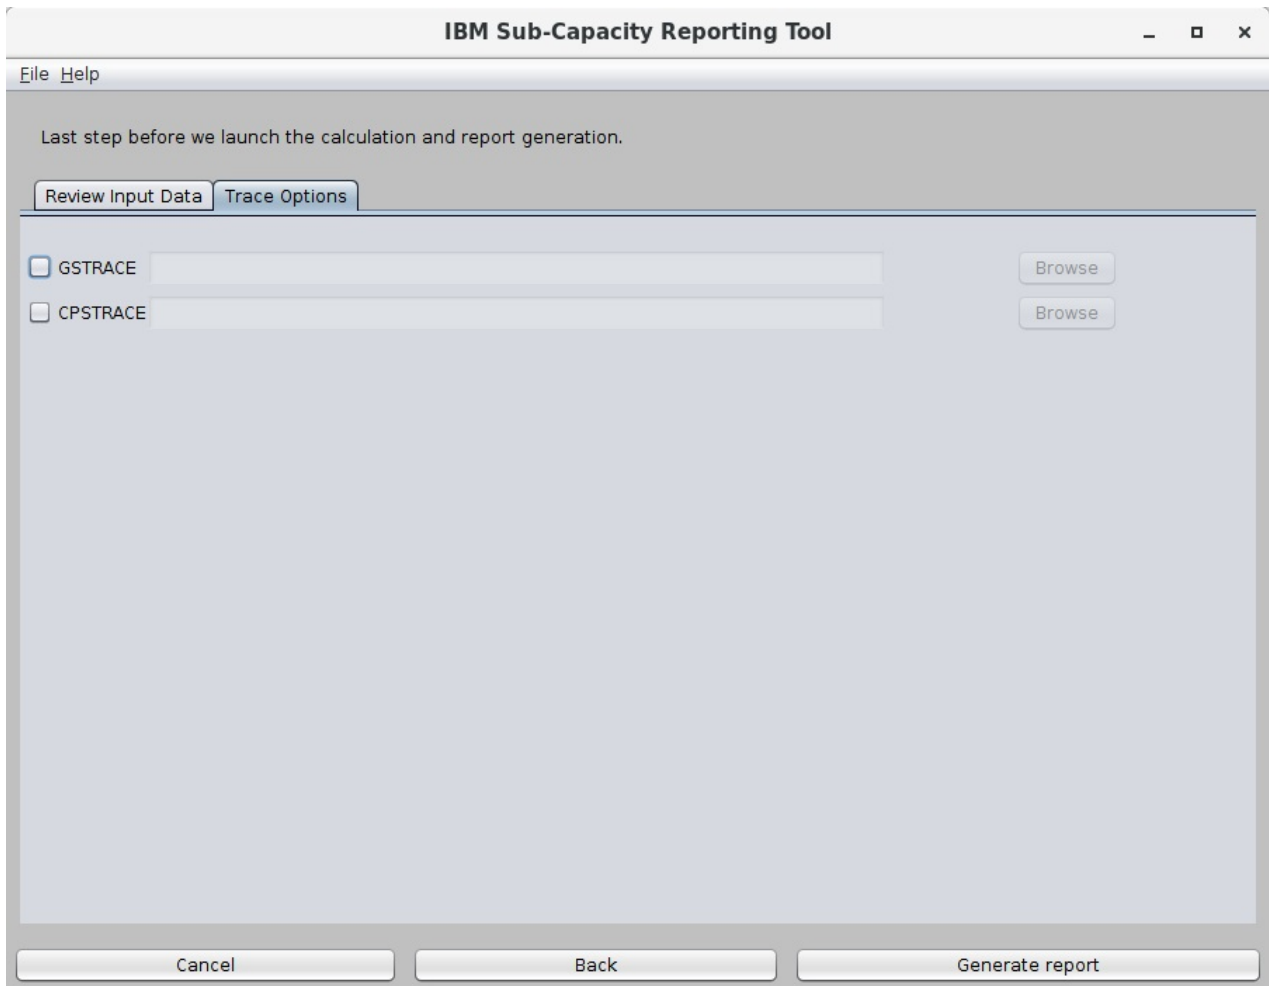


Figure 15. SCRT for Windows and Linux: Trace Options tab on the report preview panel

Click **Generate report** to begin SCRT processing. The processing time, before advancing to the next panel, depends on the volume of SMF or SCRT89 input data to be processed.

Report results panel

The report results panel appears when SCRT processing is complete and report output has been generated successfully. The panel provides information about the report generation process.

[Figure 16 on page 90](#) shows an example of the exclude panel.

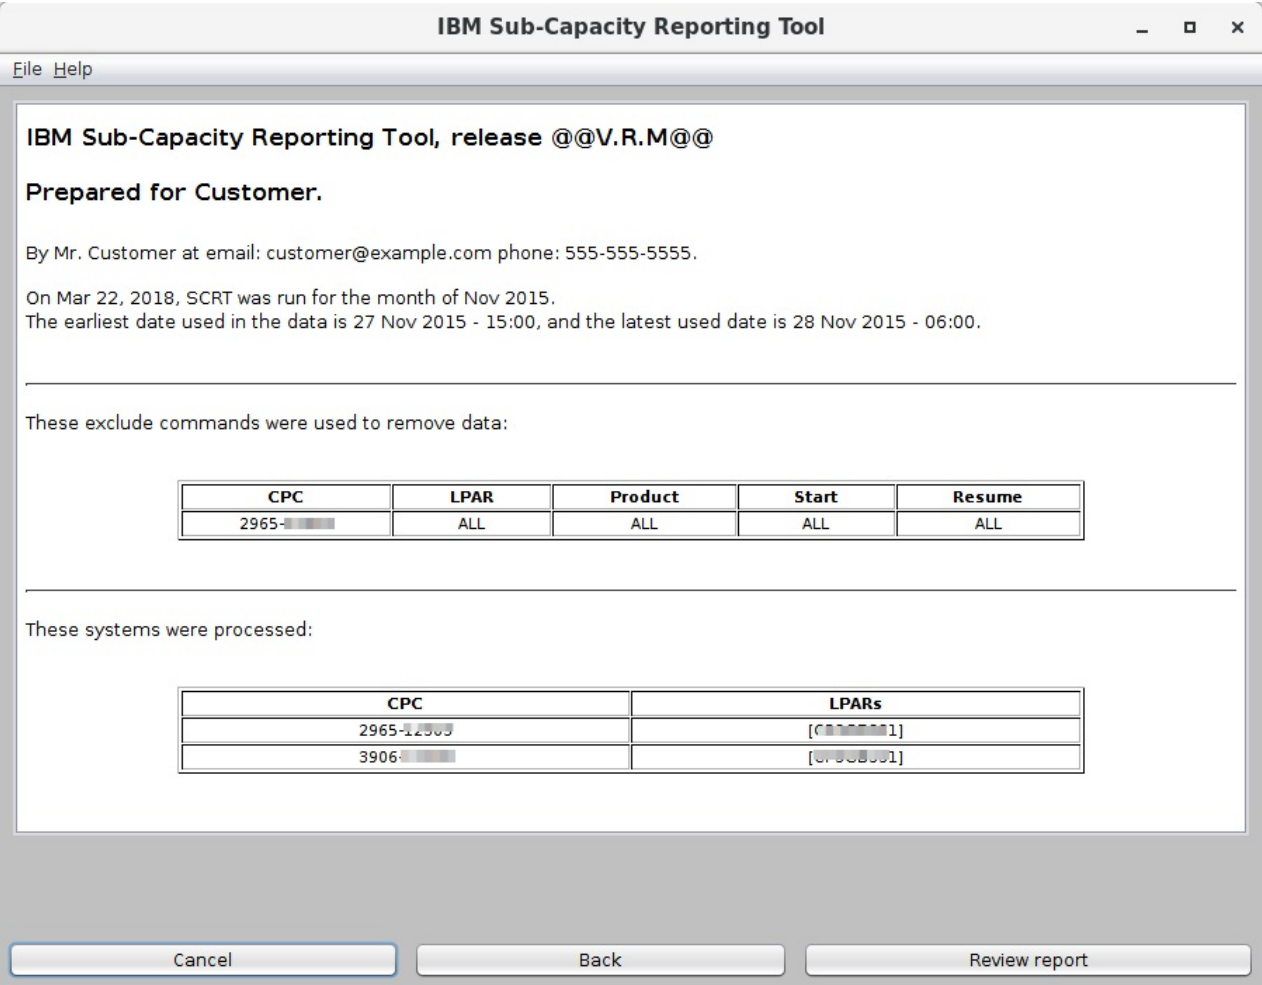


Figure 16. SCRT for Windows and Linux: Report results panel

Note: On Windows-based systems, the report results panel includes a **Review report** button that opens the generated report. This button is not present on Linux-based systems.

The SCRT Wizard is now concluded. Click **Cancel** to exit.

Proceed with the same report submission process as for reports generated on z/OS.

Using the SCRT command line interface

SCRT supports a command line interface (CLI) on Windows and Linux platforms.

The CLI supports the same SCRT processing options as the Windows and Linux graphical user interface (GUI).

Note: The CLI is included when you install the SCRT deliverable for Windows or Linux, as described in [“Obtaining and installing SCRT for Windows and Linux”](#) on page 70.

Prerequisites for using the command line interface

Observe the following prerequisites for running the SCRT CLI:

- SCRT must be installed on a Windows or Linux system.
- SMF or SCRT89 data sets must be copied to the Windows or Linux system.

Running the SCRT command line interface

Issue the appropriate command to invoke the SCRT CLI on a Linux or Windows system.

Optionally, SCRT also provides sample scripts (a .bat file for Windows systems and a shell script for Linux systems) in the `samples` directory which you can customize and use to invoke the CLI.

Invoking the command line interface on Linux systems

1. From your local shell, navigate to the SCRT installation directory.
2. Issue the following command to invoke the command-line utility and specify the processing options, followed by one or more paths to the SMF or SCRT89 input data sets:

```
./scrtc options {SMFpaths|SCRT89paths}
```

Invoking the command line interface on Windows systems

1. From a Windows command prompt, navigate to the SCRT installation directory.
2. Issue the following command to invoke the command-line utility and specify the processing options, followed by one or more paths to the SMF or SCRT89 input data sets:

```
scrtc options {SMFpaths|SCRT89paths}
```

Specifying command line options

The SCRT command line utility accepts a number of options that provide the same SCRT processing options as the Windows and Linux GUI interface.

- Many command line options have equivalent short and long forms. You can specify either form and you can mix forms between options.
 - Long form options begin with two hyphens (- -).
 - Short form options begin with one hyphen (-).
 - If an option has no short form, you must specify the long form.
- Most command line options take an argument. Use a blank character to separate the option from the argument.
- If an argument contains blank characters or characters that have special meaning to the operating system or shell, enclose the argument in double quotation marks (").

Table 7 on page 91 describes the command line options.

Table 7. SCRT command line options

Long form	Short form	Required	Description
--customer-name	-c	yes	These options correspond to their equivalent PARS DD control statements. For descriptions of these options and syntax rules for the argument values, see “PARMS DD statement” on page 34.
--customer-number	-n	yes	
--contact-name	-u	yes	
--contact-email	-e	yes	For details about the --customer-number option, see “Specifying customer number assignments” on page 93.
--contact-phone	-p	yes	For details about the --lpar-comment option, see “Specifying LPAR comments” on page 95.
--po-number	<i>none</i>	no	
--lpar-comment	-l	no	
--report-period	-r	no	

Table 7. SCRT command line options (continued)

Long form	Short form	Required	Description
--report-time	<i>none</i>	no	Direct SCRT to report dates and times from z/OS systems on the sub-capacity report using either a special form of local time or UTC time. Valid values are SPECIAL_LOCAL_TIME or UTC_TIME.
--no89	-a	yes	Specifies an input file that contains the formatted list of NO89 products and their designations. For details, see “Specifying NO89 statements” on page 93.
--output	-o	yes	Specifies an output file for the SCRT report.
--wcsupp	<i>none</i>	no	Specifies an output file for the Workload Classification Supplement. For more information, see Chapter 22, “Workload Classification Supplement,” on page 409.
--container	<i>none</i>	no	For Tailored Fit Pricing for IBM Z, associates an LPAR with a container solution ID. For details, see “Associating an LPAR with a solution ID” on page 95.
--country-multiplex-pricing	-m	no	Requests the use of the Country Multiplex Pricing (CMP) option.
--cpstrace	<i>none</i>	no	For Tailored Fit Pricing for IBM Z, designates an output file for the container pricing detail report.
--enterprise-tailored-fit-pricing	<i>none</i>	no	Requests an Enterprise Tailored Fit report. For details, see “Requesting an Enterprise Tailored Fit report” on page 96.
--exclude	-x	no	Specifies one or more Exclude control statements to exclude processing of certain data that is reported in SMF or SCRT89 records. For details, see “Specifying Exclude statements” on page 94.
--include	<i>none</i>	no	Specifies one or more INCLUDE control statements to specify a subset of images to be reported on by SCRT. For details, see “Specifying INCLUDE statements” on page 94.
--feature	<i>none</i>	no	Requests to enable an SCRT feature.
--generate-detailed-data	-d	no	Specifies to generate a detailed report under the conditions described in “GENERATE_DETAILED_DATA control statement” on page 52.
--gstrace	<i>none</i>	no	Designates an output file for the IBM Getting Started Sub-Capacity Pricing for z/OS IPLA software (GSSP) trace report.
--ignore-container	<i>none</i>	no	For Tailored Fit Pricing for IBM Z, designates a period of time for which SCRT is to ignore a container. For details, see “Discarding container-related information” on page 96.
--inputcsv	-i	no	Designates an input file that contains data for Mobile Workload Pricing (MWP), IBM Z Workload Pricing for Cloud (zWPC), or IBM Z Collocated Application Pricing (zCAP). For details, see “Specifying input data for special offerings” on page 94.
--integrated-workload-pricing	<i>none</i>	no	Requests Integrated Workload Pricing.
--ipla-all	<i>none</i>	no	Requests that SCRT report an IPLA family line even when only one product in the product family is in use.
--isv	<i>none</i>	no	For ISV reporting, specifies an ISV library file and causes SCRT to generate an ISV-specific report. For details, see “Generating ISV reports” on page 96.
--isv-show-containers	<i>none</i>	no	For ISV reporting, requests container processing. By default, SCRT does not process containers in ISV mode. For details, see “Processing containers in ISV mode” on page 96.

Table 7. SCRT command line options (continued)

Long form	Short form	Required	Description
--iwptrace	<i>none</i>	no	Designates an output file for the Integrated Workload Pricing trace report.
--mlc-all	<i>none</i>	no	Requests that SCRT report an MLC family line even when only one product in the product family is in use.
--update-container	<i>none</i>	no	For Tailored Fit Pricing for IBM Z, updates the attributes of a container, such as the container name. For details, see “Updating container properties” on page 95 .
--update-solution-id	<i>none</i>	no	For Tailored Fit Pricing for IBM Z, replaces the specified solution ID with a newly specified solution ID. For details, see “Replacing one solution ID with another” on page 96 .
--ztpf-transformation-engines-offering	<i>none</i>	no	Requests that SCRT report on the z/TPF Transformation Engines Offering.
--zwpc	<i>none</i>	no	Specifies the WLM classification to be assigned to the IBM Z Workload Pricing for Cloud (zWPC) offering. For details, see “Specifying the WLM classification for IBM Z Workload Pricing for Cloud” on page 94 .
--version	-v	no	Displays the SCRT version.
--help	-h	no	Displays help information for the SCRT command line interface.

Specifying NO89 statements

Specify your NO89 statements in a separate input file and use the **--no89** (or **-a**) option to identify that file to SCRT. The syntax for the NO89 statements in the NO89 input file is the same as for the statements in the NO89 DD in the SCRT JCL. For details, see [“NO89 DD statement” on page 61](#).

The following example shows a sample of some statements in a NO89 input file:

```
* TIVOLI NETVIEW FOR Z/OS V6
5697-NV6=*NONE
* TIVOLI NETVIEW FOR Z/OS V5
5697-ENV=*NONE
:
```

For an example of the use of the **--no89** option on the command line, see [“Sample SCRT command line invocation” on page 97](#).

Specifying customer number assignments

You must specify at least one customer number to be used as a default. Optionally, if you require that different CPCs be associated with different customer numbers, you can use the following syntax to override the default customer number on a per-CPC basis:

```
--customer-number cc-snnnnnnnn dd-pmmmmmmmm,CPC=tttt-sssss
```

where:

cc-snnnnnnnn

The default country code and customer number. For more information about customer numbers, see [“PARMS DD statement” on page 34](#).

dd-pmmmmmmmm

The override country code and customer number. For more information about customer numbers, see [“PARMS DD statement” on page 34](#).

tttt

The machine type of the CPU to which the override customer number is to apply.

SSSSS

The serial number of the CPU to which the override customer number is to apply.

Example: The following example shows the specification of a default customer number (US-S0000000000) and an override customer number (US-S123456789) for a specific CPC (2614-1234):

```
--customer-number US-S0000000000 US-S123456789,CPC=2614-1234
```

Specifying input data for special offerings

Use the **--inputcsv** option to specify the path of a file or files that contain input data for offerings such as IBM Mobile Workload Pricing (MWP), IBM Z Workload Pricing for Cloud (zWPC), or IBM Z Collocated Application Pricing (zCAP).

To specify multiple input files, repeat this option as necessary.

Example:

```
--inputcsv /path/to/input.csv
```

Specifying the WLM classification for IBM Z Workload Pricing for Cloud

Use the **--zwpc** option to request WLM-based zWPC and specify the appropriate WLM classification.

Example:

```
--zwpc "CATEGORYA"
```

Specifying Exclude statements

Use the **--exclude** (or **-x**) option to specify **Exclude** statements. The syntax for the statements is the same as for the **Exclude** statement in the SPECIAL DD in the SCRT JCL but must be enclosed in double quotation marks. For details, see [“EXCLUDE control statement” on page 44](#).

You can specify multiple statements by separating them with white space within a single argument or by specifying the **--exclude** option multiple times.

Example: The following example shows the specification of an **--exclude** option. The option is entered as part of a single, continuous command line input; the line breaks in the example are for formatting purposes only.

```
--exclude "CPC=2817-00000,IMAGE_ID=SYS2,PRODUCT_ID=5605-DB2,  
START=2017/01/01/00,RESUME=2017/12/31/23"
```

For an example of the use of the **--exclude** option on the command line, see [“Sample SCRT command line invocation” on page 97](#).

Specifying INCLUDE statements

Use the **--include** option to specify **INCLUDE** statements. The syntax for the statements is the same as for the **INCLUDE** statement in the SPECIAL DD in the SCRT JCL but must be enclosed in double quotation marks. For details, see [“INCLUDE control statement” on page 54](#).

You can specify multiple statements by separating them with white space within a single argument or by specifying the **--include** option multiple times.

Example: The following example shows the specification of an **--include** option. The option is entered as part of a single, continuous command-line input; any line breaks in the example are for formatting purposes only.

```
--include "CPC=2817-00000,IMAGE_ID=SYS1"
```

Specifying LPAR comments

Use the **--lpar-comment** (or **-l**) option to specify LPAR comments. The format of the LPAR comments the same as for those in the SPECIAL DD statement in the SCRT JCL but must be enclosed in double quotes. You can specify multiple comments by separating them with white space within a single argument or by specifying the **--lpar-comment** option multiple times.

Note that the format of the LPAR comment requires the comment text to be enclosed in either single or double quotation marks. Under Windows, you must use single quotation marks. Under Linux, you can use single quotation marks or properly escaped double quotation marks.

Example (Windows): Use single quotation marks around the LPAR comment as in the following example:

```
--lpar-comment "CPC=2817-00000,LPAR=SYS1,COMMENT='this is my comment'"
```

Example (Linux): Under most Linux shells, you can use a backslash character (\) to escape the double quotation marks around the comment text within the argument as in the following example:

```
--lpar-comment "CPC=2817-00000,LPAR=SYS1,COMMENT=\"this is my comment\""
```

Examples of multiple LPAR comments:

- **Option 1:** Specify the **--lpar-comment** option multiple times:

```
--lpar-comment "CPC=2817-00000,LPAR=SYS1,COMMENT='this is my comment'"
--lpar-comment "CPC=2817-00000,LPAR=SYS2,COMMENT='this is my comment'"
```

- **Option 2:** Separate multiple comments with white space:

```
--lpar-comment "CPC=2817-00000,LPAR=SYS1,COMMENT='this is my
comment'" "CPC=2817-00000,LPAR=SYS2,COMMENT='this is my comment'"
```

For an example of the use of the **--lpar-comment** option on the command line, see [“Sample SCRT command line invocation”](#) on page 97.

Associating an LPAR with a solution ID

Use the **--container** option to associate an LPAR with a solution ID for Tailored Fit Pricing for IBM Z.

Example:

```
--container CPC=3906-00000,IMAGE_ID=SYS1,  
ID=Z1000000-00000000-0000000000-00000000-0000-0000-0000-0000000-00000000
```

Updating container properties

Use the **--update-container** option to associate the LMS-provided solution name with a Tailored Fit Pricing solution based on the associated solution ID and to designate NO89 products that run in the container. The syntax for the option is the same as for the **UPDATE CONTAINER** control statement in the SPECIAL DD in the SCRT JCL.

Examples:

- Specifying the name of a container:

```
--update-container ID=D9111111-N4ED9AD-DB3D215CAF-NTZZZZZZ-07F5-4EDF-AF34-E062EB-0ADCD7,  
SET NAME="Container-Name"
```

- Specifying a single NO89 product:

```
--update-container ID=D9111111-N4ED9AD-DB3D215CAF-NTZZZZZZ-07F5-4EDF-AF34-E062EB-0ADCD7,  
SET NAME="Container-Name",N089 PRODUCT ID=nnnn-nnn
```

- Specifying multiple NO89 products:

```
--update-container ID=D9111111-N4ED9AD-DB3D215CAF-NTZZZZZZ-07F5-4EDF-AF34-E062EB-0ADCD7,
SET_NAME="Container-Name",N089_PRODUCT_ID=(nnnn-nnn,mmm-mmm,...)
```

The container name must be enclosed in double quotation marks. Depending on your operating system shell or environment, you might need to escape the quotation marks.

IBM requires that the specified container name match the container name that was specified when the solution ID was assigned in the LMS web portal.

Replacing one solution ID with another

Use the **--update-solution-id** option to replace one solution ID with another for Tailored Fit Pricing for IBM Z. The syntax for the option is the same as for the **UPDATE SOLUTION_ID** control statement in the SPECIAL DD in the SCRT JCL.

Example:

```
--update-solution-id OLD=Z1nnnnnn-nnnnnnnn-nnnnnnnnnn-nnnnnnnn-nnnn-nnnn-nnnn-nnnnnn-nnnnnn,
NEW=Z1mmmmmm-mmmmmmm-mmmmmmmmmmm-mmmmmmmmm-mmmm-mmmm-mmmm-mmmmmmm-mmmmmmm
```

You can use this option to correct an incorrect solution ID in the input data.

Discarding container-related information

Use the **--ignore-container** option to instruct SCRT to ignore container-related information for the specified solution ID and time interval. The syntax for the option is the same as for the **IGNORE CONTAINER** control statement in the SPECIAL DD in the SCRT JCL.

Examples:

- To ignore data for a container for a specified time interval:

```
--ignore-container ID=Z1nnnnnn-nnnnnnnn-nnnnnnnnnn-nnnnnnnn-nnnn-nnnn-nnnn-nnnnnn-nnnnnn,
START=yyyy/mm/dd/hh,RESUME=yyyy/mm/dd/hh
```

- To ignore all containers for the entire reporting period:

```
--ignore-container ID=*ALL
```

When specifying ID=*ALL, the **START** and **RESUME** keywords are not required; if specified, they are ignored.

Requesting an Enterprise Tailored Fit report

Use the **--enterprise-tailored-fit-pricing** option to generate an Enterprise Tailored Fit report. Enterprise Tailored Fit reports are for use by customers with a Tailored Fit Pricing Enterprise Solution agreement or contract. For more information, see [Chapter 20, “Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution,” on page 401.](#)

Generating ISV reports

To generate ISV reports using the command line, specify the **--isv** option and provide the path to the ISV library file.

Example:

```
--isv /path/to/isvlib.xml
```

Processing containers in ISV mode

By default, SCRT does not process nor report on container-related information when running in ISV mode. To include container information in ISV reports, use the **--isv-show-containers** option.

Specifying SMF or SCRT89 input data files

You must specify one or more SMF or SCRT89 input files at the end of the SCRT command line.

- Names of SMF input data files must have a .smf extension.
- Names of SCRT89 input data files must have a .bin, .scrt89, or .vse extension.
- You can specify relative or fully-qualified path names.
- If a path contains blank characters, enclose the path with double quotation marks.
- If the last option you specify accepts multiple arguments, such as the **--exclude** or **--lpar-comment** options, you must precede your SMF or SCRT89 paths by a double hyphen surrounded by white space (--) so that those paths are not interpreted as additional arguments for the preceding option.

For an example of how to specify SMF or SCRT89 input data files on the command line, see [“Sample SCRT command line invocation” on page 97](#).

Sample SCRT command line invocation

The following example shows a sample command line invocation of SCRT on a Linux system. The command is to be entered as a single, continuous line; the line breaks in this example are for formatting purposes only.

```
$ ./scrtc --customer-name "John Customer" --customer-number 897000001
--contact-name "John Smith" --contact-email "smithj@customer.com"
--contact-phone "800-555-1234" --no89 /path/to/no89data.txt
--exclude "CPC=2817-00000,IMAGE_ID=SYS1,PRODUCT_ID=ALL,START=ALL,RESUME=ALL"
--exclude "CPC=2817-00000,IMAGE_ID=SYS2,PRODUCT_ID=ALL,START=ALL,RESUME=ALL"
--lpar-comment "CPC=2817-00000,LPAR=SYS1,COMMENT=\"This is my comment\""
--output report.csv
/path/to/smf/data.smf /path/to/more/smf/data.smf
```

The example specifies the following information:

- The required options
- A NO89 input file (/path/to/no89data.txt)
- The required report output file (report.csv)
- Two SMF input data files (/path/to/smf/data.smf and /path/to/more/smf/data.smf)

Using the SCRT CLI sample scripts

SCRT provides a sample `scrt-cli.bat` batch script (for Windows systems) and a `scrt-cli.sh` shell script (for Linux systems) that you can customize and use to invoke the CLI. These scripts are intended to resemble the SCRT JCL that is used on z/OS and z/VSE systems and may provide a more familiar interface than using the command line interface directly.

The sample scripts are included in the `samples` subdirectory of the SCRT installation directory. Each script contains a number of variables that you must set before you can run the script. You must enter values for all of the required options, such as the customer information and the location of your SMF or SCRT89 input files. See the comments in each file for details.

Using the SCRT CLI batch script on Windows systems

Use a text editor to customize the `scrt-cli.bat` file and observe the following guidelines to customize the file:

- Enclose variables that contain white space in double quotation marks (").
- Do not include new line characters or line breaks on any variable.

- For the **--output** parameter, specify an absolute path or a relative path to the report file. Any subdirectories in the path must already exist in the file system. You must also specify the report file name. For instance:

```
set SCRTARGS=--output="c:\path\to\report.csv"
```

- To set a parameter value, specify the value after the equal sign (=) at the end of the line. For instance:

```
set SCRTARGS=%SCRTARGS% --customer-name="ABC Corp"
```

- To specify multiple SMF or SCRT89 input files, use double quotation marks around each input file and use white space to separate each file. For example:

```
set SMF="c:\path\to\file1.smf" "c:\path\to\file2.smf"
```

- To specify LPAR comments, use double quotation marks around the entire argument and single quotation marks around the comment keyword value. You can specify multiple LPAR comments by repeating the line. For example:

```
set SCRTARGS=%SCRTARGS% --lpar-comment "CPC=2817-00000,LPAR=SYS1,COMMENT='this is my comment'"
set SCRTARGS=%SCRTARGS% --lpar-comment "CPC=2817-00000,LPAR=SYS2,COMMENT='this is my comment'"
```

- As with LPAR comments, you can specify multiple **--exclude** parameters by repeating the line. For example:

```
set SCRTARGS=%SCRTARGS% --exclude "CPC=2817-00000,IMAGE_ID=SYS1,PRODUCT_ID=5605-DB2,START=2017/01/01/00,RESUME=2017/12/31/23"
set SCRTARGS=%SCRTARGS% --exclude "CPC=2817-00000,IMAGE_ID=SYS2,START=2017/01/01/00,RESUME=2017/12/31/23"
```

- The NO89 information is specified inline in the .bat file. Append your NO89 specification to the appropriate entries in the file without otherwise modifying the entries. For example:

```
rem NETVIEW PERF MONITOR V2
echo 5655-043=★NONE
:
```

After customizing the file, you can open a Windows command prompt and invoke the batch script:

```
scrt-cli.bat
```

The script prints the generated command line and runs the CLI utility. The exit value from the script reflects the exit value from the CLI utility.

Using the SCRT CLI shell script on Linux systems

Use a text editor to customize the `scrt-cli.sh` file and observe the following guidelines to customize the file:

- Enclose all variables in double quotation marks (").
- For the **--output** parameter, specify an absolute path or a relative path to the report file. Any subdirectories in the path must already exist in the file system. You must also specify the report file name. For instance:

```
OUTPUT="/path/to/report.csv"
```

- To set a parameter value, specify the value after the equal sign (=) at the end of the line. For instance:

```
CUSTOMER_NAME="ABC Corp"
```

- When specifying multiple SMF or SCRT89 input files, multiple LPAR comments, or multiple Exclude statements, enclose the entire variable value in double quotation marks and enter each value on a separate line. Within individual values, do not include new line characters or line breaks. For example:

```
SMF="
/path/to/file1.smf
```



```

/path/to/file2.bin
"
:
LPAR_COMMENT="
CPC=2817-00000,LPAR=SYS1,COMMENT='this is my comment'
CPC=2817-00000,LPAR=SYS2,COMMENT='this is my comment'
"
:
EXCLUDE="
CPC=2817-00000,IMAGE_ID=SYS1,PRODUCT_ID=5605-
DB2,START=2017/01/01/00,RESUME=2017/12/31/23
CPC=2817-00000,IMAGE_ID=SYS2,START=2017/01/01/00,RESUME=2017/12/31/23
"
:

```

- The NO89 information is specified inline in the shell script. Append your NO89 specification to the appropriate entries in the file without otherwise modifying the entries. For example:

```

*NETVIEW PERF MONITOR V2
5655-043=*NONE
:

```

After customizing the shell script, you can invoke it as:

```
./scrt-cli.sh
```

The script prints the generated command line and runs the CLI utility. The exit value from the script reflects the exit value from the CLI utility.

Part 3. Understanding and interpreting sub-capacity reports

This part helps you understand and interpret the sub-capacity report output. It provides a detailed description of the contents of a basic sub-capacity report, as well as sub-capacity reports for various other special pricing metrics, by providing sample report output and comprehensive explanations of the information that appears in each report.

The information in this part applies to sub-capacity reports generated on all platforms on which SCRT runs.

This part contains the following chapters:

- [Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103](#), describes a basic sub-capacity report and provides a foundation for understanding the sub-capacity reports for all of the other pricing metrics.

Many of the report sections and fields described here are common to the sub-capacity reports for other pricing metrics. Unless there are significant differences, the other sub-capacity reporting chapters do not describe these common report sections and fields and, instead, refer back to this chapter.

- [Chapter 6, “Sub-capacity reporting for z/OS systems that are z/VM guests,” on page 133](#)
- [Chapter 7, “Sub-capacity reporting for System z New Application License Charges \(zNALC\),” on page 145](#)
- [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#)
- [Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing \(IWP\),” on page 183](#)
- [Chapter 10, “Sub-capacity reporting for Mobile Workload Pricing \(MWP\),” on page 199](#)
- [Chapter 11, “Sub-capacity reporting for IBM Z Collocated Application Pricing \(zCAP\),” on page 229](#)
- [Chapter 12, “Sub-capacity reporting for IBM Z Workload Pricing for Cloud \(zWPC\),” on page 251](#)
- [Chapter 13, “Sub-capacity reporting for z/TPF systems,” on page 281](#)
- [Chapter 14, “Reporting for the z/TPF Transformation Engines Offering,” on page 299](#)
- [Chapter 15, “Sub-capacity reporting for z/VSE systems,” on page 305](#)
- [Chapter 16, “Sub-capacity reporting for Country Multiplex Pricing \(CMP\),” on page 343](#)
- [Chapter 17, “Sub-capacity reporting and Multi-Version Measurement \(MVM\),” on page 367](#)
- [Chapter 18, “Sub-capacity reporting for products that report a non-MSU metric,” on page 371](#)
- [Chapter 19, “Reporting for Tailored Fit Pricing for IBM Z,” on page 375](#)
- [Chapter 20, “Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution,” on page 401](#)
- [Chapter 21, “Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution,” on page 405](#)
- [Chapter 23, “Reporting for independent software vendor programs,” on page 413](#)

Chapter 5. Sub-capacity reporting for z/OS native systems

This chapter describes the tasks that the person in your organization responsible for controlling software billing (sometimes called the software asset manager) must perform to analyze the output produced by the SCRT. While this chapter will describe a sub-capacity report created for a CPC where only native z/OS systems (not z/OS guest systems running under z/VM) were running, the descriptions in this chapter also apply to z/OS guest, z/TPF and z/VSE environments as well. This chapter contains a full description of all of the fields in the sample sub-capacity report.

For information about SCRT execution and reporting that is unique to:

- z/OS systems that are z/VM guests, see [Chapter 6, “Sub-capacity reporting for z/OS systems that are z/VM guests,”](#) on page 133
- z/OS systems that include zNALC products, see [Chapter 7, “Sub-capacity reporting for System z New Application License Charges \(zNALC\),”](#) on page 145
- z/OS systems that include Getting Started Sub-Capacity Pricing products, see [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),”](#) on page 165
- z/OS systems that include Integrated Workload Pricing, see [Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing \(IWP\),”](#) on page 183
- z/OS systems that include Mobile Workload Pricing, see [Chapter 10, “Sub-capacity reporting for Mobile Workload Pricing \(MWP\),”](#) on page 199.
- z/OS systems that include IBM Z Collocated Application Pricing, see [Chapter 11, “Sub-capacity reporting for IBM Z Collocated Application Pricing \(zCAP\),”](#) on page 229.
- z/OS systems that include Country Multiplex Pricing, see [Chapter 16, “Sub-capacity reporting for Country Multiplex Pricing \(CMP\),”](#) on page 343.
- z/OS systems that include z/TPF, see [Chapter 13, “Sub-capacity reporting for z/TPF systems,”](#) on page 281
- z/OS systems that include z/VSE, see [Chapter 15, “Sub-capacity reporting for z/VSE systems,”](#) on page 305.

See [Chapter 24, “Viewing and editing the sub-capacity report,”](#) on page 419 for a description of how to view and edit a sub-capacity report.

Using the scenario with the sample reports: This chapter presents a customer scenario that does not necessarily represent a realistic customer configuration; the elements in it were chosen to illustrate a variety of possible SCRT outputs.

Identities reported for z/OS native systems

SCRT assigns identities to z/OS systems in order to associate each system with its matching information in the report.

When a single SYSID runs on an LPAR during the reporting period, SCRT identifies the system by the LPAR name. When multiple systems (that is, systems with different SYSIDs) run on an LPAR during the reporting period, SCRT identifies each system using the format *lparname/sysid*.

Advanced image naming

SCRT uses several formats to identify system images in a report.

Several factors influence image names, including the LPAR name, the operating system (z/OS, z/VSE, or z/TPF), the z/OS SYSID, whether the system is running under z/VM, and whether the image is running under zNALC pricing.

SCRT JCL control statements, CLI options, and other externals, such as the NO89 DD, generally support identifying images by LPAR name only (or by SYSID if the system is running under z/VM). This restriction sometimes makes it impossible to uniquely identify an image in these externals. Some examples of ambiguous naming include:

- An LPAR that hosts multiple z/OS images with unique SYSIDs at different times during the reporting period.
- An LPAR that runs under both traditional and zNALC pricing at different times during the reporting period.

If you need to uniquely identify images that have ambiguous names, you can enable advanced image naming support by specifying the following SPECIAL DD control statement in your SCRT JCL:

```
FEATURE ADVANCED_IMAGE_NAMING
```

If you use the Windows or Linux CLI, you can enable advanced image naming support by specifying the following option:

```
--feature advanced_image_naming
```

When advanced image naming is enabled, most SCRT externals will accept the full SCRT identity in order to uniquely identify images. The full SCRT identity is the identity used by SCRT in the report. Advanced image naming affects only commands and control statements that are used as input to SCRT and does not impact or change the names in the SCRT report. The externals that accept full SCRT identities are:

- **CONTAINER** control statement
- **EXCLUDE** control statement
- **INCLUDE** control statement
- **LPAR_COMMENT** control statement
- **NO89** DD statement

Sample base z/OS single-version configuration

The sample configuration, which is the basis for the sub-capacity reports described in the topics that follow, consists of a 2097 Model 405 CPC with a serial number of 02-12345. During the reporting period, the CPC was configured to support five LPARs: three are running z/OS natively on the LPAR, one is running Linux systems under z/VM, and another is supporting CMS users under z/VM. Prior to the reporting period, LPARSP was defined and used as a test LPAR for system programmer use; however, LPARSP was deactivated prior to the start of the reporting period. The z/OS V1 systems in LPAR1, LPAR2 and LPAR3 have z/OS SYSIDs of SYS1, SYS2 and SYS3. The z/OS systems are running a combination of the following z/OS middleware products:

- DB2 UDB V8
- CICS TS V3
- MQSeries V6
- IMS V8
- Cobol V3
- File Manager V6
- WebSphere Application Server V6

- WebSphere MQ Workflow V3

Table 8 on page 105 shows the sample hardware and software configuration.

CPC: 2097-405 (serial = 02-12345)					
LPAR1	LPAR2	LPAR3	LPARVM1	LPARVM2	LPARSP
z/OS V1 (SYSID = SYS1)	z/OS V1 (SYSID = SYS2)	z/OS V1 (SYSID = SYS3)	z/VM	z/VM	z/OS V1 (SYSID = SYSP)
CICS TS V3	DB2 UDB V8 CICS TS V3 MQSeries V6	DB2 UDB V8 CICS TS V3 MQSeries V6	Linux	CMS	
IMS V8	IMS V8	IMS V8 Cobol V3			
File Manager V6	File Manager V6 WebSphere Application Server V6	WebSphere Application Server V6			
	WebSphere MQ Workflow V3	WebSphere MQ Workflow V3			

Sample sub-capacity report for single-version z/OS native systems

This sample sub-capacity report illustrates the case where only z/OS V1 systems were running.

Note: The format of the report sections would be the same if only z/OS V2 or z/OS 3 systems were running, except that all references to the z/OS V1 product (5694-A01) would instead refer to the z/OS V2 product (5650-ZOS) or the z/OS 3 product (5655-ZOS), respectively. The formats for some report sections would differ for the case where multiple versions of z/OS systems were running during the same reporting period, as described in “Sub-capacity reporting for multiple versions of z/OS” on page 127. For purposes of these reporting examples, the distinguishing factor is not the particular version of z/OS that was running but whether a single version or multiple versions of z/OS were running on a given CPC during the reporting period.

The following sections show a sample sub-capacity report that contains two major sections:

- The billing-related section contains data that IBM needs to reconcile your company's sub-capacity bills.
- The customer verification section contains information that is useful to you when verifying the information in the sub-capacity report. The customer verification section starts after the section L5, Detail Data Sections - For Customer Analysis Purposes Only.

Billing-related section

The billing-related section of the sub-capacity report displays customer information, tool information, special conditions, product summary information, and the detail data collection. The customer and tool information appears in [Figure 17 on page 107](#), the special conditions section appears in [Figure 18 on page 108](#) and [Figure 19 on page 108](#), and the product summary information appears in [Figure 22 on page 112](#). The detail data collection appears in [Figure 31 on page 120](#).

Customer information

Figure 17 on page 107 shows the customer information, which contains the following fields. The fields marked with an asterisk (*) are generated by SCRT based on information supplied in the PARMS DD statement (as described in “PARMS DD statement” on page 34).

*Name of Person Submitting the Report

Name of the person who IBM should contact for billing-related issues

*E-mail Address of Report Submitter

Email address of the person who IBM should contact for billing-related issues

***Phone Number of Report Submitter**

Telephone number of the person who IBM should contact for billing-related issues

***Customer Name**

Name of your company

***Customer Number**

Your two-character ISO country or region code followed by a dash (-) and your ten-digit IBM site ID. For more information, see the description of the **Customer_Number** parameter in [“PARMS parameters that apply to all CPCs” on page 34](#).

Machine Serial Number

Serial number of the CPC to be analyzed in the format *xx-yyyyy*, where *xx* is the plant code and *yyyyy* is the serial number. This is the serial number to which IBM software is licensed.

In rare instances, when the serial number of the CPC differs from the serial number of the CPUs used by a z/OS system, the value in this field will be displayed as *xx-yyyyy (zzzz)*, where *zzzz* is the serial number of the CPUs (from the SMF70SER field in the SMF type 70 records). The CPU serial number is listed because there are cases where SCRT requires you to use the CPU serial number, instead of the CPC serial number, on an SCRT control statement that refers to a z/OS system. However, the CPU serial number is not listed for z/TPF and z/VSE systems, even when the CPU serial number differs from the CPC serial number, because SCRT always accepts the CPC serial number for all SCRT control statements that refer to z/TPF and z/VSE systems. For more information, see the description of the **CPC** parameter in [“EXCLUDE control statement” on page 44](#).

Machine Type/Model

CPC type and model

The message `Machine Model change observed` is displayed in the field following the **Machine Type/Model** value whenever the CPC model changes one or more times during the reporting period. In such a case, SCRT displays the highest machine model observed in the reporting period. Whenever the machine model change message is displayed, IBM requests that you provide information about that change. For details, see [“Machine software model changes” on page 441](#).

Machine Rated Capacity (MSUs)

MSU capacity for the CPC

Machine Peak Utilization

The highest observed 4-hour rolling average MSU value for the CPC based on all systems reported by SCRT. This value corresponds to the CPC value reported in the Detail LPAR Data section (N5).

Machine MSU Consumed

The total MSU consumption for the CPC based on all systems reported by SCRT. This value corresponds to the CPC value reported in the Detail LPAR Usage Data section (N7).

***Purchase Order Number**

The purchase order number assigned by your company's procurement department.

Your company may order software through a purchase order, which represents a commitment from your procurement department that a software order is authorized and will be paid for. If your company requires a purchase order, you must submit your sub-capacity reports with your purchase order number assigned by your company's procurement department. If you submitted software orders through a purchase order before you began using sub-capacity pricing, you need to enter a purchase order number on your sub-capacity reports. Contact your procurement department for more information.

The sample sub-capacity report shows the result of not assigning a purchase order number through the **PO_Number** PARMS control statement. Whenever a purchase order number is not provided, SCRT displays (optional) for the missing purchase order value, which indicates that you can also provide a purchase order value after the sub-capacity report has been generated, either by using the LMS web application or by using a spreadsheet application.

Customer Comments

(Optional) Field where you can specify, by editing the report, general comments about the report. Replace the (optional) placeholder text with your comment. Your comment must not exceed 255 characters in length. You can also use this field for these purposes:

- To provide IBM with information for a product that appears in the Product Summary Information section but for which SCRT does not provide any placeholders in the Customer Comments column for that product
- To describe why a processor model change occurred during the reporting period. For details, see [“Machine software model changes” on page 441](#)

```

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====
Run Date/Time                03 Nov 2024 - 09:17
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter:  444-999-9999

Customer Name                ABC Corporation
Customer Number              8971234567
Machine Serial Number        02-12345
Machine Type and Model       2097-405
Machine Rated Capacity (MSUs) 118
Purchase Order Number        (optional)
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust
the billable MSUs in inventory for all MLC Products listed under the MLC Product Name
column on this report. In accordance with our agreement, IBM will treat a change in product
licensed capacity as an order. If the MSUs have changed since the last report, software billing
based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to
bill those IPLA products listed under the IPLA Product Name column in this report which
exceed your entitled capacity. In accordance with our agreement, IBM will treat the use
of a product in excess of its entitled capacity as an order and you will be billed for
the amount in excess of your entitlement.

==C5=====
TOOL INFORMATION

Tool Release                  30.1.0
Reporting Period              2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

```

Figure 17. Sample sub-capacity report: Customer Information and Tool Information sections

Tool Information section

Figure 17 on page 107 also shows the Tool Information section, which contains the following fields:

Tool Release

Version, release, and modification level of SCRT that was used to produce this report.

Reporting Period

The start and end date of the reporting period, which includes the second day of one month to the first day of the next month.

Special Conditions section

The optional Special Conditions section is only generated when one or more special conditions must be reported.

One set of special conditions that causes the Special Conditions section to appear is when you request special processing options. For example, if you request Integrated Workload Pricing, a message appears in this section of the report, as shown in [Figure 18 on page 108](#).

```
==D5=====
SPECIAL CONDITIONS

IWP Control Statement
Integrated_Workload_Pricing
```

Figure 18. Sample sub-capacity report: Optional Special Conditions section

Special Conditions section for the Exclude control statement

The Special Conditions section also appears if any SCRT **Exclude** control statements applied to the CPC for which the sub-capacity report is being generated. Figure 19 on page 108 shows an example of the information that would be listed based on the **Exclude** control statements used in Example 3 in “EXCLUDE control statement” on page 44.

==D5=====						
SPECIAL CONDITIONS						
SMF/SCRT89 Exclusion Control Statement	Image ID	Product ID	Start	Resume	Customer Comments (255 chars max)	
Exclude Control Statement	BANKS	*ALL	15 Oct 2024 - 19:00	15 Oct 2024 - 20:00	(required)	
Exclude Control Statement	BANKS	*ALL	20 Oct 2024 - 17:00 UTC	20 Oct 2024 - 20:00 UTC	(required)	
SMF/SCRT89 Exclusion Results	Image ID	Product ID	Start	Resume	Hours	
SMF/SCRT89 records excluded	BANKS	*ALL	15 Oct 2024 - 23:00 UTC	16 Oct 2024 - 00:00 UTC	1	
SMF/SCRT89 records excluded	BANKS	*ALL	20 Oct 2024 - 17:00 UTC	20 Oct 2024 - 20:00 UTC	3	

Figure 19. Sample sub-capacity report: Optional Special Conditions section containing Exclude control statements

The Special Conditions section of the sub-capacity report indicates the number of hours for which SMF or SCRT89 records were read by SCRT and then excluded based on the criteria set by the corresponding **Exclude** control statement. For instance, if the range of hours specified by the **Start** and **Resume** parameters covered ten hours, and SCRT only read SMF or SCRT89 records for two of those hours, then SCRT would report two hours as being excluded for that **Exclude** control statement. If you use an **Exclude** control statement, be sure to remove that **Exclude** control statement when running SCRT for the next month because SCRT accepts and processes **Exclude** control statements whose **Start** and **Resume** range are outside of the reporting period, and reports on the number of hours that were excluded even when those hours are outside of the reporting period.

IBM will accept the exclusion provided that the number of excluded hours within the reporting period is small and a valid explanation has been given in the **Customer Comments** field for each **Exclude** control statement, as indicated by the (required) placeholder. You can enter such comments by editing the .csv report file or by using the LMS application. (Any comments—that is, SPECIAL DD statements that start with an asterisk (*) in column 1—that you use to document the **Exclude** statements in the JCL are ignored by SCRT and will not appear in the Special Conditions section.)

Note: IBM will *not* accept a sub-capacity report for which a large number of hours have been excluded; such reports will result in IBM charging the full CPC capacity against all of the products on that report.

The rows in the Special Conditions section that are labeled **Exclude Control Statement** list the values provided by each **Exclude** statement that was processed for this CPC. The last column of each of these rows provides for customer comments. IBM will not accept a sub-capacity report that does not contain an acceptable explanation for each **Exclude** statement.

The rows labeled **SMF/SCRT89 records excluded** report the results for the requests that are listed in the rows labeled **Exclude Control Statement**, which appear earlier in this section. The **Start** and **Resume** values reflect the dates and times that were provided in the original **Exclude** request. The **Hours** column is a count of the number of hours that were actually excluded from processing for the specified images and product IDs. For instance, if SCRT only found SMF and SCRT89 records for a single hour, then a value of 1 hour would be reported, even if the requested **Start** and **Resume** values on the **Exclude** control

statement include the entire reporting period. For the case where no **Product_Id** values are provided and the **Hours** value is reported as zero, no SMF or SCRT89 records met all of the requested criteria and none of the records were excluded. For the case where **Product_Id** values are provided and the **Hours** value is reported as zero, none of the requested products were found to be running on the specified image IDs during the selected date and time range.

In Example 3 in “[EXCLUDE control statement](#)” on page 44, a time zone change occurred (that is, a change in the UTC offset) during the range of hours when the records are to be excluded on 20 October. To ensure that SCRT knows precisely which hours to exclude, SCRT requires that the **Start** and **Resume** values on that **Exclude** control statement be UTC date/time values expressed in the UTC format (*yyyy/mm/dd/hh/UTC*). However, since no time zone change occurred on 15 October, SCRT accepts the **Start** and **Resume** values either as local date/time values in local format (*yyyy/mm/dd/hh*) or as UTC date/time values in UTC format (*yyyy/mm/dd/hh/UTC*). In the example, the **Exclude** control statement for 15 October specifies the date/time values in local format while the **Exclude** control statement for 20 October specifies the date/time values in the required UTC format.

The results for the **Exclude** control statements in Example 3 in “[EXCLUDE control statement](#)” on page 44 are shown in [Figure 19](#) on page 108. The SMF/SCRT89 Exclusion Control Statements section lists the **Exclude** control statements in the same local or UTC date/time format that was specified on each of the **Exclude** control statements. Thus, the **Start** and **Resume** values for the **Exclude** control statement for 15 October are in local format and the values for the **Exclude** control statement for 20 October are in UTC format.

Note, however, that the results for each of those **Exclude** control statements in the **SMF/SCRT89 Exclusion Results** section are shown in the same date/time format as that which is used in all other sections of the sub-capacity report that contain date/time values. If the date/time values are shown in local format in the other sections of the sub-capacity report, then the date/time values will also be shown in local format in this section of the report. The same is true whenever the date/time values for the other sections of the sub-capacity report are shown in UTC format. SCRT shows the date/time values in UTC format in all sections of the sub-capacity report, except for the **SMF/SCRT89 Exclusion Control Statements** section, whenever two or more LPARs on the same CPC have different UTC offsets during the reporting period.

In the example, the date/time values for the other sections of the sub-capacity report appear in UTC format, so the **Start** and **Resume** date/time values for both of the **Exclude** control statements appear in UTC format. Note that the local date/time values for the **Exclude** control statement for 15 October were converted from local date/time values to UTC date/time values using the UTC offset value of -4, which appeared in the corresponding SMF records. Also note that the conversion to UTC time caused a date change in the **Resume** date/time value from 15 October to 16 October.

Special Conditions section for the INCLUDE control statement

The Special Conditions section also appears if any SCRT **INCLUDE** control statements applied to the CPC for which the sub-capacity report is being generated. [Figure 20](#) on page 109 shows an example of the information that would be listed.

```

=====
SPECIAL CONDITIONS

Include Control Statements  Image ID          Customer Comments (255 chars max)
Include Control Statement  SYSA,SYSB      (required)

Include Results            Removed Image IDs
Images not included        SYSC,SYSD,SYSE
  
```

Figure 20. Sample sub-capacity report: Optional Special Conditions section containing INCLUDE control statements

The Special Conditions section indicates images that were present in the input data but not included via an **INCLUDE** control statement.

The rows in the Special Conditions section that are labeled **Include Control Statements** list the values provided by each **INCLUDE** control statement. The last column of each of these rows provides for customer comments. A valid explanation for including only a subset of images must be provided in the **Customer Comments** field for each **INCLUDE** control statement.

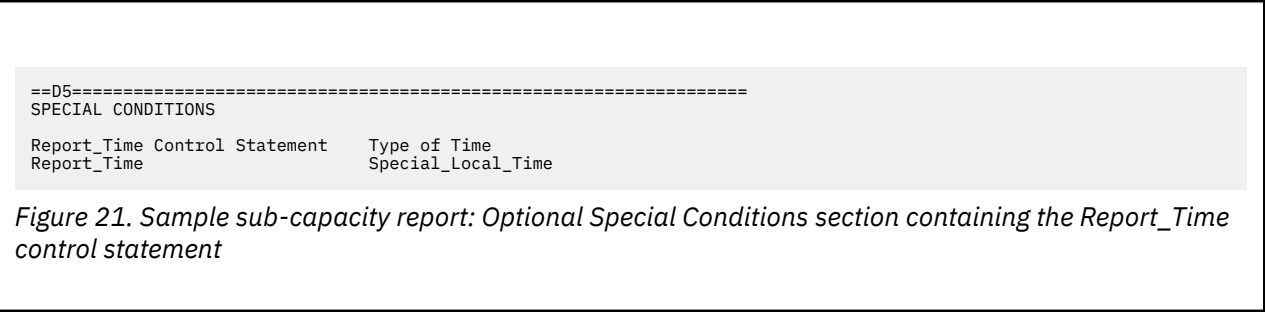
The rows labeled **Include Results** report the images that were not included based on the supplied **INCLUDE** control statements. These are images that were found in the SMF or SCRT89 data but were not explicitly included.

Special Conditions section for the Integrated_Workload_Pricing control statement

The Special Conditions section also appears to report on the **Integrated_Workload_Pricing** control statement whenever that control statement has been used. See [Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing \(IWP\),” on page 183](#) for an example of such a Special Conditions section and for more information about Integrated Workload Pricing.

Special conditions section for the Report_Time control statement

The Special Conditions section is also used to report on the optional **Report_Time** control statement whenever it is used. See [“REPORT_TIME control statement” on page 56](#) for a description and examples. [Figure 21 on page 110](#) shows how the **Report_Time** control statement used in Example 1 in [“REPORT_TIME control statement” on page 56](#) would appear in the Special Conditions section.



Product Summary Information section

The Product Summary Information section of the sub-capacity report contains the data for the current reporting period. [Figure 22 on page 112](#) shows the Product Summary Information section, which contains the following fields:

MLC Product Name

Listing, by name, of all supported, sub-capacity MLC products that are active in these LPARs during this reporting period. For information about sub-capacity reports that report on both z/OS V1 and z/OS V2, see [“Sub-capacity reporting for multiple versions of z/OS” on page 127](#). For information about zNALC, see [Chapter 7, “Sub-capacity reporting for System z New Application License Charges \(zNALC\),” on page 145](#).

Whenever the Quality Partnership Program designation, (QPP), appears after a product's name and you are participating in that QPP, IBM will treat that product in accordance with the terms and conditions specified in the IBM International License Agreement for Early Release of Programs (used for QPP programs) or equivalent agreement. If you are not participating in the QPP, IBM will treat the product in accordance with the standard IBM International License Agreement.

Note that older releases of SCRT may continue to use the QPP designation for products which are no longer under QPP but newer releases of SCRT will drop the QPP designation for QPP products which have become generally available. The QPP designation will also be used in the Product Max Contributors section, the Product Grid Snapshot section, and the Defined Capacity LPARs section.

MLC Product ID

Listing, by product ID, of all supported MLC products that are active in these LPARs during this period.

Tool MSUs

The combined LPAR 4-hour rolling average utilization values, in MSUs, for each sub-capacity eligible product shown on the report. If you have set up a defined capacity for an LPAR in which a native z/OS system is running, and the 4-hour rolling average utilization value is greater than the defined capacity set for the LPAR, then SCRT reports the defined capacity value. (See [“Capping individual LPARs” on page 497](#) for more information.) If the LPAR in which a native z/OS 1.8 or later system is running is part of an LPAR group and the LPARs in the group exceed the capacity set for the group, then SCRT reports the group capacity value. (See [“Group capacity” on page 498](#) for more information.) Note that measurement of the 4-hour rolling average utilization is for workloads executing on general-purpose CPs only. Workloads executing on zAAP and zIIP processors do not contribute to the 4-hour rolling average utilization.

Although the **Tool MSUs** value is used as the basis for billing IBM MLC products, those values may be increased by the IBM billing system in order to meet certain minimum capacity values that are part of the specific pricing option (for instance, WLC, MWLC), or those values may be adjusted downward, as can be the case for zNALC values. For more information, see [“Where to find more information” on page xvi](#) and see the zNALC explanation in [“Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads” on page 145](#).

Customer MSUs

(Optional) Whenever this field contains the (optional) placeholder, you can use this field to override the value displayed in the **Tool MSUs** field with a whole number of MSUs with a maximum of four digits. Customer MSU overrides are only supported for NO89 products. You must use the SCRT **Exclude** control statement to change the MSU values for all of the products that are tracked by SMF89 or SCRT89 records (that is, the products that have no placeholder in the **Customer MSUs** column). Any MSU override value provided in a **Customer Comments** field that does not have an (optional) placeholder will be ignored by the IBM business systems.

Customer Comments

Whenever SCRT places an (optional) placeholder in the **Customer MSUs** field, SCRT also places a (conditional) placeholder in the corresponding **Customer Comments** field. The (conditional) placeholder indicates that a comment is required whenever you have provided a value in the **Customer MSUs** field. Your comment must not exceed 255 characters in length.

Footnotes

(Optional) When present, a footnote provides additional information about the **Tool MSUs** value for a given product. A short description for each footnote appears in the rows immediately after the last product in the Product Summary Information section. For more information about footnote (iv), see Chapter 8, [“Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\)” on page 165](#). For more information about footnote (viii), see Chapter 9, [“Sub-capacity reporting for Integrated Workload Pricing \(IWP\)” on page 183](#).

IPLA Product Name

Listing, by name, of all supported IPLA products that are active in these LPARs during this period.

IPLA Product ID

Listing, by product ID, of all supported IPLA products active in these LPARs during this period.

```

=====
PRODUCT SUMMARY INFORMATION

```

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS V1	5694-A01	98		
DB2 UDB for z/OS V8	5625-DB2	71		
CICS TS for z/OS V3	5655-M15	98		
MQSeries for z/OS V6	5655-L82	71		
IMS V8	5655-C56	98		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	39	(optional)	(conditional)
IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
File Manager for z/OS V6	5655-P17	60	(optional)	(conditional)
IBM WebSphere Application Server for z/OS V6	5655-N01	71		
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	71		

Figure 22. Sample sub-capacity report: Product Summary Information section

Virtual products: The following virtual product entries may appear in this and other report sections:

z/OS Unreduced

The z/OS Unreduced special product entry reflects z/OS usage prior to reductions from IBM pricing programs, such as IBM Z Collocated Application Pricing and Tailored Fit Pricing for IBM Z. Note that the z/OS Unreduced peak hour may differ from the z/OS peak hour after reductions.

IPLA z/OS-Based

The IPLA z/OS-Based virtual product entry reflects simultaneous usage of all versions of a/OS, including Traditional and zNALC, when present.

IPLA CICS-Based

The IPLA CICS-Based virtual product entry reflects simultaneous usage of MLC and VUE editions of all CICS products. This product entry is reported when more than one version of CICS is active.

IPLA COBOL-Based

The IPLA COBOL-Based virtual product entry reflects simultaneous usage of MLC and VUE editions of all COBOL products. This product entry is reported when more than one version of COBOL is active.

IPLA DB2-Based

The IPLA DB2-Based virtual product entry reflects simultaneous usage of MLC and VUE editions of all Db2 products. This product entry is reported when more than one version of Db2 is active.

IPLA IMS-Based

The IPLA IMS-Based virtual product entry reflects simultaneous usage of MLC and VUE editions of all IMS products. This product entry is reported when more than one version of IMS is active.

IPLA MQ-Based

The IPLA MQ-Based virtual product entry reflects simultaneous usage of MLC and VUE editions of all MQ products. This product entry is reported when more than one version of MQ is active.

IPLA PL/I-Based

The IPLA PL/I-Based virtual product entry reflects simultaneous usage of MLC and VUE editions of all PL/I products. This product entry is reported when more than one version of PL/I is active.

Missing LPAR Data Details section

The Missing LPAR Data Details section is generated to notify you that SCRT detected missing LPAR data for the CPC. The missing data notifications fall into the following categories, as indicated by the notification messages:

- LPARs with unmatched SMF70 and SMF89 records

SCRT detects that, for a given interval (indicated by the **Detected** and **Resolved** fields), it has either SMF type 70 or type 89 records for the named LPAR. The time period for which the mismatch was detected and resolved is listed per incident.

- Active LPARs with no SMF/SCRT89 records

SCRT detects that SMF or SCRT89 data was not provided for an LPAR that was active during the time interval indicated by the **Detected** and **Resolved** fields. This data is listed per incident. An example of this condition is shown in Figure 23 on page 113.

When either of these conditions are detected, the **Justification for low data collection** field is also generated. For each LPAR that has either of these conditions, you must explain why the condition occurred. You must either provide your explanation by editing the multiplex report using a spreadsheet application or by entering your explanation in the LMS web application. The (required) placeholder appears in the field under the **Justification for low data collection** column when an explanation is necessary. Your explanation should replace the (required) placeholder. No explanation is necessary when the (required) placeholder does not appear.

When SCRT detects that data is missing, it generates an estimated peak MSU value for the primary operating system product. SCRT considers the primary operating system product to be the most recent z/OS release or encompassing family product. On machines without any z/OS images, the primary operating system may be a z/VSE or z/TPF system. Note the following information about the estimated peak MSU value:

- The estimated value reflects what the MSU value for the operating system would have been if all of the missing data had been present and attributed to that product. This value can be compared to the MSU value for the operating system as reported in section E5 in standard sub-capacity reports or section B5 in multiplex reports.
- Estimated values are derived from information reported by other LPARs running on the CPC during the data collection incidents and might not match the actual values that would have been recorded in SMF type 70 or 89 records. Estimated values also do not include capping considerations nor input from special pricing programs (such as IBM Z Collocated Application Pricing).
- IBM does not bill based on the estimated MSU value but may use it as a gauge to determine the relevance of the data that is missing from the report.

If SCRT determines that all expected data has been provided, the following notification message appears and no customer justification is needed:

All expected data collected.

Figure 23 on page 113 shows an example of the Missing LPAR Data Details section.

==H4=====

MISSING LPAR DATA DETAILS

Active LPARs with no SMF/SCRT89 records	Detected	Resolved (missing hours)	Est Peak MSU	Justification for low data collection (255 chars max)							
P00	25 Nov 2024 - 23:00	27 Nov 2024 - 14:00 (39 hours)	23	(required)							
P01	25 Nov 2024 - 23:00	27 Nov 2024 - 14:00 (39 hours)	5	(required)							
P02	25 Nov 2024 - 23:00	27 Nov 2024 - 14:00 (39 hours)	19	(required)							
P03	25 Nov 2024 - 23:00	27 Nov 2024 - 14:00 (39 hours)	19	(required)							
Total Missing Hours		156									
OS Product Name	OS Product ID	Multiplex	Peak MSU	Est Peak MSU	Date/Time	C00	C09	P00	P01	P02	P03
z/OS V2	5650-Z0S		20780	5009	26 Nov 2024 - 03:00	1972	2975	21	4	19	18

Figure 23. Sample sub-capacity report: Missing LPAR Data Details section

The Missing LPAR Data Details section contains the following fields:

First column

Lists the identities of the LPARs for which missing or unmatched SMF type 70 and type 89 records were detected.

Detected

The date and time of the start of the measurement interval for which missing or unmatched SMF type 70 and type 89 records were detected.

Resolved (missing hours)

The date and time of the start of the measurement interval when the missing data incident was resolved. The total length of the incident, in hours, is also reported.

Est Peak MSU

The estimated peak MSU value of the LPAR during the data collection incident. This value is computed from the effective dispatch time of the LPAR, as reported in the SMF records of other LPARs on the CPC.

Justification for low data collection

SCRT generates a placeholder of (required) in this field. Prior to submitting the sub-capacity report, you must replace the placeholder with an appropriate explanation to justify the missing data.

One-hour data gaps: For one-hour data gaps, SCRT automatically provides the following comment in the **Justification for low data collection** field (instead of the (required) placeholder):

```
1 hour missing data (no SMF89 data collected)
```

SCRT provides this comment as a convenience to you. You may keep this comment when you submit the sub-capacity report, or you may provide a more specific explanation. Regardless, your contractual agreement still requires that you collect and preserve complete SMF data. IBM reserves the right to question any data collection incident and to reject any reports that seem to indicate habitual or suspicious data collection incidents.

Total Missing Hours

The total number of hours that had missing or mismatched records.

OS Product Name

The name of the operating system product (for instance, z/OS V2) for which an estimated peak MSU value is computed.

OS Product ID

The product ID of the operating system product for which an estimated peak MSU value is computed.

Multiplex Peak MSU

An estimated multiplex peak MSU value, had all of the missing data had been included and attributed to the indicated operating system product. This column only appears in multiplex reports.

Est Peak MSU

An estimated peak MSU value, had all of the missing data been included and attributed to the indicated operating system product. In multiplex reports, this value is the machine contribution to the multiplex peak MSU value.

Date/Time

The measurement interval when the estimated peak MSU value would have occurred.

LPAR/Group contribution to estimated peak

Each column represents an LPAR and group's contribution to the estimated peak MSU value. The value is the actual MSU value if data was provided for the LPAR or group during the estimated peak hour, or is an estimated MSU value if data was not provided for the LPAR during the estimated peak hour.

Missing CPC Data Details section

Figure 24 on page 115 shows the Missing CPC Data Details section.


```

=====
MISSING CPC DATA DETAILS

CPC with no SMF/SCRT89      Detected      Resolved (missing hours)      Justification for low data
data from any LPAR                                     collection (255 chars max)

2964-840F7                  02 Feb 2024 - 00:00          08 Feb 2024 - 11:00 (155 hours)  (required)
2964-840F7                  08 Feb 2024 - 20:00          02 Mar 2024 - 00:00 (508 hours)  (required)

Total Missing Hours                                     663

```

Figure 24. Sample sub-capacity report: Missing CPC Data Details section

The Missing CPC Data Details section contains the following fields:

First column

This column identifies the CPCs for which no data was found in the SMF (or SCRT89) input data.

- In standard sub-capacity reports, each CPC is formatted as *tttt-sssss*, where *tttt* is the machine type and *sssss* is the machine serial number.
- In multiplex reports, this column contains the machine identifier tokens that SCRT assigned to the CPCs. Each machine identifier token is formatted as *MxCy*, where *x* is a unique machine number assigned to the CPC within the multiplex, and *y* is a number that differentiates between the customer numbers assigned to the CPCs within the multiplex.

Detected

The date and time of the start of the measurement interval for which no data for any LPAR on the CPC was detected.

Resolved (*missing hours*)

The date and time of the start of the measurement interval when the missing data incident was resolved. The total length of the incident, in hours, is also reported.

Justification for low data collection

SCRT generates a placeholder of (required) in this field. Prior to submitting the sub-capacity report, you must replace the placeholder with an appropriate explanation to justify the missing data.

Total Missing Hours

The total number of hours for which no data was received from any LPAR.

Product Features section

The Product Features section shows, for select IBM products, both the priced (or "entitled") and unpriced features that were active during the reporting period.

Figure 25 on page 116 shows an example of the Product Features section.

```

==F5=====
PRODUCT FEATURES

Product Name      Product ID      Entitlement ID      Feature Name
z/OS V2           5650-Z0S        S017291             z/OS V2 DFSMS dsshsm
z/OS V2           5650-Z0S        S017292             z/OS V2 DFSMS imm
z/OS V2           5650-Z0S        S017293             z/OS V2 DFSMStvs
z/OS V2           5650-Z0S        S01729C             z/OS V2 RMF
z/OS V2           5650-Z0S        S01729D             z/OS V2 SDSF
z/OS V2           5650-Z0S        S01729F             z/OS V2 Security Server
z/OS V2           5650-Z0S        S01780D             z/OS V2 z Enterprise Data Compression (zEDC)

OTHER REPORTABLE FEATURES

Product Name      Product ID      Feature Name
z/OS V2           5650-Z0S        C/C++/DEBUG
z/OS V2           5650-Z0S        JES2
z/OS V2           5650-Z0S        RACF
z/OS V2           5650-Z0S        TCP/IP BASE

```

Figure 25. Sample sub-capacity report: Product Features section

Capacity Changes section

Figure 26 on page 116 shows the Capacity Changes section.

```

==G5=====
CAPACITY CHANGES

Date/Time          Active Model    Temporary Model    Permanent Model
03 Mar 2019 - 00:00  2964-701        2964-701          2964-701
09 Mar 2019 - 21:00  2964-706        2964-701          2964-701
11 Mar 2019 - 00:00  2964-701        2964-701          2964-701

```

Figure 26. Sample sub-capacity report: Capacity Changes section

The Capacity Changes section shows the hours at which the CPC underwent a capacity change to any of the active, temporary, or permanent capacity models. The first hour reported indicates the initial capacity of the CPC for the reporting period, and subsequent hours indicate the new capacity information.

This section does not appear in the report if there were no capacity changes during the reporting period.

SCRT Softcap Enforcement section

Figure 27 on page 116 shows the SCRT Softcap Enforcement section.

```

==G6=====
SCRT SOFTCAP ENFORCEMENT - HOURS

Softcap Enforced    Image/Group    100%-110%    110%-125%    125%-150%    150%-200%    200% or more    Total SCRT Enforced
WLA Softcap Enforced  LPAR1          82           24           1            0            0              107

z/OS Uncapped Peak    Highest        Date/Time    LPAR1    LPAR1    LPAR3    LPAR4    LPAR5
                        805           24 Mar 2019 - 05:00    245     311      6      239      4

```

Figure 27. Sample sub-capacity report: SCRT Softcap Enforcement section

The SCRT Softcap Enforcement section reports, for the z/OS operating system, the number of hours that SCRT enforced either the defined capacity value for an individual LPAR or the group capacity for the LPARs in an LPAR group. Enforcement hours are grouped according to the ratio of actual utilization to the softcap value in effect.

For example, if the long term average CPU utilization of LPAR1 in a given hour is 200 MSUs and the defined capacity value of the LPAR for that hour is 175 MSU, the hour would be included in the **110%-125%** column, as 200 is 114 percent of 175.

This section also reports the **z/OS Uncapped Peak** value, which represents what the peak 4-hour rolling average would have been for z/OS if SCRT did not apply the soft capacity value in hours in which the actual utilization exceeded the soft capacity value.

In a Country Multiplex Pricing report, if the sum of the reported CPCs' **z/OS Uncapped Peak** values exceeds the summed reported peak values, then all CPCs in the report will include section G6 that shows their corresponding **z/OS Uncapped Peak** value for the uncapped peak hour.

This section does not appear in the report if SCRT did not enforce the soft capacity value during the reporting period, or if all occurrences of SCRT enforced soft capping were less than 110 percent of utilization to soft capacity.

System Recovery Boost Usage section

Figure 28 on page 117 shows the System Recovery Boost Usage section.

```

==G8=====
HOURS WITH BOOST ACTIVITY

LPAR1      IPL      SHUTDOWN  RECOVERY PROCESS
LPAR2      13       6         11
           0       2         1
  
```

Figure 28. Sample sub-capacity report: System Recovery Boost Usage section

The System Recovery Boost Usage section reports, for each system, the number of reporting period hours in which a boost event occurred.

zIIP Usage Data section

The zIIP Usage Data section reports information about overall zIIP usage and Java zIIP usage for the CPC. For each hour, the section reports the following metrics:

- zIIP eligible service units that executed on CPs
- zIIP eligible service units that executed on zIIPs
- Java zIIP eligible service units that executed on CPs
- Java zIIP eligible service units that executed on zIIPs
- The number of online CPs
- The number of online zIIPs

Note: These metrics are expressed in service units, not MSUs, and cannot be converted to nor compared to MSUs.

This section is only generated if at least one z/OS system on the CPC has the required system maintenance (PTFs) to collect the necessary data in the SMF type 70 records. If some systems on the CPC have the required maintenance and others do not, this section will indicate via a footnote that the data is incomplete.

Figure 29 on page 118 shows the zIIP Usage Data section.

```
==I8=====
zIIP USAGE DATA
```

	zIIP on CP	zIIP on zIIP	Java zIIP on CP	Java zIIP on zIIP	CPs	zIIPs	Model	Footnotes
2024-10-17T10:00:00	1217	43414351	1212	43329278	75	20	8561-775	
2024-10-17T11:00:00	4812	65161873	4753	65010608	75	20	8561-775	
2024-10-17T12:00:00	1450	60581553	1409	16571402	75	20	8561-775	
2024-10-17T13:00:00	494	65674030	466	209318	75	20	8561-775	

Figure 29. Sample sub-capacity report: zIIP Usage Data section

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the product summary section. In a simple report, it consists of the following:

- SMF / SCRT Input Data Statistics section, shown in [“SMF / SCRT89 Input Data Statistics section” on page 118](#).
- Detail LPAR Data Section, shown in [Figure 31 on page 120](#).
- Product Max Contributors section, shown in [Figure 33 on page 123](#).
- Product Grid Snapshot section, shown in [Figure 34 on page 124](#).
- In some cases, the report includes the Sub-Capacity Capping Enforced by SCRT section and the Group Capacity LPARs section. This section is also shown in [Figure 34 on page 124](#) and is described in [“Sub-Capacity Capping Enforced by SCRT section” on page 124](#) and [“Group Capacity LPARs section” on page 125](#).

SMF / SCRT89 Input Data Statistics section

This section contains data statistics which are similar to those in the Detail Data Collection section, except that the statistics in this section are for *all* of the SMF and SCRT89 records processed by SCRT, including the records outside of the reporting period and includes the reporting of all LPARs, not just the LPARs that were busy during the reporting period, as is the case for the Detail Data Collection section.

[Figure 30 on page 119](#) shows the SMF / SCRT89 Date Input Statistics section, which contains the following fields:

First column

Shows the names of the LPARs that were busy according to the SMF records that were processed. For cases where z/OS guest systems were running, SCRT shows the z/OS SYSID followed by the (guest) suffix, instead of the LPAR name. Whenever native z/OS systems with different SYSIDs run in the same LPAR within the reporting period, SCRT identifies those systems using the format *lparname/sysid* and reports on each such combination separately.

Example: If z/OS operating systems with SYSIDs of SYS1 and SYSA ran in LPAR1 at different times during the same reporting period, the SMF / SCRT89 Input Data Statistics section would report the statistics for LPAR1/SYS1 and LPAR1/SYSA separately, each on its own line.

Whenever the SYSID changes within an hour such that the same LPAR has two SYSIDs, SCRT reports all of the SMF statistics for that hour under the SYSID that ran for the longest time during the hour. If both SYSIDs ran for the same amount of time during the hour, SCRT reports all of the SMF statistics for that hour under the SYSID that had the highest 4-hour rolling average utilization for the hour. Note that this choice of the SYSID does not affect the value of the 4-hour rolling average utilization that is assigned to that hour; it simply affects which SYSID is assigned the 4-hour rolling average utilization for the hour. If both SYSIDs ran for the same amount of time and have the same 4-hour rolling average utilization, the SYSID which is first, in alphabetic order, will be assigned the 4-hour rolling average utilization value for the hour.

If you are running z/OS as a guest of z/VM, refer to the sample reports shown in [Chapter 6, “Sub-capacity reporting for z/OS systems that are z/VM guests,” on page 133](#).

SYSID

The ID of a z/OS system running in this LPAR. No SYSIDs are available for LPARs that were busy during the reporting period but provided no detailed SMF data.

Input Data Start Date

The date and time of the start of the earliest *measurement interval* for which records were found. This date and time is for the start of the measurement interval, not the date and time when the SMF or SCRT89 record was written. The dates and times are listed by LPAR or by guest system. This date might be outside the reporting period.

Input Data End Date

The date and time of the end of the last *measurement interval* for which records were found. This date and time is calculated by SCRT and does not appear on an SMF or SCRT89 record. The dates and times are listed by LPAR or by guest system. This date might be outside the reporting period.

In the sample report, LPAR1, LPAR2, and LPAR3 are LPARs that were busy during the reporting period and provided detailed SMF data. LPARVM1 and LPARVM2 are LPARs that were busy during the reporting period but provided no detailed SMF data. Both of these LPARs are also reported in the Detail Data Collection section. LPARSP is an example of an LPAR that was busy prior to the reporting period but was not busy during the reporting period and provided no detailed SMF data. LPARSP is not reported in the Detail Data Collection section.

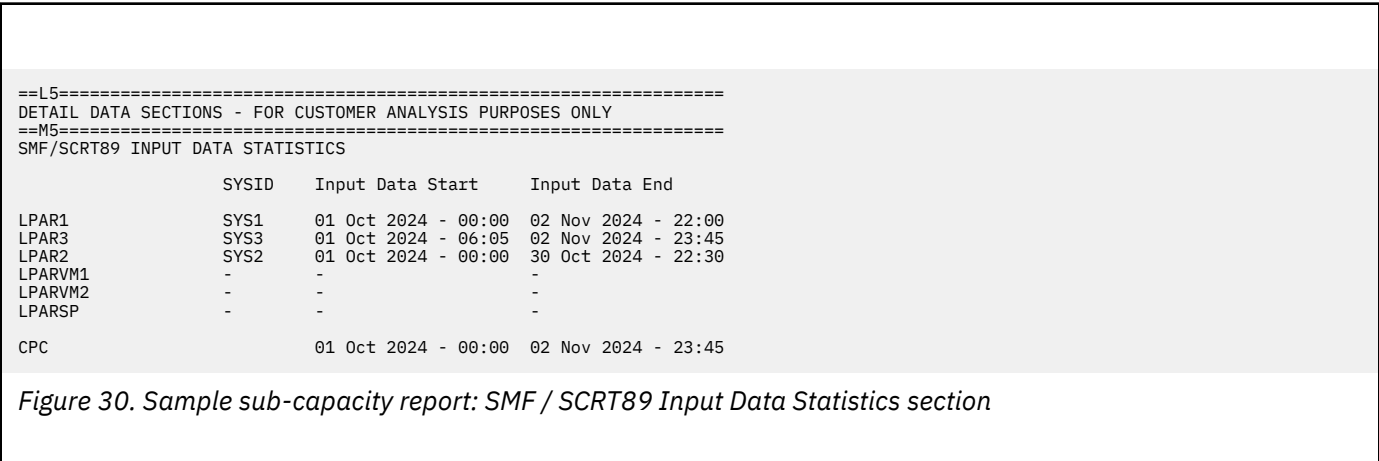


Figure 30. Sample sub-capacity report: SMF / SCRT89 Input Data Statistics section

Detail LPAR Data section

The Detail LPAR Data section of the sub-capacity report shows, for each LPAR on a CPC, the interval (date and hour) during the reporting period when the LPAR reached a maximum 4-hour rolling average utilization value (in MSUs), and the interval when the LPAR reached its second-highest 4-hour rolling average utilization value. It also shows the interval in which the sum of all the LPARs reached a maximum 4-hour rolling average value , and the interval in which the sum reached the second-highest 4-hour rolling average utilization value. If the SCRT encounters a maximum (for an individual LPAR or a sum of LPARs) in more than one interval, the first interval is shown.

When different LPARs defined on one CPC are set to different time zones, the SCRT normalizes these times to coordinated universal time, known as UTC, and reports all LPAR times as UTC. When this is the case, all times associated with SMF data are followed by the UTC suffix.

Whenever the Detail LPAR Data section is for a CPC that ran in power save mode any time during the reporting period, the report will also show additional columns which categorize the hours used in power save mode. See [“Static power save mode” on page 506](#) for a description of power save mode and a sample Detail LPAR Data section with the additional columns.

[Figure 31 on page 120](#) shows an example of the Detail LPAR Data section.

```
==N5=====
Detail LPAR Data section
```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR1	32	1	03 Oct 2024 - 07:00	30	5	03 Oct 2024 - 06:00
LPAR3	35	6	03 Oct 2024 - 14:00	34	11	03 Oct 2024 - 13:00
LPAR2	41	4	04 Oct 2024 - 11:00	40	7	04 Oct 2024 - 10:00
CPC	107	5	03 Oct 2024 - 12:00	106	14	03 Oct 2024 - 11:00

Figure 31. Sample sub-capacity report: Detail LPAR Data section

The Detail LPAR Data section contains the following fields:

First column

The names of the LPARs appear in the first (leftmost) column; the CPC entry appears last and represents the MSU values for the busiest and second busiest hours for the entire CPC. The CPC values may be less than the sum of the MSU values for the individual LPARs, except for the case where the entire machine consists of a single LPAR.

Highest

Highest observed 4-hour rolling average MSU value for the LPAR.

Notes:

1. Whenever the highest 4-hour rolling average MSU value is capped to the defined capacity value by SCRT, the defined capacity value will be used for the highest value. However, this value does not reflect any product capping that SCRT enforces as a result of LPAR group capping because such capping takes place in aggregate across all of the LPARs in the group and does not require SCRT to make specific judgments about the capping of specific LPARs in the group.
2. This value does not reflect any MSU reductions resulting from Mobile Workload Pricing or similar offerings.
3. This value is taken prior to the application of the Tailored Fit Pricing for IBM Z algorithm.

See “Group capacity” on page 498 for a more detailed explanation of how SCRT deals with group capping.

Hour Count

Number of hours during which the highest 4-hour rolling average MSU value occurred.

Date/Time

Date and time of the earliest measurement interval when the highest 4-hour rolling average MSU value occurred.

2nd Highest

Second-highest observed 4-hour rolling average MSU value for the LPAR.

Hour Count

Number of hours during which the second-highest 4-hour rolling average MSU value occurred.

Date/Time

Date and time of the earliest measurement interval when the second-highest 4-hour rolling average MSU value occurred.

Detail LPAR Usage Data section

The Detail LPAR Usage Data section of the sub-capacity report shows, for each LPAR on a CPC, the total MSUs consumed during the reporting period, the hour of peak MSU consumption, and the number of MSUs consumed during the hour of peak consumption. The Detail LPAR Usage Data section shows native LPARs and z/VM guests with the RMF Monitor I VMGUEST option.

Note: For Tailored Fit Pricing for IBM Z, the highest hour depends on when a product ran on each of the tenant resource groups that specify the container solution ID, along with any dedicated active LPARs on the CPC that are configured with the **CONTAINER** command as part of the Tailored Fit Pricing solution.

Figure 32 on page 121 shows an example of the Detail LPAR Data section.

```

==N7=====
DETAIL LPAR USAGE DATA SECTION

          Total MSU Consumed   Peak Hour Consumption   Date/Time           OS
D0                957             108   02 Sep 2024 - 08:00   z/OS
D2                117             12    02 Sep 2024 - 03:00   z/OS
D4                125             19    02 Sep 2024 - 08:00   z/OS
D6                105              8    02 Sep 2024 - 00:00   z/OS
CPC                1304            159   02 Sep 2024 - 07:00   z/OS

```

Figure 32. Sample sub-capacity report: Detail LPAR Usage Data section

The Detail LPAR Usage Data section contains the following fields:

First column

The names of the LPARs appear in the first (leftmost) column; CPC appears last and represents the total MSU consumption for the entire CPC.

Total MSU Consumed

The total MSU consumption value for the LPAR.

To calculate this value, SCRT uses the effective dispatch time of the LPAR, as reported in the SMF70EDT field of the SMF type 70 subtype 1 records, and converts this value to MSUs on a per-hour basis (using an hour interval) for each applicable LPAR.

Note: The MSU values displayed here do not reflect any MSU adjustments or reductions.

Peak Hour Consumption

The number of MSUs consumed during the peak hour.

Date/Time

The hour of peak consumption.

OS

The operating system running on the LPAR.

Product Max Contributors section

The Product Max Contributors section of the sub-capacity report shows the highest observed utilization values for each of the listed products that were active in your configuration. Whenever the highest product value occurs for multiple hours during the reporting period, the date and time for the first of the highest hours is listed. Note that the highest hour depends on when the product was running on each of the active LPARs on the CPC.

Only products that always run at the same time that z/OS runs and in every LPAR in which z/OS runs will match the z/OS value for highest MSUs and the z/OS value for the earliest date and time at which the value occurred. In the sample report, CICS V3 and IMS V8 always ran in the same LPARs and at the same times that z/OS ran, and are shown having the identical **Highest** MSU value, **Date/Time** value, and individual LPAR and group MSU values as z/OS. All of the other products ran in a subset of the LPARs in which z/OS ran or a subset of the dates/times during which z/OS ran.

Figure 33 on page 123 shows the Product Max Contributors section, which contains the following fields:

Product Name

Listing, by name, of all active, supported products

Product ID

Listing, by product ID, of all active, supported products

Highest

Highest combined 4-hour rolling average MSU value for all the LPARs in which the product runs. This is the value that your cost for the product is based on. For information about the contents of this field for z/OS systems at V1R6 and earlier levels that are z/VM guests, see [“Summary of MSUs reported for sub-capacity eligible products running on z/OS guest systems under z/VM” on page 133.](#)

Date/Time

Date and time when the highest 4-hour rolling average value was first reached.

LPAR names and values

Names of the LPARs that contributed to the highest 4-hour rolling average MSU values and the amount of the contribution from each in MSUs. Whenever z/OS systems with different SYSIDs run in the same LPAR within the reporting period, SCRT identifies those systems using the format LPARNAME/SYSID and reports on each LPARNAME/SYSID combination separately. For example, if z/OS operating systems with SYSIDs of SYS1 and SYSA ran in LPAR1 at different times during the same reporting period, the Product Max Contributors section would report the statistics for LPAR1/SYS1 and LPAR1/SYSA separately, each on its column.

Whenever the SYSID changes within an hour such that the same LPAR has two SYSIDs, SCRT reports all of the SMF statistics for that hour under the SYSID that ran for the longest time during the hour. If both SYSIDs ran for the same amount of time during the hour, SCRT reports all of the SMF statistics for that hour under the SYSID that had the highest 4-hour rolling average utilization for the hour. Note that this choice of the SYSID does not affect the value of the 4-hour rolling average utilization that is assigned to that hour; it simply affects which SYSID is assigned the 4-hour rolling average utilization for the hour. If both SYSIDs ran for the same amount of time and have the same 4-hour rolling average utilization, the SYSID which is first, in alphabetic order, will be assigned the 4-hour rolling average utilization value for the hour.

Note that this section may report higher product MSUs than are reported for z/OS for specific LPARs. This only happens for products that ran in a subset of the LPARs in which z/OS ran and/or ran during a subset of the dates/times during which z/OS ran. While it might appear that the product MSUs are higher than the z/OS MSUs for the same date and hour, that is not the case. The MSU values being reported for that LPAR are for a different **Date/Time** value than was reported for z/OS. While z/OS was assigned the same MSUs as were reported for that product for the same date and time for that LPAR, that was not the date and time during which z/OS reported its **Highest** value across all of the LPARs on the CPC. The MSUs reported for z/OS for that LPAR were for a different date and time than the MSUs reported for that product, so the z/OS MSUs reported for that LPAR may be different. In the sample report, DB2 V8, MQ V6, WebSphere Application Server V6, and WebSphere MQ Workflow V3 are examples of products whose reported MSUs are higher for LPAR3 than the MSUs reported for z/OS, due to the differences in **Date/Time** values.

Group names and values

Names of the LPAR groups that were defined during the reporting period. An MSU value is listed under a group, instead of under the LPARs in the group, whenever the earliest peak hour for the associated product has been capped by SCRT, because group capping did not react fast enough to prevent this group of LPARs from exceeding the group cap. This column is displayed for every group encountered during the reporting period, regardless of whether SCRT was required to cap a group of LPARs during the reporting interval or not.

In the sample report, the earliest peak hour for every product listed in the Product Max Contributors section did not require SCRT to perform any group capping, since no values were reported under the **GROUP** column. See [“Group capacity” on page 498](#) for examples of Product Max Contributors sections where SCRT did cap the MSU values for a group of LPARs.


```
==P5=====
PRODUCT MAX CONTRIBUTORS
```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	GROUP LGRPA
z/OS V1	5694-A01	98	06 Oct 2024 - 11:00	27	32	39	0
DB2 UDB for z/OS V8	5625-DB2	71	03 Oct 2024 - 14:00	0	34	37	0
CICS TS for z/OS V3	5655-M15	98	06 Oct 2024 - 11:00	27	32	39	0
MQSeries for z/OS V6	5655-L82	71	03 Oct 2024 - 14:00	0	34	37	0
IMS V8	5655-C56	98	06 Oct 2024 - 11:00	27	32	39	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	39	04 Oct 2024 - 10:00	0	0	39	0
File Manager for z/OS V6	5655-P17	60	09 Oct 2024 - 15:00	29	31	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	71	03 Oct 2024 - 14:00	0	34	37	0
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	71	03 Oct 2024 - 14:00	0	34	37	0

Figure 33. Sample sub-capacity report: Product Max Contributors section

Product Grid Snapshot section

The Product Grid Snapshot section of the sub-capacity report shows what products ran in each LPAR during the time period covered by the report and what percentage of that time each product is active in each LPAR. (Note that products supported by the SCRT NO89 control statement are not listed in the Product Grid Snapshot section because this section is intended to report on the times when products actually ran.) Whenever z/OS systems with different SYSIDs run in the same LPAR within the reporting period, SCRT identifies those systems using the format LPARNAME/SYSID and reports on each LPARNAME/SYSID combination separately.

Example: If z/OS operating systems with SYSIDs of SYS1 and SYSA ran in LPAR1 at different times during the same reporting period, the product grid section would report the statistics for LPAR1/SYS1 and LPAR1/SYSA separately, each in its own column.

Whenever the SYSID changes within an hour such that the same LPAR has two SYSIDs, SCRT counts that hour under both SYSIDs.

SCRT outputs one decimal place for the percentages (although your spreadsheet preferences may be set to display two decimal places for these fields.)

Figure 34 on page 124 shows the Product Grid Snapshot section.

```
==Q5=====
PRODUCT GRID SNAPSHOT
```

Product Name	Product ID	LPAR1	LPAR3	LPAR2
z/OS V1	5694-A01	92.4%	92.7%	91.5%
DB2 UDB for z/OS V8	5625-DB2		92.7%	91.5%
CICS TS for z/OS V3	5655-M15	92.4%	92.7%	91.5%
MQSeries for z/OS V6	5655-L82		92.7%	91.5%
IMS V8	5655-C56	92.4%	92.7%	91.5%
IBM WebSphere Application Server for z/OS V6	5655-N01		92.7%	91.5%
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		92.7%	91.5%

```
==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS
```

Product Name	Product ID	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	GROUP LGRPA
z/OS V1	5694-A01	4	8	7	
DB2 UDB for z/OS V8	5625-DB2		8	7	
CICS TS for z/OS V3	5655-M15	4	8	7	
MQSeries for z/OS V6	5655-L82		8	7	
IMS V8	5655-C56	4	8	7	
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53			7	
File Manager for z/OS V6	5655-P17	4	8	7	
IBM WebSphere Application Server for z/OS V6	5655-N01		8	7	
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		8	7	

```
==U5=====
GROUP CAPACITY LPARS
```

```
Group Name      LPARS
LGRPA           LPAR3(iii)(vii), LPAR2(iii)(vi)
```

```
(iii) LPAR was not a member of the group for the entire reporting interval
(vi) LPAR was hardcapped during a portion of the reporting period while it was defined to the group
(vii) LPAR was hardcapped during the entire reporting period while it was defined to the group
```

```
==W5=====
This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.
```

Figure 34. Sample sub-capacity report: Product Grid Snapshot, Sub-Capacity Capping Enforced by SCRT - Hours, and Group Capacity LPARs sections

Sub-Capacity Capping Enforced by SCRT section

The Sub-Capacity Capping Enforced by SCRT section is only generated, if necessary, to indicate when the capped value was exceeded and SCRT itself had to cap the MSU value set for a specific LPAR, or SCRT had to cap the sum of the capacities of all of the LPARs defined to a group.

The value shown is the number of hours when SCRT capped a specific LPAR or the sum of the capacities of the LPARs in an LPAR group. SCRT will only cap the reported MSUs when the system function performing the capping is unable to react fast enough to enforce the capped value that you set. This value is *not* the number of hours when PR/SM along with WLM were successful at capping an LPAR or group of LPARs to the requested capped value. It is also *not* the capped MSU value that was used. A blank value is shown in this section if SCRT was not required to cap an LPAR or a group of LPARs.

This section of the sub-capacity report is only generated when either or both of the following conditions are met:

- SCRT had to cap one or more LPARs to the specific capping value set for each of those LPARs during the reporting period.
- SCRT had to cap the LPARs in one or more LPAR groups during the reporting period.

Whenever z/OS systems with different SYSIDs run in the same LPAR within the reporting period, SCRT identifies those systems using the format *lparname/sysid* and reports on each *lparname/sysid* combination separately.

Example: If z/OS operating systems with SYSIDs of SYS1 and SYSA ran in LPAR1 at different times during the same reporting period, the Sub-Capacity Capping Enforced by SCRT section would report the statistics for LPAR1/SYS1 and LPAR1/SYSA separately, each in its own column.

Whenever the SYSID changes within an hour that has been capped by SCRT, SCRT counts that hour under both SYSIDs.

Figure 34 on page 124 shows the Sub-Capacity Capping Enforced by SCRT section. In the sample report, SCRT was required to cap the capacity in four hours in LPAR1, in seven hours in LPAR2, and in eight hours in LPAR3. SCRT was not required to provide any group capping in this reporting period.

For more information about capping individual LPARs, see [“Capping individual LPARs” on page 497](#).

Group Capacity LPARs section

The Group Capacity LPARs section of the sub-capacity report, also shown in [Figure 34 on page 124](#), displays the group name and LPARs assigned to that group. This section of the sub-capacity report is intended to assist you in understanding the effects of using the group capacity hardware and software function to set the MSU cap for all of the LPARs within a group of LPARs. This section is only generated when any LPAR is defined as a member of an LPAR group.

Footnote (iii) indicates that SCRT found that the LPAR was also either not defined to a group for periods of time during the reporting period or defined to a different group for periods of time during the reporting period. Note that the absence of footnote (iii) does not guarantee that the LPAR was always a member of the indicated LPAR group. It simply means that the LPAR was defined to the same group in all of the SMF records that were processed by SCRT for that reporting period.

Footnote (vi) indicates that the LPAR was hard capped for a portion of the reporting period while that LPAR was also defined as a member of the indicated LPAR group. Such LPARs are not eligible for group capping for the hours during which they were hard capped. This footnote helps you understand why group capping was not as effective as you might have expected it to be. In the sample report, LPAR2 was a member of an LPAR group for part of the reporting period and was also hard capped for one or more hours while it was a member of LPAR group LGRPA, based on the SMF records processed by SCRT.

Footnote (vii) indicates that the LPAR was hard capped at the same time when that LPAR was also defined as a member of the indicated LPAR group. Such LPARs are not eligible for group capping. This footnote helps you understand why group capping was not as effective as you might have expected it to be. In the sample report, LPAR3 was a member of an LPAR group for part of the reporting period and was also hard capped while it was a member of LPAR group LGRPA, based on the SMF records processed by SCRT.

See [“Group capacity” on page 498](#) for more information on group capacity.

Sub-hour Data Collection Details section

The optional Sub-hour Data Collection Details section of the sub-capacity report shows the number of unreported minutes of SMF type 70 data on an hour-by-hour basis. This section does not include hours that have been excluded or hours that were completely missing as reported in the Missing LPAR Data Details (H4) or Missing CPC Data Details (H6) sections. Hours during which there are no missing minutes for any LPAR are suppressed. This optional section only appears when you specify the **Generate_Detailed_Data** control statement in the SPECIAL DD (or the CLI or GUI equivalent).

SCRT reports missing minutes whenever the aggregate time for an hour across all SMF type 70 records is less than 60.

[Figure 35 on page 126](#) shows an example of the Sub-Hour Data Collection Details section.

```

==V6=====
SUB-HOUR DATA COLLECTION DETAILS - MINUTES MISSING

Date/Time          D0          D2
21 Aug 2024 - 10:00      26          30
27 Aug 2024 - 23:00      56          45

```

Figure 35. Sample sub-capacity report: Sub-Hour Data Collection Details section

The Detail LPAR Usage Data section contains the following fields:

Date/Time

The hours during which SCRT detected unreported minutes for one or more LPARs.

LPAR columns

The number of unreported minutes that SCRT detected for the listed hours. If an IPL or a shutdown occurred in or near the hour, the value is indicated as $x(I)$ and the (I) footnote appears.

Footnote

If an IPL or shutdown occurred in any reported hours in this section, the following footnote is reported:

(I) indicates that an IPL or a shutdown occurred during or near this hour.

In Figure 35 on page 126, system D0 had 26 unreported minutes during the 10:00 hour on 21 Aug and 56 unreported minutes during the 23:00 hour on 27 August. Likewise, system D2 had 30 and 45 unreported minutes for those hours.

System Recovery Boost Usage Details section

The System Recovery Boost Usage Details section reports System Recovery Boost usage on an hour-by-hour basis. This optional section only appears when you specify the **Generate_Detailed_Data** control statement in the SPECIAL DD (or the CLI or GUI equivalent).

Figure 36 on page 126 shows an example of the System Recovery Boost Usage Details section.

```

==V8=====
HOURS WITH BOOST ACTIVITY DETAIL

          LPAR1          SHUTDOWN  RECOVERY PROCESS          LPAR2          SHUTDOWN  RECOVERY PROCESS
          IPL            -            -            IPL            -            -
10 Jun 2020 - 10:00      -            -            0            0            1
11 Jun 2020 - 09:00      1            1            0            -            -
11 Jun 2020 - 10:00      0            1            -            -            -
11 Jun 2020 - 13:00      -            -            0            1            0
12 Jun 2020 - 11:00      1            1            -            -            -
12 Jun 2020 - 12:00      1            1            -            -            -
12 Jun 2020 - 13:00      1            0            -            -            -
15 Jun 2020 - 14:00      1            0            -            -            -
15 Jun 2020 - 15:00      1            0            -            -            -
16 Jun 2020 - 09:00      1            0            -            -            -
16 Jun 2020 - 12:00      -            -            0            1            0
16 Jun 2020 - 13:00      1            0            -            -            -
16 Jun 2020 - 14:00      1            0            -            -            -
17 Jun 2020 - 11:00      1            0            -            -            -
17 Jun 2020 - 12:00      1            1            -            -            -
17 Jun 2020 - 13:00      1            1            -            -            -
17 Jun 2020 - 14:00      0            0            -            -            -
18 Jun 2020 - 09:00      1            0            -            -            -

```

Figure 36. Sample sub-capacity report: System Recovery Boost Usage Details section

The System Recovery Boost Usage Details section contains the following fields:

(column 1)

The hours during which SCRT detected System Recovery Boost usage on at least one system.

System / Boost class columns

For each system, there is a column for each boost class (IPL, Shutdown, and Recovery Process). A non-zero value indicates that at least one boost event for the corresponding system and boost class

occurred. Dashes appear under the boost class columns for a system if no boost event occurred on that system for the given hour.

Sub-capacity reporting for multiple versions of z/OS

Whenever multiple versions of z/OS—that is, z/OS V1, z/OS V2, or z/OS 3—run on the same CPC during the same reporting period, the format of the sub-capacity report is extended to report on each product.

SCRT considers each version of z/OS to be a separate sub-capacity product and reports on each of them independently of one another. However, in addition to reporting separate sub-capacity values for these products, SCRT also reports an additional sub-capacity value which represents the concurrent sub-capacity peak value for all active versions of z/OS, as if they were treated as a single product. The sub-capacity report entry that lists this additional value shows a product name of z/OS (without any version designation) and a product ID of (A11). This entry appears prior to the z/OS V1 and z/OS V2, and z/OS 3 entries in the following sections of the sub-capacity report, but only when multiple versions of z/OS appear on the same sub-capacity report.

- Product Summary Information
- Product Max Contributors
- Product Grid Snapshot
- Sub-capacity Capping Enforced by SCRT

IBM uses the sub-capacity value reported for the z/OS (A11) entry to determine the billing for the z/OS V1, z/OS V2, and z/OS 3 products whenever multiple versions of z/OS are reported in the same sub-capacity report. This occurs under Multi-Version Measurement (MVM) rules, which replace Migration Pricing Option (MPO) and Single Version Charging (SVC) rules, where applicable. For a detailed explanation of MVM, go to [Multi-Version Measurement \(MVM\) \(www.ibm.com/downloads/cas/YPWEG7QL\)](http://www.ibm.com/downloads/cas/YPWEG7QL).

See “Sample sub-capacity report for z/OS V1 and z/OS V2 running traditional and zNALC workloads” on page 149 for a description of the special reporting used for z/OS systems where a combination of z/OS traditional and zNALC workloads were running.

The sub-capacity values reported for IBM middleware products running on the z/OS systems continue to be based on the 4-hour rolling average utilization of the z/OS native system or guest system in which those products run, regardless of which version of z/OS is running. The sub-capacity values reported for middleware products that are eligible for Getting Started Sub-Capacity Pricing or Integrated Workload Pricing are also not affected by the version of z/OS that is running in that z/OS system image.

The special reporting for multiple versions of z/OS described here is only used when multiple versions of z/OS are running on the same CPC during the same reporting period; it is not used when only a single version of z/OS is running on a CPC. For the latter case, the sub-capacity report will show either z/OS V1, z/OS V2, or z/OS 3 depending on which one ran during the reporting period, as described in “Sample sub-capacity report for single-version z/OS native systems” on page 105.

The following topics show an example of a sub-capacity report where multiple versions of z/OS (in this case, z/OS V1 and z/OS V2) are running.

Sample base z/OS V1 and z/OS V2 configuration

This example shows a migration scenario where one of three z/OS V1 systems was migrated to a z/OS V2 system several days after the reporting period started. The migrated z/OS system used the same z/OS SYSID for both the z/OS V1 and V2 levels of that z/OS system image. The other z/OS V1 systems remained as z/OS V1 systems for the entire reporting period.

Except for the levels of z/OS being used, this sample configuration is identical to the configuration shown in “Sample base z/OS single-version configuration” on page 104. This makes it easier to clearly describe the differences between the sub-capacity report sections that are generated for this scenario and those generated for the base z/OS V1 scenario.

The sample hardware consists of a 2097 Model 405 CPC with a serial number of 02-12345. During the reporting period, the CPC was configured to support five LPARs: three are running z/OS natively on the LPAR, one is running Linux systems under z/VM, and another is supporting CMS users under z/VM. (A sixth LPAR, LPARSP, was deactivated prior to the start of the reporting period and, thus, is not shown in this sample configuration.) The z/OS V1 systems in LPAR1, LPAR2, and LPAR3 use z/OS SYSIDs of SYS1, SYS2, and SYS3. The z/OS V1 system running on LPAR3 with a SYSID of SYS3 was replaced by a z/OS V2 system using the same SYSID on 5 October. The z/OS systems are running a combination of the following z/OS middleware products:

- DB2 UDB V8
- CICS TS V3
- MQSeries V6
- IMS V8
- Cobol V3
- File Manager V6
- WebSphere Application Server V6
- WebSphere MQ Workflow V3

Table 9 on page 128 shows the sample hardware and software configuration, along with the dates when z/OS V1 and z/OS V2 ran in LPAR3.

CPC: 2097-405 (serial = 02-12345)				
LPAR1	LPAR2	LPAR3	LPARVM1	LPARVM2
z/OS V1	z/OS V1	z/OS V1 (02 Oct - 05 Oct) z/OS V2 (05 Oct - 02 Nov)	z/VM	z/VM
(SYSID = SYS1)	(SYSID = SYS2)	(SYSID = SYS3)	Linux	CMS
CICS TS V3	DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8	DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8 Cobol V3		
IMS V8	File Manager V6	WebSphere Application Server V6		
File Manager V6	WebSphere Application Server V6	WebSphere MQ Workflow V3		
	WebSphere MQ Workflow V3			

Sample sub-capacity report for z/OS V1 and z/OS V2 native systems

Most of the report sections in this sample sub-capacity report are identical to those shown in “Sample sub-capacity report for single-version z/OS native systems” on page 105. Therefore, this example only describes those report sections that differ for this scenario.

The Customer Information section and Tool Information section are the same as those shown in “Billing-related section” on page 105. These sections are not listed here. There are no error conditions and no SPECIAL SCRT control statements were used for this example, so no Special Conditions section was generated.

Product Summary Information section

Figure 37 on page 129 shows the Product Summary Information section for this sample scenario.

```

=====
PRODUCT SUMMARY INFORMATION

```

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS	(A11)	98		
z/OS V2	5650-ZOS	38		
z/OS V1	5694-A01	72		
DB2 UDB for z/OS V8	5625-DB2	71		
CICS TS for z/OS V3	5655-M15	98		
MQSeries for z/OS V6	5655-L82	71		
IMS V8	5655-C56	98		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	39	(optional)	(conditional)

IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
File Manager for z/OS V6	5655-P17	60	(optional)	(conditional)
IBM WebSphere Application Server for z/OS V6	5655-N01	71		
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	71		

Figure 37. Sample sub-capacity report for z/OS V1 and z/OS V2: Product Summary Information section

Since z/OS V1 and z/OS V2 were running during the reporting period, both z/OS V1 (5694-A01) and z/OS V2 (5650-ZOS) are reported in the Product Summary Information section, along with an additional z/OS (A11) product entry. The additional z/OS (A11) entry reports the peak concurrent sub-capacity value for both z/OS V1 and z/OS V2, as if they were a single product. .

Starting with the April 2017 reporting period, IBM charges for z/OS based on the sub-capacity value reported for the z/OS (A11) entry, which is 98 MSUs in this example. Note that this charge is for the latest program family version, which is z/OS V2. This occurs under MVM rules, which replace SVC. Note that under MVM rules, there is no time limit for running multiple versions of a products.

The IBM middleware products continue to be charged based on their reported peak sub-capacity values, which are independent of the z/OS version on which those middleware products run. See [“Sub-capacity reporting for multiple versions of z/OS” on page 127](#) for more information.

Note that the z/OS (A11) entry is reported whenever both z/OS V1 and z/OS V2 are running on the same CPC at any time during the reporting period.

You can find more information about the fields in this section and other sections of the sub-capacity report in [“Sample sub-capacity report for single-version z/OS native systems” on page 105](#).

The SMF / SCRT89 Input Data Statistics section and Detail LPAR Data section are the same as the sections shown in [Figure 30 on page 119](#) and [Figure 31 on page 120](#), respectively, because the configurations in both examples are the same and these sections do not report on any product-specific information. These sections are not shown here.

Product Max Contributors section

[Figure 38 on page 130](#) shows the Product Max Contributors section for this sample scenario.

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	GROUP LGRPA
z/OS	(A11)	98	06 Oct 2024 - 11:00	27	32	39	0
z/OS V2	5650-ZOS	38	05 Oct 2024 - 10:00	0	38	0	0
z/OS V1	5694-A01	72	07 Oct 2024 - 13:00	30	0	42	0
DB2 UDB for z/OS V8	5625-DB2	71	03 Oct 2024 - 14:00	0	34	37	0
CICS TS for z/OS V3	5655-M15	98	06 Oct 2024 - 11:00	27	32	39	0
MQSeries for z/OS V6	5655-L82	71	03 Oct 2024 - 14:00	0	34	37	0
IMS V8	5655-C56	98	06 Oct 2024 - 11:00	27	32	39	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	39	04 Oct 2024 - 10:00	0	0	39	0
File Manager for z/OS V6	5655-P17	60	09 Oct 2024 - 15:00	29	31	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	71	03 Oct 2024 - 14:00	0	34	37	0
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	71	03 Oct 2024 - 14:00	0	34	37	0

Figure 38. Sample sub-capacity report for z/OS V1 and z/OS V2: Product Max Contributors section

The Product Max Contributors section reports the z/OS (A11) entry along with the z/OS V1 and z/OS V2 product entries whenever both z/OS V1 and z/OS V2 run on the same CPC anytime during the reporting period. The z/OS (A11) entry lists the concurrent peak sub-capacity value for both z/OS V1 and z/OS V2 products, as if they were a single product. The **Date/Time** values are the earliest concurrent peak values for each product or for each combination of product versions, such as z/OS (A11).

For the migration case where a z/OS V1 system is replaced by a z/OS V2 system using the same SYSID in the same LPAR in the same hour, SCRT assigns the sub-capacity value for that entire hour to the z/OS V2 system and does not assign any sub-capacity value to the z/OS V1 system in order to prevent charging both z/OS V1 and z/OS V2 products for the sub-capacity value for to that hour.

The only difference between the configuration for this sample report and the sample report in Figure 33 on page 123 is that one of the z/OS V1 systems was replaced by a z/OS V2 system. Accordingly, the **Date/Time** value and all of the sub-capacity values reported for the z/OS (A11) entry in the Product Max Contributors section of this sample report match the corresponding values reported for the z/OS V1 entry in Figure 33 on page 123.

You can find more information about the fields in this section and other sections of the sub-capacity report in “Sample sub-capacity report for single-version z/OS native systems” on page 105.

Product Grid Snapshot section

Figure 39 on page 130 shows the Product Grid Snapshot section for this sample scenario.

```

==Q5=====
PRODUCT GRID SNAPSHOT

```

Product Name	Product ID	LPAR1	LPAR3	LPAR2
z/OS	A11	92.4%	92.7%	91.5%
z/OS V2	5650-ZOS		81.4%	
z/OS V1	5694-A01	92.4%	11.3%	91.5%
DB2 UDB for z/OS V8	5625-DB2		92.7%	91.5%
CICS TS for z/OS V3	5655-M15	92.4%	92.7%	91.5%
MQSeries for z/OS V6	5655-L82		92.7%	91.5%
IMS V8	5655-C56	92.4%	92.7%	91.5%
IBM WebSphere Application Server for z/OS V6	5655-N01		92.7%	91.5%
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		92.7%	91.5%

Figure 39. Sample sub-capacity report for z/OS V1 and z/OS V2: Product Grid Snapshot section

The Product Grid Snapshot section reports the z/OS (A11) entry and the z/OS V1 and z/OS V2 product entries whenever both z/OS V1 and z/OS V2 run on the same CPC anytime during the reporting period.

You can find more information about the fields in this section and other sections of the sub-capacity report in “Sample sub-capacity report for single-version z/OS native systems” on page 105.

The Group Capacity LPARs section is the same as the corresponding section shown in Figure 34 on page 124 because the configurations in both sample reports are the same. This section will not be shown here.

Sub-capacity Capping Enforced by SCRT section

Figure 40 on page 131 shows the Sub-capacity Capping Enforced by SCRT section for this sample scenario.

```

=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS

```

Product Name	Product ID	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	GROUP LGRPA
z/OS	(A11)	4	8	7	
z/OS V2	5650-Z0S		8		
z/OS V1	5694-A01	4		7	
DB2 UDB for z/OS V8	5625-DB2		8	7	
CICS TS for z/OS V3	5655-M15	4	8	7	
MQSeries for z/OS V6	5655-L82		8	7	
IMS V8	5655-C56	4	8	7	
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53			7	
File Manager for z/OS V6	5655-P17	4	8	7	
IBM WebSphere Application Server for z/OS V6	5655-N01		8	7	
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		8	7	

Figure 40. Sample sub-capacity report for z/OS V1 and z/OS V2: Sub-capacity Capping Enforced by SCRT section

The Sub-Capacity Capping Enforced by SCRT section reports the z/OS (A11) entry and the z/OS V1 and z/OS V2 product entries whenever both z/OS V1 and z/OS V2 run on the same CPC anytime during the reporting period. Note that the capping results for the z/OS (A11) entry match those of the z/OS V1 product entry in the same report section in Figure 34 on page 124 because the z/OS (A11) entry reported all of the z/OS activity (for both z/OS V1 and z/OS V2) that was reported only as z/OS V1 activity before the migration of the z/OS V1 system to z/OS V2. Also note that the capping statistics in this sample report are also listed for each version of z/OS because each version of z/OS is treated as a separate product.

You can find more information about the fields in this section and other sections of the sub-capacity report in “Sample sub-capacity report for single-version z/OS native systems” on page 105.

The sub-capacity values reported for Getting Started Sub-Capacity Pricing are unaffected by a simple change in the version of z/OS under which such products run. The sub-capacity report sections related to Getting Started Sub-Capacity Pricing do not apply to the configuration in this example, so those report sections are not shown here.

The Group Capacity LPARs section is the same as the corresponding section shown in Figure 34 on page 124 because the configurations in both sample reports are the same. This section will not be shown here.

Chapter 6. Sub-capacity reporting for z/OS systems that are z/VM guests

This chapter describes the differences between the SCRT support for z/OS systems that run as guest systems under z/VM and the support for z/OS native systems described in Chapter 4. These differences involve the MSU calculations for guest systems and the identities that SCRT uses to report on guest systems. This chapter also shows sample sub-capacity reports for a CPC on which different levels of z/OS guests are running and describes the fields in those reports.

Prior to z/OS V1.7, SCRT reported product MSUs for z/OS systems that were z/VM guests based on the maximum capacity of the LPAR in which the z/OS guest ran. With z/OS V1.7, SCRT reports the 4-hour rolling average MSU value for z/OS guest systems and the products that run under them.

To take advantage of the support provided by z/OS V1.7 (and later), all z/OS guest systems on a CPC must be at z/OS V1.7 or later. All z/OS V1.7, V1.8, and V1.9 systems must have SMF APAR OA20314 applied. If the z/OS guests are all running z/OS V1.6 or earlier levels, or there is a mixed environment with systems at both z/OS V1.6 or earlier and z/OS V1.7 or later, the product MSUs will be based on the maximum capacity of the LPARs in which the z/OS guest systems ran. For example, if LPAR1 has a guest running z/OS V1.6 and LPAR2 has guests running z/OS V1.7, then both LPARs are reported based on their maximum capacity. For an illustration of this example, see [Table 17 on page 136](#).

The maximum capacity of an LPAR is equal to the MSU capacity of the machine divided by the number of CPs corresponding to that machine's software model, times the number of CPs configured for that LPAR.

Note: If your z/VM LPAR is configured as part of a z/VM Single System Image (SSI) cluster, you must specify VMRELOCATE OFF in the virtual machine directory entry for each z/OS guest.

Summary of MSUs reported for sub-capacity eligible products running on z/OS guest systems under z/VM

Table 10 on page 133 shows a summary of how the sub-capacity MSU value for sub-capacity eligible z/OS products is calculated for all combinations of mixed native z/OS systems and guest z/OS systems at different release levels on the same CPC.

Table 10. Calculating MSUs for z/OS products in configurations with z/VM guest systems

Is z/OS running natively on this CPC?	And is z/OS running as a guest of z/VM, with any guest at V1.6 or earlier?	And is z/OS running as a guest of z/VM, with any guest at V1.7?	Then MSUs are calculated for sub-capacity-eligible programs as...
Yes	Yes	Yes	The highest combination of:
Yes	Yes	No	<ul style="list-style-type: none"> The 4-hour rolling average utilization across all z/OS native LPARs, and The maximum capacity of all z/OS guest LPARs.
Yes	No	Yes	The highest combined 4-hour rolling average utilization across all z/OS native LPARs and z/OS guest LPARs.
Yes	No	No	
No	No	Yes	
No	Yes	Yes	The highest combined maximum capacity of all z/OS guest LPARs.
No	Yes	No	

Capacity examples for z/OS guest systems

The examples show how the SCRT calculates MSUs for products running under z/OS guest systems in various configurations.

Example 1: Single LPAR with all guests at z/OS V1.6 or lower

This example shows a configuration in which there is a single LPAR (LPAR1) with two z/OS guests (VM Guest 1 and VM Guest 2), both running z/OS V1.6.

Note: All the examples in this chapter assume that all the products listed are running 100 percent of the time in the indicated LPAR and that the LPARs shown all have their peak 4-hour rolling averages in the same hour. These assumptions, while unlikely in practice, make the examples clearer.

Table 11. Example of a single LPAR with all guests at z/OS V1.6 or lower	
LPAR1, maximum capacity = 208 MSUs	
z/VM	
VM Guest 1	VM Guest 2
z/OS V1.6 Db2 CICS	z/OS V1.6 Db2 WebSphere Application Server for z/OS V5

How MSUs are calculated for example 1

Table 12. How MSUs are calculated for example 1: Single LPAR with all guests z/OS V1.6 or lower			
Product	VM Guest 1 MSUs	VM Guest 2 MSUs	Product MSUs
Db2	maximum capacity of LPAR1 (208)	maximum capacity of LPAR1 (208)	208
WebSphere Application Server	not applicable	maximum capacity of LPAR1 (208)	208
CICS	maximum capacity of LPAR1 (208)	n/a	208
z/OS	maximum capacity of LPAR1 (208)	maximum capacity of LPAR1 (208)	208

The product MSUs for the sub-capacity products Db2, CICS, WebSphere Application Server, and z/OS are based on the maximum capacity of LPAR1 in which both VM Guest 1 and VM Guest 2 ran, which is 208 MSUs.

Example 2: Single LPAR with all guests at z/OS V1.7 or higher

This example shows a configuration in which there is a single LPAR (LPAR1) with two z/OS guests (VM Guest 1 and VM Guest 2), both running z/OS V1.7.

Note: All the examples in this chapter assume that all the products listed are running 100 percent of the time in the indicated LPAR and that the LPARs shown all have their peak 4-hour rolling averages in the same hour. These assumptions, while unlikely in practice, make the examples clearer.

Table 13. Example of a single LPAR with all guests at z/OS V1.7 or higher	
LPAR1, maximum capacity = 208 MSUs	
z/VM	
VM Guest 1	VM Guest 2
z/OS V1.7 Db2 CICS 4-hour rolling average value = 75 MSUs	z/OS V1.7 Db2 WebSphere Application Server for z/OS V5 4-hour rolling average value = 80 MSUs

How MSUs are calculated for example 2

The product MSUs for the sub-capacity products Db2, CICS, WebSphere Application Server, and z/OS are based on the maximum 4-hour rolling average value for the guest system or combination of guest systems in which each product runs.

Table 14. How MSUs are calculated for example 2: Single LPAR with all guests z/OS V1.7 or higher			
Product	VM Guest 1 MSUs	VM Guest 2 MSUs	Product MSUs
Db2	75	80	155
WebSphere Application Server	not applicable	80	80
CICS	75	not applicable	75
z/OS	75	80	155

Example 3: Single LPAR with mixed z/OS guest systems

This example shows a configuration in which there is a single LPAR (LPAR1) with two z/OS guests (VM Guest 1 and VM Guest 2). VM Guest 1 is running z/OS V1.6 and VM Guest 2 is running z/OS V1.7.

Note: All the examples in this chapter assume that all the products listed are running 100 percent of the time in the indicated LPAR and that the LPARs shown all have their peak 4-hour rolling averages in the same hour. These assumptions, while unlikely in practice, make the examples clearer.

Table 15. Example of a single LPAR with mixed guest systems	
LPAR1, maximum capacity = 208 MSUs	
z/VM	
VM Guest 1	VM Guest 2
z/OS V1.6 Db2 CICS	z/OS V1.7 Db2 WebSphere Application Server for z/OS V5

How MSUs are calculated for example 3

The product MSUs for the sub-capacity products Db2, CICS, WebSphere Application Server, and z/OS are based on the maximum capacity of LPAR1 in which both VM Guest 1 and VM Guest 2 ran, which is 208 MSUs.

Table 16. How MSUs are calculated for example 3: Single LPAR with mixed guest systems

Product	VM Guest 1 MSUs	VM Guest 2 MSUs	Product MSUs
Db2	maximum capacity of LPAR1 (208)	maximum capacity of LPAR1 (208)	208
WebSphere Application Server	not applicable	maximum capacity of LPAR1 (208)	208
CICS	maximum capacity of LPAR1 (208)	not applicable	208
z/OS	maximum capacity of LPAR1 (208)	maximum capacity of LPAR1 (208)	208

Example 4: Multiple LPARs on the same CPC; one with mixed z/OS guests and one with only z/OS V1.7 guests

This example shows a configuration in which there are two LPARs. LPAR1 has two z/OS guests (VM Guest 1 and VM Guest 2), with a mix of z/OS levels running in these VM guests (VM Guest 1 is running z/OS V1.6 and VM Guest 2 is running z/OS V1.7.) LPAR2 has only one z/OS guest running z/OS V1.7.

Note: All the examples in this chapter assume that all the products listed are running 100 percent of the time in the indicated LPAR and that the LPARs shown all have their peak 4-hour rolling averages in the same hour. These assumptions, while unlikely in practice, make the examples clearer.

Table 17. Example of two LPARs with mixed guest z/OS systems

LPAR1 max. capacity = 156 MSUs		LPAR2 max. capacity = 56 MSUs
z/VM		z/VM
VM Guest 1	VM Guest 2	VM Guest 3
z/OS V1.6 Db2 CICS	z/OS V1.7 Db2 WebSphere Application Server for z/OS V5	z/OS V1.7 Db2 CICS

How MSUs are calculated for example 4

The product MSUs for products running in LPAR2, which is running z/OS V1.7, are based on the maximum capacity of LPAR2, since there is at least one z/OS guest running a pre-z/OS V1.7 level on this same machine (VM Guest 1 is running z/OS V1.6 in LPAR1).

Table 18. How MSUs are calculated for example 4: Two LPARs with mixed guest z/OS systems

Product	VM Guest 1 MSUs	VM Guest 2 MSUs	VM Guest 3 MSUs	Product MSUs
Db2	maximum capacity of LPAR1 (156)	maximum capacity of LPAR1 (156)	maximum capacity of LPAR2 (56)	212
WebSphere Application Server	not applicable	maximum capacity of LPAR1 (156)	not applicable	156
CICS	maximum capacity of LPAR1 (156)	not applicable	maximum capacity of LPAR2 (56)	212
z/OS	maximum capacity of LPAR1 (156)	maximum capacity of LPAR1 (156)	maximum capacity of LPAR2 (56)	212

Example 5: Multiple LPARs on the same CPC; one with mixed z/OS guests and one with a native z/OS V1.7 system

This example shows a configuration in which there are two LPARs. LPAR1 has two z/OS guests (VM Guest 1 and VM Guest 2), with a mix of z/OS levels running in these VM guests (VM Guest 1 is running z/OS V1.6 and VM Guest 2 is running z/OS V1.7.) LPAR2 has a native z/OS system running z/OS V1.7.

Note: All the examples in this chapter assume that all the products listed are running 100 percent of the time in the indicated LPAR and that the LPARs shown all have their peak 4-hour rolling averages in the same hour. These assumptions, while unlikely in practice, make the examples clearer.

LPAR1 max. capacity = 156 MSUs			LPAR2 max. capacity = 56 MSUs
z/VM		z/OS	
VM Guest 1	VM Guest 2	Native z/OS system	
z/OS V1.6 Db2 CICS	z/OS V1.7 Db2 WebSphere Application Server for z/OS V5	z/OS V1.7 Db2 CICS 4-hour rolling average value = 40 MSUs	

How MSUs are calculated for example 5

The product MSUs for any products running in native LPARs are still based on a 4-hour rolling average, even though there are z/OS guests running pre-z/OS V1.7 levels. However, product MSUs for products running the z/OS guests in LPAR1 are based on the maximum capacity of LPAR1, since there is at least one z/OS guest running a pre-z/OS V1.7 level on this same machine (the VM Guest1 system running z/OS V1.6 in LPAR1).

Table 20. How MSUs are calculated for example 5: Multiple LPARs; one with mixed guests and one with a native z/OS system at any level

Product	VM Guest 1 MSUs	VM Guest 2 MSUs	Native z/OS system MSUs	Product MSUs
Db2	Maximum capacity of LPAR1 (156)	Maximum capacity of LPAR1 (156)	40	196
WebSphere Application Server	not applicable	Maximum capacity of LPAR1 (156)	n/a	156
CICS	Maximum capacity of LPAR1 (156)	not applicable	40	196
z/OS	Maximum capacity of LPAR1 (156)	Maximum capacity of LPAR1 (156)	40	196

Example 6: Multiple LPARs on the same CPC; one with z/OS guests at V1.7 and one with a native z/OS system at any level

This example shows a configuration in which there are two LPARs. LPAR1 has two z/OS guests (VM Guest 1 and VM Guest 2), both running z/OS V1.7. LPAR2 is running a native z/OS system at the z/OS V1.6 level.

Note: All the examples in this chapter assume that all the products listed are running 100 percent of the time in the indicated LPAR and that the LPARs shown all have their peak 4-hour rolling averages in the same hour. These assumptions, while unlikely in practice, make the examples clearer.

Table 21. Example of multiple LPARs, one with z/OS V1.7 guests and one with a native z/OS system at any level

LPAR1 max. capacity = 156 MSUs		LPAR2 max. capacity = 56 MSUs
z/VM		z/OS
VM Guest 1	VM Guest 2	Native z/OS system
z/OS V1.7 Db2 CICS 4-hour rolling average value = 75 MSUs	z/OS V1.7 Db2 WebSphere Application Server for z/OS V5 4-hour rolling average value = 80 MSUs	z/OS V1.6 Db2 CICS 4-hour rolling average value = 40 MSUs

How MSUs are calculated for example 6

The product MSU value for products running in the z/OS guests is based on a 4-hour rolling average, since both z/OS guests are running z/OS V1.7, even though the native z/OS LPAR is running z/OS V1.6.

Table 22. How MSUs are calculated for example 6: Multiple LPARs; one with z/OS V1.7 guests and one with a native z/OS system at any level

Product	VM Guest 1 MSUs	VM Guest 2 MSUs	Native z/OS MSUs	Product MSUs
Db2	75	80	40	195
WebSphere Application Server	n/a	80	n/a	80
CICS	75	n/a	40	115
z/OS	75	80	40	195

Identities reported for z/OS guest systems

SCRT assigns identities to z/OS systems in order to associate each system with its matching statistics in the following sections on the sub-capacity report:

- Detail Data Collection section
- SMF / SCRT89 Input Data Statistics Section
- Detail LPAR Data Section
- Product Max Contributors section
- Product Grid Snapshot section
- Sub-Capacity Capping Enforced by SCRT section
- Getting Started Sub-Capacity Pricing Assigned section
- Getting Started Sub-Capacity Pricing Equal to Traditional Sub-Capacity section
- Getting Started Sub-Capacity Pricing Not Assigned section

SCRT uses several formats to report the identities of z/OS guest systems based on the levels of the z/OS guest systems running on the CPC. As explained in [“Summary of MSUs reported for sub-capacity eligible products running on z/OS guest systems under z/VM” on page 133](#) and [“Capacity examples for z/OS guest systems” on page 134](#), whenever one or more z/OS V1.6 or earlier guest systems run on a CPC during the reporting period, SCRT treats all of the z/OS guest systems—including z/OS V1.7 and higher guest systems running on that same CPC—at the LPAR level, rather than at the individual guest level. Whenever this condition exists, SCRT assigns the LPAR name in which those guest systems run, as specified by the

VMGUESTS DD control statement, as the single identity for all of the guests running in that LPAR. No guest related information is used as part of the identity.

Whenever all of the guest systems running on the CPC are at z/OS V1.7 or higher, SCRT treats each z/OS guest as a separate z/OS system and assigns a guest identity using the format *lparname(guest_sysid)*. If the LPAR name is not available, SCRT assigns the guest identity using the format **lpar*(guest_sysid)*.

Whenever all of the guest systems running on the CPC are at z/OS V1.7 or higher and are running under the zNALC pricing option, SCRT treats each z/OS guest as a separate z/OS system and assigns a guest identity using the format *lparname(guest_sysid/zNALC)*. If the LPAR name is not available, SCRT assigns the guest identity using the format **lpar*(guest_sysid/zNALC)*.

Sample report showing any z/OS V1.6 or earlier guest systems

If there are *any* z/OS guest systems at V1.6 or earlier, reporting for all z/OS guest systems is at the LPAR level including z/OS V1.7 and higher level guest systems. MSUs for z/OS on that CPC are calculated as the highest combined maximum capacity of:

- All z/OS guest systems, including z/OS V1.7 and higher level guest systems, **plus**
- The 4-hour rolling average utilization across all z/OS native LPARs, if any

For z/OS V1.6 and earlier releases, the SCRT reports z/OS guest information at the LPAR level, using the LPAR names supplied to you in the VMGUESTS DD statement, as shown in the report in this section. When any z/OS guest system on a CPC is at z/OS V1.6 or an earlier release, MSUs for all guest systems are reported at the LPAR level. This also applies to any z/OS V1.7 guest systems that might also be running on the same CPC. This is different than how SCRT treats z/OS guest systems when all of the z/OS guest systems are at a z/OS V1.7 level or higher. See [“Sample sub-capacity report showing only z/OS V1.7 \(or later\) guest systems” on page 141](#) for a description of how SCRT reports on CPCs where all of the z/OS guest systems are at a z/OS V1.7 level or higher.

The following sections show the detail data sections of a sample sub-capacity report that corresponds to the machine configuration described in [Table 23 on page 140](#) (this configuration is not intended to be representative, but rather to show several possibilities in the sub-capacity report.) Only the Detail Data sections are shown because these are the sections where the differences from the reports shown in [“Sample sub-capacity report showing only z/OS V1.7 \(or later\) guest systems” on page 141](#) occur.

- LPAR1 runs a native z/OS system, SYS1
- LPAR2 runs a z/VM system which supports one z/OS V1.7 guest, V2A
- LPAR3 runs a z/VM system which supports one z/OS V1.6 guest, V3A, and one z/OS V1.7 guest, V3B
- LPARVM runs a z/VM system which supports multiple CMS users (not a z/OS system)

Because at least one of the guest systems in the sample configuration is a z/OS V1.6 guest system, the following VMGUESTS DD control statements were used to identify the LPAR in which each z/OS guest, including each z/OS V1.7 guest, ran:

```
LPAR2=V2A
LPAR3=V3A, V3B
```

Table 23. Sample configuration with z/OS V1.6 and V1.7 systems as guests of z/VM

CPC: 2097-405 (serial = 02-12345)				
LPAR1	LPAR2 (z/VM)	LPAR3 (z/VM)		LPARVM (z/VM)
z/OS V1.7 (SYSID = SYS1)	z/OS V1.7 (SYSID = V2A)	z/OS V1.6 (SYSID = V3A)	z/OS V1.7 (SYSID = V3B)	z/VM CMS
CICS TS V3	DB2 UDB V8	DB2 UDB V8	DB2 UDB V8	
IMS V8	MQSeries V6	MQSeries V6	MQSeries V6	
File Manager V6	Cobol V3	WebSphere Application Server for z/OS V6	WebSphere Application Server for z/OS V6	
		WebSphere MQ Workflow V3	WebSphere MQ Workflow V3	

==L5=====

DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY

==M5=====

SMF/SCRT89 INPUT DATA STATISTICS

	SYSID	Input Data Start		Input Data End	
LPAR1	SYS1	01 Oct 2024 - 00:00	02 Nov 2024 - 22:00		
LPAR2		01 Oct 2024 - 00:00	30 Oct 2024 - 22:30		
LPAR3		01 Oct 2024 - 06:05	02 Nov 2024 - 23:45		
LPARVM	-	-	-		
CPC		01 Oct 2024 - 00:00	02 Nov 2024 - 23:45		

==N5=====

DETAIL LPAR DATA SECTION

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR1	50	2	03 Oct 2024 - 07:00	49	5	03 Oct 2024 - 06:00
LPAR2	27	4	04 Oct 2024 - 11:00	26	7	04 Oct 2024 - 10:00
LPAR3	26	5	03 Oct 2024 - 09:00	25	13	03 Oct 2024 - 08:00
CPC	107	5	03 Oct 2024 - 12:00	106	14	03 Oct 2024 - 11:00

Figure 41. Sample sub-capacity report with z/OS guest systems at V1.6 or earlier levels: SMF / SCRT89 Input Data Statistics and Detail LPAR Data sections

```
==P5=====
PRODUCT MAX CONTRIBUTORS
```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR2	LPAR LPAR3
z/OS V1	5694-A01	101	06 Oct 2024 - 11:00	45	28	28
DB2 UDB for z/OS V8	5625-DB2	56	03 Oct 2024 - 18:00	0	28	28
CICS TS for z/OS V3	5655-M15	48	03 Oct 2024 - 07:00	48	0	0
MQSeries for z/OS V6	5655-L82	56	03 Oct 2024 - 18:00	0	28	28
IMS V8	5655-C56	48	03 Oct 2024 - 07:00	48	0	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	28	04 Oct 2024 - 11:00	0	28	0
File Manager for z/OS V6	5655-P17	48	03 Oct 2024 - 07:00	48	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	28	03 Oct 2024 - 13:00	0	0	28
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	28	03 Oct 2024 - 13:00	0	0	28

```
==Q5=====
PRODUCT GRID SNAPSHOT
```

Product Name	Product ID	LPAR1	LPAR2	LPAR3
z/OS V1	5694-A01	92.4%	92.7%	91.5%
DB2 UDB for z/OS V8	5625-DB2		92.7%	91.5%
CICS TS for z/OS V3	5655-M15	92.4%		
MQSeries for z/OS V6	5655-L82		92.7%	91.5%
IMS V8	5655-C56	92.4%		
IBM WebSphere Application Server for z/OS V6	5655-N01			91.5%
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM			91.5%

```
==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS
```

Product Name	Product ID	LPAR LPAR1	LPAR LPAR2	LPAR LPAR3
z/OS V1	5694-A01	2		
DB2 UDB for z/OS V8	5625-DB2			
CICS TS for z/OS V3	5655-M15	2		
MQSeries for z/OS V6	5655-L82			
IMS V8	5655-C56	2		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53			
File Manager for z/OS V6	5655-P17	2		
IBM WebSphere Application Server for z/OS V6	5655-N01			
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM			

Figure 42. Sample sub-capacity report with z/OS guest systems at V1.6 or earlier levels: Product Max Contributors, Product Grid Snapshot, and Sub-Capacity Capping Enforced by SCRT sections

Interpreting the Detail Data sections for V1.6 or earlier guests

This report shows how SCRT reports on CPCs where at least one of the z/OS guest systems is at a z/OS V1.6 level or earlier. SCRT continues to report on all guest systems on such CPCs at the LPAR level on an LPAR-by-LPAR basis, regardless of how many z/OS guests are running in the LPAR. This is different than how SCRT reports on CPCs where all of the guest systems are at a z/OS V1.7 level or higher. For a description of how SCRT determines the product MSUs for such CPCs, see [“Summary of MSUs reported for sub-capacity eligible products running on z/OS guest systems under z/VM” on page 133.](#)

Sample sub-capacity report showing only z/OS V1.7 (or later) guest systems

The following sample scenario is for a CPC with four LPARS named LPAR1, LPAR2, LPAR3, and LPARVM:

- LPAR1 runs a native z/OS system, SYS1
- LPAR2 runs a z/VM system which supports one z/OS V1.7 guest, V2A
- LPAR3 runs a z/VM system which supports two z/OS V1.7 guests, V3A and V3B
- LPARVM runs a z/VM system which supports multiple CMS users (not a z/OS system)

[Table 24 on page 142](#) shows this configuration.

Table 24. Sample configuration with z/OS V1.7 (or later) systems as guests z/VM

CPC: 2097-405 (serial = 02-12345)				
LPAR1	LPAR2 (z/VM)	LPAR3 (z/VM)		LPARVM (z/VM)
z/OS V1.7 (SYSID = SYS1)	z/OS V1.7 (SYSID = V2A)	z/OS V1.7 (SYSID = V3A)	z/OS V1.7 (SYSID = V3B)	z/VM CMS
CICS TS V3	DB2 UDB V8	DB2 UDB V8	DB2 UDB V8	
IMS V8	MQSeries V6	MQSeries V6	MQSeries V6	
File Manager V6	Cobol V3	WebSphere Application Server for z/OS V6	WebSphere Application Server for z/OS V6	
		WebSphere MQ Workflow V3	WebSphere MQ Workflow V3	

The following sections show a sample sub-capacity report that corresponds to the information presented in Table 24 on page 142.

The sub-capacity report contains two major sections:

- The billing-related section contains data that IBM needs to reconcile your company's sub-capacity bills.
- The customer verification section contains information that is useful to you when verifying the information in the sub-capacity report.

Billing-related fields in the sub-capacity report for guest systems

The billing-related section displays customer information, tool information, and product summary information and detail data collection. All of the sections in the billing-related section for z/OS guest systems are identical to those for native z/OS systems, as described in [“Billing-related section”](#) on page 105.

See Table 24 on page 142 for an explanation of how SCRT calculates the product MSUs for CPCs where all of the z/OS guest systems are at z/OS V1.7 or higher.

When all z/OS guest systems on a CPC are at V1.7, reporting for z/OS on that CPC is based on the highest combined 4-hour rolling average utilization across all z/OS guest systems and z/OS native LPARs, if any.

Customer verification section

The customer verification section contains information that supports and expands on the information in the product summary section. It consists of the following sections:

- SMF / SCRT89 Input Data Statistics section, Detail LPAR Data Section, and Product Max Contributors section, shown in [Figure 43 on page 143](#)
- Product Grid Snapshot, shown in [Figure 44 on page 144](#)
- In some cases, the report includes the Sub-Capacity Capping Enforced by SCRT section. This section is shown in [Figure 44 on page 144](#) and is described in [“Sub-Capacity Capping Enforced by SCRT section”](#) on page 124.

Prior to z/OS V1.7, SCRT reported z/OS guest information at the LPAR level using the LPAR name supplied in the VMGUESTS DD statement.

For CPCs where all of the z/OS guests are at z/OS V1.7 or later, SCRT reports on each guest system separately, in the format *lparname(guest_sysid)*. If the LPAR name is not available, SCRT uses the format **lpar*(guest_sysid)*.

The three z/OS guest systems in this sample scenario appear as LPAR2 (V2A), LPAR3 (V3A), and LPAR3 (V3B) in the sample report.

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
LPAR1	SYS1	01 Oct 2024 - 00:00	02 Nov 2024 - 22:00
LPAR2(V2A)	V2A	01 Oct 2024 - 00:00	30 Oct 2024 - 22:30
LPAR3(V3A)	V3A	01 Oct 2024 - 06:05	02 Nov 2024 - 23:45
LPAR3(V3B)	V3B	01 Oct 2024 - 02:25	02 Nov 2024 - 23:45
LPARVM	-	-	-
CPC		01 Oct 2024 - 00:00	02 Nov 2024 - 23:45

```

==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR1	187	1	03 Oct 2024 - 07:00	49	5	03 Oct 2024 - 06:00
LPAR2(V2A)	27	4	04 Oct 2024 - 11:00	26	7	04 Oct 2024 - 10:00
LPAR3(V3A)	13	5	03 Oct 2024 - 09:00	10	13	03 Oct 2024 - 08:00
LPAR3(V3B)	15	8	03 Oct 2024 - 14:00	14	9	03 Oct 2024 - 13:00
CPC	107	5	03 Oct 2024 - 12:00	106	14	03 Oct 2024 - 11:00

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR2(V2A)	LPAR LPAR3(V3A)	LPAR LPAR3(V3B)
z/OS V1	5694-A01	96	06 Oct 2024 - 11:00	45	27	13	11
DB2 UDB for z/OS V8	5625-DB2	51	06 Oct 2024 - 18:00	0	25	13	13
CICS TS for z/OS V3	5655-M15	48	03 Oct 2024 - 07:00	48	0	0	0
MQSeries for z/OS V6	5655-L82	51	03 Oct 2024 - 18:00	0	25	13	13
IMS V8	5655-C56	48	03 Oct 2024 - 07:00	48	0	0	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	27	04 Oct 2024 - 11:00	0	27	0	0
File Manager for z/OS V6	5655-P17	48	03 Oct 2024 - 07:00	48	0	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	27	03 Oct 2024 - 13:00	0	0	13	14
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	27	03 Oct 2024 - 13:00	0	0	13	14

Figure 43. Sample sub-capacity report showing only z/OS V1.7 (or later) guest systems: SMF / SCRT89 Input Data Statistics, Detail LPAR Data, and Product Max Contributors sections

```
==Q5=====
PRODUCT GRID SNAPSHOT
```

Product Name	Product ID	LPAR1	LPAR2(V2A)	LPAR3(V3A)	LPAR3(V3B)
z/OS V1	5694-A01	92.4%	92.7%	91.5%	91.9%
DB2 UDB for z/OS V8	5625-DB2		92.7%	91.5%	91.9%
CICS TS for z/OS V3	5655-M15	92.4%			
MQSeries for z/OS V6	5655-L82		92.7%	91.5%	91.9%
IMS V8	5655-C56	92.4%			
IBM WebSphere Application Server for z/OS V6	5655-N01			91.5%	91.9%
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM			91.5%	91.9%

```
==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS
```

Product Name	Product ID	LPAR LPAR1	LPAR LPAR2(V2A)	LPAR LPAR3(V3A)	LPAR LPAR(V3B)
z/OS V1	5694-A01	2			
DB2 UDB for z/OS V8	5625-DB2				
CICS TS for z/OS V3	5655-M15	2			
MQSeries for z/OS V6	5655-L82				
IMS V8	5655-C56	2			
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53				
File Manager for z/OS V6	5655-P17	2			
IBM WebSphere Application Server for z/OS V6	5655-N01				
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM				

Figure 44. Sample sub-capacity report showing only z/OS V1.7 (or later) guest systems: Product Grid Snapshot and Sub-Capacity Capping Enforced by SCRT sections

Chapter 7. Sub-capacity reporting for System z New Application License Charges (zNALC)

System z New Application License Charges (zNALC) pricing applies to the z/OS base feature and z/OS priced features. zNALC is available for z/OS on LPARs dedicated to qualified applications, among other requirements. See [“IBM Z New Application License Charges \(zNALC\)” on page 7](#) for details. For more information about z/OS zNALC guest systems see [“Sub-capacity reporting for z/OS zNALC guest systems” on page 153](#).

Note: All of the examples in this chapter are based on the assumption that the customer has signed a zNALC pricing contract with IBM for the CPCs being reported. Identifying zNALC systems by naming LPARs ZNALxxxx or IPLing with the LICENSE=zNALC parameter is not sufficient to qualify for zNALC pricing. zNALC pricing also requires a signed zNALC contract which describes the conditions under which zNALC pricing is available.

Whenever SCRT encounters SMF data either from a z/OS system that has been IPLed in an LPAR named ZNALxxxx, or from a z/OS system that has been IPLed with the LICENSE=zNALC parameter (with any LPAR name), SCRT will report z/OS MSUs either by listing separate values for z/OS, z/OS (Traditional), and z/OS (zNALC), or by just listing the z/OS (zNALC) value, depending on the z/OS configuration being reported. These SCRT changes occur to the following sections of the sub-capacity report:

- Product Summary Information
- Product Max Contributors
- Product Grid Snapshot
- Sub-Capacity Capping Enforced by SCRT

SCRT will also identify the LPARs in which a zNALC system was running by adding the (zNALC) suffix to those LPAR names. Whenever you use the LICENSE=zNALC IPL parameter to run z/OS traditional workloads (that is, workloads that do not qualify for zNALC pricing) and zNALC workloads in the same LPAR at different times of the month, SCRT will separate the LPAR-related statistics into statistics for the traditional z/OS LPAR and the zNALC LPAR and report on each separately. The zNALC LPAR will have a suffix of (zNALC). These changes will appear in the following sections of the sub-capacity report:

- Detail Data Collection
- SMF / SCRT89 Input Data Statistics
- Detail LPAR Data Section
- Product Max Contributors
- Product Grid Snapshot
- Sub-Capacity Capping Enforced by SCRT
- Group Capacity LPARs

Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads

Table 25 on page 146 contains the sample hardware and software configuration for a machine where native z/OS systems ran both traditional and zNALC workloads during the reporting period. LPAR1 only ran z/OS traditional workloads. LPAR3 ran z/OS traditional workloads until 17 October at 6:00. A z/OS system was then IPLed in LPAR3 with LICENSE=zNALC and ran zNALC workloads for the remainder of the month. LPARSAP ran the zNALC workload of Db2 as a database client for SAP.

Table 25. Sample configuration for native z/OS systems running traditional and zNALC workloads

CPC: 2097-405 (serial = 02-12345)				
LPAR1	LPAR3		ZNALSAP	LPARVM
2 October – 1 November	2 October – 17 October	17 October – 1 November	2 October – 1 November	2 October – 1 November
z/OS V1 (SYSID = SYS1) DB2 UDB V8 CICS TS V3 IMS V8 MQSeries V6 File Manager V6	z/OS V1 (SYSID = SYS3) DB2 UDB V8 CICS TS V3 Cobol V3 File Manager V6 WebSphere Application Server V6 WebSphere MQ Workflow V3	z/OS V1 (SYSID = SYS3) (LICENSE = ZNALC) DB2 UDB V8 Cobol V3 WebSphere Application Server V6 WebSphere MQ Workflow V3	z/OS V1 (SYSID = SAP1) DB2 UDB V8	CMS

The following figures show all of the sub-capacity report sections that were generated for the hardware and software configuration shown in [Table 25 on page 146](#). These sections show two sets of changes compared with the sub-capacity report for the base z/OS configuration, which is described in [Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103](#).

When the SMF data is from a CPC which runs both traditional z/OS workloads and zNALC workloads (either the LPAR name is ZNALxxxx or z/OS was IPLed with LICENSE=ZNALC), SCRT reports three values for z/OS:

z/OS

This value represents the MSU values for z/OS running in all LPARs on the CPC. This value is used when assigning capacity to the z/OS based IPLA sub-capacity products running on this CPC.

z/OS (Traditional)

This value represents the MSU value charged against traditional z/OS use.

z/OS (zNALC)

This value is used to determine what MSU value is charged against zNALC use.

When the difference between the z/OS value and the z/OS (Traditional) value is smaller than the zNALC value, then the difference is used, unless the difference is less than three MSUs. See [z Systems New Application License Charges \(zNALC\) \(www.ibm.com/downloads/cas/WGVAN8A2\)](#) for the rules that govern the case for when the resulting zNALC value is less than three MSUs.

SCRT will identify the LPARs in which a zNALC system was running by adding the (zNALC) suffix to those LPAR names. In the sample configuration the customer used LPAR3 to run traditional z/OS workloads until 17 October at 6:00 AM. The customer then IPLed a z/OS zNALC system using LICENSE=ZNALC and ran zNALC workloads in LPAR3 for the remainder of the month. Note that SCRT separated the LPAR3 statistics into statistics for the traditional z/OS workloads and statistics for the zNALC workloads and reported on each separately. The statistics for the traditional z/OS workloads running in LPAR3 are listed under LPAR3. The statistics for the z/OS zNALC workloads are listed under LPAR3 (zNALC). ZNALSAP ran zNALC workloads for the entire month and so all of the statistics for ZNALSAP are listed under ZNALSAP (zNALC).

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 03 Nov 2024 - 09:17
 Name of Person Submitting Report: John Customer
 E-mail Address of Report Submitter: customer@abc.com
 Phone Number of Report Submitter: 444-999-9999

Customer Name ABC Corporation
 Customer Number 8971234567
 Machine Serial Number 02-12345
 Machine Type and Model 2097-405
 Machine Rated Capacity (MSUs) 118
 Purchase Order Number (optional)
 Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

==C5=====

TOOL INFORMATION

Tool Release 30.1.0
 Reporting Period 2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

Figure 45. Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads: Customer information and Tool Information sections

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS V1	5694-A01	98		
z/OS V1 (Traditional)	5694-A01	60		
z/OS V1 (zNALC)	5694-A01	57		
DB2 UDB for z/OS V8	5625-DB2	98		
CICS TS for z/OS V3	5655-M15	60		
MQSeries for z/OS V6	5655-L82	32		
IMS V8	5655-C56	32		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	35	(optional)	(conditional)
IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
File Manager for z/OS V6	5655-P17	60	(optional)	(conditional)
IBM WebSphere Application Server for z/OS V6	5655-N01	35		
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	35		

Figure 46. Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads: Product Summary Information section

```

=====
==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
=====
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
LPAR1	SYS1	01 Oct 2024 - 00:00	02 Nov 2024 - 22:00
LPAR3	SYS3	01 Oct 2024 - 06:05	17 Oct 2024 - 06:00
LPAR3(zNALC)	SYS3	17 Oct 2024 - 06:15	02 Nov 2024 - 23:45
ZNALSAP(zNALC)	SAP1	01 Oct 2024 - 00:00	30 Oct 2024 - 22:30
LPARVM	-	-	-
CPC		01 Oct 2024 - 00:00	02 Nov 2024 - 23:45

```

=====
==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR1	32	1	03 Oct 2024 - 07:00	30	5	03 Oct 2024 - 06:00
LPAR3	35	6	03 Oct 2024 - 14:00	34	11	03 Oct 2024 - 13:00
LPAR3(zNALC)	17	2	18 Oct 2024 - 16:00	16	3	18 Oct 2024 - 15:00
ZNALSAP(zNALC)	41	4	04 Oct 2024 - 11:00	40	7	04 Oct 2024 - 10:00
CPC	107	5	03 Oct 2024 - 12:00	106	9	03 Oct 2024 - 11:00

Figure 47. Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads: SMF / SCRT89 Input Data Statistics section and Detail LPAR Data section

```

=====
==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR3	LPAR LPAR3(zNALC)	LPAR ZNALSAP(zNALC)	GROUP LGRPA
z/OS V1	5694-A01	98	06 Oct 2024 - 11:00	27	32	0	39	0
z/OS V1 (Traditional)	5694-A01	60	03 Oct 2024 - 07:00	32	28	0	0	0
z/OS V1 (zNALC)	5694-A01	57	18 Oct 2024 - 16:00	0	0	17	40	0
DB2 UDB for z/OS V8	5625-DB2	98	06 Oct 2024 - 11:00	27	32	0	39	0
CICS TS for z/OS V3	5655-M15	60	03 Oct 2024 - 07:00	32	28	0	0	0
MQSeries for z/OS V6	5655-L82	32	03 Oct 2024 - 07:00	32	0	0	0	0
IMS V8	5655-C56	32	03 Oct 2024 - 07:00	32	0	0	0	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	35	03 Oct 2024 - 14:00	0	35	0	0	0
File Manager for z/OS V6	5655-P17	60	09 Oct 2024 - 15:00	29	31	0	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	35	03 Oct 2024 - 14:00	0	35	0	0	0
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	35	03 Oct 2024 - 14:00	0	35	0	0	0

Figure 48. Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads: Product Max Contributors section

```

==Q5=====
PRODUCT GRID SNAPSHOT

```

Product Name	Product ID	LPAR1	LPAR3	LPAR3 (zNALC)	ZNALSAP (zNALC)
z/OS V1	5694-A01	92.4%	44.9%	46.7%	91.5%
z/OS V1 (Traditional)	5694-A01	92.4%	44.9%		
z/OS V1 (zNALC)	5694-A01			46.7%	91.5%
DB2 UDB for z/OS V8	5625-DB2	92.4%	44.9%	46.9%	91.5%
CICS TS for z/OS V3	5655-M15	92.4%	44.9%		
MQSeries for z/OS V6	5655-L82	92.4%			
IMS V8	5655-C56	92.4%			
IBM WebSphere Application Server for z/OS V6	5655-N01		44.9%	46.7%	
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		44.9%	46.7%	

```

==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS

```

Product Name	Product ID	LPAR LPAR1	LPAR LPAR3	LPAR LPAR3 (zNALC)	LPAR ZNALSAP (zNALC)	Group LGRPA
z/OS V1	5694-A01	2	1		4	1
z/OS V1 (Traditional)	5694-A01	2	1			
z/OS V1 (zNALC)	5694-A01				4	1
DB2 UDB for z/OS V8	5625-DB2	2	1		4	1
CICS TS for z/OS V3	5655-M15	2	1			
MQSeries for z/OS V6	5655-L82	2				
IMS V8	5655-C56	2				
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53		1			
File Manager for z/OS V6	5655-P17	2	1			
IBM WebSphere Application Server for z/OS V6	5655-N01		1			
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		1			

Figure 49. Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads: Product Grid Snapshot and Sub-Capacity Capping Enforced by SCRT sections

```

==U5=====
GROUP CAPACITY LPARS

```

Group Name	LPARS
LGRPA	LPAR3(iii), LPAR3(zNALC)(iii), ZNALSAP(iii)

(iii) LPAR was not a member of the group for the entire reporting interval

```

==W5=====
This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.

```

Figure 50. Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads: Group Capacity LPARs section

Sample sub-capacity report for z/OS V1 and z/OS V2 running traditional and zNALC workloads

Table 26 on page 150 contains the sample hardware and software configuration for a CPC with z/OS V1 traditional and zNALC workloads that is in the process of migrating both sets of workloads to z/OS V2. This example shows how SCRT reports on the combination of z/OS V1 and z/OS V2 systems running both traditional and zNALC workloads.

The LPAR names indicate both the type of workload (traditional or zNALC) and the version of the z/OS system (V1 or V2) that was running in each LPAR. (The names of the zNALC LPARs each start with ZNALxxxx, so no LICENSE=ZNALC parameter was required at IPL time.)

Table 26. Sample configuration for native z/OS V1 and z/OS V2 systems running traditional and zNALC workloads

CPC: 2097-405 (serial = 02-12345)			
TRADV1	TRADV2	ZNALCV1	ZNALCV2
z/OS V1 (SYSID = SYS1)	z/OS V2 (SYSID = SYS2)	z/OS V1 (SYSID = SYS3)	z/OS V2 (SYSID = SYS4)
DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8	DB2 UDB V8 CICS TS V3	DB2 UDB V8	DB2 UDB V8
File Manager V6	Cobol V3 File Manager V6	Cobol V3	Cobol V3
	WebSphere Application Server for z/OS V6	WebSphere Application Server for z/OS V6	WebSphere Application Server for z/OS V6
	WebSphere MQ Workflow V3	WebSphere MQ Workflow V3	WebSphere MQ Workflow V3

The Customer Information section and Tool Information section are the same as those shown in “Sample sub-capacity report for native z/OS systems running traditional and zNALC workloads” on page 145, so those sections are not repeated here. There are no error conditions and no SPECIAL SCRT control statements were used, so no Special Conditions section was generated.

Product Summary Information section

Figure 51 on page 150 shows the Product Summary Information section that was generated for the configuration in Table 26 on page 150.

```

==E5=====
PRODUCT SUMMARY INFORMATION

MLC Product Name                MLC Product ID    Tool MSUs    Customer MSUs    Customer Comments (255 chars max)

z/OS                            (A11)              98
z/OS (Traditional)              (A11)              60
z/OS (zNALC)                   (A11)              42

z/OS V2                         5650-ZOS           56
z/OS V2 (Traditional)          5650-ZOS           35
z/OS V2 (zNALC)                5650-ZOS           24

z/OS V1                         5694-A01           48
z/OS V1 (Traditional)          5694-A01           32
z/OS V1 (zNALC)                5694-A01           22

DB2 UDB for z/OS V8             5625-DB2           98
CICS TS for z/OS V3             5655-M15           60
MQSeries for z/OS V6            5655-L82           32
IMS V8                          5655-C56           32
IBM Enterprise Cobol for z/OS and OS/390 V3 5655-G53           35      (required)    (conditional)

IPLA Product Name              IPLA Product ID    Tool MSUs    Customer MSUs    Customer Comments (255 chars max)

File Manager for z/OS V6        5655-P17           60      (required)    (conditional)
IBM WebSphere Application Server for z/OS V6 5655-N01           35
IBM WebSphere MQ Workflow for z/OS V3      5655-BPM           35

```

Figure 51. Sample sub-capacity report for native z/OS V1 and z/OS V2 systems running traditional and zNALC workloads: Product Summary Information section

Since z/OS V1 and z/OS V2 were running during the reporting period, all of the z/OS V1 (5694-A01) product entries (overall, traditional, zNALC) and all of the z/OS V2 (5650-ZOS) product entries are reported in the Product Summary Information section along with an additional set of z/OS (A11) product entries corresponding to each of the z/OS overall, traditional and zNALC product entries for each version of z/OS. A z/OS entry will always be reported whenever a z/OS V1 and z/OS V2 system run during the same reporting period. A z/OS (Traditional) entry will only be reported when at least one z/OS V1 or z/OS V2 system was running a traditional workload during that reporting period; otherwise, the z/OS (Traditional) entry will not be reported. A z/OS (zNALC) entry will only be reported when at least one z/OS V1 or z/OS V2 system was running a zNALC workload during that reporting period; otherwise, the z/OS (zNALC) entry will not be reported.

The additional set of z/OS (A11) product entries (overall, traditional, zNALC) reports the peak concurrent sub-capacity value for the corresponding z/OS V1 and z/OS V2 product entries, as if those entries were a single product. If the z/OS V1 product is under Multi-Version Measurement (MVM), IBM waives the charges for the z/OS V1 product and will charge for z/OS V2 based on the traditional and zNALC sub-capacity values reported for the z/OS (A11) entries.

This example includes both z/OS V1 and z/OS V2 systems running traditional and zNALC workloads, and z/OS V1 is under MVM, so IBM charges for z/OS V2 traditional based on the z/OS (Traditional) value of 60 MSUs and will charge for z/OS V2 zNALC based on the MSUs calculated by subtracting the z/OS (Traditional) value from the z/OS value (98 - 60 = 38 MSUs). IBM does not charge for the z/OS V1 product under MVM rules. (See “[Sub-capacity reporting for multiple versions of z/OS](#)” on [page 127](#) for more information, including a detailed explanation of MVM for multiple z/OS versions.)

Note that the z/OS (A11) entry is reported whenever both z/OS V1 and z/OS V2 are running on the same CPC at any time during the reporting period.

You can find more information about the fields in this section and other sections of the sub-capacity report in “[Sample sub-capacity report for single-version z/OS native systems](#)” on [page 105](#).

There is no notable impact on the Detail Data Collection section, SMF / SCRT89 Input Data Statistics section, and Detail LPAR Data section for this scenario other than differences due to the different configuration, so these report sections are not shown here.

Product Max Contributors section

[Figure 52 on page 151](#) shows the Product Max Contributors section for this sample scenario.

```
==P5=====
PRODUCT MAX CONTRIBUTORS
```

Product Name	Product ID	Highest	Date/Time	LPAR TRADV1	LPAR TRADV2	LPAR ZNALCV1(ZNALC)	LPAR ZNALCV2(ZNALC)	GROUP LGRPA
z/OS	(A11)	98	06 Oct 2024 - 11:00	27	32	18	21	0
z/OS (Traditional)	(A11)	60	03 Oct 2024 - 07:00	32	28	0	0	0
z/OS (zNALC)	(A11)	42	18 Oct 2024 - 16:00	0	0	19	23	0
z/OS V2	5650-Z0S	56	05 Oct 2024 - 13:00	0	34	0	22	0
z/OS V2 (Traditional)	5650-Z0S	35	03 Oct 2024 - 14:00	0	35	0	0	0
z/OS V2 (zNALC)	5650-Z0S	24	12 Oct 2024 - 11:00	0	0	0	24	0
z/OS V1	5694-A01	48	04 Oct 2024 - 08:00	29	0	19	0	0
z/OS V1 (Traditional)	5694-A01	32	04 Oct 2024 - 10:00	32	0	0	0	0
z/OS V1 (zNALC)	5694-A01	22	13 Oct 2024 - 15:00	0	0	22	0	0
DB2 UDB for z/OS V8	5625-DB2	98	06 Oct 2024 - 11:00	27	32	0	39	0
CICS TS for z/OS V3	5655-M15	60	03 Oct 2024 - 07:00	32	28	0	0	0
MQSeries for z/OS V6	5655-L82	32	03 Oct 2024 - 07:00	32	0	0	0	0
IMS V8	5655-C56	32	03 Oct 2024 - 07:00	32	0	0	0	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	35	03 Oct 2024 - 14:00	0	35	0	0	0
File Manager for z/OS V6	5655-P17	60	09 Oct 2024 - 15:00	29	31	0	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	35	03 Oct 2024 - 14:00	0	35	0	0	0
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	35	03 Oct 2024 - 14:00	0	35	0	0	0

Figure 52. Sample sub-capacity report for native z/OS V1 and z/OS V2 systems running traditional and zNALC workloads: Product Max Contributors section

The entire set of z/OS (overall, traditional, zNALC) entries and the entire set of z/OS V1 and z/OS V2 product entries that are listed in the Product Summary Information section are also listed in the Product Max Contributors section in order to report each z/OS system’s individual MSU contribution to each of the sub-capacity values listed in the Product Summary Information section.

You can find more information about the fields in this section and other sections of the sub-capacity report in “[Sample sub-capacity report for single-version z/OS native systems](#)” on [page 105](#).

Product Grid Snapshot section

[Figure 53 on page 152](#) shows the Product Grid Snapshot section for this sample scenario.

```

=====
PRODUCT GRID SNAPSHOT

```

Product Name	Product ID	TRADV1	TRADV2	ZNALCV1(ZNALC)	ZNALCV2(ZNALC)
z/OS	(All)	92.4%	44.9%	46.7%	91.5%
z/OS (Traditional)	(All)	92.4%	44.9%		
z/OS (zNALC)	(All)			46.7%	91.5%
z/OS V2	5650-ZOS		44.9%		91.5%
z/OS V2 (Traditional)	5650-ZOS		44.9%		
z/OS V2 (zNALC)	5650-ZOS				91.5%
z/OS V1	5694-A01	92.4%		46.7%	
z/OS V1 (Traditional)	5694-A01	92.4%			
z/OS V1 (zNALC)	5694-A01			46.7%	
DB2 UDB for z/OS V8	5625-DB2	92.4%	44.9%	46.9%	91.5%
CICS TS for z/OS V3	5655-M15	92.4%	44.9%		
MQSeries for z/OS V6	5655-L82	92.4%			
IMS V8	5655-C56	92.4%			
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53		44.9%		
File Manager for z/OS V6	5655-P17	92.4%	44.9%		
IBM WebSphere Application Server for z/OS V6	5655-N01		44.9%	46.7%	
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		44.9%	46.7%	

Figure 53. Sample sub-capacity report for native z/OS V1 and z/OS V2 systems running traditional and zNALC workloads: Product Grid Snapshot section

The entire set of z/OS, z/OS V1, and z/OS V2 entries, and middleware product entries that are listed in the Product Summary Information section are also reported in the Product Grid Snapshot section.

You can find more information about the fields in this section and other sections of the sub-capacity report in [“Sample sub-capacity report for single-version z/OS native systems”](#) on page 105.

Sub-capacity Capping Enforced by SCRT section

Figure 54 on page 152 shows the Sub-capacity Capping Enforced by SCRT section for this sample scenario.

```

=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS

```

Product Name	Product ID	LPAR TRADV1	LPAR TRADV2	LPAR ZNALCV1(zNALC)	LPAR ZNALCV2(zNALC)	Group LGRPA
z/OS	(All)	2	1	4		1
z/OS (Traditional)	(All)	2	1			
z/OS (zNALC)	(All)			4		1
z/OS V2	5650-ZOS		1			
z/OS V2 (Traditional)	5650-ZOS		1			
z/OS V2 (zNALC)	5650-ZOS					
z/OS V1	5694-A01	2		4		1
z/OS V1 (Traditional)	5694-A01	2				
z/OS V1 (zNALC)	5694-A01			4		1
DB2 UDB for z/OS V8	5625-DB2	2	1	4		1
CICS TS for z/OS V3	5655-M15	2	1			
MQSeries for z/OS V6	5655-L82	2				
IMS V8	5655-C56	2				
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53		1			
File Manager for z/OS V6	5655-P17	2	1			
IBM WebSphere Application Server for z/OS V6	5655-N01					
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM					

Figure 54. Sample sub-capacity report for native z/OS V1 and z/OS V2 systems running traditional and zNALC workloads: Sub-capacity Capping Enforced by SCRT section

The Sub-Capacity Capping Enforced by SCRT section shows how many hours SCRT enforced either the cap for an individual LPAR or the cap for the LPARs in an LPAR group for each of the z/OS, z/OS V1, and z/OS V2 entries, and for the middleware product entries that are listed in the Product Summary Information section.

You can find more information about the fields in this section and other sections of the sub-capacity report in [“Sample sub-capacity report for single-version z/OS native systems”](#) on page 105.

Sample product summary information section for a CPC only running zNALC workloads

Table 27 on page 153 contains the sample hardware and software configuration for a CPC where native z/OS systems only ran zNALC workloads during the reporting period.

CPC: 2097-405 (serial = 02-12345)	
ZNALWAS	ZNALSAP
z/OS V1 (SYSID = WAS1) DB2 UDB V8 Cobol V3 File Manager V6 WebSphere Application Server for z/OS V6 WebSphere MQ Workflow V3	z/OS V1 (SYSID = SAP1) DB2 UDB V8

When the SMF data is from a CPC where all of the z/OS systems are running zNALC workloads (that is, the z/OS systems are running in LPARs named ZNALxxxx or have been IPLed with LICENSE=ZNALC), then SCRT reports all of the z/OS values as zNALC values and lists z/OS (zNALC) as the name in the **MLC Product Name** column of the sub-capacity report. In this case, IBM charges zNALC pricing for all of the z/OS use on the CPC. See Figure 55 on page 153. Note that the sample report applies to both native z/OS zNALC systems and to z/OS zNALC guest systems.

```
==E5=====
```

PRODUCT SUMMARY INFORMATION				
MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS V1 (zNALC)	5694-A01	98		
DB2 UDB for z/OS V8	5625-DB2	71		
CICS TS for z/OS V3	5655-M15	98		
MQSeries for z/OS V6	5655-L82	71		
IMS V8	5655-C56	98		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	39	(optional)	(conditional)
IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
File Manager for z/OS V6	5655-P17	60	(optional)	(conditional)
IBM WebSphere Application Server for z/OS V6	5655-N01	17		
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	17		

Figure 55. Sample sub-capacity report: product summary information for a CPC with zNALC only

Sub-capacity reporting for z/OS zNALC guest systems

SCRT uses the same rules to recognize zNALC systems regardless of whether those systems are native zNALC systems or zNALC guest systems running under z/VM. All zNALC native and guest systems are either running in LPARs named ZNALxxxx or have been IPLed with the LICENSE=ZNALC parameter.

Whenever SCRT encounters SMF data from zNALC systems, it reports the sub-capacity MSUs for such systems as zNALC MSUs, regardless of whether those systems are native zNALC systems or zNALC guest systems. When SCRT encounters SMF data from a CPC that has a combination of z/OS traditional workloads and zNALC workloads, it lists separate MSU values for z/OS, z/OS (Traditional), and z/OS (zNALC). The list of sub-capacity report sections, which report on all three z/OS MSU values, appears at the beginning of this chapter.

SCRT also uses the (zNALC) suffix to identify all zNALC systems (native and guest systems) in several sections of the sub-capacity report. The list of sub-capacity report sections that identify all zNALC systems using the (zNALC) suffix appears at the beginning of this chapter.

There is one exception where SCRT does not use the (zNALC) suffix. Just like it does for native z/OS systems, SCRT will report on a combination of traditional z/OS guests and zNALC guests, each running in different LPARs on the same CPC. SCRT also supports a combination of z/OS traditional guests and z/OS zNALC guests running in the same LPAR, either at different times or at the same time. However, one combination of z/OS traditional and z/OS zNALC workloads must be reported slightly differently than all of the others. Whenever, a combination of z/OS traditional guests and z/OS zNALC guests run in the same LPAR and one of those guests is running z/OS V1.6 (or earlier releases), SCRT does not report on zNALC using the (zNALC) suffix because SCRT reports all of the statistics for z/OS V1.6 LPARs at the LPAR level and does not provide a further break down within that LPAR. However, SCRT does report the z/OS, z/OS (Traditional), and z/OS (zNALC) MSU values for that case.

Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in different LPARs

Table 28 on page 154 contains the sample hardware and software configuration for a machine where z/OS V1.6 guest systems ran traditional and zNALC workloads in different LPARs during the same reporting period. One z/OS V1.6 guest (SYSID S16A) ran a traditional workload in LPAR5 and another z/OS V1.6 guest (SYSID S16B) IPLed with LICENSE=zNALC and ran a zNALC workload in LPAR6. This example also includes a native z/OS V1.7 system that ran a traditional workload in LPAR3.

CPC: 2097-406 (serial = 02-12345)			
LPAR3	LPAR5 (z/VM)	LPAR6 (z/VM)	LPARVM (z/VM)
z/OS V1.7 (SYSID = SYS3) DB2 UDB V8 CICS TS V3 IMS V8 MQSeries V6 File Manager V6	z/OS V1.6 (SYSID = S16A) DB2 UDB V8 CICS TS V3 Cobol V3 File Manager V6 WebSphere Application Server V6 WebSphere MQ Workflow V3	z/OS V1.6 (SYSID = S16B) (LICENSE = ZNALC) Cobol V3 WebSphere Application Server V6 WebSphere MQ Workflow V3	CMS

The following figures show all of the sub-capacity report sections that were generated for the hardware and software configuration shown in Table 28 on page 154. These sections show how SCRT reports on a z/OS guest environment where traditional and zNALC workloads were run in different LPARs.

Note that since at least one z/OS V1.6 guest was running on this CPC, SCRT will use the z/OS V1.6 guest format to report all guest data, including data from z/OS V1.7 (and later) guests, if present—that is, all guests will be reported under the name of the LPAR assigned to that guest by the associated VMGUESTS control statements. In this example the VMGUESTS control statements defined the guest with SYSID S16A as running in LPAR5 and the guest with SYSID S16B as running in LPAR6. If guest S16B were a z/OS V1.7 guest, instead of a z/OS V1.6 guest, the same VMGUESTS statement would be required for the z/OS V1.7 guest system because SCRT treats all guest as z/OS V1.6 guests on a CPC where at least one z/OS V1.6 guest is present.

Note, also, that SCRT will assign the maximum capacity of the LPAR to all of the products running in the guest system running on that LPAR because at least one z/OS V1.6 (or earlier) guest was running on that CPC. “Sample sub-capacity report for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads” on page 161 shows an example of how SCRT reports on a CPC where all of the z/OS guests are at z/OS V1.7 or later.

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 03 Nov 2024 - 09:17
 Name of Person Submitting Report: John Customer
 E-mail Address of Report Submitter: customer@abc.com
 Phone Number of Report Submitter: 444-999-9999

Customer Name ABC Corporation
 Customer Number 8971234567
 Machine Serial Number 02-12345
 Machine Type and Model 2097-406
 Machine Rated Capacity (MSUs) 139
 Purchase Order Number (optional)
 Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

==C5=====

TOOL INFORMATION

Tool Release 30.1.0
 Reporting Period 2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

Figure 56. Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in different LPARs: Customer information and Tool Information sections

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS V1	5694-A01	117		
z/OS V1 (Traditional)	5694-A01	89		
z/OS V1 (zNALC)	5694-A01	28		
DB2 UDB for z/OS V8	5625-DB2	117		
CICS TS for z/OS V3	5655-M15	89		
MQSeries for z/OS V6	5655-L82	34		
IMS V8	5655-C56	34		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	83	(optional)	(conditional)
IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
File Manager for z/OS V6	5655-P17	89	(optional)	(conditional)
IBM WebSphere Application Server for z/OS V6	5655-N01	83		
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	83		

Figure 57. Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in different LPARs: Product Summary Information section

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
LPAR3	SYS3	01 Oct 2024 - 06:05	02 Nov 2024 - 23:45
LPAR5	-	01 Oct 2024 - 00:00	02 Nov 2024 - 23:45
LPAR6(zNALC)	-	01 Oct 2024 - 00:00	02 Nov 2024 - 23:00
LPARVM	-	-	-
CPC		01 Oct 2024 - 00:00	02 Nov 2024 - 23:45

```

==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR3	34	6	03 Oct 2024 - 14:00	32	11	03 Oct 2024 - 06:00
LPAR5	55	715	04 Oct 2024 - 11:00			
LPAR6(zNALC)	28	725	17 Oct 2024 - 06:30			
CPC	117	5	03 Oct 2024 - 12:00	96	11	03 Oct 2024 - 13:00

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR3	LPAR LPAR5	LPAR LPAR6(zNALC)
z/OS V1	5694-A01	117	03 Oct 2024 - 14:00	34	55	28
z/OS V1 (Traditional)	5694-A01	89	03 Oct 2024 - 14:00	34	55	0
z/OS V1 (zNALC)	5694-A01	28	02 Oct 2024 - 00:00	0	0	28
DB2 UDB for z/OS V8	5625-DB2	117	03 Oct 2024 - 14:00	34	55	28
CICS TS for z/OS V3	5655-M15	89	03 Oct 2024 - 14:00	34	55	0
MQSeries for z/OS V6	5655-L82	34	03 Oct 2024 - 14:00	34	0	0
IMS V8	5655-C56	34	03 Oct 2024 - 14:00	34	0	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	83	02 Oct 2024 - 00:00	0	55	28
File Manager for z/OS V6	5655-P17	89	03 Oct 2024 - 14:00	34	55	0
IBM WebSphere Application Server for z/OS V6	5655-N01	83	02 Oct 2024 - 00:00	0	55	28
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	83	02 Oct 2024 - 00:00	0	55	28

Figure 58. Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in different LPARs: SMF / SCRT89 Input Data Statistics, Detail LPAR Data Section, and Product Max Contributors sections

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                                Product ID    LPAR3    LPAR5    LPAR6(zNALC)
z/OS V1                                    5694-A01      91.3%    92.1%      90.4%
z/OS V1 (Traditional)                     5694-A01      91.3%    92.1%
z/OS V1 (zNALC)                           5694-A01      91.3%    92.1%      90.4%
DB2 UDB for z/OS V8                       5625-DB2      91.3%    92.1%      90.4%
CICS TS for z/OS V3                       5655-M15      91.3%    92.1%
MQSeries for z/OS V6                      5655-L82      91.3%
IMS V8                                    5655-C56      91.3%
IBM WebSphere Application Server for z/OS V6 5655-N01      92.1%      90.4%
IBM WebSphere MQ Workflow for z/OS V3      5655-BPM      92.1%      90.4%

```

```

==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS

Product Name                                Product ID    LPAR    LPAR    LPAR
                                                LPAR3      LPAR5    LPAR6(zNALC)
z/OS V1                                    5694-A01      8        34
z/OS V1 (Traditional)                     5694-A01      8        34
z/OS V1 (zNALC)                           5694-A01      8        0
DB2 UDB for z/OS V8                       5625-DB2      8        34
CICS TS for z/OS V3                       5655-M15      8        34
MQSeries for z/OS V6                      5655-L82      8        34
IMS V8                                    5655-C56      8        34
IBM Enterprise Cobol for z/OS and OS/390 V3 5655-G53      8        0
File Manager for z/OS V6                  5655-P17      8        34
IBM WebSphere Application Server for z/OS V6 5655-N01      8        0
IBM WebSphere MQ Workflow for z/OS V3      5655-BPM      8        0

```

```

==W5=====

LPAR5 includes z/OS guest sysids S16A.
LPAR6 includes z/OS guest sysids S16B.

This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.

```

Figure 59. Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in different LPARs: Product Grid Snapshot and Sub-Capacity Capping Enforced by SCRT sections

Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in the same LPAR

This topic deals with the special z/OS guest case where a combination of z/OS traditional guests and z/OS zNALC guests run in the same LPAR and one of those guests is running z/OS V1.6 (or earlier releases). The sub-capacity report for this case is slightly different than the report for any of the other cases for zNALC systems because it is the only case where SCRT does not report on zNALC using the (zNALC) suffix. SCRT does not use the (zNALC) suffix because SCRT always reports all of the statistics for z/OS V1.6 LPARs at the LPAR level. Except for reporting the z/OS, z/OS (Traditional), and z/OS (zNALC) MSUs, SCRT does not provide a further breakdown between traditional and zNALC within that LPAR. That means that the (zNALC) suffix will not appear in the following sections of the sub-capacity report, even though a zNALC workload was running in one of those LPARs:

- SMF / SCRT89 Input Data Statistics
- Detail LPAR Data Section
- Product Max Contributors
- Product Grid Snapshot
- Sub-Capacity Capping Enforced by SCRT
- Group Capacity LPARs

Table 29 on page 158 contains the sample hardware and software configuration for a CPC where the customer from the previous example migrated from a z/OS V1.6 guest system running traditional workloads to a z/OS system running zNALC qualifying workloads in the same LPAR (LPAR6). The customer identified the z/OS system that ran with zNALC qualifying workloads by IPLing that system (SYSID S16B) with the LICENSE=zNALC parameter, midway through the month. This example also includes a native z/OS system which ran traditional workloads in LPAR3 and includes a z/OS V1.6 guest system (SYSID S16A) which ran a traditional workload in LPAR5.

CPC: 2097-406 (serial = 02-12345)				
LPAR3	LPAR5 (z/VM)	LPAR6 (z/VM)		LPARVM (z/VM)
2 October – 1 November	2 October – 1 November	2 October – 17 October	18 October – 1 November	2 October – 1 November
z/OS V1.7 (SYSID = SYS3)	z/OS V1.6 (SYSID = S16A)	z/OS V1.6 (SYSID = S16B)	z/OS V1.6 (SYSID = S16B) (LICENSE = ZNALC)	CMS
DB2 UDB V8	DB2 UDB V8	DB2 UDB V8	DB2 UDB V8	
CICS TS V3	CICS TS V3	CICS TS V3		
IMS V8				
MQSeries V6				
	Cobol V3	Cobol V3	Cobol V3	
File Manager V6	File Manager V6			
	WebSphere Application Server V6	WebSphere Application Server V6	WebSphere Application Server V6	
	WebSphere MQ Workflow V3	WebSphere MQ Workflow V3	WebSphere MQ Workflow V3	

The following figures show all of the sub-capacity report sections that were generated for the hardware and software configuration shown in Table 29 on page 158. These sections show how SCRT reports on a z/OS V1.6 guest environment where traditional and zNALC workloads were run in the same LPAR. Note that the same z/OS V1.6 guest reporting rules and LPAR capacity rules apply to this example as applied to the previous example.

The main difference in reporting between this example and the previous example is that the (zNALC) suffix does not appear on this report because the LPAR in which the zNALC guest ran also supported a traditional guest system during the same reporting period.

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 03 Nov 2024 - 09:17
 Name of Person Submitting Report: John Customer
 E-mail Address of Report Submitter: customer@abc.com
 Phone Number of Report Submitter: 444-999-9999

Customer Name ABC Corporation
 Customer Number 8971234567
 Machine Serial Number 02-12345
 Machine Type and Model 2097-406
 Machine Rated Capacity (MSUs) 139
 Purchase Order Number (optional)
 Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

==C5=====

TOOL INFORMATION

Tool Release 30.1.0
 Reporting Period 2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

Figure 60. Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in the same LPAR: Customer information and Tool Information sections

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS V1	5694-A01	117		
z/OS V1 (Traditional)	5694-A01	89		
z/OS V1 (zNALC)	5694-A01	28		
DB2 UDB for z/OS V8	5625-DB2	117		
CICS TS for z/OS V3	5655-M15	117		
MQSeries for z/OS V6	5655-L82	34		
IMS V8	5655-C56	34		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	83	(optional)	(conditional)

IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
File Manager for z/OS V6	5655-P17	89	(optional)	(conditional)
IBM WebSphere Application Server for z/OS V6	5655-N01	83		
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	83		

Figure 61. Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in the same LPAR: Product Summary Information section

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
LPAR3	SYS3	01 Oct 2024 - 06:05	02 Nov 2024 - 23:45
LPAR5	-	01 Oct 2024 - 00:00	02 Nov 2024 - 23:45
LPAR6	-	01 Oct 2024 - 00:00	02 Nov 2024 - 23:00
LPARVM	-	-	-
CPC		01 Oct 2024 - 00:00	02 Nov 2024 - 23:45

```

==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR3	34	6	03 Oct 2024 - 14:00	32	11	03 Oct 2024 - 13:00
LPAR5	55	715	04 Oct 2024 - 11:00			
LPAR6	28	725	17 Oct 2024 - 06:30			
CPC	117	5	03 Oct 2024 - 12:00	96	11	03 Oct 2024 - 13:00

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR3	LPAR LPAR5	LPAR LPAR6
z/OS V1	5694-A01	117	03 Oct 2024 - 14:00	34	55	28
z/OS V1 (Traditional)	5694-A01	89	03 Oct 2024 - 14:00	34	55	28
z/OS V1 (zNALC)	5694-A01	28	02 Oct 2024 - 00:00	0	0	28
DB2 UDB for z/OS V8	5625-DB2	117	03 Oct 2024 - 14:00	34	55	28
CICS TS for z/OS V3	5655-M15	117	03 Oct 2024 - 14:00	34	55	28
MQSeries for z/OS V6	5655-L82	34	03 Oct 2024 - 14:00	34	0	0
IMS V8	5655-C56	34	03 Oct 2024 - 14:00	34	0	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	83	02 Oct 2024 - 00:00	0	55	28
File Manager for z/OS V6	5655-P17	89	03 Oct 2024 - 14:00	34	55	0
IBM WebSphere Application Server for z/OS V6	5655-N01	83	02 Oct 2024 - 00:00	0	55	28
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	83	02 Oct 2024 - 00:00	0	55	28

Figure 62. Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in the same LPAR: SMF / SCRT89 Input Data Statistics, Detail LPAR Data Section, and Product Max Contributors sections

```
==Q5=====
PRODUCT GRID SNAPSHOT
```

Product Name	Product ID	LPAR3	LPAR5	LPAR6
z/OS V1	5694-A01	91.3%	92.1%	90.4%
z/OS V1 (Traditional)	5694-A01	91.3%	92.1%	
z/OS V1 (zNALC)	5694-A01			90.4%
DB2 UDB for z/OS V8	5625-DB2	91.3%	92.1%	90.4%
CICS TS for z/OS V3	5655-M15	91.3%	92.1%	90.4%
MQSeries for z/OS V6	5655-L82	91.3%		
IMS V8	5655-C56	91.3%		
IBM WebSphere Application Server for z/OS V6	5655-N01		92.1%	90.4%
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		92.1%	90.4%

```
==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS
```

Product Name	Product ID	LPAR LPAR3	LPAR LPAR5	LPAR LPAR6
z/OS V1	5694-A01	8		
z/OS V1 (Traditional)	5694-A01	8		
z/OS V1 (zNALC)	5694-A01			
DB2 UDB for z/OS V8	5625-DB2	8		
CICS TS for z/OS V3	5655-M15	8		
MQSeries for z/OS V6	5655-L82	8		
IMS V8	5655-C56	8		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53			
File Manager for z/OS V6	5655-P17	8		
IBM WebSphere Application Server for z/OS V6	5655-N01			
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM			

```
==W5=====
```

LPAR5 includes z/OS guest sysids S16A.
LPAR6 includes z/OS guest sysids S16B.

This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such Customer is solely responsible for the completeness and accuracy of information and data used to create this report. Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report. Any questions regarding the contents of this report should be directed to the Customer.

Figure 63. Sample sub-capacity report for z/OS V1.6 guest systems running traditional and zNALC workloads in the same LPAR: Product Grid Snapshot and Sub-Capacity Capping Enforced by SCRT sections

Sample sub-capacity report for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads

Whenever all of the z/OS guest systems running on a CPC are at a z/OS V1.7 level (or later), SCRT reports on each of the guest systems (identified by their SYSID) as if they were running in their own LPAR. SCRT treats these guest systems this way regardless of whether they are traditional systems or zNALC systems. SCRT reports on traditional guest systems using the format *lparname(guest_sysid)*, and reports on zNALC guest systems using the format *lparname(guest_sysid/zNALC)* or, if the LPAR name is not available, **lpar*(guest_sysid/zNALC)*. For details about the formats used for z/OS guest systems, see “Identities reported for z/OS guest systems” on page 138.

Table 30 on page 162 contains the sample hardware and software configuration for a CPC where the customer migrated from a z/OS V1.7 guest system running traditional workloads to a z/OS system running zNALC qualifying workloads in the same LPAR, LPAR7. The customer identified the z/OS system that ran with zNALC qualifying works by IPLing that system (SYSID S17B), using LICENSE=zNALC, midway through the month. This example also includes a native z/OS system which ran traditional workloads in LPAR3 and included a z/OS V1.7 guest system (SYSID S17A) which ran a traditional workload in LPAR5.

Table 30. Sample configuration for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads in the same CPC

CPC: 2097-405 (serial = 02-12345)				
LPAR3	LPAR5 (z/VM)	LPAR7 (z/VM)		LPARVM (z/VM)
2 October – 1 November	2 October – 1 November	2 October – 16 October	16 October – 1 November	2 October – 1 November
z/OS V1.6 (SYSID = SYS3)	z/OS V1.7 (SYSID = S17A)	z/OS V1.7 (SYSID = S17B)	z/OS V1.7 (SYSID = S17B) (LICENSE = ZNALC)	CMS
DB2 UDB V8	DB2 UDB V8	DB2 UDB V8	DB2 UDB V8	
CICS TS V3	CICS TS V3	CICS TS V3		
IMS V8				
MQSeries V6				
File Manager V6	Cobol V3 File Manager V6	Cobol V3	Cobol V3	
	WebSphere Application Server V6	WebSphere Application Server V6	WebSphere Application Server V6	
	WebSphere MQ Workflow V3	WebSphere MQ Workflow V3	WebSphere MQ Workflow V3	

The following figures show all of the sub-capacity report sections that were generated for the hardware and software configuration shown in Table 30 on page 162. These sections show how SCRT reports on a z/OS V1.7 guest environment where traditional and zNALC workloads were run in the same LPAR. Note that the same z/OS V1.6 guest reporting rules that applied to the previous example do not apply to this CPC because all of the z/OS guest systems are at z/OS V1.7 (or higher). The native z/OS V1.6 system running in LPAR3 plays no role in deciding whether or not the z/OS V1.6 guest rules apply.

```
==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====
```

```
Run Date/Time           03 Nov 2024 - 09:17
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter:  444-999-9999
```

```
Customer Name           ABC Corporation
Customer Number          8971234567
Machine Serial Number    02-12345
Machine Type and Model    2097-405
Machine Rated Capacity (MSUs) 118
Purchase Order Number     (optional)
Customer Comments (255 chars max) (optional)
```

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

```
==C5=====
```

TOOL INFORMATION

```
Tool Release           30.1.0
Reporting Period        2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)
```

Figure 64. Sample sub-capacity report for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads: Customer information and Tool Information sections


```

==E5=====
PRODUCT SUMMARY INFORMATION

```

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS V1	5694-A01	103		
z/OS V1 (Traditional)	5694-A01	70		
z/OS V1 (zNALC)	5694-A01	35		
DB2 UDB for z/OS V8	5625-DB2	103		
CICS TS for z/OS V3	5655-M15	103		
MQSeries for z/OS V6	5655-L82	34		
IMS V8	5655-C56	34		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	69	(optional)	(conditional)

IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
File Manager for z/OS V6	5655-P17	70	(optional)	(conditional)
IBM WebSphere Application Server for z/OS V6	5655-N01	69		
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	69		

Figure 65. Sample sub-capacity report for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads: Product Summary Information section

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
LPAR3	SYS3	01 Oct 2024 - 06:05	02 Nov 2024 - 23:45
LPAR5(S17A)	S17A	01 Oct 2024 - 00:00	02 Nov 2024 - 23:45
LPAR7(S17B)	S17B	01 Oct 2024 - 00:00	16 Nov 2024 - 04:00
LPAR7(S17B/zNALC)	S17B	16 Oct 2024 - 04:10	02 Nov 2024 - 23:00
LPARVM	-	-	-
CPC		01 Oct 2024 - 00:00	02 Nov 2024 - 23:45


```

==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR3	34	6	03 Oct 2024 - 14:00	32	11	03 Oct 2024 - 13:00
LPAR5(S17A)	38	8	04 Oct 2024 - 11:00	36	12	04 Oct 2024 - 10:00
LPAR7(S17B)	35	5	02 Oct 2024 - 06:00	34	8	02 Oct 2024 - 05:00
LPAR7(S17B/zNALC)	35	9	17 Oct 2024 - 12:00	33	15	17 Oct 2024 - 15:00
CPC	103	5	03 Oct 2024 - 14:00	101	11	03 Oct 2024 - 13:00


```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR3	LPAR LPAR5(S17A)	LPAR LPAR7(S17B)	LPAR LPAR7(S17B/zNALC)
z/OS V1	5694-A01	103	03 Oct 2024 - 14:00	34	37	32	0
z/OS V1 (Traditional)	5694-A01	70	04 Oct 2024 - 11:00	32	38	0	0
z/OS V1 (zNALC)	5694-A01	35	17 Oct 2024 - 12:00	0	0	0	35
DB2 UDB for z/OS V8	5625-DB2	103	03 Oct 2024 - 14:00	34	37	32	0
CICS TS for z/OS V3	5655-M15	103	03 Oct 2024 - 14:00	34	37	32	0
MQSeries for z/OS V6	5655-L82	34	03 Oct 2024 - 14:00	34	0	0	0
IMS V8	5655-C56	34	03 Oct 2024 - 14:00	34	0	0	0
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	69	17 Oct 2024 - 12:00	0	36	0	33
File Manager for z/OS V6	5655-P17	70	04 Oct 2024 - 11:00	32	38	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	69	17 Oct 2024 - 12:00	0	36	0	33
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	69	17 Oct 2024 - 12:00	0	36	0	33

Figure 66. Sample sub-capacity report for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads: SMF / SCRT89 Input Data Statistics, Detail LPAR Data, and Product Max Contributors sections

==Q5=====

PRODUCT GRID SNAPSHOT

Product Name	Product ID	LPAR3	LPAR5(S17A)	LPAR7(S17B)	LPAR7(S17B/zNALC)
z/OS V1	5694-A01	91.3%	92.1%	45.2%	46.4%
z/OS V1 (Traditional)	5694-A01	91.3%	92.1%		
z/OS V1 (zNALC)	5694-A01			45.2%	46.4%
DB2 UDB for z/OS V8	5625-DB2	91.3%	92.1%	45.2%	46.4%
CICS TS for z/OS V3	5655-M15	91.3%	92.1%	45.2%	
MQSeries for z/OS V6	5655-L82	91.3%			
IMS V8	5655-C56	91.3%			
IBM WebSphere Application Server for z/OS V6	5655-N01		92.1%	45.2%	46.2%
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM		92.1%	45.2%	46.2%

==S5=====

SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS

Product Name	Product ID	LPAR LPAR3	LPAR LPAR5(S17A)	LPAR LPAR7(S17B)	LPAR LPAR7(S17B/zNALC)
z/OS V1	5694-A01				
z/OS V1 (Traditional)	5694-A01				
z/OS V1 (zNALC)	5694-A01				
DB2 UDB for z/OS V8	5625-DB2	8			
CICS TS for z/OS V3	5655-M15	8			
MQSeries for z/OS V6	5655-L82	8			
IMS V8	5655-C56				
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53				
File Manager for z/OS V6	5655-P17				
IBM WebSphere Application Server for z/OS V6	5655-N01	8			
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	8			

==W5=====

This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such Customer is solely responsible for the completeness and accuracy of information and data used to create this report. Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report. Any questions regarding the contents of this report should be directed to the Customer.

Figure 67. Sample sub-capacity report for z/OS V1.7 (or later) guest systems running traditional and zNALC workloads: Product Grid Snapshot and Sub-Capacity Capping Enforced by SCRT sections

Chapter 8. Sub-capacity reporting for Getting Started Sub-Capacity Pricing (GSSP)

IBM Getting Started Sub-Capacity Pricing for z/OS IPLA software (or simply Getting Started Sub-Capacity Pricing) is a sub-capacity pricing option that applies to a specific subset of IPLA products, as described in US Software Announcement letters 208-088 and 209-162 and at [IPLA Sub-Capacity Overview \(www.ibm.com/downloads/cas/AOJR8PEA\)](http://www.ibm.com/downloads/cas/AOJR8PEA).

Note: For the most current list of IPLA products that are eligible for Getting Started Sub-Capacity Pricing, see [IPLA Execution-Based Programs \(www.ibm.com/downloads/documents/us-en/10a99803f52fd884\)](http://www.ibm.com/downloads/documents/us-en/10a99803f52fd884).

Prior to Getting Started Sub-Capacity Pricing, SCRT only implemented one algorithm that applied to all sub-capacity eligible MLC and IPLA software. Under this algorithm, the 4-hour rolling average LPAR MSU utilization of the z/OS system (usually the entire LPAR) is assigned to all execution-based sub-capacity products running on that z/OS system for that LPAR for that hour. The sub-capacity values generated by this algorithm will be referred to as *traditional sub-capacity* to distinguish them from the sub-capacity values generated by the Getting Started Sub-Capacity Pricing algorithm.

Unlike other IBM pricing options where you must specifically order the new pricing option, SCRT automatically applies Getting Started Sub-Capacity Pricing to all eligible product levels that are under the sub-capacity IPLA option. SCRT automatically assigns the traditional sub-capacity value to instances of the Getting Started Sub-Capacity Pricing products whose product levels are not eligible for Getting Started Sub-Capacity Pricing.

Considerations for Getting Started Sub-Capacity Pricing

The following special conditions affect the processing and reporting of Getting Started Sub-Capacity Pricing eligible products:

1. Using Getting Started Sub-Capacity Pricing eligible products on z/OS guest systems
2. Using a Getting Started Sub-Capacity Pricing eligible product whose prior level is not eligible for Getting Started Sub-Capacity Pricing
3. Support for power save mode
4. Support for WebSphere Transformation Extender for z/OS V8 (5655-R95)

Restrictions for z/OS 1.6 guest systems

There are restrictions on all Getting Started Sub-Capacity Pricing eligible products that run on z/OS 1.6 (and earlier) guest systems. Whenever SCRT processes SMF records from a z/OS 1.6 guest system, it sets the sub-capacity value for all execution-based sub-capacity products running on that guest to the maximum size of the LPAR. SCRT applies these same z/OS 1.6 guest rules for Getting Started Sub-Capacity Pricing products. This means that Getting Started Sub-Capacity Pricing is not available for products running on z/OS guest systems running under z/OS 1.6 guest rules. Note that z/OS guest rules also apply to z/OS 1.7 (and later) guests whenever SCRT encounters one or more z/OS 1.6 (or earlier) guest systems running on that CPC at any time during the reporting period. For a more detailed explanation and examples of the z/OS 1.6 guest rules, see [Chapter 6, “Sub-capacity reporting for z/OS systems that are z/VM guests,” on page 133](#).

Restrictions for partially supported products

In order to be eligible for Getting Started Sub-Capacity Pricing, a product must be instrumented to appear on an SMF type 89 subtype 1 record so that SCRT can obtain the required product utilization data to use in the Getting Started Sub-Capacity Pricing algorithm. In order for IBM to make Getting Started Sub-Capacity Pricing available as soon as possible, later levels of existing versions of some products were instrumented to support Getting Started Sub-Capacity Pricing. This means that the

earlier levels of those products do not support Getting Started Sub-Capacity Pricing, while the newest levels of those same products do. You must migrate to the product levels listed at [IPLA Execution-Based Programs \(www.ibm.com/downloads/documents/us-en/10a99803f52fd884\)](http://www.ibm.com/downloads/documents/us-en/10a99803f52fd884) in order to be eligible to receive Getting Started Sub-Capacity Pricing for those products. SCRT automatically assigns the traditional sub-capacity value to instances of the Getting Started Sub-Capacity Pricing products whose product levels are not eligible for Getting Started Sub-Capacity Pricing.

Restrictions for power save mode

For CPCs that support power save mode (where the CPC can reduce its power requirements but at a decrease in CPC performance), SCRT can adjust for the decrease in CPC performance when reporting the sub-capacity values for Getting Started Sub-Capacity Pricing products but only when the required APARs have been applied to the z/OS system and only when the new SMF **MAXEVENTINTRECS** parameter has been set to a value high enough to prevent an SMF type 89 record from containing busy time for both power save mode and for high performance mode (non-power save mode). See [“Static power save mode” on page 506](#) for additional details on power save mode and for the actions required to enable SCRT to perform the appropriate adjustments for Getting Started Sub-Capacity Pricing products.

Considerations for WebSphere Transformation Extender V8 (5655-R95)

WebSphere Transformation Extender V8 (5655-R95) differs from all other Getting Started Sub-Capacity Pricing products because you are always required to provide SCRT with a NO89 control statement for this product and because only parts of this product (only two of the four priced features) are eligible for Getting Started Sub-Capacity Pricing products. Because neither of these conditions applies to any other Getting Started Sub-Capacity Pricing products, we specifically describe these restrictions here.

Requirements for using the NO89 control statement for 5655-R95

All four of the WebSphere Transformation Extender V8 priced features were not instrumented to be tracked by SMF type 89 records until Service Level 2 of the product; therefore, SCRT treats all levels of the version 8 product as a NO89 product in order to support the function delivered by both the initial release and the Service Level 1 release. This means that SCRT requires you to supply a NO89 control statement for all levels of the WebSphere Transformation Extender V8 product.

Unlike other IBM products which typically consist of a base product and possibly one or more optional priced features, where the base product provides a specified set of functions and the optional features provide extensions to the base product function or provide additional functions, WebSphere Transformation Extender V8 only consists of four optional, priced features (Command Server, Launcher, Application Programming and Integration Server) with no base functions available outside of these features, though the features do share some common functions. This means that the required NO89 control statement, which is at the product level, must be constructed based on the use of those features under the various conditions described later. For a detailed description of the NO89 control statement see [“NO89 DD statement” on page 61](#).

For *any* of the following conditions, set the value of the NO89 control statement to *NONE, as shown in the following example. The NO89 control statement in the example is preceded by the corresponding comment statement containing the product’s name. All statements in the example start in column one.

- None of the WebSphere Transformation Extender V8 features have been licensed or used during any part of the reporting period.
- None of the licensed WebSphere Transformation Extender V8 features have been deployed or used during any part of the reporting period.
- All of the licensed WebSphere Transformation Extender V8 features that were used at any time during the reporting period were at Service Level 2 or higher.

```
* WEBSHERE TRANSFORMATION EXTENDER FOR Z/OS V8
5655-R95=*NONE
```

For the case where one or more of the priced features were run at a product level earlier than Service Level 2 at any time during the reporting period, customize the NO89 control statement to list the image ID (*image_id*) for each z/OS system in which any of those earlier levels of the feature ran, regardless of whether other instances of that feature or instances of other features ran at Service Level 2 or higher in the same z/OS image ID, either at the same time or at a different time during the same reporting period. When the earlier levels of two or more of the features ran in the same z/OS image ID, that image ID should only be listed once on the NO89 control statement for the same CPC. For native z/OS systems, set each *image_id* value equal to the corresponding LPAR name. For z/OS guest systems, set each *image_id* value equal to the SID (sysid) value of the guest system. For the case where all of the z/OS systems were running an early level of one or more of the features at any time during the reporting period, a value of *ALL can be used, instead of listing all of the z/OS systems.

The following set of complex conditions would require the NO89 control statement shown in the following example. All statements shown in the example start in column one.

- **CPC: 2097-12345**

- **LPAR: PROD1A**

- Command Server feature at Service Level 1 running during the first 10 days of the reporting period
 - Command Server feature at Service Level 2 running during the remainder of the reporting period
 - Launcher feature at Service Level 2 running during the reporting period

- **LPAR: PROD1B**

- Application Programming feature at Service Level 2 running during the reporting period

- **LPAR: PROD1C**

- Launcher feature at Service Level 1 running during the first 10 days of the reporting period

- **CPC: 2097-23456**

- **LPAR: PROD2A**

- Integration Server feature at Service Level 1 running during the first 10 days of the reporting period
 - Integration Server feature at Service Level 2 running during the remainder of the reporting period

- **LPAR: PROD2B**

- Application Programming feature Service Level 2 running during the reporting period

- **LPAR: PROD2C**

- Launcher feature at Service Level 2 running during the reporting period

```
* WEBSPPHERE TRANSFORMATION EXTENDER FOR Z/OS V8
5655-R95=2097-12345:PROD1A,PROD1C,2097-23456:PROD2A
```

5655-R95 feature eligibility for Getting Started Sub-Capacity Pricing

When running at Service Level 2 or higher, the Command Server feature and the Launcher feature of this product are both eligible for Getting Started Sub-Capacity Pricing. Earlier levels of these features are not eligible for Getting Started Sub-Capacity Pricing because they have not been instrumented to be tracked by any SMF89 record. SCRT automatically assigns either Getting Started Sub-Capacity Pricing or traditional sub-capacity pricing to these features based on the contents of the NO89 control statement and the contents of the SMF type 89 records being processed.

The Application Programming feature and the Integration Server feature of this product are not eligible for Getting Started Sub-Capacity Pricing, regardless of the service level at which they are running. SCRT automatically assigns traditional sub-capacity pricing to all uses of these features. For feature instances that are running at a level earlier than Service Level 2, the traditional sub-capacity value will be based on the contents of the NO89 control statement. For feature instances that are running at Service Level 2 or higher, the traditional sub-capacity value will be based on the contents of SMF type 89 subtype 2 records.

Getting Started Sub-Capacity Pricing algorithm

SCRT applies the following Getting Started Sub-Capacity Pricing algorithm to product levels that are tracked on SMF type 89-1 records:

1. SCRT uses the product-specific utilization data collected for each Getting Started Sub-Capacity Pricing eligible product to calculate the 4-hour rolling average MSU utilization for those eligible products for every hour for every LPAR.
2. SCRT subtracts the 4-hour rolling average utilization of the Db2 time, which is used by the Getting Started Sub-Capacity Pricing eligible product via the Db2 JDBC type 2 connector, from the 4-hour rolling average utilization of that product before applying the remainder of the Getting Started Sub-Capacity Pricing rules, provided that the following conditions are met:
 - The SMF records must be collected from a z/OS 1.10 system or higher with SMF APAR OA31615 applied. This APAR collects the Db2 JDBC type 2 connector time whenever the remainder of the following required conditions are met. Collecting SMF records from z/OS 1.9 or earlier systems, or collecting SMF records from z/OS 1.10 or later systems without SMF APAR OA31615 applied, means that the Db2 data required for the subtraction has not been collected.
 - DB2 UDB for OS/390 V6 or higher or DB2 V8 for z/OS Value Unit Edition or higher must be used by the Getting Started Sub-Capacity Pricing product via a Db2 JDBC type 2 connector.
 - Db2 must be configured to collect the required Db2 data that will be used for the subtraction. That Db2 configuration must be activated prior to Db2 being used by the Getting Started Sub-Capacity Pricing eligible product. The default Db2 configuration setting is not to collect the required data. You must change the default Db2 setting by configuring Db2 to use the **SMF89 YES** parameter of the DSN6SYSP macro.
 - The active SMFPRMxx member of PARMLIB must either be set to use MULCFUNC or must use the default MULCFUNC setting so that Db2 will collect the Db2 JDBC type 2 connector busy time.
 - DB2 V8 (UDB or Value Unit Edition) and higher should have APARs PM13466 and PM13467 applied. These APARs correct a problem where Db2 does not collect all of the Db2 busy time to be subtracted. Failure to apply these APARs could mean that less of the Db2 JDBC type 2 connector's busy time will be subtracted from the Getting Started Sub-Capacity Pricing eligible product.
3. SCRT multiplies the value calculated in step “2” on [page 168](#) by the Getting Started Sub-Capacity Pricing scaling factor in order to determine each product's Getting Started Sub-Capacity Pricing value for every hour for every LPAR.

SCRT uses a default scaling factor of 2, which means that SCRT reports reduced MSUs for eligible products whose 4-hour rolling average utilization is up to 50 percent of the 4-hour rolling average utilization for the corresponding hour of the z/OS system. Future Getting Started Sub-Capacity Pricing eligible products could choose to use a different scaling factor to more accurately reflect how those products are used relative to the utilization of the LPAR in which they run.

4. SCRT assigns the lesser of the product's Getting Started Sub-Capacity Pricing value and the product's traditional sub-capacity value to the product for that LPAR for that hour.
5. SCRT combines all of the LPAR values across the CPC to determine the product's value for that hour on that CPC.
6. SCRT assigns the product's highest CPC hour value found in the reporting period as the product's sub-capacity value.

Sub-capacity report externals for Getting Started Sub-Capacity Pricing

SCRT automatically applies the Getting Started Sub-Capacity Pricing rules to all of the eligible products that SCRT encounters when processing SMF records. In addition to reporting the appropriate MSU value for each eligible product (either the Getting Started Sub-Capacity Pricing based MSU value or the traditional sub-capacity MSU value or an MSU value that is a combination of one or more Getting Started

Sub-Capacity Pricing values for some LPARs and traditional MSU values for other LPARs), SCRT also does the following:

1. Uses footnotes to indicate when the MSU value shown for a Getting Started Sub-Capacity Pricing product is based totally or partially on the Getting Started Sub-Capacity Pricing MSU value
2. Reports the number of hours when the Getting Started Sub-Capacity Pricing MSU value is less than, equal to, or greater than the product's traditional sub-capacity MSU value
3. Creates trace records that indicate:
 - Why SCRT did not assign the Getting Started Sub-Capacity Pricing MSU value to the product
 - When SCRT did assign the Getting Started Sub-Capacity Pricing MSU value but that value was higher than it would have been if the appropriate actions had been taken

These externals are also discussed in [“Sample sub-capacity report for Getting Started Sub-Capacity Pricing” on page 174.](#)

Getting Started Sub-Capacity Pricing footnote

For the case where the Too1 MSUs value for one or more products that are already listed in the Product Summary Information section is based totally or partially on the product's Getting Started Sub-Capacity Pricing value, SCRT does the following on the sub-capacity report:

1. Generates an optional Footnotes column header for the IPLA section
2. Posts footnote (iv) in the footnote column for each product whose Too1 MSUs value is either:
 - Totally based on the product's Getting Started Sub-Capacity Pricing value, or
 - Partially based on the product's Getting Started Sub-Capacity Pricing value, as when that product's MSU values are based on its Getting Started Sub-Capacity Pricing value from one or more LPARs and on its traditional sub-capacity value from other LPARs.
3. Generates a footnote subsection immediately after the last IPLA product is listed. The footnote subsection contains a footnote header row followed by a description of footnote (iv).

Footnote (iv) is optional. All three footnote-related items will only be generated if one or more of the Getting Started Sub-Capacity Pricing eligible products requires such a footnote.

SCRT posts this footnote whenever any of the hours during which the peak value occurred received a benefit from the Getting Started Sub-Capacity Pricing algorithm. Prior to SCRT V20.1.0, the footnote was posted only when the earliest peak hour received a benefit from the Getting Started Sub-Capacity Pricing algorithm.

All numeric footnotes use lower case Roman numerals (i, ii, iii, ...) instead of Arabic numerals (1, 2, 3, ...) to prevent spreadsheet applications from automatically translating the footnotes into some other format.

Getting Started Sub-Capacity Pricing report sections

Whenever SCRT processes data for Getting Started Sub-Capacity Pricing eligible products, it automatically generates up to three additional report sections in the Customer verification section of the sub-capacity report, provided that the SPECIAL **Suppress_Report_Section** control statement was not used to suppress them. These report sections can be used along with the GSTRACE records to determine the effectiveness of Getting Started Sub-Capacity Pricing for the products reported in the Product Summary Information section of the same sub-capacity report.

Each of the three optional report sections only appear when there is data to report for that specific section. That means that one, two, or all three report sections can appear on a sub-capacity report whenever SCRT processes data for any Getting Started Sub-Capacity Pricing eligible products. None of these sections appear when there is no data from any Getting Started Sub-Capacity Pricing eligible product.

The Getting Started Sub-Capacity Pricing report sections are:

Getting Started Sub-Capacity Pricing Assigned

This optional sub-capacity report section lists the number of hours, by LPAR, when SCRT assigned the Getting Started Sub-Capacity Pricing value to an eligible product because that value was less than the product's traditional sub-capacity value for that same hour. SCRT automatically generates this section whenever it assigns the Getting Started Sub-Capacity Pricing value for one or more hours to any product appearing in this sub-capacity report section; otherwise, this section is not generated. Only the Getting Started Sub-Capacity Pricing eligible products that have hours to report in this section will appear here.

Getting Started Sub-Capacity Pricing Equal to Traditional Sub-capacity

This optional sub-capacity report section lists the number of hours, by LPAR, when an eligible product's Getting Started Sub-Capacity Pricing value was equal to that product's traditional sub-capacity value for that same hour. Whenever an eligible product runs in multiple LPARs that are members of the same LPAR group and the sum of that product's Getting Started Sub-Capacity Pricing capacity in the individual LPARs is equal to the LPAR group cap, SCRT reports that hour by LPAR group.

Note that the rules for posting data under an LPAR group are different for this section than for the Product Max Contributors section, Sub-Capacity Capping Enforced by SCRT section, and Getting Started Sub-capacity Pricing Not Assigned section. In this section, product data (hours) is posted under the LPAR **Group** column whenever the sum of the product's MSUs across the LPARs in a group equals the value of the LPAR group cap because, for that case, the product did not receive a reduction in MSUs solely due to Getting Started Sub-Capacity Pricing. The product would have reported the same MSU value, even without the Getting Started Sub-Capacity Pricing algorithm, due to group capping. However, the rule for reporting product data (MSUs or hours) under the LPAR **Group** column for the other sections of the sub-capacity report are different. For all of those other sections, data is reported under the LPAR **Group** column only when group capping is actually used to cap the MSU values. Data is not reported under the LPAR **Group** column for the case where the LPAR group cap value equals the sum of the product's MSUs across the LPARs in the LPAR group.

SCRT automatically generates this section whenever the Getting Started Sub-Capacity Pricing value for one or more eligible products is equal to the traditional sub-capacity value for those products for one or more hours; otherwise, this section is not generated. Only the Getting Started Sub-Capacity Pricing eligible products that have hours to report in this section will appear here.

Getting Started Sub-Capacity Pricing Not Assigned

This optional sub-capacity report section lists the number of hours, by LPAR, when an eligible product's Getting Started Sub-Capacity Pricing value was higher than that product's traditional sub-capacity value for that same hour. Whenever an eligible product runs in multiple LPARs that are members of the same LPAR group and the sum of that product's Getting Started Sub-Capacity Pricing capacity in the individual LPARs is greater than the LPAR group cap, SCRT reports that hour by LPAR group. SCRT always assigns the lesser traditional sub-capacity value, which might be the capped value for the LPAR group, to the Getting Started Sub-Capacity Pricing eligible products for the hours listed in this section. SCRT automatically generates this section whenever the Getting Started Sub-Capacity Pricing value for one or more eligible products is higher than the traditional sub-capacity value for those products for one or more hours; otherwise, this section is not generated. Only the Getting Started Sub-Capacity Pricing eligible products that have hours to report in this section will appear here.

Suppress_Report_Section control statement

SCRT supports an optional SPECIAL DD control statement:

```
Suppress_Report_Section=GSSP
```

This control statement directs SCRT not to generate all three of the Getting Started Sub-Capacity Pricing report sections described in [“Getting Started Sub-Capacity Pricing report sections” on page 169](#). The default is for SCRT to generate these report sections unless directed not to do so by this control statement (or unless there are no Getting Started Sub-Capacity Pricing eligible products on which to report).

Customers who do not benefit from Getting Started Sub-Capacity Pricing because the traditional sub-capacity values for all of their Getting Started Sub-Capacity Pricing eligible products are always equal to or less than the Getting Started Sub-Capacity Pricing values for those products can use this control statement to prevent the new report sections from appearing on the sub-capacity report.

Getting Started Sub-Capacity Pricing trace records

SCRT supports an optional GSTRACE DD JCL statement, which is used to track the reasons why SCRT did not assign the Getting Started Sub-Capacity Pricing value by CPC by LPAR or LPAR group, by product, by hour. The fields of the trace records that are written to this file are all separated by commas (that is, the file is in comma separated value format) so that the file can be downloaded in text format onto a workstation where it can be processed by a spreadsheet application, such as Lotus 1-2-3 or Microsoft Excel.

The very first record in this file is a header record which contains the column titles for the fields of all the remaining rows in the spreadsheet. The titles contained in the header record are:

CPC (tttt-sssss)

The machine type (*tttt*) and serial number (*sssss*) of the CPC where the reported product ran. Example: 2097-12345

Image Id

The Image ID of the z/OS system on which the reported product ran.

- For native z/OS systems that are not under group capping or whose group cap is higher than the product's Getting Started Sub-Capacity Pricing value, this value is the LPAR name.

Example: LPAR1

- For native z/OS systems that are being capped under group capping and the capped value is less than or equal to the sum of the values of the Getting Started Sub-Capacity Pricing products running in the capped LPARs, this value is the LPAR group name followed by the (group) suffix.

Example: LGRPA (group)

- For z/OS 1.6 and earlier guest systems, this value is the LPAR name.

Example: LPAR5

- For z/OS 1.7 and later guest systems, this value is the SID (SYSID) of the guest system.

Example: SYS1

Product

The PID number of the Getting Started Sub-Capacity Pricing eligible product that was not assigned a Getting Started Sub-Capacity Pricing value.

If the report line is for a product family root, this field contains an abbreviation that represents the family.

Example: 5655-N01

Date:Hour(YyyyMmmDdd:Hhh)

The year (*Yyyy*), month (*Mmm*), day (*Ddd*), and hour (*Hhh*) for which the product specified in the **Product #** field was not assigned a Getting Started Sub-Capacity Pricing value.

Example: Y2024M10D02:H00 represents 00:00 on 2 October 2024.

Traditional Subcap Value

The traditional sub-capacity value that was assigned to the product in the **Product #** field, instead of the Getting Started Sub-Capacity Pricing value, because the traditional sub-capacity value was equal or less.

Whenever group capping is responsible for creating the trace record, the Traditional Subcap Value is set to the group cap value. That is, when the eligible product is running in multiple LPARs that are members of the same LPAR group, the Traditional Subcap Value will be the group cap

value whenever the group cap value is less than or equal to the sum of the product's Getting Started Sub-Capacity Pricing values for the LPARs in the LPAR group.

GSSP Value

The Getting Started Sub-Capacity Pricing value that was calculated for the product in the **Product #** field for the cases where eligibility was based on the product's 4-hour rolling average utilization relative to the 4-hour rolling average utilization of the LPAR or z/OS guest system in which the product was running.

Whenever group capping is responsible for creating the trace record, the GSSP Value is set to the sum of the product's Getting Started Sub-Capacity Pricing values for the LPARs in the LPAR group. That is, when the eligible product is running in multiple LPARs that are members of the same LPAR group, the GSSP Value will be the sum of the product's Getting Started Sub-Capacity Pricing values for the LPARs in the LPAR group whenever the group cap value is less than or equal to the sum of the product's Getting Started Sub-Capacity Pricing values for the LPARs in the LPAR group.

Note that SCRT treats cases where the Getting Started Sub-Capacity Pricing value is equal to the traditional sub-capacity value as not being eligible for Getting Started Sub-Capacity Pricing. This is because, by definition, Getting Started Sub-Capacity Pricing is only eligible when the Getting Started Sub-Capacity Pricing value is less than the traditional sub-capacity pricing value.

Reason for not getting GSSP

The reason or reasons why SCRT assigned the traditional sub-capacity value to the product in the **Product #** field, instead of assigning the Getting Started Sub-Capacity Pricing value, or the reason for possibly assigning a Getting Started Sub-Capacity Pricing value that is higher than it would have been had the appropriate actions been performed. These reasons include:

NO89

The pre-Getting Started Sub-Capacity Pricing level of this product was running in the LPAR or z/OS guest system, as indicated by that product's SCRT NO89 control statement. Instances of the NO89 level of the product without instances of the SMF type 89-1 level of the product running during the same hour in the same LPAR make that product ineligible for Getting Started Sub-Capacity Pricing in those LPARs and z/OS guest systems. A value of zero MSUs is posted in the GSSP Value field for the product.

z/OS 1.6 Guest Rules

The product is ineligible for Getting Started Sub-Capacity Pricing because it is running in a z/OS guest system that is running under z/OS 1.6 guest rules. See the explanation in [“Restrictions for z/OS 1.6 guest systems” on page 165](#). A value of zero MSUs is posted in the GSSP Value field for the product.

Product level not eligible

Only the pre-Getting Started Sub-Capacity Pricing level of this product was running in the LPAR or z/OS guest system, making the product ineligible for Getting Started Sub-Capacity Pricing in this LPAR or z/OS guest system. A value of zero MSUs is posted in the GSSP Value field for the product.

Getting Started Sub-Capacity Pricing equaled subcap MSUs

The product's Getting Started Sub-Capacity Pricing value is equal to its traditional sub-capacity value. SCRT considers a product to be eligible for Getting Started Sub-Capacity Pricing only when its Getting Started Sub-Capacity Pricing value is less than its traditional sub-capacity value.

Getting Started Sub-Capacity Pricing exceeded subcap value by *nnn* MSUs

The product's Getting Started Sub-Capacity Pricing value is higher than its traditional sub-capacity value by the indicated amount (*nnn*). According to Getting Started Sub-Capacity Pricing rules, SCRT assigns the lesser traditional sub-capacity value to the product, instead of the higher Getting Started Sub-Capacity Pricing value.

The following reasons only apply to the case where the Getting Started Sub-Capacity Pricing eligible product is using Db2 via a Db2 JDBC type 2 connector. Only if that condition applies, then the following reasons could have resulted in SCRT assigning the specified GSSP product a traditional sub-capacity value instead of the GSSP value or in SCRT assigning a GSSP value that was higher than it would have been had all of the requirements listed in step 2 of [“Getting Started Sub-Capacity](#)

Pricing algorithm” on page 168 been satisfied prior to the current SCRT reporting period. However, there is no guarantee that satisfying these requirements would have resulted in a lower sub-capacity value for any GSSP product because all of the required conditions must apply for that GSSP product's peak hour.

IBM recommends that you satisfy all of the requirements listed in step 2 of “Getting Started Sub-Capacity Pricing algorithm” on page 168 in order for SCRT to report the lowest sub-capacity value for GSSP products.

Missing APAR OA31615

The required z/OS SMF APAR has not been applied. Any busy time used by Db2 via the Db2 JDBC type 2 connector, which is being counted as part of the GSSP product's busy time, was not reported on an SMF record and so could not be subtracted by SCRT as indicated in step 2 of “Getting Started Sub-Capacity Pricing algorithm” on page 168. The possible results are that SCRT assigned the specified GSSP product a traditional sub-capacity value instead of the GSSP value or that SCRT assigned a GSSP value that was higher than it would have been had all of the requirements listed in step 2 been satisfied prior to the current SCRT reporting period.

This reason is posted once per hour and uses a Product # of : : ALL : : to apply to all of the GSSP products running at the specified hour, CPC, and LPAR. No values are provided for the Traditional Subcap Value nor the GSSP Value because none apply.

DB2 data collection

Some or all of the recommended Db2 related configuration or service related actions listed in “Getting Started Sub-Capacity Pricing algorithm” on page 168 have not been implemented. Either all or part of the busy time used by Db2 via the Db2 JDBC type 2 connector, which is being counted as part of the GSSP product's busy time, was not reported on an SMF record and so could not be subtracted by SCRT as indicated in step 2. The possible results are that SCRT assigned the specified GSSP product a traditional sub-capacity value instead of the GSSP value or that SCRT assigned a GSSP value that was higher than it would have been had all of the requirements listed in step 2 been satisfied prior to the current SCRT reporting period.

No values are provided for the Traditional Subcap Value nor the GSSP Value because none apply.

You can use the GSTRACE file together with the Getting Started Sub-Capacity Pricing Assigned section, the Getting Started Sub-Capacity Pricing Equal to Traditional Sub-Capacity section, and the Getting Started Sub-Capacity Pricing Not Assigned section of the sub-capacity report to determine the effectiveness of Getting Started Sub-Capacity Pricing for the products listed in the Product Summary Information section of that report.

Sample configuration for Getting Started Sub-Capacity Pricing

Table 31 on page 174 illustrates the sample hardware and software configuration for a CPC where three Getting Started Sub-Capacity Pricing products were running during the reporting period. This is the same hardware configuration as was used for the base z/OS sub-capacity report described in Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103. The software configuration, however, differs slightly to demonstrate how SCRT reports on Getting Started Sub-Capacity Pricing products.

The sample configuration consists of a 2097 Model 405 CPC with a serial number of 02-12345. The CPC is configured to support four LPARs: three are running z/OS natively in the LPAR and one is running Linux systems under z/VM. The z/OS systems in LPAR1, LPAR2, and LPAR3 have z/OS SYSIDs of SYS1, SYS2, and SYS3. LPAR2 and LPAR3 are configured to be members of LPAR group LGRPA. The z/OS systems are running a combination of the following z/OS products:

- DB2 UDB V8
- CICS TS V3
- MQSeries V6
- IMS V8
- Lotus Domino V7

- File Manager V6
- IBM WebSphere Application Server V6
- WebSphere Message Broker for z/OS V6
- WebSphere Transformation Extender for z/OS V8

Table 31. Sample configuration with Getting Started Sub-Capacity Pricing eligible products

CPC: 2097-405 (serial = 02-12345)			
LPAR1	LPAR2	LPAR3	LPARVM
z/OS V1 (SYSID = SYS1)	z/OS V1 (SYSID = SYS2)	z/OS V1 (SYSID = SYS3)	z/VM
CICS TS V3	DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8	DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8 Cobol V3	(Linux guests)
IMS V8	File Manager V6	WebSphere Application Server V6 WebSphere Message Broker V6 WebSphere Transformation Ext. V8	
File Manager V6	WebSphere Application Server V6 WebSphere Message Broker V6 WebSphere Transformation Ext. V8		

Sample control statement input to SCRT for Getting Started Sub-Capacity Pricing

SCRT supports additional externals for reporting on Getting Started Sub-Capacity Pricing products. The sample sub-capacity report discussed in [“Sample sub-capacity report for Getting Started Sub-Capacity Pricing”](#) on page 174 is based on the following inputs to SCRT:

NO89 control statement

The following NO89 control statement was used for WebSphere Transformation Extender for z/OS V8:

```
5655-S95=★NONE
```

The value ★NONE is used because all instances of this product that were running in both LPARs were enabled for Getting Started Sub-Capacity Pricing.

SPECIAL DD control statements

None. That is, the statement `Suppress_Report_Section=GSSP` was not used so that SCRT will report on the hours where eligible products were assigned Getting Started Sub-Capacity Pricing values versus being assigned traditional sub-capacity values.

GSTRACE DD control statement

The GSTRACE DD statement was uncommented and used to generate a .csv formatted trace file.

Sample sub-capacity report for Getting Started Sub-Capacity Pricing

The sub-capacity report for Getting Started Sub-Capacity Pricing has the same basic format and report sections as the sub-capacity report for a base z/OS system. However, the sub-capacity report for Getting Started Sub-Capacity Pricing can also contain a new footnote and one to three new report sections, depending on the MSUs used by each Getting Started Sub-Capacity Pricing eligible product.

This sample report describes the differences between a sub-capacity report for a base z/OS system and a sub-capacity report that also reports on Getting Started Sub-Capacity Pricing eligible products. For a complete description of all of the fields appearing in a base sub-capacity report see [“Sample sub-capacity report for single-version z/OS native systems”](#) on page 105.

Billing-related section

The billing-related section of the sub-capacity report displays the customer information, tool information, special conditions, product summary information, and detail data collection information, which are all described in [“Billing-related section” on page 105](#).

Customer Information section

There are no differences between the Customer Information section for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report. For details about this section, see [“Customer information” on page 105](#). The sample Customer Information section appears in [Figure 68 on page 175](#).

Tool Information section

There are no differences between the Tool Information section for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report. For details about this section, see [“Tool Information section” on page 107](#). The sample Tool Information section appears in [Figure 68 on page 175](#).

Special Conditions section

There are no differences between the Special Conditions section for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report. For details about this section, see [“Special Conditions section” on page 107](#).

Based on the sample hardware and software configurations used for this sample sub-capacity report, the optional Special Conditions section was not generated and, thus, does not appear in [Figure 68 on page 175](#).

```

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time                03 Nov 2024 - 09:17
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter:  444-999-9999

Customer Name                ABC Corporation
Customer Number              8971234567
Machine Serial Number        02-12345
Machine Type and Model       2097-405
Machine Rated Capacity (MSUs) 118
Purchase Order Number        (optional)
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust
the billable MSUs in inventory for all MLC Products listed under the MLC Product Name
column on this report. In accordance with our agreement, IBM will treat a change in product
licensed capacity as an order. If the MSUs have changed since the last report, software billing
based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to
bill those IPLA products listed under the IPLA Product Name column in this report which
exceed your entitled capacity. In accordance with our agreement, IBM will treat the use
of a product in excess of its entitled capacity as an order and you will be billed for
the amount in excess of your entitlement.

==C5=====
TOOL INFORMATION

Tool Release                30.1.0
Reporting Period            2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

```

Figure 68. Sample sub-capacity report for Getting Started Sub-Capacity Pricing: Customer information and Tool Information sections

Product Summary Information section

An optional **Footnotes** column header will appear in the IPLA section whenever SCRT needs to post a footnote for one of the Getting Started Sub-Capacity Pricing eligible products that are listed in this section of the report. The **Footnotes** heading does not appear whenever there are no footnotes for this section.

The sample Product Summary Information section appears in [Figure 69 on page 176](#).

```

==E5=====
PRODUCT SUMMARY INFORMATION

```

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)	
z/OS V1	5694-A01	98			
DB2 UDB for z/OS V8	5625-DB2	71			
CICS TS for z/OS V3	5655-M15	98			
MQSeries for z/OS V6	5655-L82	71			
IMS V8	5655-C56	98			
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	39	(optional)	(conditional)	

IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)	Footnotes
File Manager for z/OS V6	5655-P17	60	(optional)	(conditional)	
IBM WebSphere Application Server for z/OS V6	5655-N01	64			(iv)
WebSphere Message Broker for z/OS V6	5655-M74	64			(iv)
WebSphere Transformation Extender for z/OS V8	5655-R95	47			(iv)

Footnotes:
(iv) Product's Tool MSUs are based on Getting Started Sub-capacity Pricing from one or more LPARs

Figure 69. Sample sub-capacity report for Getting Started Sub-Capacity Pricing: Product Summary Information section

There are no differences between the **MLC Product Name**, **MLC Product ID**, **Customer Comment**, **IPLA Product Name**, and **IPLA Product ID** fields on a basic z/OS sub-capacity report and on a Getting Started Sub-Capacity Pricing sub-capacity report. For a description of these fields see [“Product Summary Information section” on page 110](#).

Tool MSUs

SCRT applies the rules for traditional sub-capacity pricing to all of the MLC products and IPLA products that are not eligible for Getting Started Sub-Capacity Pricing. SCRT posts the traditional sub-capacity values for these products under the **Tool MSUs** column.

SCRT then automatically applies the Getting Started Sub-Capacity Pricing rules to all products that are eligible for Getting Started Sub-Capacity Pricing, as indicated at [IPLA Execution-Based Programs \(www.ibm.com/downloads/documents/us-en/10a99803f52fd884\)](http://www.ibm.com/downloads/documents/us-en/10a99803f52fd884). Then, SCRT either posts the Getting Started Sub-Capacity Pricing value for that product, posts the traditional sub-capacity value for that product, or posts some combination of both types of values for the case where the maximum product value is comprised of the Getting Started Sub-Capacity Pricing value from some LPARs and the traditional sub-capacity value from other LPARs.

The case for combining the different types of values can be more easily understood by looking at the LPAR values for WebSphere Application Server in the Product Max Contributors section of the sample sub-capacity report. For that product, the MSUs from LPAR3 are based on the lesser values due to the Getting Started Sub-Capacity Pricing algorithm, while the MSUs from LPAR2 are based on the traditional sub-capacity value because the Getting Started Sub-Capacity Pricing value for WebSphere Application Server, running in LPAR2, was higher than the traditional sub-capacity value. SCRT always posts the lesser of the two MSU values in the **Tool MSUs** column.

For additional information about Tool MSUs, see [“Product Summary Information section” on page 110](#).

Customer MSUs

You can use the **Customer MSUs** field to enter whole-number MSU override values (maximum of four digits) for the MLC products and IPLA products only when the product's **Customer MSUs** field contains the (optional) placeholder. If no placeholder appears in this field, you must use the

Exclude control statement, as described in “[EXCLUDE control statement](#)” on page 44, to adjust the **Tool MSUs** value. You must also provide an acceptable explanation in the **Customer Comments** column.

For more information about the **Customer MSUs** field, see “[Product Summary Information section](#)” on page 110.

Footnotes

SCRT posts footnote (iv) whenever the **Tool MSUs** value posted for a product is based totally or partially on the Getting Started Sub-Capacity Pricing algorithm. The **Tool MSUs** are partially based on the Getting Started Sub-Capacity Pricing algorithm whenever the posted value for the product is combined from multiple LPARs where the product’s value is based on the traditional sub-capacity value from some LPARs and on the Getting Started Sub-Capacity Pricing values from other LPARs. Whenever SCRT posts footnote (iv) in the **Footnotes** column, it also generates a row with a footnote header followed by a description for the footnote.

Note: The footnote is posted whenever any of the hours during which the peak value occurred received a benefit from the Getting Started Sub-Capacity Pricing algorithm.

For an explanation of footnote (viii), see “[Product Summary Information section](#)” on page 191.

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the Product Summary section.

SMF / SCRT89 Input Data Statistics section

There are no differences between the SMF / SCRT89 Input Data Statistics section for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report. For details about this section, see “[SMF / SCRT89 Input Data Statistics section](#)” on page 118. The sample SMF / SCRT89 Input Data Statistics section appears in [Figure 70 on page 178](#).

Detail LPAR Data Section

There are no differences between the Detail LPAR Data Section for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report. For details about this section, see “[Detail LPAR Data section](#)” on page 119. The sample Detail LPAR Data Section appears in [Figure 70 on page 178](#).

Product Max Contributors section

The MSU values that are listed in the Product Max Contributors section are the same for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report, with one subtle change. While all of the MSU values listed in this section continue to reflect the product MSUs by LPAR or LPAR Group for the earliest hour for the product’s maximum MSU value, the MSU values posted for Getting Started Sub-Capacity Pricing eligible products are based on the Getting Started Sub-Capacity Pricing algorithm whenever that value is less than the traditional sub-capacity value for that product for that hour. The result is that some of the LPAR values posted for Getting Started Sub-Capacity Pricing eligible products are less than the MSUs posted for z/OS, even when the Getting Started Sub-Capacity Pricing eligible product runs in all of the same LPARs and at the same times as z/OS.

Note: The MSUs for Getting Started Sub-Capacity Pricing eligible products are posted under an LPAR group whenever the LPAR group cap is less than the Getting Started Sub-Capacity Pricing value for the earliest hour with the maximum product value. That is not the case for this sample configuration.

If the WebSphere Transformation Extender product had been running NO89 instances of the product in either LPAR2 or LPAR3 without also running instances of the product level that are eligible for Getting Started Sub-Capacity Pricing, then SCRT would have automatically assigned the traditional sub-capacity value to that product for each LPAR listed on the NO89 statement. This is because the NO89 instances of the product make the product ineligible for Getting Started Sub-Capacity Pricing in each LPAR where

a NO89 instance of the product runs without a matching instance of the product level that is eligible for Getting Started Sub-Capacity Pricing.

For additional information about the Product Max Contributors section, see “Product Max Contributors section” on page 121. The sample Product Max Contributors section appears in Figure 70 on page 178.

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

          SYSID      Input Data Start      Input Data End
LPAR1      SYS1      01 Oct 2024 - 00:00  02 Nov 2024 - 22:00
LPAR3      SYS3      01 Oct 2024 - 06:05  02 Nov 2024 - 23:45
LPAR2      SYS2      01 Oct 2024 - 00:00  31 Oct 2024 - 22:30
LPARVM      -        -                  -
CPC                01 Oct 2024 - 00:00  02 Nov 2024 - 23:45

==N5=====
DETAIL LPAR DATA SECTION

          Highest Hour Count Date/Time          2nd Highest Hour Count Date/Time
LPAR1          32          1 03 Oct 2024 - 07:00          30          5 03 Oct 2024 - 06:00
LPAR3          35          6 03 Oct 2024 - 14:00          34          11 03 Oct 2024 - 13:00
LPAR2          41          4 04 Oct 2024 - 11:00          40          7 04 Oct 2024 - 10:00
CPC            107          5 03 Oct 2024 - 12:00          106         14 03 Oct 2024 - 11:00

==P5=====
PRODUCT MAX CONTRIBUTORS

Product Name                                Product ID Highest Date/Time          LPAR  LPAR  LPAR  GROUP
                                           LPAR1  LPAR3  LPAR2  LGRP A
z/OS V1                                5694-A01      98 06 Oct 2024 - 11:00      27   32   39
DB2 UDB for z/OS V8                    5625-DB2      71 03 Oct 2024 - 14:00       0   34   37
CICS TS for z/OS V3                     5655-M15      98 06 Oct 2024 - 11:00      27   32   39
MQSeries for z/OS V6                    5655-L82      71 03 Oct 2024 - 14:00       0   34   37
IMS V8                                  5655-C56      98 06 Oct 2024 - 11:00      27   32   39
IBM Enterprise Cobol for z/OS and OS/390 V3 5655-G53      39 04 Oct 2024 - 10:00       0    0   39
File Manager for z/OS V6                 5655-P17      60 09 Oct 2024 - 15:00      29   31    0
IBM WebSphere Application Server for z/OS V6 5655-N01      64 03 Oct 2024 - 14:00       0   27   37
WebSphere Message Broker for z/OS V6      5655-M74      64 03 Oct 2024 - 14:00       0   27   37
WebSphere Transformation Extender for z/OS V8 5655-R95      47 03 Oct 2024 - 12:00       0   23   24

```

Figure 70. Sample sub-capacity report for Getting Started Sub-Capacity Pricing: SMF / SCRT89 Input Data Statistics, Detail LPAR Data, and Product Max Contributors sections

Product Grid Snapshot section

There are no differences between the Product Grid Snapshot section for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report. For details about this section, see “Product Grid Snapshot section” on page 123. The sample Product Grid Snapshot section appears in Figure 71 on page 179.

Sub-Capacity Capping Enforced by SCRT

The MSU values that are listed in the Sub-Capacity Capping Enforced by SCRT section are the same for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report. Usually, the number of capped hours is the same for the LPARs or LPAR groups where the products are running at the same time that z/OS is running. This is not always true for Getting Started Sub-Capacity Pricing eligible products because those products might get the lesser Getting Started Sub-Capacity Pricing value, which is less than the capped value for the LPAR.

For additional information about the Sub-Capacity Capping Enforced by SCRT section, see “Sub-Capacity Capping Enforced by SCRT section” on page 124. The sample Sub-Capacity Capping Enforced by SCRT section appears in Figure 71 on page 179.

```

==Q5=====
PRODUCT GRID SNAPSHOT

```

Product Name	Product ID	LPAR1	LPAR3	LPAR2
z/OS V1	5694-A01	95.2%	92.7%	91.5%
DB2 UDB for z/OS V8	5625-DB2		92.7%	91.5%
CICS TS for z/OS V3	5655-M15	92.4%	92.7%	91.5%
MQSeries for z/OS V6	5655-L82		92.7%	91.4%
IMS V8	5655-C56	92.4%	92.5%	91.5%
IBM WebSphere Application Server for z/OS V6	5655-N01		92.7%	91.5%
WebSphere Message Broker for z/OS V6	5655-M74		92.7%	91.5%
WebSphere Transformation Extender for z/OS V8	5655-R95		74.4%	79.2%

```

==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS

```

Product Name	Product ID	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	GROUP LGRPA
z/OS V1	5694-A01			8	
DB2 UDB for z/OS V8	5625-DB2			8	
CICS TS for z/OS V3	5655-M15			8	
MQSeries for z/OS V6	5655-L82			8	
IMS V8	5655-C56			8	
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53				
File Manager for z/OS V6	5655-P17			8	
IBM WebSphere Application Server for z/OS V6	5655-N01			8	
WebSphere Message Broker for z/OS V6	5655-M74			8	
WebSphere Transformation Extender for z/OS V8	5655-R95				

Figure 71. Sample sub-capacity report for Getting Started Sub-Capacity Pricing: Product Grid Snapshot and Sub-Capacity Capping Enforced by SCRT sections

Getting Started Sub-Capacity Pricing Assigned section

This optional report section appears because the Tool MSU value for one or more products in the Product Summary Information section of this sub-capacity report is based on the Getting Started Sub-Capacity Pricing value for one or more LPARs or z/OS guest systems. In the sample configuration, all of the hours for all three of the Getting Started Sub-Capacity Pricing eligible products running in LPAR3 are assigned the Getting Started Sub-Capacity Pricing value because those products always met all of the eligibility requirements. Note that the number of hours reported for LPAR3 (and all of the other LPARs) is based on the number of hours that each product actually ran in that LPAR and not on the number of hours when z/OS may have been running in that LPAR.

Note that the WebSphere Application Server and WebSphere Message Broker products are only assigned Getting Started Sub-Capacity Pricing for 22 hours in LPAR2 because the Getting Started Sub-Capacity Pricing value for those products was less than the traditional value for only 22 hours during the reporting period. Also note that the WebSphere Transformation Extender product is assigned the Getting Started Sub-Capacity Pricing value for LPAR2 for every hour in which that product ran in LPAR2.

The sample Getting Started Sub-Capacity Pricing Assigned section appears in Figure 72 on page 180.

Getting Started Sub-Capacity Pricing Equal to Traditional Sub-Capacity section

This optional report section appears because one or more of the products in the Product Summary Information section of this sub-capacity report have Getting Started Sub-Capacity Pricing values that equal the traditional sub-capacity values for those products. In the sample configuration, that condition only occurred for both the WebSphere Application Server and WebSphere Message Broker products for 17 hours in LPAR2. The GSTRACE file lists the hours by LPAR by product when this condition occurred.

Note that the hours for Getting Started Sub-Capacity Pricing eligible products are posted under an LPAR group whenever the LPAR group cap is equal to the sum of the product's Getting Started Sub-Capacity Pricing value for the LPARs in the LPAR group. That is not the case for this sample configuration.

The sample Getting Started Sub-Capacity Pricing Equal to Traditional Sub-Capacity section appears in [Figure 72 on page 180](#).

Getting Started Sub-Capacity Pricing Not Assigned section

This optional report section appears because one or more of the products in the Product Summary Information section of this sub-capacity report have Getting Started Sub-Capacity Pricing values that are higher than the traditional sub-capacity values for those product. Note that SCRT assigned the traditional sub-capacity value to those products under those conditions. In the sample configuration, that condition occurred for both the WebSphere Application Server and WebSphere Message Broker products for 642 hours in LPAR2. The GSTRACE file lists the hours by LPAR by product when this condition occurred.

Note that the hours for Getting Started Sub-Capacity Pricing eligible products are posted under an LPAR group whenever the LPAR group cap is less than the sum of the product's Getting Started Sub-Capacity Pricing value for the LPARs in the LPAR group. That is not the case for this sample configuration.

The sample Getting Started Sub-Capacity Pricing Not Assigned section appears in [Figure 72 on page 180](#).

==S7=====					
GETTING STARTED SUB-CAPACITY PRICING ASSIGNED - HOURS					
Product Name	Product ID	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	
IBM WebSphere Application Server for z/OS V6	5655-N01		697	22	
WebSphere Message Broker for z/OS V6	5655-M74		697	22	
WebSphere Transformation Extender for z/OS V8	5655-R95		554	589	
==S8=====					
GETTING STARTED SUB-CAPACITY PRICING EQUAL TO TRADITIONAL SUB-CAPACITY - HOURS					
Product Name	Product ID	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	GROUP LGRPA
IBM WebSphere Application Server for z/OS V6	5655-N01			17	
WebSphere Message Broker for z/OS V6	5655-M74			17	
==S9=====					
GETTING STARTED SUB-CAPACITY PRICING NOT ASSIGNED - HOURS					
Product Name	Product ID	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	GROUP LGRPA
IBM WebSphere Application Server for z/OS V6	5655-N01			642	
WebSphere Message Broker for z/OS V6	5655-M74			642	

Figure 72. Sample sub-capacity report for Getting Started Sub-Capacity Pricing: Getting Started Sub-Capacity Pricing-related sections

Group Capacity LPARs section

There are no differences between the Group Capacity LPARs section for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report. For details about this section, see [“Group Capacity LPARs section” on page 125](#). The sample Group Capacity LPARs section appears in [Figure 73 on page 181](#).

```

==U5=====
GROUP CAPACITY LPARS

Group Name          LPARS
LGRPA               LPAR3(iii)(vii),LPAR2(iii)(vi)

(iii) LPAR was not a member of the group for the entire reporting interval
(vi)  LPAR was hardcapped during a portion of the reporting period while it was defined to the group
(vii) LPAR was hardcapped during the entire reporting period while it was defined to the group

==W5=====
This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.

```

Figure 73. Sample sub-capacity report for Getting Started Sub-Capacity Pricing: Group Capacity LPARs section

Chapter 9. Sub-capacity reporting for Integrated Workload Pricing (IWP)

IBM Integrated Workload Pricing (IWP) is an AWLC sub-capacity pricing option which is intended to enable qualifying z/OS customers to add a selected set of workloads to an existing LPAR without increasing the sub-capacity value of other products already running in that LPAR, as described in US Software Announcement letter 211-011. This sub-capacity pricing option is effective for SMF data collected on 2 February 2011 and later, provided that all of the required IWP conditions have been met.

Whenever the required IWP conditions have been met, SCRT will reduce the sub-capacity value that is reported for a select set of products, referred to as *IWP adjusted programs*, by subtracting the 4-hour rolling average utilization values of the general purpose processor time consumed by a different set of products, referred to as *IWP defining programs*, from the traditional sub-capacity values of each IWP adjusted program, on an hour-by-hour basis, per LPAR. For example, using IWP, a new WebSphere Application Server workload could be added to an LPAR which is already running CICS and IMS workloads, without significantly affecting the sub-capacity values of the CICS and IMS products for that LPAR.

For the case where a z/OS system runs as a guest under z/VM, each eligible z/OS guest system is treated as its own LPAR when applying any sub-capacity adjustments for IWP.

The IWP adjusted program product's sub-capacity values for the same hour are summed across all of the LPARs and z/OS guest systems in which the product runs to create an adjusted sub-capacity value for that product, for the given CPC, for each hour. The CPC hour with the product's highest sub-capacity value is used to report the product's sub-capacity value. If all of an IWP adjusted program's hourly sub-capacity values were adjusted during the SCRT reporting period, the sub-capacity value reported for that product will be a value that was adjusted by IWP. If one or more of the product's hourly sub-capacity values were not adjusted during the reporting period due to the failure to meet one or more of the IWP conditions, then the sub-capacity value reported for the product might not be a value that was adjusted by IWP. Note that the adjustment to each IWP adjusted program is made on an hour-by-hour and LPAR-by-LPAR basis, as opposed to subtracting the peak 4-hour rolling average for the CPC of each IWP defining program from the peak 4-hour rolling average for the CPC of each IWP adjusted program.

Table 32 on page 183 lists the IWP adjusted programs supported by SCRT V30.1.0. These are the products whose sub-capacity pricing values will be reduced when the required set of IWP conditions are met.

Note: For the most current list of IWP adjusted programs, see [Integrated Workload Pricing \(IWP\) Defining and Adjusted Programs \(www.ibm.com/downloads/documents/us-en/10a99803d9afd758\)](http://www.ibm.com/downloads/documents/us-en/10a99803d9afd758).

Table 32. IWP adjusted programs

Product family	Product name	Product ID
CICS¹	CICS/ESA V4	5655-018
	CICS TS for OS/390	5655-147
	CICS TS for OS/390 V2	5697-E93
	CICS TS for z/OS V3	5655-M15
	CICS TS for z/OS V4	5655-S97
	CICS TS for z/OS V5	5655-Y04

Table 32. IWP adjusted programs (continued)

Product family	Product name	Product ID
IMS²	IMS/ESA® V6	5655-158
	IMS V7	5655-B01
	IMS V8	5655-C56
	IMS V9	5655-J38
	IMS V10	5635-A01
	IMS V11	5635-A02
	IMS V12	5635-A03
	IMS V13	5635-A04
	IMS V14	5635-A05
Footnotes:		
1. The Value Unit Edition (VUE) of CICS (CICS TS VUE V5 (5722-DFJ)) is not eligible as an IWP adjusted program.		
2. The Value Unit Editions (VUE) of IMS (IMS DB VUE V12 (5655-DSQ), IMS DB VUE V13 (5655-DSM)) are not eligible as IWP adjusted programs.		

Table 33 on page 184 lists the IWP defining programs supported by SCRT V30.1.0. These are the products whose busy times will be used to reduce the sub-capacity values of the IWP adjusted programs when the required set of IWP conditions are met.

Note: For the most current list of IWP defining programs, see Integrated Workload Pricing (IWP) Defining and Adjusted Programs (www.ibm.com/downloads/documents/us-en/10a99803d9afd758).

Table 33. IWP defining programs

Product family	Product name	Product ID
IBM Business Process Manager for z/OS	IBM Business Process Manager for z/OS V7 ¹	5655-W05
	IBM Business Process Manager for z/OS V8 ²	5655-Y02
IBM Integration Bus for z/OS	IBM Integration Bus for z/OS V9	5655-IBB
	IBM Integration Bus for z/OS V10	5655-AB1
IBM Integration Bus Standard Edition for z/OS	IBM Integration Bus Standard Edition for z/OS V9	5655-IBC
	IBM Integration Bus Standard Edition for z/OS V10	5655-AB2
WebSphere Application Server	IBM WebSphere Application Server for z/OS V6 Service level 16 or higher	5655-N01
	IBM WebSphere Application Server for z/OS V7	5655-N02
	IBM WebSphere Application Server for z/OS V8	5655-W65
	IBM WebSphere Application Server for z/OS V9	5655-WAS
WebSphere Process Server	WebSphere Process Server for z/OS V6	5655-N53
WebSphere Message Broker	WebSphere Message Broker for z/OS V6 Level 6.1.0.2 or higher	5655-M74
	WebSphere Message Broker for z/OS V7	5655-V60
	WebSphere Message Broker for z/OS V8	5697-P44
WebSphere Enterprise Service Bus	WebSphere Enterprise Service Bus for z/OS V6	5655-R15
	WebSphere Enterprise Service Bus for z/OS V7	5655-W09
Footnotes:		
1. Prior to Release 7.5, this product was named WebSphere Process Server for z/OS V7.		
2. SCRT V22.2.0 is the first SCRT release to support this product as an IWP defining program.		

Qualifications for Integrated Workload Pricing

The following conditions must be met before IBM will accept and process sub-capacity reports where Integrated Workload Pricing is being reported via footnote (viii). Some of these conditions are also described in more detail elsewhere in this chapter but are consolidated here for ease of reference.

1. The CPC from which the SMF data is collected must be under AWLC or AEWLC sub-capacity terms and conditions. Because these terms and conditions require that products be licensed to IBM zEnterprise System servers or later (IBM z17, z16®, z15®, z14, z13, z13s, zBC12, zEC12, z114, and z196 CPCs), the SMF data must be collected from one or more z/OS systems running on zEnterprise CPCs. SCRT will only apply the IWP algorithm to SMF data collected from zEnterprise CPCs. zEnterprise CPCs that are in a qualified Parallel Sysplex with non-zEnterprise CPCs and whose IBM products are being charged under VWLC due to the terms of the AWLC Transition Charges for Sysplexes offering are still eligible for Integrated Workload Pricing.
2. The SMF data must be collected from a z/OS V1.10 or higher system that has SMF APAR OA31615 applied prior to collecting the SMF data. SCRT will not apply the IWP algorithm to SMF data collected from z/OS V1.9 or earlier z/OS systems. For zEnterprise CPCs with a mixture of z/OS V1.9 or earlier systems and z/OS V1.10 or later systems, SCRT will only apply the IWP algorithm to the SMF data collected from eligible z/OS 1.10 and later systems. For the z/OS 1.10 and later systems, SCRT will only apply the IWP algorithm to SMF data collected for the hours during which the system ran with APAR OA31615 applied. The result could be that SCRT applies the IWP algorithm to the SMF data collected only from a subset of the z/OS systems running on the CPC and only for a subset of the hours that a z/OS system was running.
3. For z/OS guest systems running under z/VM, all z/OS guest systems must be at z/OS V1.10 or higher. Integrated Workload Pricing is not supported on any z/OS guest system whenever at least one z/OS V1.6 or earlier guest system is running on the CPC anytime during that SCRT reporting period. If a z/OS V1.6 guest system is running on the CPC, then none of the products running on any of the z/OS guests, including z/OS V1.10 and higher guests with APAR OA31615 applied, are eligible for Integrated Workload Pricing on that CPC. This restriction is consistent with SCRT's existing set of z/OS V1.6 guest rules, where all sub-capacity products are assigned the maximum capacity of the LPAR in which they execute. Note that native z/OS systems, including z/OS V1.6, are not affected by SCRT's z/OS V1.6 guest rules. Products running on native z/OS V1.10 and higher systems are eligible for Integrated Workload Pricing when all of the other IWP conditions are met, even when z/OS V1.6 or earlier guest systems are also running in the same CPC. The guest restriction only applies to z/OS guest systems and has no effect on native z/OS systems.

When using z/OS guest systems, either z/VM V5.4 with APAR VM64798 or z/VM V6.1 with APAR VM64879 are required.

4. The traditional sub-capacity values of IWP adjusted programs are only adjusted when one or more IWP defining programs were running in the same z/OS system (LPAR) during the same 4-hour window that was used to determine the traditional sub-capacity value of the IWP adjusted product. IWP adjusted programs running in one z/OS system (LPAR) are not affected by IWP defining programs running during the same hour in a different z/OS system (LPAR).
5. Only SMF data collected on 2 February 2011 and later is eligible for Integrated Workload Pricing, as set forth in the Integrated Workload Pricing announcement letter. Although SCRT will generate sub-capacity reports based on the IWP algorithm for SMF data collected earlier than 2 February 2011, the IBM business process with reject those reports.
6. You must generate and submit a monthly sub-capacity report that contains IWP adjustments to IBM during the normal sub-capacity submission window (business as usual).
7. You must use SCRT V19.2.0 or later to generate the sub-capacity report that is submitted for Integrated Workload Pricing. Note that since this is a new release of SCRT (rather than a new version), IBM will also accept sub-capacity reports generated by either SCRT V19.10 or SCRT V19.2.0, but only SCRT V19.2.0 reports adjustments based on the IWP algorithm.
8. You must use the IWP control statement to request that SCRT generate a sub-capacity report using the IWP algorithms whenever all of the required IWP conditions are met. SCRT's default processing does

not use the IWP algorithms. For details, see [“INTEGRATED_WORKLOAD_PRICING control statement” on page 55](#).

9. You must agree to the IWP Addendum prior to submitting a sub-capacity report where footnote (viii) appears in the Product Information Summary section, since footnote (viii) indicates that the sub-capacity value for one or more products were adjusted based on IWP. Failure to formally agree to the IWP Addendum prior to submitting a sub-capacity report with footnote (viii) will cause the IBM business process to reject that report. The rejection will contain an explanation and request that you generate a replacement sub-capacity report without using the IWP control statement, and submit it to IBM in order to continue to receive traditional sub-capacity pricing for the IBM products on that CPC.

SCRT externals for Integrated Workload Pricing

SCRT provides the following externals for Integrated Workload Pricing. All of these externals are specific to IWP and are treated as optional when IWP is not being requested or does not apply to the SMF data being processed by SCRT.

1. SCRT supports an optional IWP control statement, **Integrated_Workload_Pricing**, which must be used to request that SCRT apply the IWP algorithm to the SMF data from all CPCs being processed by that execution of SCRT.
2. SCRT reports the IWP control statement, when used, in the Special Conditions section of the sub-capacity report.
3. SCRT uses a footnote, (viii), in the Product Summary Information section of the sub-capacity report to indicate that the Tool MSUs value for the corresponding product was adjusted by the IWP algorithm.
4. SCRT uses an optional IWPTRACE DD to track the reasons why an IWP defining program was not eligible for use by the IWP algorithm for a given hour and LPAR.

Integrated_Workload_Pricing control statement

The default SCRT processing does not support Integrated Workload Pricing. You must specify the **Integrated_Workload_Pricing** control statement to specifically request that SCRT use the IWP algorithm when generating sub-capacity reports.

Note: IBM authorizes the use of the IWP control statement only when you have signed or otherwise agreed to the terms of the IWP Addendum. IBM will reject all sub-capacity reports that show an IWP benefit if the IWP contractual terms and conditions are not in place.

See [“INTEGRATED_WORKLOAD_PRICING control statement” on page 55](#) for the format of the IWP control statement and for operational and IWP Addendum details.

When you specify the IWP control statement, SCRT applies the IWP algorithm only to the hours for the LPARs and z/OS guest systems that meet all of the IWP requirements. There is no guarantee that the resulting sub-capacity report will show any benefit due to IWP, especially when the z/OS release level or SMF APAR requirements have not been met for the entire SCRT reporting period. See [“Sample sub-capacity report for Integrated Workload Pricing” on page 189](#) for examples of the conditions under which SCRT does and does not apply the IWP algorithm.

Addition to the Special Conditions section

Whenever SCRT encounters the IWP control statement, SCRT displays a Special Conditions header for the control statement, followed by the IWP control statement itself, in the Special Conditions section of each sub-capacity report generated by that execution of SCRT. SCRT displays the IWP control statement in the Special Conditions section, even if SCRT never applied the IWP algorithm to any of the SMF data processed for that sub-capacity report because none of that data for that CPC met all of the IWP requirements. The control statement appears in the Special Conditions section to indicate that Integration Workload Pricing was requested for every CPC whose data was processed by that execution of SCRT, even if some or none of those CPCs ended up receiving any benefit from IWP.

Integrated Workload Pricing footnote

For the case where the Tool MSUs value for one or more products listed in the Product Summary Information section have been adjusted for IWP, SCRT applies the following changes to the Product Summary Information Section:

1. Generates an optional Footnotes column header for the MLC section.
2. Posts footnote (viii) in the Footnotes column for each IWP adjusted product whose peak sub-capacity value was adjusted for IWP. SCRT posts footnote (viii) whenever any of the hours at which the peak value for the product occurred received a benefit from the Getting Started Sub-Capacity Pricing algorithm.
3. Generates a footnote subsection immediately after the last IPLA product is listed. The footnote subsection contains a footnote header row followed by a description of footnote (viii).

Footnote (viii) is optional. All three footnote-related items will only be generated if one or more of the IWP adjusted programs requires such a footnote.

All numeric footnotes use lower case Roman numerals (i, ii, iii, ...) instead of Arabic numerals (1, 2, 3, ...) to prevent spreadsheet applications from automatically translating the footnotes into some other format.

Integrated Workload Pricing trace records

SCRT supports an optional IWPTRACE DD JCL statement, which is used to track the reasons why an IWP defining program was not eligible to be used for an IWP algorithm on a CPC, by image ID, by IWP defining program, by hour basis. The fields of the trace records that are written to this file are all separated by commas (that is, the file is in comma separated value format) so that the file can be downloaded in text format onto a workstation where it can be processed by a spreadsheet application, such as Microsoft Excel.

The very first record in this file is a header record which contains the column titles for the fields of all the remaining rows in the spreadsheet. The titles contained in the header record follow, along with an explanation of the corresponding fields in the actual trace records.

CPC (tttt-sssss)

The machine type (*tttt*) and serial number (*sssss*) of the CPC where the reported product ran.

Example: 2817-12345

Image Id

The Image ID of the z/OS system on which the reported product ran.

- For native z/OS systems that are not under group capping or whose group cap is higher than the product's IWP value, this value is the LPAR name.

Example: LPAR1

- For native z/OS systems that SCRT must cap under group capping, this value is the LPAR group name followed by the (Group) suffix.

Example: LGRPA (group)

- For z/OS guest systems, this value is the SID (SYSID) of the guest system.

Example: SYS1

Product

The product ID (PID) number of the IWP defining program that was not used by the IWP algorithm. Only IWP defining programs will be reported because they are the only products whose busy times are used when calculating the adjustments for the IWP adjusted programs.

Example: 5655-N01

Date:Hour(YyyyMmmDdd:Hhh)

The year (*Yyyy*), month (*Mmm*), day (*Ddd*), and hour (*Hhh*) for which the product specified in the **Product #** field was not used by the IWP algorithm.

Example: Y2024M02D02:H00 represents 00:00 on 2 February 2024.

Reason for not contributing to IWP

The reason why the busy time for the IWP defining program listed in the **Product #** field was not eligible for use by the IWP algorithm for adjusting the sub-capacity values of the IWP adjusted programs. The possible reasons for SCRT not using the IWP algorithm are listed below in priority order. Whenever multiple reasons apply, only the first applicable reason will be listed.

CPC

None of the IWP defining programs on the matching sub-capacity report are eligible for use by the IWP algorithm because those products did not run on an IBM z13, zEC12, z196, z13s, zBC12, or z114 CPC, which is one of the requirements for IWP. When this condition occurs, SCRT will only create a single trace entry for all of the image IDs, products, and dates for that CPC, since the same eligibility problem would apply to all of those entries. This trace entry will use the value `: : ALL : :` for the **Image Id**, **Product #** and **Date** values of that trace entry.

Missing APAR OA31615

None of the IWP defining programs for the specified CPC, image ID, and date are eligible for use by the IWP algorithm because the required SMF APAR, OA31615, has not been applied and the SMF type 89 records have not collected all of the data needed to support IWP. This trace entry will use the value `: : ALL : :` for the **Product #** field, as it applies to all eligible IWP defining programs. Note that this trace entry will also be created for z/OS V1.9 and earlier releases because SMF APAR OA31615 only supports z/OS V1.10 and later releases. This trace entry will also be created for the case where no IWP defining programs were running during the 4-hour window that ends with the hour listed in the trace entry.

Product level not eligible

The specified IWP defining program was not at the minimum required level, so the busy time for that product was not collected on SMF type 89 records and could not be used by the IWP algorithm. The minimum level for every IWP defining product that has such a requirement is listed in [Table 33 on page 184](#).

Not first product to register

The specified IWP defining program was not the first product to register in the address space where one or more other IWP defining programs also registered and ran. When that happens, the busy time for an IWP defining program that did not register first is not used as part of the IWP algorithm, because using that busy time would have the effect of using the same address space busy time twice in the IWP algorithm, once for the IWP defining program that registered first and again for the other IWP defining program that registered later in that same address space.

No IWP defining products

None of the eligible IWP defining programs were running during the 4-hour rolling average utilization window that would be used to adjust the sub-capacity value of any IWP adjusted programs running during the specified hour. Note that a value of `: : ALL : :` is listed for the **Product #** because all of the IWP defining programs were absent from that LPAR for that 4-hour window.

Sample configuration for Integrated Workload Pricing

[Table 34 on page 189](#) illustrates the sample hardware and software configuration for a CPC where two IWP adjusted programs were running during the reporting period. While the mix of software products and versions used in this sample configuration might not be realistic, it serves to illustrate some of the finer points of the IWP algorithm.

The sample configuration consists of a 2817 Model 404 CPC with a serial number of 02-12345. The CPC is configured to support four LPARs: three are running z/OS natively, each in its own LPAR, and the fourth LPAR is running Linux systems under z/VM. The z/OS systems in LPAR1, LPAR2, and LPAR3 have z/OS SYSIDs of SYS1, SYS2, and SYS3. LPAR2 and LPAR3 are configured to be members of LPAR group LGRPA. The z/OS systems are running a combination of the following z/OS products (listed in the order that they will appear on the sub-capacity report):

- DB2 V9
- CICS TS V4

- WebSphere MQ for z/OS V7
- IMS V11
- IMS V10
- IBM WebSphere Application Server V7
- IBM WebSphere Application Server V6 (service level 15)
- IBM WebSphere Application Server V6 (service level 16)
- WebSphere Message Broker for z/OS V7
- WebSphere Portal Enable for z/OS V6

CPC: 2817-404 (serial = 02-12345)			
LPAR1	LPAR2 (LPAR group: LGRPA)	LPAR3 (LPAR group: LGRPA)	LPARVM
z/OS V1.10 (SYSID = SYS1) DB2 V9 CICS TS V4 WebSphere MQ V7 IMS V10 WebSphere Application Server V7 WebSphere Application Server V6 (service level 16) WebSphere Message Broker V6 WebSphere Portal V6	z/OS V1.9 (SYSID = SYS2) DB2 V9 CICS TS V4 WebSphere MQ V7 IMS V10 WebSphere Application Server V7 WebSphere Application Server V6 (service level 16) WebSphere Message Broker V6 WebSphere Portal V6	z/OS V1.10 (SYSID = SYS3) DB2 V9 CICS TS V4 WebSphere MQ V7 IMS V11 WebSphere Application Server V6 (service level 15) WebSphere Portal V6	z/VM (Linux guests)

Sample control statement input to SCRT for Integrated Workload Pricing

SCRT supports additional externals for IWP. The sample sub-capacity report discussed in “[Sample sub-capacity report for Integrated Workload Pricing](#)” on page 189 is based on the following SCRT control statement and JCL inputs to SCRT:

SPECIAL DD control statements

The following SCRT control statement was used to request that SCRT use the IWP algorithm whenever all of the IWP conditions have been satisfied. See “[INTEGRATED_WORKLOAD_PRICING control statement](#)” on page 55 for a detailed description and restrictions on its use.

```
Integrated_Workload_Pricing
```

The following SCRT control statement was used to request that SCRT suppress the reporting of Getting Started Sub-Capacity Pricing statistics, as this chapter focuses on IWP, not GSSP. This control statement has no effect on whether SCRT uses the IWP algorithm. (See “[SUPPRESS_REPORT_SECTION control statement](#)” on page 58 for a detailed description.)

```
Suppress_Report_Section=GSSP
```

Sample sub-capacity report for Integrated Workload Pricing

The sub-capacity report for Integrated Workload Pricing has the same basic format and report sections as the sub-capacity report for a base z/OS system. However, the sub-capacity report for IWP always contains a Special Conditions section to list the IWP control statement and may also contain a footnote to indicate which IWP adjusted programs benefitted from the IWP algorithm.

The following sample report sections describe the differences between a sub-capacity report for a base z/OS system and a sub-capacity report that uses the IWP control statement. For a complete description

of all of the fields appearing in a base sub-capacity report see [“Sample sub-capacity report for single-version z/OS native systems”](#) on page 105.

Billing-related section

The billing-related section of the sub-capacity report displays the customer information, tool information, special conditions, product summary information, and detail data collection information, which are all described in [“Billing-related section”](#) on page 105.

Customer Information section

There are no differences between the Customer Information section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the IWP control statement. For details about this section, see [“Customer information”](#) on page 105. The sample Customer Information section appears in [Figure 74](#) on page 190.

Tool Information section

There are no differences between the Tool Information section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the IWP control statement. For details about this section, see [“Tool Information section”](#) on page 107. The sample Tool Information section appears in [Figure 74](#) on page 190.

```

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time                02 Mar 2024 - 09:17
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter:  444-999-9999

Customer Name                ABC Corporation
Customer Number              8971234567
Machine Serial Number        02-12345
Machine Type and Model       2817-404
Machine Rated Capacity (MSUs) 110
Purchase Order Number        (optional)
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust
the billable MSUs in inventory for all MLC Products listed under the MLC Product Name
column on this report. In accordance with our agreement, IBM will treat a change in product
licensed capacity as an order. If the MSUs have changed since the last report, software billing
based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to
bill those IPLA products listed under the IPLA Product Name column in this report which
exceed your entitled capacity. In accordance with our agreement, IBM will treat the use
of a product in excess of its entitled capacity as an order and you will be billed for
the amount in excess of your entitlement.

==C5=====
TOOL INFORMATION

Tool Release                30.1.0
Reporting Period            2 Feb, 2024 - 1 Mar, 2024 inclusive (28 days)

```

Figure 74. Sample sub-capacity report for Integrated Workload Pricing: Customer information and Tool Information sections

Special Conditions section

One of the major differences between the Special Conditions section for a basic z/OS sub-capacity report and a sub-capacity report generated using the IWP control statement is that the latter will always contain a Special Conditions section, which will list the IWP control statement. The IWP control statement is listed to reflect the fact that it was provided to SCRT to request the use of the IWP algorithm when

generating this sub-capacity report. If an Exclude control statement is also used for this execution of SCRT for this CPC, the information related to the exclusion information is also listed in the Special Conditions section. For general details about this section, see [“Special Conditions section” on page 107](#). The sample Special Conditions section appears in [Figure 75 on page 191](#).

```
==D5=====
SPECIAL CONDITIONS

IWP Control Statement
Integrated_Workload_Pricing
```

Figure 75. Sample sub-capacity report for Integrated Workload Pricing: Special Conditions section

Product Summary Information section

An optional Footnotes column header appears in the MLC section whenever SCRT needs to post a footnote for an IWP adjusted program listed in this section of the report. The Footnotes heading does not appear whenever there are no footnotes for this section.

There are no differences between the MLC Product Name, MLC Product ID, Customer Comment, IPLA Product Name, and IPLA Product ID fields on a basic z/OS sub-capacity report and on a sub-capacity report generated using the IWP control statement. For a description of these fields see [“Product Summary Information section” on page 110](#). The sample Product Summary Information section appears in [Figure 76 on page 192](#).

Tool MSUs

The values listed in the Tool MSUs column represent the peak sub-capacity values calculated by SCRT for each product using the appropriate sub-capacity rules for that product. SCRT uses the traditional 4-hour rolling average calculation for MLC products that are not IWP adjusted programs and for IPLA products that are not Getting Started Sub-Capacity Pricing products. SCRT uses the IWP algorithm for IWP adjusted programs for those hours when all of the IWP requirements have been met, and it uses the traditional 4-hour rolling average calculation for IWP adjusted programs for those hours when all of the IWP requirements have not been met. SCRT uses the Getting Started Sub-Capacity Pricing algorithm for IPLA products that are eligible for Getting Started Sub-Capacity Pricing.

See [“Product Max Contributors section” on page 192](#) for a more detailed discussion of how the IWP algorithm affects the IWP adjusted programs that are listed in this section of the sub-capacity report. For additional general information about Tool MSUs, see [“Product Summary Information section” on page 110](#).

Customer MSUs

There are no differences between the Customer MSUs field on a basic z/OS sub-capacity report and on a sub-capacity report generated using the IWP control statement. For additional information about Customer MSUs, see [“Product Summary Information section” on page 110](#).

Footnotes

SCRT posts footnote (viii) whenever the Tool MSUs value posted for a product is based totally or partially on adjustments by the IWP algorithm. The Tool MSUs are partially based on the IWP algorithm whenever the posted value for the product is combined from multiple LPARs where the product's value is based on the traditional sub-capacity value from some LPARs and on the IWP adjusted values from other LPARs. Whenever SCRT posts footnote (viii) in the Footnotes column, it also generates a row with a footnote header followed by a description for the footnote.

The Product Summary Information section in [Figure 76 on page 192](#) uses footnote (viii) to indicate that one or more of the peak sub-capacity values for IWP adjusted programs CICS TS V4 and IMS V10 were adjusted by the IWP algorithm. However, this was not true for the IWP adjusted program IMS V11. See [“Product Max Contributors section” on page 192](#) for a more detailed discussion of why footnote (viii) was posted for certain IWP adjusted programs and not for others.

SCRT also posts footnote (iv) whenever the Tool MSUs value posted for a product is based totally or partially on having been adjusted by the Getting Started Sub-Capacity Pricing algorithm. See [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#) for a detailed description of Getting Started Sub-Capacity Pricing.

SCRT posts the footnote whenever any of the hours at which the peak value occurred received a benefit from the Getting Started Sub-Capacity Pricing algorithm.

```

=====
PRODUCT SUMMARY INFORMATION

```

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)	Footnotes
z/OS V1	5694-A01	98			
DB2 V9 for z/OS	5635-DB2	98			
CICS TS for z/OS V4	5655-S97	82			(viii)
WebSphere MQ for z/OS V7	5655-R36	98			
IMS V11	5635-A02	41			
IMS V10	5635-A01	47			(viii)

IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)	Footnotes
IBM WebSphere Application Server for z/OS V7	5655-N02	24			(iv)
IBM WebSphere Application Server for z/OS V6	5655-N01	42			(iv)
WebSphere Message Broker for z/OS V7	5655-V60	14			(iv)
IBM WebSphere Portal Enable for z/OS V6	5655-R17	42			(iv)

Footnotes:
 (iv) Product's Tool MSUs are based on Getting Started Sub-capacity Pricing from one or more LPARs
 (viii) Product's Tool MSUs are based on Integrated Workload Pricing from one or more LPARs

Figure 76. Sample sub-capacity report for Integrated Workload Pricing: Product Summary Information section

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the Product Summary Information section.

SMF / SCRT89 Input Data Statistics section

There are no differences between the SMF / SCRT89 Input Data Statistics section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the IWP control statement. For details about this section, see [“SMF / SCRT89 Input Data Statistics section” on page 118](#). The sample SMF / SCRT89 Input Data Statistics section appears in [Figure 77 on page 196](#).

Detail LPAR Data Section

There are no differences between the Detail LPAR Data Section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the IWP control statement. For details about this section, see [“Detail LPAR Data section” on page 119](#). The sample Detail LPAR Data Section appears in [Figure 77 on page 196](#).

Product Max Contributors section

The MSU values that are listed in the Product Max Contributors section are the same for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report, with one subtle change. While all of the MSU values listed in this section continue to reflect the product MSUs by LPAR or LPAR Group for the earliest hour for the product's maximum MSU value, the MSU values posted for Getting Started Sub-Capacity Pricing eligible products are based on the Getting Started Sub-Capacity Pricing algorithm whenever that value is less than the traditional sub-capacity value for that product for that hour. The result is that some of the LPAR values posted for Getting Started Sub-Capacity Pricing eligible products are less than the MSUs posted for z/OS, even when the Getting Started Sub-Capacity Pricing eligible product runs in all of the same LPARs and at the same times as z/OS.

Note: The MSUs for Getting Started Sub-Capacity Pricing eligible products are posted under an LPAR group whenever the LPAR group cap is less than the Getting Started Sub-Capacity Pricing value for the earliest hour with the maximum product value. That is not the case for this sample configuration.

This section provides an additional layer of detail behind the sub-capacity values reported for the products listed in the Product Summary Information section by reporting the date/time that the sub-capacity value occurred together with the product's sub-capacity contribution from each LPAR in which that product was running at that date and time. When the *same* peak sub-capacity value occurs for multiple hours, the Product Max Contributors section lists the data for the *earliest* hour. The impact of SCRT reporting a product's sub-capacity value based on the earliest peak hour whenever that peak value occurs for multiple hours is important enough to reiterate in the following explanations because the IWP conditions for the other peak hours can differ from the IWP conditions for the earliest peak hour.

Given the level of detail in this section, the sample Product Max Contributors section shown in [Figure 77 on page 196](#) will be used to explain Integrated Workload Pricing in more detail. However, this explanation assumes you have a general understanding of the contents and use of each field in the Product Max Contributors section, as described in [“Product Max Contributors section” on page 121](#).

How SCRT implements Integrated Workload Pricing

It is important to understand that SCRT uses the appropriate algorithm to calculate the sub-capacity value for each product for each hour for each LPAR, independently of what is happening with that product on other LPARs or guest systems during that same hour. Then, SCRT sums each of the product's sub-capacity values across all of the LPARs for the same hour to calculate a sub-capacity value for the product for the CPC for that hour. After performing these sub-capacity calculations for the product across the CPC for all of the hours during the reporting period, SCRT selects the earliest hour across the CPC with the product's peak sub-capacity value and reports the date and time for that hour, along with the individual sub-capacity values for the product for each LPAR for that same hour, in the Product Max Contributors section.

For products like z/OS, whose sub-capacity is entirely based on the 4-hour rolling average MSU utilization of the LPAR, the 4-hour rolling average MSU utilization is always reported as the product's sub-capacity value. In the sample Product Max Contributors section, the peak z/OS sub-capacity value across the CPC was 98 MSUs and that peak first occurred on 03 February 2024 at 12:00. The z/OS product running in LPAR1, LPAR2, and LPAR3 contributed 27, 32, and 39 MSUs, respectively, to the z/OS peak sub-capacity value of 98 MSUs for the CPC.

Based on the configuration (in [“Sample configuration for Integrated Workload Pricing” on page 188](#)), DB2 V9 and WebSphere MQ V7 are the only other products whose sub-capacity values are based entirely on the 4-hour rolling average utilization of the LPARs in which they run. Because these products ran in the same LPARs as z/OS and at the same time that z/OS ran, both the earliest peak hour and the sub-capacity values for Db2 and WebSphere MQ are identical to those for z/OS.

CICS TS V4, IMS V10, and IMS V11 are IWP adjusted programs, which means that SCRT applies the IWP algorithm to each of those products on an hour-by-hour, LPAR-by-LPAR basis, whenever all of the IWP required conditions have been met for that hour on that LPAR. Whenever all of the required IWP conditions have not been met, SCRT uses the traditional 4-hour rolling average MSU utilization for that hour on that LPAR when reporting that product's sub-capacity value.

A note about IWP adjustments

The MSU values used for IWP adjustments are calculated internally by SCRT and do not appear in, nor can they be derived directly from, the Product Max Contributors section of the sub-capacity report. The Product Max Contributors section is only intended to show the contribution of each LPAR to the specified product's peak sub-capacity value for earliest peak hour.

Results for CICS TS V4: After SCRT applied the appropriate algorithm hour by hour, LPAR by LPAR, and then summed the sub-capacity values for the same hour across all of the LPARs in which the product ran during the same hour, SCRT reported 03 February 2024 at 11:00 as the earliest hour with the peak sub-capacity value for CICS TS V4, as indicated in the sample Product Max Contributors section. Note that the earliest peak hour for CICS TS V4 is different than the earliest peak hour for all of the other products in the sample Product Max Contributors section. SCRT reported a value of 10 MSUs for LPAR1, based on the IWP algorithm for CICS TS V4 for that hour on that LPAR because all of the following IWP conditions were met:

1. The SCRT **Integrated_Workload_Pricing** control statement was used to request IWP (otherwise, SCRT would not use the IWP algorithm and would not issue any IWPTRACE records).

2. The CPC was an IBM z/16, z15, z14, z13, zEC12, z196, z13s, zBC12, or z114 (otherwise, the IWPTRACE file would have contained a single trace entry indicating that the CPC was not a supported machine type).
3. z/OS V1.10 was running with the required APAR OA31615 applied for that hour (otherwise, the IWPTRACE file would have contained a trace entry indicating that APAR OA31615 had not been applied for that LPAR for that hour).
4. One or more eligible IWP defining programs that met the minimum eligibility level required for IWP were running during that hour.

The result is that SCRT used the 4-hour rolling average busy times for 03 February 2024 at 11:00 for each of the IWP defining programs listed above (with values of 5, 6, and 4 MSUs, respectively) to calculate an adjustment value of 15 MSUs for CICS TS V4 for LPAR1. SCRT then subtracted the 15 MSU adjustment value from the 4-hour rolling average LPAR utilization for that hour for LPAR1, which was 25 MSUs, giving a result of 10 MSUs to be reported as the sub-capacity value of CICS TS V4 for LPAR1 for that hour.

Because the sub-capacity value for CICS TS V4 for that hour for LPAR1 was adjusted based on the IWP algorithm, SCRT posted footnote (viii) for CICS TS V4 in the Product Summary Information section, regardless of whether the sub-capacity value for CICS TS V4 was adjusted on any other LPAR during the same hour. Note that no IWPTRACE entries would have been created for this hour and LPAR because all of the required IWP conditions had been met. Also note that the values of 5, 6, and 4 MSUs are not the sub-capacity values that will be reported for the products that had those busy times. Instead, SCRT will use the Getting Started Sub-Capacity Pricing algorithm (as described in [“Getting Started Sub-Capacity Pricing algorithm”](#) on page 168) when reporting the sub-capacity values for each of those products.

For CICS TS V4 on LPAR2, applying the same process for the same hour, 03 February 2024 at 11:00, produced different results than for LPAR1 because one of the required IWP conditions was not met on LPAR2. The z/OS system running in LPAR2 during that hour did not have required APAR OA31615 applied, which made that hour on LPAR2 ineligible for the IWP algorithm. Note that APAR OA31615 is only available for z/OS V1.10 and later systems. Instead of using the IWP algorithm, SCRT reported the 4-hour rolling average LPAR utilization value of 33 MSUs for that hour as the sub-capacity value for CICS TS V4 for LPAR2. SCRT would have generated an IWPTRACE entry with a reason of Missing APAR OA31615 for that hour for LPAR2.

For CICS TS V4 on LPAR3, applying the same process for the same hour, 03 February 2024 at 11:00, produced different results than were produced for LPAR1 or LPAR2 because a different IWP condition was not met on LPAR3. The z/OS V1.10 system with APAR OA31615 applied and running in LPAR3 during that hour did not have any eligible IWP defining programs running during any of the 4-hour windows that make up the 4-hour rolling average for that hour. The only IWP defining program running on LPAR3 was WebSphere Application Server V6 at service level 15, which ran at that level for the entire reporting period. Unfortunately, that service level does not collect the product busy time information that IWP requires. The result is that the IWP adjustment will always be zero MSUs, which means that SCRT will essentially be reporting the 4-hour rolling average LPAR utilization for that hour, 39 MSUs, as the sub-capacity of CICS TS V4 for LPAR3. SCRT would have also generated an IWPTRACE entry for WebSphere Application Server V6 with a reason of Product level not eligible for that hour for LPAR3.

Results for IMS V11: SCRT applied the same IWP process as described for CICS TS V4 to IMS V11. Since IMS V11 only ran on LPAR3, SCRT only applied the IWP process to IMS V11 on LPAR3. The result was that IMS V11 did not receive an IWP adjustment for the same reason that CICS TS V4 did not receive an IWP adjustment for LPAR3—that is, the only IWP defining program that was running on LPAR3 during the IMS V11 peak hour was not running at the required product service level and so there was no adjustment to apply to IMS V11. SCRT, therefore, reported the highest 4-hour rolling average LPAR3 utilization of 41 MSUs as the IMS V11 sub-capacity value in the sample Product Max Contributors section for both LPAR3 and for the CPC, since IMS V11 did not run on any of the other LPARs on the CPC. Based on the sample Product Max Contributors section, the earliest peak hour occurred at 04 February 2024 at 11:00, which is a different earliest peak hour than was reported for any of the other products in this section.

Even though the peak sub-capacity value for IMS V11 may have occurred during multiple (different) hours, footnote (viii) was not posted for IMS V11 in the Product Summary information section because the peak IMS V11 value was never adjusted for any of those peak hours based on the busy time of one or

more IWP adjusted programs. Note that the IWPTRACE entry for WebSphere Application Server V6 with a reason of Product level not eligible for that hour for LPAR3 is only generated once for the LPAR, even when multiple IWP adjusted programs are affected during the same hour and LPAR, since the same IWPTRACE entry would apply to both CICS TS V4 and IMS V11.

Results for IMS V10: SCRT applied the same IWP process as described for CICS TS V4 to IMS V10, but only for LPAR1 and LPAR2 because IMS V10 only ran in those two LPARs. Since IMS V10 ran in LPAR1 and LPAR2, but CICS TS V4 ran in LPAR1, LPAR2, and LPAR3, both the peak sub-capacity value and earliest peak hour for IMS V10 are different than those for CICS TS V4. Based on the sample Product Max Contributors section, the earliest peak hour for IMS V10 was 03 February 2024 at 15:00.

Since the IWP conditions in each LPAR remained the same for the sample configuration for the entire reporting period, LPAR1 passed all of the IWP conditions and LPAR2 failed the same IWP condition as described for CICS TS V4. Using the IWP algorithm for IMS V10 running on LPAR1 for that peak hour, SCRT used the 4-hour rolling average busy times for WebSphere Application Server V7, WebSphere Application Server V6, and WebSphere Message Broker V6 (with values of 3, 4, and 8 MSUs, respectively) to calculate that IMS V10 should be adjusted by 15 MSUs. Subtracting 15 MSUs from the 4-hour rolling average LPAR utilization value of 28 MSUs for that hour resulted in SCRT reporting 13 MSUs for IMS V10 for LPAR1. Since the sub-capacity value for IMS V10 for LPAR1 for this earliest peak hour was adjusted based on the IWP algorithm, SCRT posted footnote (viii) for IMS V10 in the Product Summary Information section, regardless of whether the sub-capacity value of IMS V10 is adjusted on any other LPAR during the same hour. No IWPTRACE entry was generated for LPAR1 for this hour, as all of the IWP defining programs passed all of the IWP conditions.

For IMS V10 on LPAR2, applying the same process for IMS V10 for the same hour, 03 February 2024 at 15:00, produced results similar to those for CICS TS V4 for LPAR2. The z/OS system running in LPAR2 during that hour did not have required APAR OA31615 applied, which made that hour on LPAR2 ineligible for the IWP algorithm for any IWP adjusted program, including IMS V10, running during that hour on that LPAR. Instead of using the IWP algorithm, SCRT reported the 4-hour rolling average LPAR utilization value of 34 MSUs for that hour as the sub-capacity value for IMS V10 for LPAR2. SCRT would have generated a single IWPTRACE entry with a reason of Missing APAR OA31615 for that hour for LPAR2, regardless of the number of IWP adjusted programs that were affected. In the sample configuration, both CICS TS V4 and IMS V10 would be covered by the same IWPTRACE entry for LPAR2 for the same specified hour.

The remaining products listed in the sample Product Max Contributors section are all Getting Started Sub-Capacity Pricing eligible products. For an explanation of how SCRT determines the sub-capacity values for Getting Started Sub-Capacity Pricing eligible products, see [“Getting Started Sub-Capacity Pricing algorithm”](#) on page 168.

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
LPAR1	SYS1	02 Feb 2024 - 00:00	02 Mar 2024 - 22:00
LPAR3	SYS3	02 Feb 2024 - 00:00	02 Mar 2024 - 23:45
LPAR2	SYS2	02 Feb 2024 - 00:00	28 Feb 2024 - 22:30
LPARVM	-	-	-
CPC		02 Feb 2024 - 00:00	02 Mar 2024 - 23:45

```

==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR1	32	1	03 Feb 2024 - 07:00	30	5	03 Feb 2024 - 06:00
LPAR3	35	6	03 Feb 2024 - 14:00	34	11	03 Feb 2024 - 13:00
LPAR2	41	4	04 Feb 2024 - 11:00	40	7	04 Feb 2024 - 10:00
CPC	98	5	03 Feb 2024 - 12:00	97	14	03 Feb 2024 - 11:00

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR3	LPAR LPAR2	GROUP LGRPA
z/OS V1	5694-A01	98	06 Feb 2024 - 12:00	27	32	39	
DB2 V9 for z/OS	5635-DB2	98	03 Feb 2024 - 12:00	27	32	39	
CICS TS for z/OS V4	5655-S97	82	06 Feb 2024 - 11:00	10	33	39	
WebSphere MQ for z/OS V7	5655-R36	98	03 Feb 2024 - 12:00	27	32	39	
IMS V11	5635-A01	41	06 Feb 2024 - 11:00	0	0	41	
IMS V10	5635-A01	47	04 Feb 2024 - 15:00	13	34	0	
IBM WebSphere Application Server for z/OS V7	5655-N02	24	09 Feb 2024 - 14:00	10	14	0	
IBM WebSphere Application Server for z/OS V6	5655-N01	42	03 Feb 2024 - 13:00	12	16	14	
WebSphere Message Broker for z/OS V7	5655-V60	14	03 Feb 2024 - 14:00	8	6	0	
IBM WebSphere Portal Enable for z/OS V6	5655-R17	42	03 Feb 2024 - 13:00	12	16	14	

Figure 77. Sample sub-capacity report for Integrated Workload Pricing: SMF / SCRT89 Input Data Statistics, Detail LPAR Data, and Product Max Contributors sections

Product Grid Snapshot section

There are no differences between the Product Grid Snapshot section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the IWP control statement. For details about this section, see [“Product Grid Snapshot section” on page 123](#). The sample Product Grid Snapshot section appears in [Figure 78 on page 197](#).

Sub-Capacity Capping Enforced by SCRT section

The MSU values that are listed in the Sub-Capacity Capping Enforced by SCRT section are the same for a basic z/OS sub-capacity report and for a sub-capacity report generated using the IWP control statement. Usually, the number of capped hours is the same for the LPARs or LPAR groups where the products are running at the same time that z/OS is running. This is not always true for IWP adjusted programs because those products might get the lower IWP adjusted value, which is less than the capped value for the LPAR.

For additional information about the Sub-Capacity Capping Enforced by SCRT section, see [“Sub-Capacity Capping Enforced by SCRT section” on page 124](#). The sample Sub-Capacity Capping Enforced by SCRT section appears in [Figure 78 on page 197](#).

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                                Product ID    LPAR1    LPAR3    LPAR2
z/OS V1                                    5694-A01      92.4%    92.7%    91.5%
DB2 V9 for z/OS                           5635-DB2      92.4%    92.7%    91.5%
CICS TS for z/OS V4                       5655-S97      92.4%    92.7%    91.5%
WebSphere MQ for z/OS V7                  5655-R36      92.4%    92.7%    91.5%
IMS V11                                    5635-A02      92.4%    92.7%    91.5%
IMS V10                                    5635-A01      92.4%    92.7%    91.5%
IBM WebSphere Application Server for z/OS V7 5655-N02      92.4%    92.7%    91.5%
IBM WebSphere Application Server for z/OS V6 5655-N01      92.4%    92.7%    91.5%
WebSphere Message Broker for z/OS V7       5655-V60      92.4%    92.7%    91.5%
IBM WebSphere Portal Enable for z/OS V6     5655-R17      92.4%    92.7%    91.5%

```

```

==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS

Product Name                                Product ID    LPAR    LPAR    LPAR    GROUP
                                                LPAR1      LPAR2    LPAR3    LGRPA
z/OS V1                                    5694-A01      8
DB2 V9 for z/OS                           5635-DB2      8
CICS TS for z/OS V4                       5655-S97      8
WebSphere MQ for z/OS V7                  5655-R36      8
IMS V11                                    5635-A02      8
IMS V10                                    5635-A01      8
IBM WebSphere Application Server for z/OS V7 5655-N02      8
IBM WebSphere Application Server for z/OS V6 5655-N01      8
WebSphere Message Broker for z/OS V7       5655-V60      8
IBM WebSphere Portal Enable for z/OS V6     5655-R17      8

```

Figure 78. Sample sub-capacity report for Integrated Workload Pricing: Product Grid Snapshot and Sub-Capacity Capping Enforced by SCRT sections

Group Capacity LPARs section

There are no differences between the Group Capacity LPARs section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the IWP control statement. For details about this section, see [“Group Capacity LPARs section” on page 125](#). The sample Group Capacity LPARs section appears in [Figure 79 on page 197](#).

```

==U5=====
GROUP CAPACITY LPARS

Group Name          LPARS
LGRPA              LPAR2(iii)(vii),LPAR3(iii)(vi)

(iii) LPAR was not a member of the group for the entire reporting interval
(vi)  LPAR was hardcapped during a portion of the reporting period while it was defined to the group
(vii) LPAR was hardcapped during the entire reporting period while it was defined to the group

==W5=====
This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.

```

Figure 79. Sample sub-capacity report for Integrated Workload Pricing: Group Capacity LPARs section

Chapter 10. Sub-capacity reporting for Mobile Workload Pricing (MWP)

IBM Mobile Workload Pricing (MWP) is a sub-capacity pricing option under Advanced Workload License Charges (AWLC), Advanced Entry Workload License Charges (AEWLC), or System z New Application License Charges (zNALC). It is intended to enable qualifying z/OS customers to mitigate the cost of growth for mobile transactions processed by select MWP defining programs, specifically in cases where the higher mobile transaction volumes may cause a spike in machine utilization (as described in US Software Announcement Letter 214-223). This sub-capacity pricing option is effected for SMF data collected on 2 June 2014 and later, provided that all of the required MWP conditions have been met.

Overview of Mobile Workload Pricing

IBM Mobile Workload Pricing (MWP) provides an enhanced way of reporting z/OS system utilization, which can improve price/performance for sub-capacity eligible programs when running in the same LPAR as mobile workloads processed by select MWP defining programs. Mobile workloads are defined as the transactions processed by a named mobile application, or transactions that can be identified as originating from a mobile device.

Table 35 on page 199 lists the MWP defining programs supported by SCRT V23.10.00 and higher. These are the products whose sub-capacity pricing values will be reduced when the required set of MWP conditions are met.

Note: For the most current list of MWP defining programs, see [Mobile Workload Pricing \(MWP\) Defining Programs \(www.ibm.com/downloads/documents/us-en/107a02e95b48f61c\)](http://www.ibm.com/downloads/documents/us-en/107a02e95b48f61c).

Table 35. MWP defining programs

Product family	Product name	Product ID
CICS	IBM CICS Transaction Server for z/OS, V4	5655-S97
	CICS Transaction Server for z/OS, V5	5655-Y04
	CICS Value Unit Edition (VUE) V5	5722-DFJ
Db2	IBM DB2 V9 for z/OS	5635-DB2
	Db2 V10 for z/OS	5605-DB2
	Db2 V11 for z/OS	5615-DB2
	Db2 12 for z/OS	5650-DB2
	DB2 VUE V9	5697-P12
	Db2 10 VUE	5697-P31
	Db2 11 VUE	5697-P43
	Db2 12 VUE	5770-AF3
IMS	IBM IMS V11	5635-A02
	IMS V12	5635-A03
	IMS V13	5635-A04
	IMS Database VUE V12	5655-DSQ
	IMS Database VUE V13	5655-DSM

Table 35. MWP defining programs (continued)

Product family	Product name	Product ID
MQ	IBM WebSphere MQ for z/OS, V6	5655-L82
	WebSphere MQ for z/OS, V7	5655-R36
	WebSphere MQ for z/OS, V8	5655-W97
	IBM MQ for z/OS V9	5655-MQ9
	WebSphere MQ VUE V7	5655-VUE
	WebSphere MQ VUE V8	5655-VU8
	IBM MQ for z/OS VUE V9	5655-VU9
WebSphere	IBM WebSphere Application Server for z/OS, V7	5655-N02
	WebSphere Application Server for z/OS, V8	5655-W65
	IBM WebSphere Application Server for z/OS V9	5655-WAS

Mobile Workload Pricing requires that you identify mobile transactions. You can use two different mechanisms to accomplish this:

- Mobile workload reporting using an INPUTCSV file
- Mobile workload reporting using z/OS workload management (WLM) workload classification

Note: You can use both the input file method and the WLM mobile workload classification method to pass mobile transaction input data to SCRT in the same run. However, you must ensure that the data passed by each method represents data for a unique mobile workload; you cannot pass data for the same mobile workload using both input methods at the same time.

Mobile workload reporting using an INPUTCSV file

You must be able to track and report the general purpose processor time (CPU time) for mobile transactions and report those values in a pre-defined format to IBM each month in order to take advantage of this pricing offering.

Whenever the required MWP conditions have been met, SCRT reduces the sub-capacity value that is reported for all products by subtracting 60 percent of the 4-hour rolling average utilization of the reported mobile transaction general purpose processor time consumed by a set of products referred to as *MWP defining programs*, which are listed in [Table 35 on page 199](#).

For the case where a z/OS system runs as a guest under z/VM, each eligible z/OS guest system is treated as its own LPAR when applying any sub-capacity adjustments for MWP.

SCRT calculates the 4-hour rolling average of the reported mobile transaction general purpose processor time consumed by the MWP defining programs and subtracts 60 percent of those values from the traditional sub-capacity MSUs for all sub-capacity eligible programs running in the same LPAR(s) as the mobile workloads, on an hour-by-hour basis, per LPAR. The program values for the same hour are summed across all of the LPARs in which the program runs to create an adjusted sub-capacity value for the program, for the given machine, for each hour. SCRT determines the billable MSU peak for a given program on a machine using the adjusted MSU values.

Important: In order to calculate the correct 4-hour rolling average for the first hour in the report period, the INPUTCSV data must contain entries for the 3 hours prior to the start of the report period. For instance, for a report period that begins at 00:00 on 2 August, the INPUTCSV data must contain mobile transaction data beginning at 21:00 on 1 August.

Mobile workload reporting using WLM workload classification

SCRT supports the use of WLM mobile workload classification.

The WLM mobile workload classification support allows you to specify within the WLM service definition those transactions that are eligible for Mobile Workload Pricing. Workloads classified as mobile must exploit products in the set of products referred to as MWP *defining programs*, which are listed in [Table 35 on page 199](#).

For more information about the WLM support for mobile workload classification, see [z/OS MVS Planning: Workload Management](#).

In order to use this offering, you must have the WLM support applied on your z/OS systems and be able to classify mobile transactions. Once implemented and activated on a z/OS image, WLM and RMF report the 4-hour rolling average of mobile CPU consumption for that image in the SMF70LACM field of SMF type 70 records. Because the data is reported in the same SMF type 70 records that SCRT already requires, you need not perform any additional data tracking or collection.

Whenever the required MWP conditions are met, SCRT reduces the sub-capacity value that is reported for all products by subtracting 60 percent of the 4-hour rolling average utilization, as reported in the SMF70LACM field.

For the case where a z/OS system runs as a guest under z/VM, each eligible z/OS guest system is treated as its own LPAR when applying any sub-capacity adjustments for MWP.

SCRT subtracts 60 percent of the SMF70LACM value from the traditional sub-capacity MSUs for all sub-capacity eligible programs running in the same LPARs as the mobile workloads, on an hour-by-hour basis, per LPAR. The program values for the same hour are summed across all of the LPARs in which the program runs to create an adjusted sub-capacity value for the program, for the given CPC, for each hour.

SCRT uses the adjusted MSU values to determine the billable MSU peak for a given program on the CPC.

Qualifications for Mobile Workload Pricing

The following conditions must be met before IBM will accept and process sub-capacity reports where Mobile Workload Pricing is being reported via footnote (d). Some of these conditions are also described in more detail elsewhere in this chapter but are consolidated here for ease of reference.

1. The CPC from which the SMF data is collected must be licensed under AWLC, AEWLC, or zNALC sub-capacity terms and conditions. The MWP terms and conditions require that defining products be installed on an IBM z13, z13s, zEC12, or zBC12 CPC that has implemented sub-capacity pricing. Alternatively, you may qualify for MWP by installing authorized MWP defining programs on an IBM z196 or z114 CPC if a z13, z13s, zEC12, or zBC12 CPC is being operated in your enterprise. Note that the SMF data must be collected from one or more z/OS systems running on zEnterprise or later CPCs; SCRT only applies the MWP algorithm when processing SMF data collected from zEnterprise or later CPCs.
2. The traditional sub-capacity values of programs are only adjusted when one or more MWP defining programs (for which data was submitted via the INPUTCSV DD statement) were running in the same z/OS system (LPAR) during the same 4-hour window that was used to determine the traditional sub-capacity value of the MWP defining program. Programs running in one z/OS system (LPAR) are not affected by MWP defining programs running during the same hour in a different z/OS system (LPAR).
3. Only SMF data collected on 2 June 2014 and later is eligible for Mobile Workload Pricing, as set forth in the Mobile Workload Pricing announcement letter. Although SCRT will generate sub-capacity reports based on the MWP algorithm for SMF data collected earlier than 2 June 2014, the IBM business process will reject those reports.
4. You must generate and submit to IBM a monthly sub-capacity report that contains MWP adjustments during the normal sub-capacity submission window (business as usual).
5. You must use SCRT V23.10.00 or later to use the INPUTCSV DD statement to generate a sub-capacity report that is submitted for Mobile Workload Pricing. You must use SCRT V23.13.0 or later to use WLM mobile workload classification to generate a sub-capacity report that is submitted for Mobile Workload Pricing.
6. When you intend to provide Mobile Workload Pricing data for a mobile defining product to SCRT in CSV format, you must use the INPUTCSV DD statement to request that SCRT generate a sub-capacity

report using the MWP algorithms whenever all of the required MWP conditions are met. The default SCRT processing does not use the MWP algorithms. For details, see [“INPUTCSV DD statement for Mobile Workload Pricing”](#) on page 203.

When you provide Mobile Workload Pricing data to SCRT via WLM mobile workload classification, SCRT always uses the MWP algorithms, unless you specify the **IGNORE_WLM_MWP_Data** control statement in the SPECIAL DD.

7. You must agree to the IBM Mobile Workload Pricing Addendum prior to submitting a sub-capacity report where footnote (d) appears in the Product Information Summary section, since footnote (d) indicates that the sub-capacity values for one or more products were adjusted based on MWP. Failure to formally agree to the MWP Addendum prior to submitting a sub-capacity report with footnote (d) will cause the IBM business process to reject that report. The rejection will contain an explanation and request that you generate a replacement sub-capacity report without using the INPUTCSV DD statement to submit MWP data, and submit it to IBM in order to continue to receive traditional sub-capacity pricing for the IBM products on that CPC.

Data collection requirements for Mobile Workload Pricing

The data collection requirements for Mobile Workload Pricing are:

- You must collect SMF type 70 and type 89 records for all LPARs on all CPCs each month, and you must retain these records for six months after the billing period for auditing purposes.

You must also collect data about mobile transactions using the following mechanism:

- Manual mobile transaction reporting via CSV file input:
 - You must collect the source data (CPU time) for the mobile transactions that will be used in monthly reporting, and you must retain this data (CPU time and CSV files) for six months after the billing period for auditing purposes.
 - You must preprocess your mobile transaction data into a predefined format to be loaded into SCRT for each sub-capacity reporting period. The data must consist of general purpose processor CPU seconds for mobile transactions, summarized by hour, by LPAR, for all CPCs that process mobile transactions.
- Mobile transaction reporting via WLM mobile workload classification:
 - You must classify mobile workloads as described in [z/OS MVS Planning: Workload Management](#).
 - Workloads that are classified as mobile in the WLM service class do not require any manual aggregation or preprocessing of data.

MWP monthly reporting process

When using Mobile Workload Pricing, you must complete the following requirements each month:

- For programs that have mobile transactions tracked manually and provided to SCRT in CSV format:
 1. Track mobile transactions, including CPU seconds, per program, on an hourly basis, per LPAR.
 2. Produce a file in the comma-separated value (CSV) format that contains the mobile CPU consumption data for each program. The format of the data in this file is described in [“Format of the INPUTCSV input data for Mobile Workload Pricing”](#) on page 204.
 3. Load the resulting data files along with the SMF type 70 and type 89 records into SCRT.
 4. Run SCRT and submit the resulting sub-capacity reports to IBM for each reporting period.
- For programs that are classified in WLM as a mobile workload:
 1. Load the SMF type 70 and type 89 records into SCRT.
 2. Run SCRT and submit the resulting sub-capacity reports to IBM for each reporting period.

Notes:

- This monthly process requires an initial setup action in WLM to classify one or more workloads as mobile workloads.
- Note that no explicit action is required for SCRT to apply the Mobile Workload Pricing algorithm for workloads that are classified through WLM. If SCRT detects that a WLM classification-based mobile workload was active during the reporting period, its default processing is to apply the MWP algorithm.

SCRT externals for Mobile Workload Pricing

The SCRT externals for Mobile Workload Pricing differ depending on whether you are supplying mobile workload data using the INPUTCSV DD or via WLM workload classification.

SCRT externals for INPUTCSV-based Mobile Workload Pricing

SCRT provides the following externals for Mobile Workload Pricing using input data in CSV format. Where noted, these externals are specific to MWP and are treated as optional when MWP is not being requested or does not apply to the SMF data being processed by SCRT.

- SCRT uses the optional DD statement, INPUTCSV, which is used to provide input CPU consumption data in a predefined, CSV format for mobile workload pricing.
- SCRT uses footnote (d) in the Product Summary Information section of the sub-capacity report to indicate that the Tool MSUs value for the corresponding product was adjusted by the MWP algorithm.
- SCRT uses an optional **Generate_Detailed_Data** control statement to track the application of mobile and IBM Z Collocated Application Pricing (zCAP) reductions per hour, per LPAR.

INPUTCSV DD statement for Mobile Workload Pricing

The default SCRT processing does not support Mobile Workload Pricing. You must provide your mobile transaction data as input on the INPUTCSV DD statement. The data in the data set pointed to by the INPUTCSV DD statement must be in a predefined format, as described in [“Format of the INPUTCSV input data for Mobile Workload Pricing”](#) on page 204.

Note: IBM authorizes the use of the INPUTCSV DD statement for providing mobile CPU consumption data to SCRT only when you have signed or otherwise agreed to the terms of the Mobile Workload Pricing Addendum.

For more information about this statement, see [“INPUTCSV DD statement”](#) on page 60.

Mobile Workload Pricing footnote

For the case where the Tool MSUs value for one or more products listed in the Product Summary Information section have been adjusted for MWP, SCRT applies the following changes to the Product Summary Information Section:

- Generates an optional **Footnotes** column header for the MLC section.
- Posts footnote (d) in the **Footnotes** column for each MWP defining product for which mobile workload pricing data was supplied via the INPUTCSV DD statement.
- Generates a footnote subsection immediately after the last IPLA product is listed. The footnote subsection contains a footnote header row followed by a description of footnote (d).

Footnote (d) is optional. All three footnote-related items will only be generated if one or more MWP defining programs requires such a footnote.

Generate_Detailed_Data control statement

SCRT supports an optional **Generate_Detailed_Data** SPECIAL control statement. This is used to track the application of mobile and z Systems Collocated Application Pricing (zCAP) reductions per hour, per LPAR.

When specified, this control statement causes the Mobile/zCAP Pricing Detailed Data section (V5) to be generated in the Sub-capacity report.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

Mobile/zCAP Pricing Detailed Data section

The Mobile/zCAP Pricing Detailed Data section (V5) of the sub-capacity report shows, for each product, the hour-by-hour application of mobile and zCAP MSU reductions.

See [“Sample sub-capacity report for INPUTCSV-based Mobile Workload Pricing” on page 209](#) for a detailed description of the output in the Mobile/zCAP Pricing Detailed Data section.

Format of the INPUTCSV input data for Mobile Workload Pricing

For Mobile Workload Pricing, the INPUTCSV DD statement must include the data sets that contain data formatted as described in this topic.

The data contained in each Mobile Workload Pricing input file must consist of general purpose processor CPU seconds for mobile transactions. These are summarized by hour, by LPAR, for all CPCs that process mobile transactions.

Each MWP defining program product family (CICS, Db2, MQ, IMS, and WebSphere) must be contained in a separate CSV data set or data set member.

Note the following points about the INPUTCSV data set:

- Any completely blank row in the file is ignored.
- Any row that starts with a double quotation mark immediately followed by an asterisk ("*) is considered a comment and is ignored.
- The INPUTCSV data uses LPAR names, *not* system IDs (SYSIDs).

Figure 80 on page 204 shows an example of the format for the contents of the mobile transaction CSV file. The field separator character may be a comma (,) or a semicolon (;), but must be consistent throughout.

```
"Product", "5655-S97", "5655-M15"
""
"Date", "2827-0F4D7:SYSF", "2827-0F4D7:SYSG", "2827-0F4D7:SYSH"
"2014-03-02T00:00:00", 3026.23, 4103.88, 229.41
"2014-03-02T01:00:00", 1954.01, 3582.93, 257.33
"2014-03-02T02:00:00", 2444.22, 3876.60, 262.84
```

Figure 80. Example of the contents of a file in the INPUTCSV DD

Note: The double quotation marks shown in [Figure 80 on page 204](#) are part of the data and are required. You must use the keyboard double quotation mark character (EBCDIC X'7F' or ASCII X'22'), **not** typographical opening and closing double quotation marks.

The contents of the INPUTCSV data set consist of a product row (line), a header row, and one or more data rows.

Product row: The first row in the file is the product row; it must contain the "Product" keyword as the first token.

Following the "Product" keyword, you must list the product ID (or IDs) for all versions of a particular product family (for instance, the product IDs associated with the CICS product family). If you are using multiple versions of a particular product, you must specify multiple product IDs. (A typical reason for this would be if you are migrating from one version of the product to another.) Each product ID must

be enclosed in double quotation marks. The example in [Figure 80 on page 204](#) shows two products: "5655-S97" and "5655-M15".

Header row: The first non-blank, non-comment row after the product row is the header row.

The first column of the header row must contain the date header information in one of the following formats:

- The "Date" token, which indicates that all times listed are local in ISO-8601 format.
- The "DateUTC" token, which indicates that all times listed are in UTC (ISO-8601 format).

Note: The use of Hypervisor offsets (either LPAR Sysplex Timer offset or z/VM virtual clock, VTOD) requires that you input the CSV time as though these values are already applied to the data. SCRT applies Hypervisor offset values as the SMF records are read and expects the CSV time values to reflect the Hypervisor offset. Normal z/OS time zone offsets do not require this adjustment.

For example, if the Hypervisor offset is used to run the system one month in the future, the CSV file should have that one month in the future date, and not the current day's date and times.

The subsequent columns of the header row must contain the LPAR names in one of the following formats:

- The LPAR name (*not* the system ID), enclosed in double quotation marks

Example: "SYSF"

- Both the type-serial and LPAR name (*not* the system ID), separated by a colon, in the format *type-serial:LPARname*, enclosed in double quotation marks.

If you have (or plan to have) duplicate LPAR names on different CPCs, you must use this format.

Example: "2827-0F4D7:SYSF"

Data rows: Any non-blank, non-comment rows after the header row must contain the mobile transaction data.

The first column of a mobile transaction data row must contain the date. The date must be in ISO-8601 format and enclosed in double quotation marks.

Example: A date value of "2014-03-21T03:00:00" indicates 3 a.m. on 21 March 2014.

The subsequent cell values of the mobile transaction data row must contain the CPU seconds for the hour specified in the first column of that row. The CPU second values must be placed in the column corresponding to the LPAR on which they were gathered. The CPU seconds are specified as a floating-point number and must not be enclosed in double quotation marks.

For more information about the INPUTCSV DD statement, see ["INPUTCSV DD statement" on page 60](#).

SCRT externals for WLM classification-based Mobile Workload Pricing

SCRT provides the following externals for Mobile Workload Pricing using WLM classification. Where noted, these externals are specific to MWP and are treated as optional when MWP does not apply to the SMF data being processed by SCRT.

- SCRT uses the MWP_Input_From_WLM message in report section D5 to indicate when mobile workload data from WLM has been detected.
- SCRT uses footnote (d) in the Product Summary Information section of the sub-capacity report to indicate the mobile defining program.
- SCRT uses an optional **IGNORE_WLM_MWP_Data** control statement to indicate that the Mobile Workload Pricing algorithm should not be applied to mobile workload data from WLM.
- SCRT uses an optional **Generate_Detailed_Data** control statement to track the application of mobile and z Systems Collocated Application Pricing (zCAP) reductions per hour, per LPAR.
- SCRT uses an optional **DIAGMSG=WLMMOB** control statement to provide additional diagnostic messaging to help determine why an expected WLM mobile adjustment does not appear on the sub-capacity report.

MWP_Input_From_WLM message

For the case where the **Tool MSUs** value for one or more products listed in the Product Summary Information section have been adjusted for MWP based on input from a WLM classification, SCRT generates an optional D5 report section. The MWP_Input_From_WLM message appears in section D5, along with any other messages applicable to SCRT execution.

Mobile Workload Pricing footnote

For the case where the Tool MSUs value for one or more products listed in the Product Summary Information section have been adjusted for MWP, SCRT applies the following changes to the Product Summary Information Section:

- Generates an optional **Footnotes** column header for the MLC section.
- Posts footnote (d) in the **Footnotes** column for the highest z/OS operating system version.
- Generates a footnote subsection immediately after the last IPLA product is listed. The footnote subsection contains a footnote header row followed by a description of footnote (d).

Footnote (d) is optional. All three footnote-related items will only be generated if the SMF type 70 data contains a valid, non-zero mobile 4-hour rolling average value for at least one interval in the reporting period.

IGNORE_WLM_MWP_DATA control statement

SCRT supports an optional **IGNORE_WLM_MWP_DATA** control statement in the SPECIAL DD.

SCRT applies the Mobile Workload Pricing algorithm as a default when at least one SMF type 70 interval record in the reporting period contains a non-zero value. You can use this SPECIAL control statement to disable this default processing when it is not desired.

Guideline: You must have an *IBM Addendum for System z Mobile Workload Pricing* agreement in place to submit a sub-capacity report with a Mobile Workload Pricing adjustment to IBM. If you have classified one or more workloads as mobile within WLM but do not have an MWP agreement in place, you must specify the **IGNORE_WLM_MWP_DATA** control statement in the SPECIAL DD before running SCRT to generate sub-capacity reports intended for submission to IBM.

Generate_Detailed_Data control statement

SCRT supports an optional **Generate_Detailed_Data** SPECIAL control statement. This is used to track the application of mobile and z Systems Collocated Application Pricing (zCAP) reductions per hour, per LPAR.

When specified, this control statement causes the Mobile/zCAP Pricing Detailed Data section (V5) to be generated in the Sub-capacity report.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

DIAGMSG=WLM MOB control statement

SCRT supports an optional **DIAGMSG=WLM MOB** SPECIAL control statement. This control statement is used to provide additional diagnostic messaging to help determine why an expected WLM mobile adjustment does not appear in the sub-capacity report.

When specified, this control statement causes SCRT to generate the SCRTTOOL610 or SCRTTOOL611 messages, when appropriate.

The SCRTTOOL610 message indicates that the specified LPAR does not appear to have the required z/OS maintenance level to support the new SMF70LACM field in the SMF type 70 record.

The SCRTTOOL611 message indicates that the specified LPAR did not report a non-zero value during the reporting period for the SMF field containing the mobile workload 4-hour rolling average CPU consumption.

Mobile/zCAP Pricing Detailed Data section

The Mobile/zCAP Pricing Detailed Data section (V5) of the sub-capacity report shows, for each product, the hour-by-hour application of mobile and zCAP MSU adjustments.

See [“Sample sub-capacity report for WLM classification-based Mobile Workload Pricing”](#) on page 219 for a detailed description of the output in the Mobile/zCAP Pricing Detailed Data section.

Reporting sample for INPUTCSV-based Mobile Workload Pricing

This topic provides a detailed example of sub-capacity reporting for INPUTCSV-based Mobile Workload Pricing.

Sample configuration for INPUTCSV-based Mobile Workload Pricing

Table 36 on page 207 illustrates the sample hardware and software configuration for a CPC where an input CSV file was provided for two MWP defining programs that were running during the reporting period. While the mix of software products and versions used in this sample configuration might not be realistic, it serves to illustrate some of the finer points of the MWP algorithm.

The sample configuration consists of a 2964 Model 796 CPC with a serial number of 02-12345. The CPC is configured with two LPARs, both running z/OS natively. The z/OS systems in LPAR1 and LPAR2 have z/OS SYSIDs of SYS1 and SYS2 respectively.

The z/OS systems are running a combination of the following z/OS products:

- Db2 11 for z/OS
- CICS TS for z/OS V5
- IBM MQ for z/OS V8
- IMS V13
- IBM Enterprise COBOL for z/OS V4
- IBM WebSphere Application Server for z/OS V8
- IBM Multi-site Workload Lifeline V2

CPC: 2964-796 (serial = 02-12345)	
LPAR1	LPAR2
z/OS V2.1 (SYSID = SYS1) Db2 11 for z/OS CICS TS for z/OS V5 IBM MQ for z/OS V8 IMS V13 IBM Enterprise COBOL for z/OS V4 IBM WebSphere Application Server for z/OS V8	z/OS V2.1 (SYSID = SYS2) Db2 11 for z/OS CICS TS for z/OS V5 IBM MQ for z/OS V8 IMS V13 IBM Multi-site Workload Lifeline V2

Sample control statement input to SCRT for INPUTCSV-based Mobile Workload Pricing

SCRT supports additional externals for MWP. The sample sub-capacity report discussed in [“Sample sub-capacity report for INPUTCSV-based Mobile Workload Pricing”](#) on page 209 is based on the following SCRT control statement and JCL inputs to SCRT:

INPUTCSV DD statement

The following SCRT DD statement was used to specify data sets containing mobile workload pricing data for the reporting period for the LPARs and SYSIDs described in [“Sample configuration for INPUTCSV-based Mobile Workload Pricing”](#) on page 207:

```
//INPUTCSV DD    DISP=SHR,DSN=MY.DSN.ID(CICSMWP)
//              DD    DISP=SHR,DSN=MY.DSN.ID(IMSMWP)
```

Specifying the INPUTCSV statement and supplying CPU consumption data for mobile transactions in the appropriate format in the INPUTCSV data sets is the method to request that SCRT use the MWP algorithm whenever all of the MWP conditions have been met. See [“INPUTCSV DD statement for Mobile Workload Pricing”](#) on page 203 for a detailed description and restrictions on its use.

In the sample configuration, the following mobile consumption data is stored in MY.DSN.ID(CICSMWP), which is the first data set referenced on the sample INPUTCSV DD statement:

```
"Product", "5655-Y04"
"
"Date", "2964-12345:SYS1"
"2015-05-04T14:00:00", 2000.00
"2015-05-04T15:00:00", 1000.00
"2015-05-04T16:00:00", 1500.00
"2015-05-04T17:00:00", 1250.00
"2015-05-04T18:00:00", 1150.00
```

The CICSMWP member contains mobile workload pricing CPU consumption data for CICS TS for z/OS V5 (5655-Y04), in the format required for MWP INPUTCSV data.

In the sample configuration, the following mobile consumption data is stored in MY.DSN.ID(IMSMWP), which is the second data set referenced on the sample INPUTCSV DD statement:

```
"Product", "5635-A04"
"
"Date", "2964-12345:SYS2"
"2015-05-04T14:00:00", 350.00
"2015-05-04T15:00:00", 750.00
"2015-05-04T16:00:00", 500.00
"2015-05-04T17:00:00", 750.00
```

The IMSMWP member contains mobile workload pricing CPU consumption data for IMS V13 (5635-A04), in the format required for MWP INPUTCSV data.

For more information about the format of the INPUTCSV data set, see [“Format of the INPUTCSV input data for Mobile Workload Pricing”](#) on page 204.

SPECIAL DD control statement

The following SCRT control statement was specified in the SPECIAL DD to request that SCRT generate the Mobile/zCAP Detailed Data section of the sub-capacity report. This control statement has no effect on whether SCRT uses the MWP algorithm to apply reductions based on the data supplied via the INPUTCSV DD statement.

```
Generate_Detailed_Data
```

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being

generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

Sample sub-capacity report for INPUTCSV-based Mobile Workload Pricing

The sub-capacity report for Mobile Workload Pricing has the same basic format and report sections as the sub-capacity report for a base z/OS system. However, the sub-capacity report for MWP always contains a footnote to indicate the case when mobile CPU consumption data was provided for MWP defining programs.

The following sample report sections describe the differences between a sub-capacity report for a base z/OS system and a sub-capacity report that uses Mobile Workload Pricing. For a complete description of all of the fields appearing in a base sub-capacity report see [“Sample sub-capacity report for single-version z/OS native systems”](#) on page 105.

Billing-related section

The billing-related section of the sub-capacity report displays the customer information, tool information, special conditions, product summary information, and detail data collection information, which are described in [“Billing-related section”](#) on page 105.

Figure 81 on page 209 shows an example of the Customer Information section and the Tool Information section.

```

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time                24 Jun 2024 - 15:19
Name of Person Submitting Report:    John Customer
E-mail Address of Report Submitter:  customer@abc.com
Phone Number of Report Submitter:    800-555-0123

Customer Name                ABC Corporation
Customer Number              US-S000000000
Machine Serial Number        02-12345
Machine Type and Model       2964-796
Machine Rated Capacity (MSUs) 9783
Purchase Order Number        (optional)
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust
the billable MSUs in inventory for all MLC Products listed under the MLC Product Name
column on this report. In accordance with our agreement, IBM will treat a change in product
licensed capacity as an order. If the MSUs have changed since the last report, software billing
based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to
bill those IPLA products listed under the IPLA Product Name column in this report which
exceed your entitled capacity. In accordance with our agreement, IBM will treat the use
of a product in excess of its entitled capacity as an order and you will be billed for
the amount in excess of your entitlement.

==C5=====
TOOL INFORMATION

Tool Release                30.1.0
Reporting Period            2 May, 2024 - 1 Jun, 2024 inclusive (31 days)

```

Figure 81. Sample sub-capacity report for Mobile Workload Pricing: Customer information and Tool Information sections

Customer Information section

There are no differences between the Customer Information section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Customer information” on page 105](#).

Tool Information section

There are no differences between the Tool Information section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Tool Information section” on page 107](#).

Special Conditions section

There are no differences between the Special Conditions section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP.

Product Summary Information section

An optional Footnotes column header appears in the MLC section whenever SCRT needs to post a footnote for an MWP defining program listed in this section of the report. The Footnotes heading does not appear whenever there are no footnotes for this section.

There are no differences between the MLC Product Name, MLC Product ID, Customer Comment, IPLA Product Name, IPLA Product ID, and Customer MSUs fields on a basic z/OS sub-capacity report and on a sub-capacity report generated using MWP. For a description of these fields see [“Product Summary Information section” on page 110](#).

The sample Product Summary Information section appears in [Figure 82 on page 211](#). The following fields differ for MWP reporting:

Tool MSUs

The values listed in the Tool MSUs column represent the peak sub-capacity values calculated by SCRT for each product using the appropriate sub-capacity rules for that product. SCRT uses the MWP algorithm and the reported mobile transaction CPU data to adjust the 4-hour rolling average sub-capacity MSU for sub-capacity eligible programs on a given CPC.

See [“Product Max Contributors section” on page 212](#) for a more detailed discussion of how the MWP algorithm affects the sub-capacity programs that are listed in this section of the sub-capacity report. For additional general information about Tool MSUs, see [“Product Summary Information section” on page 110](#).

Footnotes

SCRT posts footnote (d) for each MWP defining program for which mobile workload pricing data was supplied via the INPUTCSV DD statement.

The Tool MSUs values are partially based on the MWP algorithm whenever the posted value for the product is combined from multiple LPARs where the product's value is based on the traditional sub-capacity values from some LPARs and on the MWP adjusted values from other LPARs.

Whenever SCRT posts footnote (d) in the Footnotes column, it also generates a row with a footnote header followed by a description for the footnote.

The Product Summary Information section in [Figure 82 on page 211](#) uses footnote (d) to indicate that data was specified on the INPUTCSV DD statement for MWP defining programs CICS TS V5 and IMS V13.

SCRT also posts footnote (iv) whenever the Tool MSUs value posted for a product is based totally or partially on having been adjusted by the Getting Started Sub-Capacity Pricing algorithm. See [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#) for a detailed description of Getting Started Sub-Capacity Pricing.


```

=====
PRODUCT SUMMARY INFORMATION
MLC Product Name          MLC Product ID    Tool MSUs    Customer MSUs    Customer Comments (255 chars max)    Footnotes
z/OS V2                   5650-Z0S           3514
DB2 11 for z/OS           5615-DB2           3514
CICS TS for z/OS V5       5655-Y04           3514
IBM MQ for z/OS V8        5655-W97           3514
IMS V13                   5635-A04           3514
IBM Enterprise Cobol for z/OS V4  5655-S71           3514    (optional)    (conditional)
IPLA Product Name        IPLA Product ID    Tool MSUs    Customer MSUs    Customer Comments (255 chars max)    Footnotes
File Export for z/OS V1    5697-I12           3514    (optional)    (conditional)
IBM WebSphere Application Server for z/OS V8  5655-W65           21
Migration Utility for z/OS V4  5655-MGU           3514    (optional)    (conditional)
IBM Multi-site Workload Lifeline V2  5655-UM4           2693
Footnotes:
(iv) Product's Tool MSUs are based on Getting Started Sub-capacity Pricing from one or more LPARs
(d) Mobile Workload Pricing transaction data from this product was included from one or more LPARs

```

Figure 82. Sample sub-capacity report for Mobile Workload Pricing: Product Summary Information section

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the Product Summary Information section.

SMF / SCRT89 Input Data Statistics section

There are no differences between the SMF / SCRT89 Input Data Statistics section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“SMF / SCRT89 Input Data Statistics section” on page 118](#).

The sample SMF / SCRT89 Input Data Statistics section appears in [Figure 83 on page 211](#).

```

=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
=====
SMF/SCRT89 INPUT DATA STATISTICS
SYSID    Input Data Start    Input Data End
LPAR1     SYS1    04 May 2024 - 12:59  04 May 2024 - 17:59
LPAR2     SYS2    04 May 2024 - 12:59  04 May 2024 - 17:59
CPC              04 May 2024 - 12:59  04 May 2024 - 17:59

```

Figure 83. Sample sub-capacity report for Mobile Workload Pricing: SMF / SCRT89 Input Data Statistics section

Detail LPAR Data Section

There are no differences between the Detail LPAR Data Section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Detail LPAR Data section” on page 119](#).

The sample Detail LPAR Data Section appears in [Figure 84 on page 212](#).

```

==N5=====
DETAIL LPAR DATA SECTION

              Highest Hour Count Date/Time              2nd Highest Hour Count Date/Time
LPAR1          2701           1 04 May 2024 - 14:00      2244           1 04 May 2024 - 15:00
LPAR2           822           1 04 May 2024 - 14:00      803           1 04 May 2024 - 15:00
CPC            3523           1 04 May 2024 - 14:00      3047           1 04 May 2024 - 15:00

```

Figure 84. Sample sub-capacity report for Mobile Workload Pricing: Detail LPAR Data section

Product Max Contributors section

The MSU values that are listed in the Product Max Contributors section are the same for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report, with one subtle change. While all of the MSU values listed in this section continue to reflect the product MSUs by LPAR or LPAR Group for the earliest hour for the product's maximum MSU value, the MSU values posted for Getting Started Sub-Capacity Pricing eligible products are based on the Getting Started Sub-Capacity Pricing algorithm whenever that value is less than the traditional sub-capacity value for that product for that hour. The result is that some of the LPAR values posted for Getting Started Sub-Capacity Pricing eligible products are less than the MSUs posted for z/OS, even when the Getting Started Sub-Capacity Pricing eligible product runs in all of the same LPARs and at the same times as z/OS.

Note: The MSUs for Getting Started Sub-Capacity Pricing eligible products are posted under an LPAR group whenever the LPAR group cap is less than the Getting Started Sub-Capacity Pricing value for the earliest hour with the maximum product value. That is not the case for this sample configuration.

This section provides an additional layer of detail behind the sub-capacity values reported for the products listed in the Product Summary Information section by reporting the date/time that the sub-capacity value occurred together with the product's sub-capacity contribution from each LPAR in which that product was running at that date and time. When the *same* peak sub-capacity value occurs for multiple hours, the Product Max Contributors section lists the data for the *earliest* hour. The impact of SCRT reporting a product's sub-capacity value based on the earliest peak hour whenever that peak value occurs for multiple hours is important enough to reiterate in the following explanations because the MWP conditions for the other peak hours can differ from the MWP conditions for the earliest peak hour.

Given the level of detail in this section, the sample Product Max Contributors section shown in Figure 85 on page 212 will be used to explain Mobile Workload Pricing in more detail. However, this explanation assumes you have a general understanding of the contents and use of each field in the Product Max Contributors section, as described in "Product Max Contributors section" on page 121.

```

==P5=====
PRODUCT MAX CONTRIBUTORS

Product Name              Product ID Highest Date/Time              LPAR LPAR Mobile MSU
                           SYS1  SYS2  Reduction
z/OS V2                   5650-Z0S   3514 04 May 2024 - 14:00 2693 821      9
DB2 11 for z/OS           5615-DB2   3514 04 May 2024 - 14:00 2693 821      9
CICS TS for z/OS V5       5655-Y04   3514 04 May 2024 - 14:00 2693 821      9
IBM MQ for z/OS V8        5655-W97   3514 04 May 2024 - 14:00 2693 821      9
IMS V13                   5635-A04   3514 04 May 2024 - 14:00 2693 821      9
IBM Enterprise Cobol for z/OS V4 5655-S71   3514 04 May 2024 - 14:00 2693 821      9
File Export for z/OS V1   5697-I12   3514 04 May 2024 - 14:00 2693 821      9
IBM WebSphere Application Server for z/OS V8 5655-W65   21 04 May 2024 - 17:00 21      0
Migration Utility for z/OS V4 5655-MGU   3514 04 May 2024 - 14:00 2693 821      9
IBM Multi-site Workload Lifeline V2 5655-UM4   2693 04 May 2024 - 14:00 2693 0        8

```

Figure 85. Sample sub-capacity report for Mobile Workload Pricing: Product Max Contributors section

How SCRT implements Mobile Workload Pricing

It is important to understand that SCRT uses the appropriate algorithm to calculate the sub-capacity value for each product for each hour for each LPAR, independently of what is happening with that product on other LPARs or guest systems during that same hour. Then, SCRT sums each of the product's sub-capacity values across all of the LPARs for the same hour to calculate a sub-capacity value for the product for the CPC for that hour. After performing these sub-capacity calculations for the product across the CPC for all of the hours during the reporting period, SCRT selects the earliest hour across the CPC with the product's peak sub-capacity value and reports the date and time for that hour, along with the individual sub-capacity values for the product for each LPAR for that same hour, in the Product Max Contributors section.

For products like z/OS, whose sub-capacity is entirely based on the 4-hour rolling average MSU utilization of the LPAR, the 4-hour rolling average MSU utilization is always reported as the product's sub-capacity value. In the sample Product Max Contributors section, the peak z/OS sub-capacity value across the CPC was 3514 MSUs and that peak first occurred on 04 May 2024 at 14:00. The z/OS product running in LPAR1 and LPAR2 contributed 2693 and 821 MSUs, respectively, to the z/OS peak sub-capacity value of 3514 MSUs for the CPC.

One key difference between a basic z/OS sub-capacity report and a sub-capacity report generated using MWP is the **Mobile MSU Reduction** column in the Product Max Contributors section. This column is generated whenever mobile pricing data was provided via the INPUTCSV DD statement, as was the case in the configuration (see [“Sample configuration for INPUTCSV-based Mobile Workload Pricing”](#) on page 207).

The **Mobile MSU Reduction** column shows the adjustment that SCRT applied to products' sub-capacity values using the MWP algorithm. In the Product Max Contributors section for the sample configuration, the peak z/OS sub-capacity value and the value of all products based entirely on the 4-hour rolling average MSU utilization of the LPAR were adjusted by SCRT using the value shown in the **Mobile MSU Reduction** column.

In the Product Max Contributors section for the sample configuration, note that the MSU values shown for LPAR SYS1 and LPAR SYS2 already reflect the contribution of MWP defining programs that ran in each LPAR during the peak hour. Also note that the **Mobile MSU Reduction** column reflects the total reduction across all LPARs for that hour.

See [“Mobile/zCAP Pricing Detailed Data section for MWP”](#) on page 214 for more information about the reporting results for the sample configuration.

A note about MWP adjustments: The MSU values used for MWP adjustments are derived from the mobile transaction general purpose processor time values reported using the INPUTCSV DD statement. These adjustments are calculated internally by SCRT as specified in the IBM Addendum for System z Mobile Workload Pricing. The Product Max Contributors section is only intended to show the contribution of each LPAR to the specified product's peak sub-capacity value for earliest peak hour.

Product Grid Snapshot section

There are no differences between the Product Grid Snapshot section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Product Grid Snapshot section”](#) on page 123.

The sample Product Grid Snapshot section appears in [Figure 86 on page 214](#).

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                                Product ID    SYS1    SYS2
z/OS V2                                    5650-Z05      0.5%    0.5%
DB2 11 for z/OS                           5615-DB2      0.5%    0.5%
CICS TS for z/OS V5                       5655-Y04      0.5%    0.5%
IBM MQ for z/OS V8                        5655-W97      0.5%    0.5%
IMS V13                                   5635-A04      0.5%    0.5%
IBM WebSphere Application Server for z/OS V8 5655-W65      0.5%    0.5%
IBM Multi-site Workload Lifeline V2       5655-UM4      0.5%    0.5%

```

Figure 86. Sample sub-capacity report for Integrated Workload Pricing: Product Grid Snapshot section

Getting Started Sub-Capacity Pricing Assigned Hours section

There are no differences between the Getting Started Sub-Capacity Pricing Assigned Hours section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see “Getting Started Sub-Capacity Pricing Assigned section” on page 179.

The sample Getting Started Sub-Capacity Pricing Assigned Hours section appears in [Figure 87 on page 214](#).

```

==S7=====
GETTING STARTED SUB-CAPACITY PRICING ASSIGNED - HOURS

Product Name                                Product ID    LPAR    LPAR
                                           SYS1      SYS2
IBM WebSphere Application Server for z/OS V8 5655-W65      4

```

Figure 87. Sample sub-capacity report for Mobile Workload Pricing: Getting Started Sub-Capacity Pricing Assigned Hours section

Mobile/zCAP Pricing Detailed Data section for MWP

The optional Mobile/zCAP Pricing Detailed Data section appears when the optional **Generate_Detailed_Data** SPECIAL DD control statement is specified. This section shows the mobile MSU reduction applied by SCRT for each MWP defining program, per LPAR. This information is shown for each sub-capacity product on an hour-by-hour basis.

As shown in [Figure 88 on page 216](#), the optional Mobile/zCAP Pricing Detailed Data section reports several columns for each product:

- Date and time of the adjusted interval.
- CPC Type-model.
- CPC rated machine capacity in MSU.
- Adjusted MSU values for each LPAR. (Note that the MWP algorithm has already been applied.)
- For each mobile defining product with INPUTCSV DD data specified, the MWP reduction applied per defining product, per partition.
- Total Mobile MSU reduction applied across all MWP defining products.
- CPC Total MSU, which is the sum of the adjusted LPAR MSU values across all partitions for the adjusted product.

There are several additional values that can appear in section V5 of a sub-capacity report:

–

Indicates that there is no data for this interval.

dd(rr)

Indicates that the defined capacity limit for the LPAR was reached during this interval. The first number (*dd*) represents the defined capacity value used for this hour. The second number (*rr*) represents the 4-hour rolling average value prior to the capping.

G(rr)

Indicates that the sum of the partitions exceeded the group capacity limit; the group capacity was used for this hour for that set of LPARs. The *rr* value represents the 4-hour rolling average value prior to the capping.

P(rr)

Indicates that the sum of the partitions exceeded the processor capacity; the processor capacity was used. The *rr* value represents the 4-hour rolling average value prior to the capping.

(GSSP)

Indicates that a WebSphere program eligible for Getting Started Sub-Capacity Pricing received the GSSP value so there was no additional adjustment.

nnn(CPS)

Indicates that an LPAR was assigned to a Tailored Fit Pricing solution for this hour. In this case, the LPAR does not contribute to the product sub-capacity value. The *nnn* value represents what the LPAR sub-capacity value would have been if it were part of the standard (non-container) environment.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

The following descriptions of the Mobile Workload Pricing algorithm are based on the sample configuration in [“Sample configuration for INPUTCSV-based Mobile Workload Pricing”](#) on page 207. Note that only a subset of the products running in the environment are evaluated.

Results for z/OS V2R1

After SCRT applied the appropriate algorithm hour by hour, LPAR by LPAR, and then summed the sub-capacity values for the same hour across all of the z/OS LPARs, SCRT reported 04 May 2024 at 14:00 as the earliest hour with the peak sub-capacity value for z/OS, as indicated in the sample Product Max Contributors section.

For 04 May 2024 at 14:00, SCRT applied the appropriate algorithms and calculated a sub-capacity value of 2693 MSUs for LPAR SYS1. As shown in the CICS Mobile Partition column for SYS1, this includes an 8 MSU MWP reduction derived from the CICS mobile transaction general purpose processor time that was specified as input in the sample configuration in [“Sample configuration for INPUTCSV-based Mobile Workload Pricing”](#) on page 207.

For partition SYS2, SCRT calculated a sub-capacity value of 821 MSUs, which includes a 1 MSU MWP reduction derived from the IMS mobile transaction general purpose processor time that was specified as input in the sample configuration.

[Figure 88 on page 216](#) illustrates these results.

```

==V5=====
MOBILE/ZCAP PRICING DETAILED DATA

MLC Products

5650-ZOS                z/OS V2

Date Time of RMF Interval  Type      CPC MSU  Partition  CICS Mobile  IMS Mobile  Total
                          Type      MSU      SYS1      SYS2  Partition  Partition  Mobile MSU  CPC Total
                          Type      MSU      SYS1      SYS2  Partition  Partition  Reduction
04 May 2024 - 14:00      2964-796  9783    2693    821    8          1          9          3514
04 May 2024 - 15:00      2964-796  9783    2231    798    13         5         18         3029
04 May 2024 - 16:00      2964-796  9783    1864    749    19         7         26         2613
04 May 2024 - 17:00      2964-796  9783    1679    676    24        10         34         2355

```

Figure 88. Sample sub-capacity report for Mobile Workload Pricing: Mobile/zCap Pricing Detailed Data section – z/OS results

Results for IBM MQ for z/OS V8

SCRT applied the process as described in “Results for z/OS V2R1” on page 215. Like z/OS, IBM MQ for z/OS V8 is based on the 4-hour rolling average MSU utilization of the LPAR.

Since IBM MQ for z/OS V8 ran on both z/OS LPARs during 04 May 2024 at 14:00, the application of the MWP algorithm mirrors that previously described for z/OS.

This same process applies for all sub-capacity programs that reported a peak MSU value of 3514 with 04 May 2024 at 14:00 as the earliest hour with the peak sub-capacity value.

Figure 89 on page 216 illustrates the results.

```

5655-W97                IBM MQ for z/OS V8

Date Time of RMF Interval  Type      CPC MSU  Partition  CICS Mobile  IMS Mobile  Total
                          Type      MSU      SYS1      SYS2  Partition  Partition  Mobile MSU  CPC Total
                          Type      MSU      SYS1      SYS2  Partition  Partition  Reduction
04 May 2024 - 14:00      2964-796  9783    2693    821    8          1          9          3514
04 May 2024 - 15:00      2964-796  9783    2231    798    13         5         18         3029
04 May 2024 - 16:00      2964-796  9783    1864    749    19         7         26         2613
04 May 2024 - 17:00      2964-796  9783    1679    676    24        10         34         2355

```

Figure 89. Sample sub-capacity report for Mobile Workload Pricing: Mobile/zCap Pricing Detailed Data section – IBM MQ for z/OS results

Results for IBM Multi-site Workload Lifeline V2

SCRT applied the process as described in “Results for z/OS V2R1” on page 215 and “Results for IBM MQ for z/OS V8” on page 216. Like those products, IBM Multi-site Workload Lifeline V2 is based on the 4-hour rolling average MSU utilization of the LPAR.

However, IBM Multi-site Workload Lifeline V2 is only running in partition SYS1. As such, it is not eligible for the reduction derived from the IMS mobile defining program, as mobile transaction general purpose processor time was only specified for IMS on SYS2. As a result, SCRT calculated an MSU value of 2693 for IBM Multi-site Workload Lifeline V2, as illustrated in Figure 90 on page 217.

5655-UM4		IBM Multi-site Workload Lifeline V2				
Date Time of RMF Interval	Type	CPC MSU	Partition SYS1	CICS Mobile Partition SYS1	Total Mobile MSU Reduction	CPC Total
04 May 2024 - 14:00	2964-796	9783	2693	8	8	2693
04 May 2024 - 15:00	2964-796	9783	2231	13	13	2231
04 May 2024 - 16:00	2964-796	9783	1864	19	19	1864
04 May 2024 - 17:00	2964-796	9783	1679	24	24	1679

Figure 90. Sample sub-capacity report for Mobile Workload Pricing: Mobile/zCap Pricing Detailed Data section – IBM Multi-site Workload Lifeline results

Results for IBM WebSphere Application Server for z/OS V8

SCRT applied the Getting Started Sub-Capacity Pricing for z/OS IPLA Software process to IBM WebSphere Application Server for z/OS V8.

Since the terms and conditions of GSSP were met in the sample configuration, and the GSSP values were less than the traditional sub-capacity values for each hour of the reporting period, SCRT assigned IBM WebSphere Application Server for z/OS V8 the GSSP value for each hour in the reporting period.

When the GSSP value is used for a GSSP eligible product, the Mobile/zCAP Pricing Detailed Data section places the (GSSP) token in the MWP defining program partition reduction column, as shown in Figure 91 on page 217. In accordance with the Mobile Workload Pricing addendum, the MWP reduction is not applied when GSSP pricing is used instead of the traditional sub-capacity value.

In the sample configuration, the peak MSU value used for IBM WebSphere Application Server for z/OS V8 is 21, which is achieved during the 04 May 2024 17:00 hour.

For information about the GSSP process, see [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),”](#) on page 165.

5655-W65		IBM WebSphere Application Server for z/OS V8					
Date Time of RMF Interval		Type	CPC MSU	Partition SYS1	CICS Mobile Partition SYS1	Total Mobile MSU Reduction	CPC Total
04 May 2024 - 14:00		2964-796	9783	5	(GSSP)	0	5
04 May 2024 - 15:00		2964-796	9783	11	(GSSP)	0	11
04 May 2024 - 16:00		2964-796	9783	16	(GSSP)	0	16
04 May 2024 - 17:00		2964-796	9783	21	(GSSP)	0	21

Figure 91. Sample sub-capacity report for Mobile Workload Pricing: Mobile/zCap Pricing Detailed Data section – IBM WebSphere Application Server for z/OS results

Reporting sample for WLM classification-based Mobile Workload Pricing

This topic provides a detailed example of sub-capacity reporting for WLM classification-based Mobile Workload Pricing.

Sample configuration for WLM classification-based Mobile Workload Pricing

Table 37 on page 218 illustrates the sample hardware and software configuration for a CPC where mobile workloads are indicated using WLM classification during the reporting period. Although you must define a given workload as mobile using the appropriate mechanism in WLM, SCRT is not aware of the mobile contributors under this approach.

While the mix of software products and versions used in this sample configuration might not be realistic, it serves to illustrate some of the finer points of the MWP algorithm.

The sample configuration consists of a 2817 Model 760 CPC with a serial number of 02-12345. The CPC is configured with two LPARs, both running z/OS natively. The z/OS systems in LPAR1 and LPAR2 have z/OS SYSIDs of SYS1 and SYS2, respectively.

The z/OS systems are running a combination of the following z/OS products:

- Db2 11 for z/OS
- CICS TS for z/OS V5
- IBM MQ for z/OS V8
- IMS V13
- IBM Enterprise COBOL for z/OS V4
- IBM WebSphere Application Server for z/OS V8
- File Export for z/OS V1

CPC: 2817-760 (serial = 02-12345)	
LPAR1	LPAR2
z/OS V2.1 (SYSID = SYS1) Db2 11 for z/OS CICS TS for z/OS V5 IBM MQ for z/OS V8 IMS V13 IBM Enterprise COBOL for z/OS V4 File Export for z/OS V1 IBM WebSphere Application Server for z/OS V8	z/OS V2.1 (SYSID = SYS2) Db2 11 for z/OS CICS TS for z/OS V5 IBM MQ for z/OS V8 IMS V13 IBM Enterprise COBOL for z/OS V4 File Export for z/OS V1 IBM WebSphere Application Server for z/OS V8

Sample control statement input to SCRT for WLM classification-based Mobile Workload Pricing

SCRT supports additional externals for WLM classification-based MWP. The sample sub-capacity report discussed in [“Sample sub-capacity report for WLM classification-based Mobile Workload Pricing”](#) on page 219 is based on the following SCRT control statement and JCL inputs to SCRT

Note: No input is required for SCRT to apply the MWP algorithm. When SCRT detects that a WLM classification-based mobile workload was active during the reporting period, its default processing applies the MWP algorithm.

SPECIAL DD control statement

The following SCRT control statement was specified in the SPECIAL DD to request that SCRT generate the Mobile/zCAP Detailed Data section of the sub-capacity report. This control statement has no effect on whether SCRT uses the MWP algorithm to apply reductions based on the WLM classification-based workloads found in the SMF data.

```
Generate_Detailed_Data
```

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

Sample sub-capacity report for WLM classification-based Mobile Workload Pricing

The sub-capacity report for Mobile Workload Pricing has the same basic format and report sections as the sub-capacity report for a base z/OS system. However, the sub-capacity report for MWP always contains a Special Condition message when mobile data was provided via WLM workload classification. Additionally, a mobile footnote is always present on the highest z/OS operating system level in use on the CPC to indicate the case when mobile CPU consumption data was provided via WLM.

The following sample report sections describe the differences between a sub-capacity report for a base z/OS system and a sub-capacity report that uses WLM classification-based Mobile Workload Pricing. For a complete description of all of the fields appearing in a base sub-capacity report, see [“Sample sub-capacity report for single-version z/OS native systems” on page 105](#).

Billing-related section

The billing-related section of the sub-capacity report displays the customer information, tool information, special conditions, product summary information, and detail data collection information, which are described in [“Billing-related section” on page 105](#).

Figure 92 on page 219 shows an example of the Customer Information section and the Tool Information section.

```

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time                03 Feb 2024 - 14:00
Name of Person Submitting Report: John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter: 800-555-0123

Customer Name                ABC Corp
Customer Number              US-S000000000
Machine Serial Number        02-12345
Machine Type and Model       2817-760
Machine Rated Capacity (MSUs) 5001
Purchase Order Number        (optional)
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust
the billable MSUs in inventory for all MLC Products listed under the MLC Product Name
column on this report. In accordance with our agreement, IBM will treat a change in product
licensed capacity as an order. If the MSUs have changed since the last report, software billing
based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to
bill those IPLA products listed under the IPLA Product Name column in this report which
exceed your entitled capacity. In accordance with our agreement, IBM will treat the use
of a product in excess of its entitled capacity as an order and you will be billed for
the amount in excess of your entitlement.

==C5=====
TOOL INFORMATION

Tool Release                30.1.0
Reporting Period            2 Jan, 2024 - 1 Feb, 2024 inclusive (31 days)

```

Figure 92. Sample sub-capacity report for Mobile Workload Pricing: Customer information and Tool Information sections

Customer Information section

There are no differences between the Customer Information section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Customer information” on page 105](#).

Tool Information section

There are no differences between the Tool Information section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Tool Information section” on page 107](#).

Special Conditions section

The key difference between the Special Conditions section for a basic z/OS sub-capacity report and a sub-capacity report for WLM classification-based MWP is the presence of the following message:

```
MWP_Input_From_WLM
```

This message is always generated when SCRT detects and applies valid WLM mobile workload data in the SMF input data.

In the sample configuration, Integrated Workload Pricing is requested using the **Integrated_Workload_Pricing** control statement in the SPECIAL DD. For more information about Integrated Workload Pricing, see [Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing \(IWP\),” on page 183](#).

The sample Special Conditions section appears in [Figure 93 on page 220](#).

```
==D5=====
SPECIAL CONDITIONS

IWP Control Statement
Integrated_Workload_Pricing

MWP_Input_From_WLM
```

Figure 93. Sample sub-capacity report for Mobile Workload Pricing: Special Conditions section

Product Summary Information section

An optional Footnotes column header appears in the MLC section whenever SCRT needs to post a footnote for an MWP defining program listed in this section of the report. The Footnotes heading does not appear whenever there are no footnotes for this section.

There are no differences between the MLC Product Name, MLC Product ID, Customer Comment, IPLA Product Name, IPLA Product ID, and Customer MSUs fields on a basic z/OS sub-capacity report and on a sub-capacity report generated using MWP. For a description of these fields see [“Product Summary Information section” on page 110](#).

The sample Product Summary Information section appears in [Figure 94 on page 221](#). The following fields differ for MWP reporting:

Tool MSUs

The values listed in the Tool MSUs column represent the peak sub-capacity values calculated by SCRT for each product using the appropriate sub-capacity rules for that product. SCRT uses the MWP algorithm and the reported mobile transaction CPU data to adjust the 4-hour rolling average sub-capacity MSU for sub-capacity eligible programs on a given CPC.

See [“Product Max Contributors section” on page 212](#) for a more detailed discussion of how the MWP algorithm affects the sub-capacity programs that are listed in this section of the sub-capacity report. For additional general information about Tool MSUs, see [“Product Summary Information section” on page 110](#).

Footnotes

When mobile workload data is provided via WLM classification, SCRT posts footnote (d) for the first z/OS operating system version that appears in the sub-capacity report.

The Tool MSUs values are partially based on the MWP algorithm whenever the posted value for the product is combined from multiple LPARs where the product's value is based on the traditional sub-capacity values from some LPARs and on the MWP adjusted values from other LPARs.

Whenever SCRT posts footnote (d) in the Footnotes column, it also generates a row with a footnote header followed by a description for the footnote.

The Product Summary Information section in [Figure 94 on page 221](#) uses footnote (d) to indicate that mobile data is based on WLM mobile workload classification.

SCRT also posts footnote (iv) whenever the Tool MSUs value posted for a product is based totally or partially on having been adjusted by the Getting Started Sub-Capacity Pricing algorithm. See [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#) for a detailed description of Getting Started Sub-Capacity Pricing.

SCRT also posts footnote (viii) whenever the Tool MSUs value posted for a product is adjusted by the IWP algorithm. For more information, see [Chapter 9, “Sub-capacity reporting for Integrated Workload Pricing \(IWP\),” on page 183](#).

```

=====
PRODUCT SUMMARY INFORMATION

MLC Product Name                MLC Product ID    Tool MSUs    Customer MSUs    Customer Comments (255 chars max)    Footnotes
z/OS V2                        5650-Z0S          3021
DB2 11 for z/OS                5615-DB2          3021
CICS TS for z/OS V5            5655-Y04          3001
IBM MQ for z/OS V8             5655-W97          3021
IMS V13                        5635-A04          3001
IBM Enterprise Cobol for z/OS V4 5655-S71          3021    (optional)    (conditional)

IPLA Product Name              IPLA Product ID    Tool MSUs    Customer MSUs    Customer Comments (255 chars max)    Footnotes
File Export for z/OS V1        5697-I12          3021
IBM WebSphere Application Server for z/OS V8 5655-W65          20    (optional)    (conditional)

Footnotes:
(iv) Product's Tool MSUs are based on Getting Started Sub-capacity Pricing from one or more LPARs
(viii) Product's Tool MSUs are based on Integrated Workload Pricing from one or more LPARs
(d) Mobile Workload Pricing transaction data from this product was included from one or more LPARs

```

Figure 94. Sample sub-capacity report for Mobile Workload Pricing: Product Summary Information section

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the Product Summary Information section.

SMF / SCRT89 Input Data Statistics section

There are no differences between the SMF / SCRT89 Input Data Statistics section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“SMF / SCRT89 Input Data Statistics section” on page 118](#).

The sample SMF / SCRT89 Input Data Statistics section appears in [Figure 95 on page 221](#).

```

=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
=====
SMF/SCRT89 INPUT DATA STATISTICS

SYSID    Input Data Start    Input Data End
LPAR1    SYS1    11 Jan 2024 - 22:59  12 Jan 2024 - 23:59
LPAR2    SYS2    11 Jan 2024 - 22:59  12 Jan 2024 - 23:59
CPC      11 Jan 2024 - 22:59  12 Jan 2024 - 23:59

```

Figure 95. Sample sub-capacity report for Mobile Workload Pricing: SMF / SCRT89 Input Data Statistics section

Detail LPAR Data Section

There are no differences between the Detail LPAR Data Section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Detail LPAR Data section” on page 119](#).

The sample Detail LPAR Data Section appears in [Figure 96 on page 222](#).

```

=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR1	1611	1	12 Jan 2024 - 16:00	1594	1	12 Jan 2024 - 17:00
LPAR2	1619	1	12 Jan 2024 - 17:00	1607	1	12 Jan 2024 - 18:00
CPC	3213	1	12 Jan 2024 - 17:00	3200	1	12 Jan 2024 - 16:00

Figure 96. Sample sub-capacity report for Mobile Workload Pricing: Detail LPAR Data section

Product Max Contributors section

The MSU values that are listed in the Product Max Contributors section are the same for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report, with one subtle change. While all of the MSU values listed in this section continue to reflect the product MSUs by LPAR or LPAR Group for the earliest hour for the product's maximum MSU value, the MSU values posted for Getting Started Sub-Capacity Pricing eligible products are based on the Getting Started Sub-Capacity Pricing algorithm whenever that value is less than the traditional sub-capacity value for that product for that hour. The result is that some of the LPAR values posted for Getting Started Sub-Capacity Pricing eligible products are less than the MSUs posted for z/OS, even when the Getting Started Sub-Capacity Pricing eligible product runs in all of the same LPARs and at the same times as z/OS.

Note: The MSUs for Getting Started Sub-Capacity Pricing eligible products are posted under an LPAR group whenever the LPAR group cap is less than the Getting Started Sub-Capacity Pricing value for the earliest hour with the maximum product value. That is not the case for this sample configuration.

This section provides an additional layer of detail behind the sub-capacity values reported for the products listed in the Product Summary Information section by reporting the date/time that the sub-capacity value occurred together with the product's sub-capacity contribution from each LPAR in which that product was running at that date and time. When the *same* peak sub-capacity value occurs for multiple hours, the Product Max Contributors section lists the data for the *earliest* hour. The impact of SCRT reporting a product's sub-capacity value based on the earliest peak hour whenever that peak value occurs for multiple hours is important enough to reiterate in the following explanations because the MWP conditions for the other peak hours can differ from the MWP conditions for the earliest peak hour.

Given the level of detail in this section, the sample Product Max Contributors section shown in [Figure 97 on page 223](#) will be used to explain Mobile Workload Pricing in more detail. However, this explanation assumes you have a general understanding of the contents and use of each field in the Product Max Contributors section, as described in [“Product Max Contributors section” on page 121](#).

```

=====
==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR SYS1	LPAR SYS2	Mobile MSU Reduction
z/OS V2	5650-Z0S	3021	12 Jan 2024 - 17:00	1499	1522	192
DB2 11 for z/OS	5615-DB2	3021	12 Jan 2024 - 17:00	1499	1522	192
CICS TS for z/OS V5	5655-Y04	3021	12 Jan 2024 - 17:00	1499	1508	192
IBM MQ for z/OS V8	5655-W97	3021	12 Jan 2024 - 17:00	1499	1522	192
IMS V13	5635-A04	3021	12 Jan 2024 - 17:00	1499	1508	192
IBM Enterprise Cobol for z/OS V4	5655-S71	3021	12 Jan 2024 - 17:00	1499	1522	192
File Export for z/OS V1	5697-I12	3021	12 Jan 2024 - 17:00	1499	1522	192
IBM WebSphere Application Server for z/OS V8	5655-W65	20	12 Jan 2024 - 20:00	16	4	0

Figure 97. Sample sub-capacity report for Mobile Workload Pricing: Product Max Contributors section

How SCRT implements Mobile Workload Pricing

It is important to understand that SCRT uses the appropriate algorithm to calculate the sub-capacity value for each product for each hour for each LPAR, independently of what is happening with that product on other LPARs or guest systems during that same hour. Then, SCRT sums each of the product's sub-capacity values across all of the LPARs for the same hour to calculate a sub-capacity value for the product for the CPC for that hour. After performing these sub-capacity calculations for the product across the CPC for all of the hours during the reporting period, SCRT selects the earliest hour across the CPC with the product's peak sub-capacity value and reports the date and time for that hour, along with the individual sub-capacity values for the product for each LPAR for that same hour, in the Product Max Contributors section.

For products like z/OS, whose sub-capacity is entirely based on the 4-hour rolling average MSU utilization of the LPAR, the 4-hour rolling average MSU utilization is always reported as the product's sub-capacity value. In the sample Product Max Contributors section, the peak z/OS sub-capacity value across the CPC was 3021 MSUs and that peak first occurred on 12 Jan 2024 at 17:00. The z/OS product running in LPAR1 and LPAR2 contributed 1499 and 1522 MSUs, respectively, to the z/OS peak sub-capacity value of 3021 MSUs for the CPC.

One key difference between a basic z/OS sub-capacity report and a sub-capacity report generated using MWP is the Mobile MSU Reduction column in the Product Max Contributors section. This column is generated whenever mobile pricing data was provided via the INPUTCSV DD statement or was detected in the SMF type 70 data passed as input to SCRT, as was the case in the configuration (see [“Sample configuration for INPUTCSV-based Mobile Workload Pricing” on page 207](#)).

The Mobile MSU Reduction column shows the adjustment that SCRT applied to products' sub-capacity values using the MWP algorithm. In the Product Max Contributors section for the sample configuration, the peak z/OS sub-capacity value and the value of all products based entirely on the 4-hour rolling average MSU utilization of the LPAR were adjusted by SCRT using the value shown in the Mobile MSU Reduction column.

In the Product Max Contributors section for the sample configuration, note that the MSU values shown for LPAR SYS1 and LPAR SYS2 already reflect the contribution of MWP defining programs that ran in each LPAR during the peak hour. Also note that the Mobile MSU Reduction column reflects the total reduction across all LPARs for that hour.

See [“Mobile/zCAP Pricing Detailed Data section for MWP” on page 214](#) for more information about the reporting results for the sample configuration.

A note about MWP adjustments: The MSU values used for MWP adjustments are derived from the mobile transaction general purpose processor time values reported using WLM mobile workload classification. These adjustments are calculated internally by SCRT as specified in *IBM Addendum for System z Mobile Workload Pricing*. The Product Max Contributors section is only intended to show the contribution of each LPAR to the specified product's peak sub-capacity value for earliest peak hour.

Product Grid Snapshot section

There are no differences between the Product Grid Snapshot section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Product Grid Snapshot section” on page 123](#).

The sample Product Grid Snapshot section appears in [Figure 98 on page 224](#).

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                                Product ID    SYS1    SYS2
z/OS V2                                    5650-Z0S      3.2%    3.2%
DB2 11 for z/OS                           5615-DB2      3.2%    3.2%
CICS TS for z/OS V5                       5655-Y04      3.3%    3.3%
IBM MQ for z/OS V8                        5655-W97      3.3%    3.3%
IMS V13                                   5635-A04      3.3%    3.3%
IBM WebSphere Application Server for z/OS V8 5655-W65      3.3%    3.2%

```

Figure 98. Sample sub-capacity report for Integrated Workload Pricing: Product Grid Snapshot section

Getting Started Sub-Capacity Pricing Assigned Hours section

There are no differences between the Getting Started Sub-Capacity Pricing Assigned Hours section for a basic z/OS sub-capacity report and for a sub-capacity report generated using MWP. For details about this section, see [“Getting Started Sub-Capacity Pricing Assigned section” on page 179](#).

The sample Getting Started Sub-Capacity Pricing Assigned Hours section appears in [Figure 99 on page 224](#).

```

==S7=====
GETTING STARTED SUB-CAPACITY PRICING ASSIGNED - HOURS

Product Name                                Product ID    LPAR    LPAR
                                           SYS1      SYS2
IBM WebSphere Application Server for z/OS V8 5655-W65      25      24

```

Figure 99. Sample sub-capacity report for Mobile Workload Pricing: Getting Started Sub-Capacity Pricing Assigned Hours section

Mobile/zCAP Pricing Detailed Data section for MWP

The optional Mobile/zCAP Pricing Detailed Data section appears when the optional **Generate_Detailed_Data** SPECIAL DD control statement is specified. This section shows the mobile MSU reduction applied by SCRT for each MWP defining program, per LPAR. This information is shown for each sub-capacity product on an hour-by-hour basis.

As shown in [Figure 100 on page 226](#), the optional Mobile/zCAP Pricing Detailed Data section reports several columns for each product:

- Date and time of the adjusted interval.
- CPC Type-model.
- CPC rated machine capacity in MSU.
- Adjusted MSU values for each LPAR. (Note that the MWP algorithm has already been applied.)
- For each mobile defining product with INPUTCSV DD data specified, the MWP reduction applied per defining product, per partition.
- Total Mobile MSU reduction applied across all MWP defining products.

- CPC Total MSU, which is the sum of the adjusted LPAR MSU values across all partitions for the adjusted product.

There are several additional values that can appear in section V5 of a sub-capacity report:

–

Indicates that there is no data for this interval.

dd(rr)

Indicates that the defined capacity limit for the LPAR was reached during this interval. The first number (*dd*) represents the defined capacity value used for this hour. The second number (*rr*) represents the 4-hour rolling average value prior to the capping.

G(rr)

Indicates that the sum of the partitions exceeded the group capacity limit; the group capacity was used for this hour for that set of LPARs. The *rr* value represents the 4-hour rolling average value prior to the capping.

P(rr)

Indicates that the sum of the partitions exceeded the processor capacity; the processor capacity was used. The *rr* value represents the 4-hour rolling average value prior to the capping.

(GSSP)

Indicates that a WebSphere program eligible for Getting Started Sub-Capacity Pricing received the GSSP value so there was no additional adjustment.

nnn(CPS)

Indicates that an LPAR was assigned to a Tailored Fit Pricing solution for this hour. In this case, the LPAR does not contribute to the product sub-capacity value. The *nnn* value represents what the LPAR sub-capacity value would have been if it were part of the standard (non-container) environment.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

The following descriptions of the Mobile Workload Pricing algorithm are based on the sample configuration in “Sample configuration for INPUTCSV-based Mobile Workload Pricing” on [page 207](#). Note that only a subset of the products running in the sample environment are included here.

Results for z/OS V2R1

After SCRT applied the appropriate algorithm hour by hour, LPAR by LPAR, and then summed the sub-capacity values for the same hour across all of the z/OS LPARs, SCRT reported 12 Jan 2024 at 17:00 as the earliest hour with the peak sub-capacity value for z/OS, as indicated in the sample Product Max Contributors section.

For 12 Jan 2024 at 17:00, SCRT applied the appropriate algorithms and calculated a sub-capacity value of 1499 MSUs for LPAR1. As shown in the **WLM Mobile Partition** column for LPAR1, this includes a 95 MSU MWP reduction derived from the WLM mobile classification 4-hour rolling average data found in SMF type 70 records that are passed as input to SCRT in the sample configuration in “Sample configuration for WLM classification-based Mobile Workload Pricing” on [page 217](#).

For partition SYS2, SCRT calculated a sub-capacity value of 1522 MSUs, which includes a 97 MSU MWP reduction derived from the WLM mobile classification 4-hour rolling average data that was specified as input to SCRT in the sample configuration.

[Figure 100 on page 226](#) illustrates these results.


```
==V5=====
```

MOBILE/ZCAP PRICING DETAILED DATA

MLC Products

5650-ZOS

z/OS V2

Date Time of RMF Interval	Type	CPC MSU	Partition LPAR1	LPAR2	WLM Mobile Partition LPAR1	LPAR2	Total Mobile MSU Reduction	CPC Total
11 Jan 2024 - 23:00	2817-760	5001	-	-	-	-	0	0
12 Jan 2024 - 00:00	2817-760	5001	877	1001	56	64	120	1878
12 Jan 2024 - 01:00	2817-760	5001	973	955	62	61	123	1928
12 Jan 2024 - 02:00	2817-760	5001	1106	953	70	61	131	2059
12 Jan 2024 - 03:00	2817-760	5001	1227	926	78	58	136	2153
12 Jan 2024 - 04:00	2817-760	5001	1223	902	78	57	135	2125
12 Jan 2024 - 05:00	2817-760	5001	1164	976	74	62	136	2140
12 Jan 2024 - 06:00	2817-760	5001	1081	1050	68	67	135	2131
12 Jan 2024 - 07:00	2817-760	5001	1022	1130	65	72	137	2152
12 Jan 2024 - 08:00	2817-760	5001	978	1266	62	80	142	2244
12 Jan 2024 - 09:00	2817-760	5001	878	1342	56	85	141	2220
12 Jan 2024 - 10:00	2817-760	5001	818	1400	52	89	141	2218
12 Jan 2024 - 11:00	2817-760	5001	765	1464	49	93	142	2229
12 Jan 2024 - 12:00	2817-760	5001	861	1479	55	94	149	2340
12 Jan 2024 - 13:00	2817-760	5001	1080	1481	68	94	162	2561
12 Jan 2024 - 14:00	2817-760	5001	1256	1478	80	94	174	2734
12 Jan 2024 - 15:00	2817-760	5001	1425	1477	91	94	185	2902
12 Jan 2024 - 16:00	2817-760	5001	1514	1494	97	95	192	3008
12 Jan 2024 - 17:00	2817-760	5001	1499	1522	95	97	192	3021
12 Jan 2024 - 18:00	2817-760	5001	1401	1511	89	96	185	2912
12 Jan 2024 - 19:00	2817-760	5001	1274	1491	81	95	176	2765
12 Jan 2024 - 20:00	2817-760	5001	1099	1416	70	90	160	2515
12 Jan 2024 - 21:00	2817-760	5001	985	1293	62	82	144	2278
12 Jan 2024 - 22:00	2817-760	5001	928	1191	59	76	135	2119
12 Jan 2024 - 23:00	2817-760	5001	896	1128	57	71	128	2024

Figure 100. Sample sub-capacity report for Mobile Workload Pricing: Mobile/zCap Pricing Detailed Data section — z/OS results

Results for Db2 11 for z/OS

SCRT applied the process as described in “Results for z/OS V2R1” on [page 225](#). Like z/OS, Db2 11 for z/OS is based on the 4-hour rolling average MSU utilization of the LPAR.

This same process applies for all sub-capacity programs that reported a peak MSU value of 3021 with 12 Jan 2024 at 17:00 as the earliest hour with the peak sub-capacity value.

[Figure 101 on page 227](#) illustrates the results.

5615-DB2		DB2 11 for z/OS				CICS Mobile Partition		Total Mobile MSU Reduction	CPC Total
Date Time of RMF Interval	Type	CPC MSU	Partition LPAR1	LPAR2	LPAR1	LPAR2			
11 Jan 2024 - 23:00	2817-760	5001	-	-	-	-		0	0
12 Jan 2024 - 00:00	2817-760	5001	877	1001	56	64		120	1878
12 Jan 2024 - 01:00	2817-760	5001	973	955	62	61		123	1928
12 Jan 2024 - 02:00	2817-760	5001	1106	953	70	61		131	2059
12 Jan 2024 - 03:00	2817-760	5001	1227	926	78	58		136	2153
12 Jan 2024 - 04:00	2817-760	5001	1223	902	78	57		135	2125
12 Jan 2024 - 05:00	2817-760	5001	1164	976	74	62		136	2140
12 Jan 2024 - 06:00	2817-760	5001	1081	1050	68	67		135	2131
12 Jan 2024 - 07:00	2817-760	5001	1022	1130	65	72		137	2152
12 Jan 2024 - 08:00	2817-760	5001	978	1266	62	80		142	2244
12 Jan 2024 - 09:00	2817-760	5001	878	1342	56	85		141	2220
12 Jan 2024 - 10:00	2817-760	5001	818	1400	52	89		141	2218
12 Jan 2024 - 11:00	2817-760	5001	765	1464	49	93		142	2229
12 Jan 2024 - 12:00	2817-760	5001	861	1479	55	94		149	2340
12 Jan 2024 - 13:00	2817-760	5001	1080	1481	68	94		162	2561
12 Jan 2024 - 14:00	2817-760	5001	1256	1478	80	94		174	2734
12 Jan 2024 - 15:00	2817-760	5001	1425	1477	91	94		185	2902
12 Jan 2024 - 16:00	2817-760	5001	1514	1494	97	95		192	3008
12 Jan 2024 - 17:00	2817-760	5001	1499	1522	95	97		192	3021
12 Jan 2024 - 18:00	2817-760	5001	1401	1511	89	96		185	2912
12 Jan 2024 - 19:00	2817-760	5001	1274	1491	81	95		176	2765
12 Jan 2024 - 20:00	2817-760	5001	1099	1416	70	90		160	2515
12 Jan 2024 - 21:00	2817-760	5001	985	1293	62	82		144	2278
12 Jan 2024 - 22:00	2817-760	5001	928	1191	59	76		135	2119
12 Jan 2024 - 23:00	2817-760	5001	896	1128	57	71		128	2024

Figure 101. Sample sub-capacity report for Mobile Workload Pricing: Mobile/zCap Pricing Detailed Data section – Db2 11 for z/OS results

Results for IBM WebSphere Application Server for z/OS V8

SCRT applied the Getting Started Sub-Capacity Pricing for z/OS IPLA Software process to IBM WebSphere Application Server for z/OS V8. Since the terms and conditions of GSSP were met in the sample configuration, and the GSSP values were less than the traditional sub-capacity values for each hour of the reporting period, SCRT assigned IBM WebSphere Application Server for z/OS V8 the GSSP value for each hour in the reporting period.

When the GSSP value is used for a GSSP eligible product, the Mobile/zCAP Pricing Detailed Data section places the (GSSP) token in the MWP defining program partition reduction column, as shown in [Figure 102 on page 228](#). In accordance with the Mobile Workload Pricing addendum, the MWP reduction is not applied when GSSP pricing is used instead of the traditional sub-capacity value.

In the sample configuration, the peak MSU value used for IBM WebSphere Application Server for z/OS V8 is 20, which is achieved during the 12 Jan 2024 20:00 hour.

For information about the GSSP process, see [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#).

5655-W65		IBM WebSphere Application Server for z/OS V8							
Date Time of RMF Interval		Type	CPC MSU	Partition		WLM Mobile Partition		Total Mobile MSU Reduction	CPC Total
				LPAR1	LPAR2	LPAR1	LPAR2		
12 Jan 2024	- 00:00	2817-760	5001	2	1	(GSSP)	(GSSP)	0	3
12 Jan 2024	- 01:00	2817-760	5001	4	1	(GSSP)	(GSSP)	0	5
12 Jan 2024	- 02:00	2817-760	5001	8	2	(GSSP)	(GSSP)	0	10
12 Jan 2024	- 03:00	2817-760	5001	10	3	(GSSP)	(GSSP)	0	13
12 Jan 2024	- 04:00	2817-760	5001	10	3	(GSSP)	(GSSP)	0	13
12 Jan 2024	- 05:00	2817-760	5001	11	3	(GSSP)	(GSSP)	0	14
12 Jan 2024	- 06:00	2817-760	5001	12	3	(GSSP)	(GSSP)	0	15
12 Jan 2024	- 07:00	2817-760	5001	14	4	(GSSP)	(GSSP)	0	18
12 Jan 2024	- 08:00	2817-760	5001	14	4	(GSSP)	(GSSP)	0	18
12 Jan 2024	- 09:00	2817-760	5001	13	4	(GSSP)	(GSSP)	0	17
12 Jan 2024	- 10:00	2817-760	5001	12	4	(GSSP)	(GSSP)	0	16
12 Jan 2024	- 11:00	2817-760	5001	11	4	(GSSP)	(GSSP)	0	15
12 Jan 2024	- 12:00	2817-760	5001	11	4	(GSSP)	(GSSP)	0	15
12 Jan 2024	- 13:00	2817-760	5001	11	4	(GSSP)	(GSSP)	0	15
12 Jan 2024	- 14:00	2817-760	5001	11	4	(GSSP)	(GSSP)	0	15
12 Jan 2024	- 15:00	2817-760	5001	11	4	(GSSP)	(GSSP)	0	15
12 Jan 2024	- 16:00	2817-760	5001	12	4	(GSSP)	(GSSP)	0	16
12 Jan 2024	- 17:00	2817-760	5001	13	4	(GSSP)	(GSSP)	0	17
12 Jan 2024	- 18:00	2817-760	5001	14	4	(GSSP)	(GSSP)	0	18
12 Jan 2024	- 19:00	2817-760	5001	14	4	(GSSP)	(GSSP)	0	18
12 Jan 2024	- 20:00	2817-760	5001	16	4	(GSSP)	(GSSP)	0	20
12 Jan 2024	- 21:00	2817-760	5001	15	4	(GSSP)	(GSSP)	0	19
12 Jan 2024	- 22:00	2817-760	5001	15	4	(GSSP)	(GSSP)	0	19
12 Jan 2024	- 23:00	2817-760	5001	15	4	(GSSP)	(GSSP)	0	19

Figure 102. Sample sub-capacity report for Mobile Workload Pricing: Mobile/zCap Pricing Detailed Data section — IBM WebSphere Application Server for z/OS results

Chapter 11. Sub-capacity reporting for IBM Z Collocated Application Pricing (zCAP)

IBM Z Collocated Application Pricing (zCAP) is a sub-capacity pricing option under Advanced Workload License Charges (AWLC), Advanced Entry Workload License Charges (AEWLC), or System z New Application License Charges (zNALC). It is intended to enable qualifying z/OS customers to improve the cost of deploying new z/OS applications. zCAP allows new applications to be deployed in existing LPARs, yet be priced as if running in dedicated LPARs.

This enhancement to sub-capacity reporting is intended to enable you to essentially pay for the new application's direct usage of the IBM middleware that manages the processing of the new application, such as CICS or WebSphere Application Server, while mitigating the impact on the reported MSUs for other programs running in the same partition.

IBM Z Collocated Application Pricing is effected for SMF data collected on 2 April 2015 and later, provided that all of the required zCAP conditions have been met.

Overview of IBM Z Collocated Application Pricing

IBM Z Collocated Application Pricing provides an enhanced way of reporting z/OS system utilization when a new application is deployed in an existing LPAR. A new z/OS application is defined by sub-capacity eligible programs when running in the same LPAR as mobile workloads processed by select zCAP defining programs. Mobile workloads are defined by zCAP as an application not currently running on any IBM Z server, leveraging one of the zCAP *defining programs*.

Table 38 on page 229 lists the zCAP defining programs supported by SCRT V23.10.00 and later. These are the products whose sub-capacity pricing values will be reduced when the required set of zCAP conditions are met.

Note: For the most current list of zCAP defining programs, see [IBM Z Collocated Application Pricing \(zCAP\) Defining Programs \(www.ibm.com/downloads/documents/us-en/10a99803f4afd859\)](http://www.ibm.com/downloads/documents/us-en/10a99803f4afd859).

Table 38. zCAP defining programs

Product family	Product name	Product ID
CICS	IBM CICS Transaction Server for z/OS, V4	5655-S97
	CICS Transaction Server for z/OS, V5	5655-Y04
	CICS Value Unit Edition (VUE) V5	5722-DFJ
Db2	IBM DB2 V9 for z/OS	5635-DB2
	Db2 V10 for z/OS	5605-DB2
	Db2 V11 for z/OS	5615-DB2
	DB2 VUE V9	5697-P12
	Db2 10 VUE	5697-P31
	Db2 11 VUE	5697-P43

Table 38. zCAP defining programs (continued)

Product family	Product name	Product ID
IMS	IBM IMS V11	5635-A02
	IMS V12	5635-A03
	IMS V13	5635-A04
	IMS V14	5635-A05
	IMS Database VUE V12	5655-DSQ
	IMS Database VUE V13	5655-DSM
	IMS Transaction Manager VUE V12	5655-TM1®
	IMS Transaction Manager VUE V13	5655-TM2
MQ	IBM WebSphere MQ for z/OS, V6	5655-L82
	WebSphere MQ for z/OS, V7	5655-R36
	WebSphere MQ for z/OS, V8	5655-W97
	IBM MQ for z/OS V9	5655-MQ9
	WebSphere MQ VUE V7	5655-VUE
	WebSphere MQ VUE V8	5655-VU8
	IBM MQ for z/OS VUE V9	5655-VU9
WebSphere	WebSphere Application Server for z/OS, V7	5655-N02
	WebSphere Application Server for z/OS, V8	5655-W65
	IBM WebSphere Application Server for z/OS V9	5655-WAS

Depending on your zCAP usage characteristics, you may need to track and report the general purpose processor time (CPU time) for the new application's zCAP defining program and report those values in a pre-defined format to IBM each month in order to take advantage of this pricing offering.

Whenever the required zCAP conditions have been met, SCRT reduces the sub-capacity values for other programs (non-zCAP defining programs) using the following methodology:

1. Calculate the 4-hour rolling average for general purpose processor time for the reported zCAP defining program.
2. Subtract 100 percent of those values from the traditional sub-capacity MSUs for all other sub-capacity eligible middleware programs on that CPC on an hour-by-hour basis, per LPAR.
3. Subtract 50 percent of those values from the traditional sub-capacity MSUs for z/OS on an hour-by-hour basis, per LPAR.
4. Sum the adjusted MSU values for the same hour across all of the LPARs and the z/OS guest systems running under z/VM (if any) in which the program runs to create an adjusted sub-capacity value for the program for each hour for that CPC.

SCRT determines the billable MSU peak value for a given program on a CPC using the adjusted MSU values.

For any IPLA reference-based programs with a zCAP defining program as their parent which are purchased in support of an approved zCAP application and are not running elsewhere in the CPC or sysplex, the required license capacity will be referenced only to the MSUs for the zCAP defining program associated with the zCAP application.

For the case where a z/OS system runs as a guest under z/VM, each eligible z/OS guest system is treated as its own LPAR when applying any sub-capacity adjustments for zCAP.

The methodology that SCRT uses to determine the 4-hour rolling average utilization of the zCAP defining program depends on whether the zCAP defining program is new to the LPAR or is already running on the LPAR. SCRT supports both of these use cases.

Use case 1: Reporting methodology when the zCAP defining program is new to an LPAR

When a zCAP defining program is deployed in an LPAR where that same program is not already running, you must create a simple comma-separated value (CSV) file that provides the program number for the defining program, the machine type and serial number, and the name of the LPAR or LPARs where that program runs. The CSV file must be loaded into SCRT each month.

SCRT uses the SMF type 89 record data for the zCAP defining program for reporting, as follows:

1. **Reporting for the zCAP defining program in the LPAR:** The MSUs are retrieved from the SMF type 89 subtype 1 records for the defining program. These records contain the CPU usage time for the program, not the LPAR utilization value.
2. **Reporting for the other middleware in the LPAR:** 100 percent of the MSUs from the SMF type 89 records for the zCAP defining program are removed from the LPAR utilization values for the middleware products, each hour, as part of the calculation of the 4-hour rolling average utilization.
3. **Reporting for z/OS in the LPAR:** 50 percent of the MSUs from the SMF type 89 records for the zCAP defining program are removed from the LPAR utilization values for z/OS, each hour, as part of the calculation of the 4-hour rolling average utilization.

Note that if you make changes to the configuration of the LPARs which cause the zCAP defining program deployment scenario to align with use case 2 instead of use case 1, you must meet the reporting requirements described in use case 2, including the capture and submission of the CPU time attributable only to the new application. This could occur, for instance, if a non-zCAP workload that also uses the zCAP defining program is added to the same LPAR.

Use case 2: Reporting methodology when the zCAP defining program already runs in an LPAR

When a new instance of a zCAP defining program is deployed in an LPAR where that same program already runs, you must create a CSV file to capture only the CPU time used by the defining program in support of the new application. This is similar to the tracking requirements for IBM Mobile Workload Pricing (MWP), but applies only to new applications, not existing transactions. This may require the use of one or more dedicated address spaces for the new application in order to be able to capture and record the CPU time that is only attributable to that new application. SMF type 30 records provide a means of tracking address space CPU time.

SCRT uses the following methodology for this use case:

1. **Reporting for the zCAP defining program in the LPAR:** The MSUs for pricing will be the total reported MSU peak for the defining program (according to the standard sub-capacity LPAR utilization value) including the new application CPU time, as the new application simply represents incremental MSUs to an existing program.
2. **Reporting for the other middleware in the LPAR:** 100 percent of the incremental MSUs for the zCAP defining program, as provided by you in the CSV file, are removed from the LPAR utilization values for the middleware products, each hour, as part of the calculation of the 4-hour rolling average utilization.
3. **Reporting for z/OS in the LPAR:** 50 percent of the incremental MSUs for the zCAP defining program, as provided by you in the CSV file, are removed from the LPAR utilization values for z/OS, each hour, as part of the calculation of the 4-hour rolling average utilization.

Similar to Mobile Workload Pricing, this ensures that the monthly MSU peaks for other middleware programs do not increase simply because a new application was deployed in the same LPAR. However, the MSU peaks for other middleware programs may increase as a result of interaction with the new application.

Qualifications for IBM Z Collocated Application Pricing

The following conditions must be met before IBM will accept and process sub-capacity reports where IBM Z Collocated Application Pricing is being reported via footnote (c) or (o). Some of these conditions are also described in more detail elsewhere in this chapter but are consolidated here for ease of reference.

1. The CPC from which the SMF data is collected must be licensed under AWLC, AEWLC, or zNALC sub-capacity terms and conditions. The zCAP terms and conditions require that defining programs be installed on an IBM z13, z13s, zEC12, zBC12, z196, or z114 CPC (or follow-on CPC) that has implemented sub-capacity pricing. Note that the SMF data must be collected from one or more z/OS systems running on zEnterprise or later CPCs; SCRT only applies the zCAP algorithm when processing SMF data collected from zEnterprise or later CPCs.
2. The SMF data must be collected from a z/OS V1.10 or later system that has SMF APAR OA31615 applied prior to collecting the SMF data. SCRT does not apply the zCAP algorithm to SMF data collected from z/OS V1.9 or earlier z/OS systems. For eligible CPCs with a mixture of z/OS V1.9 or earlier systems and z/OS V1.10 or later systems, SCRT only applies the zCAP algorithm to the SMF data collected from eligible z/OS 1.10 and later systems. For the z/OS 1.10 and later systems, SCRT only applies the zCAP algorithm to SMF data collected for the hours during which the system ran with APAR OA31615 applied. The result could be that SCRT applies the zCAP algorithm to the SMF data collected only from a subset of the z/OS systems running on the CPC and only for a subset of the hours that a z/OS system was running.
3. For z/OS guest systems running under z/VM, all z/OS guest systems must be at V1.10 or later. IBM Z Collocated Application Pricing is not supported on any z/OS guest system whenever at least one z/OS V1.6 or earlier guest system is running on the CPC anytime during that SCRT reporting period. If a z/OS V1.6 guest system is running on the CPC, none of the products running on any of the z/OS guests, including z/OS V1.10 and later guests with APAR OA31615 applied, are eligible for IBM Z Collocated Application Pricing on that CPC. This restriction is consistent with the existing SCRT rules for z/OS V1.6 guests, where all sub-capacity products are assigned the maximum capacity of the LPAR in which they execute. Note that native z/OS systems, including z/OS V1.6, are not affected by the SCRT rules for z/OS V1.6 guests. Products running on native z/OS V1.10 and later systems are eligible for IBM Z Collocated Application Pricing when all of the other zCAP conditions are met, even when z/OS V1.6 or earlier guest systems are also running in the same CPC. The guest restriction only applies to z/OS guest systems and has no effect on native z/OS systems. When using z/OS guest systems, either z/VM V5.4 with APAR VM64798 or z/VM V6.1 with APAR VM64879 or later are required.
4. The traditional sub-capacity values of programs are only adjusted when one or more zCAP defining programs (for which data was submitted via the INPUTCSV DD statement) were running in the same z/OS system (LPAR) during the same 4-hour window that was used to determine the traditional sub-capacity value of the zCAP defining program. Programs running in one z/OS system (LPAR) are not affected by zCAP defining programs running during the same hour in a different z/OS system (LPAR).
5. Only SMF data collected on 2 April 2015 and later is eligible for IBM Z Collocated Application Pricing, as set forth in the IBM Z Collocated Application Pricing announcement letter. Although SCRT will generate sub-capacity reports based on the zCAP algorithm for SMF data collected earlier than 2 April 2015, the IBM business process with reject those reports.
6. You must generate and submit to IBM a monthly sub-capacity report that contains zCAP adjustments during the normal sub-capacity submission window (business as usual).
7. You must use SCRT V23.10.00 or later to generate the sub-capacity report that is submitted for IBM Z Collocated Application Pricing.
8. You must use the INPUTCSV DD statement to request that SCRT generate a sub-capacity report using the zCAP algorithms whenever all of the required zCAP conditions are met. The default SCRT processing does not use the zCAP algorithms. For details, see [“INPUTCSV DD statement for zCAP” on page 234](#).
9. You must agree to the IBM Addendum for IBM Z Collocated Application Pricing (Z126-6861) prior to submitting a sub-capacity report where footnote (c) or (o) appears in the Product Information Summary section, since footnote (c) or (o) indicates that the sub-capacity values for one or more products were adjusted based on zCAP. Failure to formally agree to the zCAP Addendum prior to

submitting a sub-capacity report with footnote (c) or (o) will cause the IBM business process to reject that report. The rejection will contain an explanation and request that you generate a replacement sub-capacity report without using the INPUTCSV DD statement to submit zCAP data, and submit it to IBM in order to continue to receive traditional sub-capacity pricing for the IBM products on that CPC.

Data collection requirements for IBM Z Collocated Application Pricing

The data collection requirements for IBM Z Collocated Application Pricing are:

- You must collect SMF type 70 and type 89 records for all LPARs on all CPCs each month, and you must retain these records for six months after the billing period for auditing purposes.
- You must collect the source data (CPU time) for the zCAP defining programs that will be used in monthly reporting, and you must retain this data (CPU time and CSV files) for six months after the billing period for auditing purposes.
- You must preprocess your zCAP defining program CPU time into a predefined format to be loaded into SCRT for each sub-capacity reporting period. The data must consist of general purpose processor CPU seconds for the zCAP defining program, summarized by hour, by LPAR, for all CPCs that process zCAP data.

zCAP monthly reporting process

When using IBM Z Collocated Application Pricing, you must complete the following requirements each month:

- When the zCAP defining program is new to the LPAR:
 - Produce a file in the comma-separated value (CSV) format that contains the program number of the defining program, the machine type and serial number, and the name of the LPAR or LPARs where the program runs. The format of the data in this file is described in [“Format of the INPUTCSV input data for IBM Z Collocated Application Pricing”](#) on page 235.
 - Load the resulting data files along with the SMF type 70 and type 89 records into SCRT.
 - Run SCRT and submit the resulting Sub-capacity reports to IBM for each reporting period.
- When the zCAP defining program represents a new instance of the same program already running in the LPAR:
 - Track the CPU time for the zCAP defining program which is attributable only to the new application, on an hourly basis, per LPAR.
 - Produce a file in the comma-separated value (CSV) format that contains the CPU seconds for the zCAP defining program by hour and LPAR. The format of the data in this file is described in [“Format of the INPUTCSV input data for IBM Z Collocated Application Pricing”](#) on page 235.
 - Load the resulting data files along with the SMF type 70 and type 89 records into SCRT.
 - Run SCRT and submit the resulting Sub-capacity reports to IBM for each reporting period.

SCRT externals for IBM Z Collocated Application Pricing

SCRT provides the following externals for IBM Z Collocated Application Pricing. Where noted, these externals are specific to zCAP and are treated as optional when zCAP is not being requested or does not apply to the SMF data being processed by SCRT.

1. SCRT uses the optional DD statement, INPUTCSV, which is used to provide input CPU consumption data in a predefined format for zCAP.
2. SCRT uses footnotes in the Product Summary Information section of the sub-capacity report to indicate that the Too1 MSUs value for the corresponding product was adjusted by the zCAP algorithm, as follows:

- Footnote (c) indicates that the product was adjusted using the zCAP algorithm when the INPUTCSV DD statement was used to specify zCAP data for that product.
 - Footnote (o) appears in conjunction with footnote (c) in cases where the INPUTCSV DD contains the ONLYZCAP or ONLYZCAPALL token. This token indicates that the zCAP defining program is only used for zCAP eligible workloads.
3. SCRT uses an optional **Generate_Detailed_Data** control statement to track the application of mobile and zCAP reductions per hour, per LPAR.

INPUTCSV DD statement for zCAP

The default SCRT processing does not support IBM Z Collocated Application Pricing. You must provide the CPU time data for your zCAP defining program as input on the INPUTCSV DD statement. The data in the data set pointed to by the INPUTCSV DD statement must be in a predefined format, as described in [“Format of the INPUTCSV input data for IBM Z Collocated Application Pricing” on page 235.](#)

Important: In order to calculate the correct 4-hour rolling average for the first hour in the report period, the INPUTCSV data must contain entries for the 3 hours prior to the start of the report period. For instance, for a report period that begins at 00:00 on 2 August, the INPUTCSV data must contain mobile transaction data beginning at 21:00 on 1 August.

Note: IBM authorizes the use of the INPUTCSV DD statement for providing zCAP defining program CPU time data to SCRT only when you have signed or otherwise agreed to the terms of the IBM Addendum for IBM Z Collocated Application Pricing.

For more information about this statement, see [“INPUTCSV DD statement” on page 60.](#)

IBM Z Collocated Application Pricing footnote

For the case where the Too1 MSUs value for one or more products listed in the Product Summary Information section has been adjusted for zCAP, SCRT applies the following changes to the Product Summary Information section:

1. Generates an optional Footnotes column header for the MLC section.
2. Posts footnote (c) in the Footnotes column for each zCAP defining product for which workload data was supplied via the INPUTCSV DD statement.
3. Posts footnote (o) in the Footnotes column for each zCAP defining product where the CSV file for that product contains the ONLYZCAP or ONLYZCAPALL token, indicating that the zCAP defining program is only in use by the zCAP workload.
4. Generates a footnote subsection immediately after the last IPLA product is listed. The footnote subsection contains a footnote header row followed by a description of the footnotes.

Footnotes (c) and (o) are optional. Footnote-related items are only be generated if one or more zCAP defining programs requires such a footnote.

Generate_Detailed_Data control statement

SCRT supports an optional **Generate_Detailed_Data** SPECIAL control statement. This is used to track the application of mobile and zCAP reductions per hour, per LPAR.

When specified, this control statement causes the Mobile/zCAP Pricing Detailed Data section (V5) to be generated in the sub-capacity report.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

Mobile/zCAP Pricing Detailed Data section

The Mobile/zCAP Pricing Detailed Data section (V5) of the sub-capacity report shows, for each product, the hour-by-hour application of mobile and zCAP MSU reductions.

See [“Sample sub-capacity report for IBM Z Collocated Application Pricing” on page 239](#) for a detailed description of the output in the Mobile/zCAP Pricing Detailed Data section.

Format of the INPUTCSV input data for IBM Z Collocated Application Pricing

For IBM Z Collocated Application Pricing, the INPUTCSV DD statement must include the data sets that contain data formatted as described in this topic.

Each version of a zCAP defining program *must* be contained in a separate CSV data set or data set member (that is, one for each individual product ID). (This differs from the CSV file for Mobile Workload Pricing, which requires that data from all of the versions of a product within the product family be included in the same CSV file.)

Note that the format of the data contained in each zCAP input file depends on the use case of the zCAP defining program. The following use cases are supported:

- **Use case 1: Reporting methodology when the zCAP defining program is new to an LPAR**

When a zCAP defining program is deployed in an LPAR where that same program is not already being run, you must supply a CSV file that provides the program number for the defining program, the CPC machine type and serial number, and the name of the LPAR or LPARs where that program runs. [Figure 103 on page 235](#) shows an example of a zCAP CSV data set input for this use case.

```
"Product-zCAP", "5635-A04"
"ONLYZCAPALL", "2964-23456:LPAR1"
```

Figure 103. Example of the contents of a file in the INPUTCSV DD for zCAP use case 1

- **Use case 2: Reporting methodology when the zCAP defining program already runs in the LPAR**

When a new instance of a zCAP defining program is deployed in an LPAR where that same program already runs, you must supply a CSV file that provides the general purpose processor CPU seconds used by the defining program only in support of the new application. These must be summarized by hour, by LPAR, for all CPCs that process zCAP workload. [Figure 104 on page 235](#) shows an example of a zCAP CSV data set input for this use case.

```
"Product-zCAP", "5655-Y04"
"Date", "2964-23456:LPAR2"
"2015-02-08T15:00:00", 500.00
"2015-02-08T16:00:00", 500.00
"2015-02-08T17:00:00", 500.00
```

Figure 104. Example of the contents of a file in the INPUTCSV DD for zCAP use case 2

Note the following points about the INPUTCSV data set:

- Any completely blank row in the file is ignored.
- Any row that starts with a double quotation mark immediately followed by an asterisk ("*) is considered a comment and is ignored.
- The field separator character may be a comma (,) or a semicolon (;), but must be consistent throughout.
- The double quotation marks (as shown in [Figure 103 on page 235](#) and [Figure 104 on page 235](#)) are required. You must use the keyboard double quotation mark character (EBCDIC X'7F' or ASCII X'22'), *not* typographical opening and closing double quotation marks.

- The INPUTCSV data uses LPAR names, *not* system IDs (SYSIDs).

The contents of the INPUTCSV data set consist of a product row (line), a header row, and, depending on the use case, zero or more data rows.

Product row: The first row in the INPUTCSV data set is the product row; it must contain the "Product-zCAP" keyword as the first token.

Following the "Product-zCAP" keyword, you must list the product ID for the zCAP defining program version for which data is being submitted by this INPUTCSV data set. Each program version must be contained in a separate INPUTCSV data set.

Header row: The first non-blank, non-comment row after the product row is the header row. The format of the header row and any subsequent rows depends on the use case of the zCAP defining program, as follows:

- For use case 1, where the zCAP defining program in each LPAR on the CPC (where the program is in use) is only processing zCAP workloads, the "ONLYZCAPALL" token instructs SCRT to use the program usage data from the SMF type 89 records, as was shown in [Figure 103 on page 235](#):

```
"Product-zCAP", "5635-A04"  
"ONLYZCAPALL", "2964-23456:LPAR1"
```

No further input is necessary in this INPUTCSV data set.

- For use case 2, the zCAP defining program is processing a mix of traditional and zCAP workloads. This requires that you provide the general purpose processor CPU time used by the zCAP defining program in support of the new application. The format (which is similar to that used for Mobile Workload Pricing) is as follows:

The first column of the header row must contain the date header information in one of the following formats:

- The "Date" token, which indicates that all times listed are local in ISO-8601 format.
- The "DateUTC" token, which indicates that all times listed are in UTC (ISO-8601 format).

Note: The use of Hypervisor offsets (either LPAR Sysplex Timer offset or z/VM virtual clock, VTOD) requires that you input the CSV time as though these values are already applied to the data. SCRT applies Hypervisor offset values as the SMF records are read and expects the CSV time values to reflect the Hypervisor offset. Normal z/OS time zone offsets do not require this adjustment.

For example, if the Hypervisor offset is used to run the system one month in the future, the CSV file should have that one month in the future date, and not the current day's date and times.

The subsequent columns of the header row must contain the LPAR names in one of the following formats:

- The LPAR name (*not* the system ID), enclosed in double quotation marks

Example: "LPAR2"

- Both the type-serial and LPAR name (*not* the system ID), separated by a colon, in the format *type-serial:LPARname*, enclosed in double quotation marks.

If you have (or plan to have) duplicate LPAR names on different CPCs, you must use this format.

Example: "2964-23456:LPAR2"

Data rows: Any non-blank, non-comment rows after the header row must contain the zCAP workload data.

The first column of a zCAP workload data row must contain the date. The date must be in ISO-8601 format and enclosed in double quotation marks.

Example: A date value of "2014-03-21T03:00:00" indicates 3 a.m. on 21 March 2014.

Note: If the "DateUTC" token appears in the header row, all times are assumed to be in UTC.

The subsequent cell values of the mobile transaction data row must contain the CPU seconds for the hour specified in the first column of that row. The CPU second values must be placed in the column corresponding to the LPAR on which they were gathered. The CPU seconds are specified as a floating-point number and must not be enclosed in double quotation marks.

The following example illustrates this format:

```
"Product-zCAP", "5655-Y04"
"Date", "2964-23456:LPAR2", "2964-23456:LPAR3"
"2015-02-08T15:00:00", 1500.23, 500.00
"2015-02-08T16:00:00", 2500.54, 550.45
"2015-02-08T17:00:00", 2700.83, 475.21
```

- A third use case is possible in which the zCAP defining program is only processing zCAP workloads in some LPARs, but is processing a mix of traditional and zCAP workloads in other LPARs on the CPC. In such a case, you can specify the "ONLYZCAP" token in place of the zCAP general purpose processor CPU time for LPARs that only process zCAP workloads for that zCAP defining program. To do this, follow the same formatting rules as in use case 2. However, for LPARs that process only zCAP workloads, specify the "ONLYZCAP" token instead of a CPU time value in the column for that LPAR.

Example: LPAR3 is processing only zCAP workloads for the zCAP defining program, 5635-A04, while LPAR4 is processing a mix of zCAP and traditional workloads for that same program. The INPUTCSV file specifies the "ONLYZCAP" token in the column for LPAR3 and specifies the zCAP general purpose processor CPU seconds attributable to the zCAP workload in the column for LPAR4, as follows:

```
"Product-zCAP", "5635-A04"
"DATE", "2964-23456:LPAR3", "2964-23456:LPAR4"
"2015-05-04T12:00:00", "ONLYZCAP", 335.00
"2015-05-04T13:00:00", "ONLYZCAP", 755.45
"2015-05-04T14:00:00", "ONLYZCAP", 844.23
"2015-05-04T15:00:00", "ONLYZCAP", 1110.44
"2015-05-04T16:00:00", "ONLYZCAP", 1200.33
"2015-05-04T17:00:00", "ONLYZCAP", 901.55
```

For more information about the INPUTCSV DD statement, see [“INPUTCSV DD statement”](#) on page 60.

Sample configuration for IBM Z Collocated Application Pricing

Table 39 on page 238 illustrates the sample hardware and software configuration for a CPC where an input CSV file was provided for two zCAP defining programs that were running during the reporting period. While the mix of software products and versions used in this sample configuration might not be realistic, it serves to illustrate some of the finer points of the zCAP algorithm.

The sample configuration consists of a 2964 Model 796 CPC with a serial number of 02-12345. The CPC is configured with two LPARs, both running z/OS natively. The z/OS systems in LPAR1 and LPAR2 have z/OS SYSIDs of SYS1 and SYS2, respectively.

The z/OS systems are running a combination of the following z/OS products:

- Db2 11 for z/OS
- CICS TS for z/OS V5
- IBM MQ for z/OS V8
- IMS V13
- IBM Enterprise COBOL for z/OS V4
- IBM WebSphere Application Server for z/OS V8
- IBM Multi-site Workload Lifeline V2

Table 39. Sample configuration for IBM Z Collocated Application Pricing

CPC: 2964-796 (serial = 02-12345)	
LPAR1	LPAR2
z/OS V2.1 (SYSID = SYS1) Db2 11 for z/OS CICS TS for z/OS V5 IBM MQ for z/OS V8 IMS V13 IBM Enterprise COBOL for z/OS V4 IBM WebSphere Application Server for z/OS V8	z/OS V2.1 (SYSID = SYS2) Db2 11 for z/OS CICS TS for z/OS V5 IBM MQ for z/OS V8 IMS V13 IBM Multi-site Workload Lifeline V2

Sample control statement input to SCRT for IBM Z Collocated Application Pricing

SCRT supports additional externals for zCAP. The sample sub-capacity report discussed in [“Sample sub-capacity report for IBM Z Collocated Application Pricing”](#) on page 239 is based on the following SCRT control statement and JCL inputs to SCRT:

INPUTCSV DD statement

The following SCRT DD statement was used to specify data sets containing zCAP workload data for the reporting period for the LPARs and SYSIDs described in [“Sample configuration for IBM Z Collocated Application Pricing”](#) on page 237:

```
//INPUTCSV DD DISP=SHR,DSN=MY.DSN.ID(CICSZCAP)
//          DD DISP=SHR,DSN=MY.DSN.ID(IMSZCAP)
```

Specifying the INPUTCSV statement and supplying CPU consumption data for zCAP workloads in the appropriate format in the INPUTCSV data sets is the method to request that SCRT use the zCAP algorithm whenever all of the zCAP conditions have been met. See [“INPUTCSV DD statement for zCAP”](#) on page 234 for a detailed description and restrictions on its use.

In the sample configuration, the following zCAP workload data is stored in MY.DSN.ID(CICSZCAP), which is the first data set referenced on the sample INPUTCSV DD statement:

```
"Product-zCAP", "5655-Y04"
"ONLYZCAPALL", "2964-12345:LPAR1"
```

The CICSZCAP member contains the ONLYZCAPALL token. CICS TS for z/OS V5 (5655-Y04) is only processing a zCAP new workload on LPAR1; it is not processing any traditional workloads. Notice that LPAR2 is not specified. In this example, CICS TS for z/OS V5 is running on LPAR2, but is only processing traditional workloads. Thus, it is not within the scope of the zCAP data collection requirement.

In the sample configuration, the following zCAP workload data is stored in MY.DSN.ID(IMSZCAP), which is the second data set referenced on the sample INPUTCSV DD statement:

```
"Product-zCAP", "5635-A04"
"Date", "2964-12345:LPAR1", "2964-12345:LPAR2"
"2015-05-04T14:00:00", 900.45, "ONLYZCAP"
"2015-05-04T15:00:00", 1300.00, "ONLYZCAP"
"2015-05-04T16:00:00", 1600.78, "ONLYZCAP"
"2015-05-04T17:00:00", 1450.53, "ONLYZCAP"
```

In the sample environment, IMS V13 (5635-A04) is processing a mix of traditional and zCAP workloads on LPAR1. For LPAR1, the IMSZCAP member contains the general purpose processor CPU seconds used by IMS (the zCAP defining program) in support of the new application. On LPAR2, IMS is processing only zCAP workloads, so the IMSZCAP member specifies the ONLYZCAP token in place of general purpose

processor CPU seconds. This instructs SCRT to retrieve the usage data from the SMF type 89 records for use in zCAP calculations for LPAR2.

For more information about the format of the INPUTCSV data set, see [“Format of the INPUTCSV input data for IBM Z Collocated Application Pricing” on page 235](#).

SPECIAL DD control statement

The following SCRT control statement was specified in the SPECIAL DD to request that SCRT generate the Mobile/zCAP Detailed Data section (V5) of the sub-capacity report. This control statement has no effect on whether SCRT uses the zCAP algorithm to apply reductions based on the data supplied via the INPUTCSV DD statement.

```
Generate_Detailed_Data
```

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

Sample sub-capacity report for IBM Z Collocated Application Pricing

The sub-capacity report for IBM Z Collocated Application Pricing has the same basic format and report sections as the sub-capacity report for a base z/OS system. However, the sub-capacity report for zCAP always contains a footnote to indicate the case where data was provided for zCAP defining programs.

The following sample report sections describe the differences between a sub-capacity report for a base z/OS system and a sub-capacity report that uses IBM Z Collocated Application Pricing. For a complete description of all of the fields appearing in a base sub-capacity report see [“Sample sub-capacity report for single-version z/OS native systems” on page 105](#).

Billing-related section

The billing-related section of the sub-capacity report displays the customer information, tool information, special conditions, product summary information, and detail data collection information, which are described in [“Billing-related section” on page 105](#).

[Figure 105 on page 240](#) shows an example of the Customer Information section and the Tool Information section.

```

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time              10 Jul 2024 - 10:08
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter:  800-555-0123

Customer Name              ABC Corporation
Customer Number            US-S000000000
Machine Serial Number      02-12345
Machine Type and Model     2964-796
Machine Rated Capacity (MSUs) 9783
Purchase Order Number      (optional)
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust
the billable MSUs in inventory for all MLC Products listed under the MLC Product Name
column on this report. In accordance with our agreement, IBM will treat a change in product
licensed capacity as an order. If the MSUs have changed since the last report, software billing
based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to
bill those IPLA products listed under the IPLA Product Name column in this report which
exceed your entitled capacity. In accordance with our agreement, IBM will treat the use
of a product in excess of its entitled capacity as an order and you will be billed for
the amount in excess of your entitlement.

==C5=====
TOOL INFORMATION

Tool Release              30.1.0
Reporting Period          2 May, 2024 - 1 Jun, 2024 inclusive (31 days)

```

Figure 105. Sample sub-capacity report for IBM Z Collocated Application Pricing: Customer information and Tool Information sections

Customer Information section

There are no differences between the Customer Information section for a basic z/OS sub-capacity report and for a sub-capacity report generated using zCAP. For details about this section, see [“Customer information” on page 105](#).

The sample Customer Information section appears in [Figure 105 on page 240](#)

Tool Information section

There are no differences between the Tool Information section for a basic z/OS sub-capacity report and for a sub-capacity report generated using zCAP. For details about this section, see [“Tool Information section” on page 107](#).

The sample Tool Information section appears in [Figure 105 on page 240](#)

Special Conditions section

One of the major differences between the Special Conditions section for a basic z/OS sub-capacity report and for a sub-capacity report generated using zCAP is that the latter will contain a Special Conditions section when the ONLYZCAP or ONLYZCAPALL tokens are specified, as described in [“Format of the INPUTCSV input data for IBM Z Collocated Application Pricing” on page 235](#).

Note: The special condition message that is currently generated uses the keyword ONLYICAP instead of ONLYZCAP. IBM intends to update this message to say ONLYZCAP in a future release.

For general details about the Special Conditions section, see [“Special Conditions section” on page 107](#).

The sample Special Conditions section appears in [Figure 106 on page 241](#).

```

==D5=====
SPECIAL CONDITIONS
IBM Collocated Application Pricing ONLYICAP parameter specified  CICS      LPAR1
IBM Collocated Application Pricing ONLYICAP parameter specified  IMS       LPAR2

```

Figure 106. Sample sub-capacity report for IBM Z Collocated Application Pricing: Special Conditions section

Product Summary Information section

An optional Footnotes column header appears in the MLC section whenever SCRT needs to post a footnote for an zCAP defining program listed in this section of the report. The Footnotes heading does not appear whenever there are no footnotes for this section.

There are no differences between the MLC Product Name, MLC Product ID, Customer Comment, IPLA Product Name, IPLA Product ID, and Customer MSUs fields on a basic z/OS sub-capacity report and on a sub-capacity report generated using zCAP. For a description of these fields see [“Product Summary Information section”](#) on page 110.

The sample Product Summary Information section appears in [Figure 107 on page 242](#). The following fields differ for zCAP reporting:

Tool MSUs

The values listed in the **Tool MSUs** column represent the peak sub-capacity values calculated by SCRT for each product using the appropriate sub-capacity rules for that product. SCRT uses the zCAP algorithm and the supplied zCAP defining program CPU time (as described in [“Overview of IBM Z Collocated Application Pricing”](#) on page 229) to adjust the 4-hour rolling average sub-capacity MSUs for sub-capacity eligible programs on a given CPC.

See [“Product Max Contributors section”](#) on page 243 for a more detailed discussion of how the zCAP algorithm affects the sub-capacity programs that are listed in this section of the sub-capacity report. For additional general information about **Tool MSUs**, see [“Product Summary Information section”](#) on page 110.

Footnotes

SCRT posts footnote (c) for each zCAP defining program for which zCAP workload data was supplied via the INPUTCSV DD statement. This occurs for either the ONLYZCAP or ONLYZCAPALL token or for zCAP general purpose processor CPU seconds in support of a new application.

SCRT also posts footnote (o) for each zCAP defining product for which the ONLYZCAP or ONLYZCAPALL token was supplied via the INPUTCSV DD statement.

The **Tool MSUs** values are partially based on the zCAP algorithm whenever the posted value for the product is combined from multiple LPARs where the product’s value is based on the traditional sub-capacity values from some LPARs and on the zCAP adjusted values from other LPARs.

Whenever SCRT posts footnote (c) or (o) in the **Footnotes** column, it also generates a row with a footnote header followed by a description for the footnote.

The Product Summary Information section in [Figure 107 on page 242](#) uses footnote (c) and (o) to indicate that data was specified on the INPUTCSV DD statement for zCAP defining programs CICS TS V5 and IMS V13. Both footnotes are used in the sample because each INPUTCSV data set contains either the ONLYZCAP or ONLYZCAPALL token.

SCRT also posts footnote (iv) whenever the **Tool MSUs** value posted for a product is based totally or partially on having been adjusted by the Getting Started Sub-Capacity Pricing algorithm. See [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),”](#) on page 165 for a detailed description of Getting Started Sub-Capacity Pricing.


```

=====
PRODUCT SUMMARY INFORMATION

```

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)	Footnotes
z/OS V2	5650-Z0S	3429			
DB2 11 for z/OS	5615-DB2	3335			
CICS TS for z/OS V5	5655-Y04	836			(c) (o)
IBM MQ for z/OS V8	5655-W97	3335			
IMS V13	5635-A04	2687			(c) (o)
IBM Enterprise Cobol for z/OS V4	5655-S71	3335	(optional)	(conditional)	
IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)	Footnotes
File Export for z/OS V1	5697-I12	3335	(optional)	(conditional)	
IBM WebSphere Application Server for z/OS V8	5655-W65	21			(iv)
Migration Utility for z/OS V4	5655-MGU	3335	(optional)	(conditional)	
IBM Multi-site Workload Lifeline V2	5655-UM4	2597			

Footnotes:

- (iv) Product's Tool MSUs are based on Getting Started Sub-capacity Pricing from one or more LPARs
(c) IBM Collocated Application Pricing transaction data from this product was included from one or more LPARs
(o) IBM Collocated Application Pricing ONLYICAP parameter specified as reported in section D5

Figure 107. Sample sub-capacity report for IBM Z Collocated Application Pricing: Product Summary Information section

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the Product Summary Information section.

SMF / SCRT89 Input Data Statistics section

There are no differences between the SMF / SCRT89 Input Data Statistics section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the zCAP algorithm. For details about this section, see [“SMF / SCRT89 Input Data Statistics section” on page 118](#).

The sample SMF / SCRT89 Input Data Statistics section appears in [Figure 108 on page 242](#).

```

=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
LPAR1	SYS1	04 May 2024 - 12:59	04 May 2024 - 17:59
LPAR2	SYS2	04 May 2024 - 12:59	04 May 2024 - 17:59
CPC		04 May 2024 - 12:59	04 May 2024 - 17:59

Figure 108. Sample sub-capacity report for IBM Z Collocated Application Pricing: SMF / SCRT89 Input Data Statistics section

Detail LPAR Data Section

There are no differences between the Detail LPAR Data Section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the zCAP algorithm. For details about this section, see [“Detail LPAR Data section” on page 119](#).

The sample Detail LPAR Data Section appears in [Figure 109 on page 243](#).


```

=====
==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR1	2701	1	04 May 2024 - 14:00	2244	1	04 May 2024 - 15:00
LPAR2	822	1	04 May 2024 - 14:00	803	1	04 May 2024 - 15:00
CPC	3523	1	04 May 2024 - 14:00	3047	1	04 May 2024 - 15:00

Figure 109. Sample sub-capacity report for IBM Z Collocated Application Pricing: Detail LPAR Data section

Product Max Contributors section

The MSU values that are listed in the Product Max Contributors section are the same for a basic z/OS sub-capacity report and for a Getting Started Sub-Capacity Pricing sub-capacity report, with one subtle change. While all of the MSU values listed in this section continue to reflect the product MSUs by LPAR or LPAR Group for the earliest hour for the product's maximum MSU value, the MSU values posted for Getting Started Sub-Capacity Pricing eligible products are based on the Getting Started Sub-Capacity Pricing algorithm whenever that value is less than the traditional sub-capacity value for that product for that hour. The result is that some of the LPAR values posted for Getting Started Sub-Capacity Pricing eligible products are less than the MSUs posted for z/OS, even when the Getting Started Sub-Capacity Pricing eligible product runs in all of the same LPARs and at the same times as z/OS.

Note: The MSUs for Getting Started Sub-Capacity Pricing eligible products are posted under an LPAR group whenever the LPAR group cap is less than the Getting Started Sub-Capacity Pricing value for the earliest hour with the maximum product value. That is not the case for this sample configuration.

This section provides an additional layer of detail behind the sub-capacity values reported for the products listed in the Product Summary Information section by reporting the date/time that the sub-capacity value occurred together with the product's sub-capacity contribution from each LPAR in which that product was running at that date and time. When the *same* peak sub-capacity value occurs for multiple hours, the Product Max Contributors section lists the data for the *earliest* hour. The impact of SCRT reporting a product's sub-capacity value based on the earliest peak hour whenever that peak value occurs for multiple hours is important enough to reiterate in the following explanations because the zCAP conditions for the other peak hours can differ from the zCAP conditions for the earliest peak hour.

Given the level of detail in this section, the sample Product Max Contributors section shown in [Figure 110 on page 244](#) will be used to explain IBM Z Collocated Application Pricing in more detail. However, this explanation assumes you have a general understanding of the contents and use of each field in the Product Max Contributors section, as described in [“Product Max Contributors section” on page 121](#).

```

=====
==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR2	zCAP MSU Reduction
z/OS V2	5650-ZOS	3429	04 May 2024 - 14:00	2649	780	94
DB2 11 for z/OS	5615-DB2	3335	04 May 2024 - 14:00	2597	738	188
CICS TS for z/OS V5	5655-Y04	836	04 May 2024 - 14:00	98	738	90
IBM MQ for z/OS V8	5655-W97	3335	04 May 2024 - 14:00	2597	738	188
IMS V13	5635-A04	2687	04 May 2024 - 14:00	2603	84	98
IBM Enterprise Cobol for z/OS V4	5655-S71	3335	04 May 2024 - 14:00	2597	738	188
File Export for z/OS V1	5697-I12	3335	04 May 2024 - 14:00	2597	738	188
IBM WebSphere Application Server for z/OS V8	5655-W65	21	04 May 2024 - 17:00	21	0	406
Migration Utility for z/OS V4	5655-MGU	3335	04 May 2024 - 14:00	2597	738	188
IBM Multi-site Workload Lifeline V2	5655-UM4	2597	04 May 2024 - 14:00	2597	0	104

zCAP Defining Program MSUs	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR2
IMS V13	5635-A04	301	04 May 2024 - 17:00	37	264
CICS TS for z/OS V5	5655-Y04	369	04 May 2024 - 17:00	369	0

Figure 110. Sample sub-capacity report for IBM Z Collocated Application Pricing: Product Max Contributors section

How SCRT implements IBM Z Collocated Application Pricing

It is important to understand that SCRT uses the appropriate algorithm to calculate the sub-capacity value for each product for each hour for each LPAR, independently of what is happening with that product on other LPARs or guest systems during that same hour. Then, SCRT sums each of the product's sub-capacity values across all of the LPARs for the same hour to calculate a sub-capacity value for the product for the CPC for that hour. After performing these sub-capacity calculations for the product across the CPC for all of the hours during the reporting period, SCRT selects the earliest hour across the CPC with the product's peak sub-capacity value and reports the date and time for that hour, along with the individual sub-capacity values for the product for each LPAR for that same hour, in the Product Max Contributors section.

In the sample Product Max Contributors section, the peak z/OS sub-capacity value across the CPC was 3429 MSUs and that peak first occurred on 04 May 2024 at 14:00. The z/OS product running in LPAR1 and LPAR2 contributed 2649 and 780 MSUs, respectively, to the z/OS peak sub-capacity value of 3429 MSUs for the CPC.

One key difference between a basic z/OS sub-capacity report and a sub-capacity report generated using the zCAP algorithm is the **zCAP MSU Reduction** column in the Product Max Contributors section. This column is generated whenever zCAP workload data was provided via the INPUTCSV DD statement, as was the case in the configuration (see [“Sample configuration for IBM Z Collocated Application Pricing” on page 237](#)).

The **zCAP MSU Reduction** column shows the adjustment that SCRT applied to products' sub-capacity values using the zCAP algorithm. In the Product Max Contributors section for the sample configuration, the peak z/OS sub-capacity value and the value of all products based entirely on the 4-hour rolling average MSU utilization of the LPAR were adjusted by SCRT using the value shown in the **zCAP MSU Reduction** column.

In the Product Max Contributors section for the sample configuration, note that the MSU values shown for LPAR SYS1 and LPAR SYS2 already reflect the contribution of zCAP defining programs that ran in each LPAR during the peak hour. Also note that the **zCAP MSU Reduction** column reflects the total reduction across all LPARs for that hour.

Another key difference between the basic z/OS sub-capacity report and a sub-capacity report generated using the zCAP algorithm is that an additional section of the Product Max Contributors section is generated. This new section begins with the **zCAP Defining Program MSUs** column and contains the same set of fields as in the normal Product Max Contributors section. The content of the fields differs in the following ways:

Highest

Highest combined zCAP MSU reduction for the zCAP defining program, applied for all LPARs in which the product ran as a zCAP defining program.

Date/Time

Date and time when the highest zCAP MSU reduction was first reached.

LPAR names and zCAP values

Names of the LPARs that contributed to the zCAP MSU reduction and the amount of the contribution from each, in MSUs.

In the sample configuration, the highest combined IMS V13 zCAP MSU reduction was 301 MSUs, which occurred on 04 May 2024 at 17:00. LPAR1 contributed 37 MSUs and LPAR2 contributed 264 MSUs to the total.

In the sample configuration, the highest combined CICS TS for z/OS V5 MSU reduction was 369 MSUs, which occurred on 04 May 2024 at 17:00. Only LPAR1 contributed to the zCAP reduction. This is because, in this sample configuration, CICS TS for z/OS V5 is not running a zCAP eligible workload on LPAR2.

See “[Mobile/zCAP Pricing Detailed Data section for zCAP](#)” on page 246 for more information about the reporting results for the sample configuration.

A note about zCAP adjustments: The MSU values used for zCAP adjustments are derived from either the zCAP defining program’s SMF type 89 records or the zCAP new application general purpose processor time values reported using the INPUTCSV DD statement. These adjustments are calculated internally by SCRT as specified in the IBM Addendum for IBM Z Collocated Application Pricing. The Product Max Contributors section is only intended to show the contribution of each LPAR to the specified product's peak sub-capacity value for earliest peak hour.

Product Grid Snapshot section

There are no differences between the Product Grid Snapshot section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the zCAP algorithm. For details about this section, see “[Product Grid Snapshot section](#)” on page 123.

The sample Product Grid Snapshot section appears in [Figure 111 on page 245](#).

```
==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                                Product ID    SYS1    SYS2
z/OS V2                                     5650-Z05      0.5%    0.5%
DB2 11 for z/OS                             5615-DB2      0.5%    0.5%
CICS TS for z/OS V5                         5655-Y04      0.5%    0.5%
IBM MQ for z/OS V8                          5655-W97      0.5%    0.5%
IMS V13                                     5635-A04      0.5%    0.5%
IBM WebSphere Application Server for z/OS V8 5655-W65      0.5%
IBM Multi-site Workload Lifeline V2         5655-UM4      0.5%
```

Figure 111. Sample sub-capacity report for Integrated Workload Pricing: Product Grid Snapshot section

Getting Started Sub-Capacity Pricing Assigned Hours section

There are no differences between the Getting Started Sub-Capacity Pricing Assigned Hours section for a basic z/OS sub-capacity report and for a sub-capacity report generated using the zCAP algorithm. For details about this section, see “[Getting Started Sub-Capacity Pricing Assigned section](#)” on page 179.

The sample Getting Started Sub-Capacity Pricing Assigned Hours section appears in [Figure 112 on page 246](#).

```

==S7=====
GETTING STARTED SUB-CAPACITY PRICING ASSIGNED - HOURS
Product Name                                Product ID    LPAR    LPAR
                                SYS1      SYS2
IBM WebSphere Application Server for z/OS V8  5655-W65      4

```

Figure 112. Sample sub-capacity report for IBM Z Collocated Application Pricing: Getting Started Sub-Capacity Pricing Assigned Hours section

Mobile/zCAP Pricing Detailed Data section for zCAP

The optional Mobile/zCAP Pricing Detailed Data section appears when the optional **Generate_Detailed_Data** SPECIAL DD control statement is specified. This section shows the zCAP MSU reduction applied by SCRT for each zCAP defining program, per LPAR. This information is shown for each sub-capacity product on an hour-by-hour basis.

As shown in [Figure 113 on page 247](#), the optional Mobile/zCAP Pricing Detailed Data section reports several columns for each product:

- Date and time of the adjusted interval.
- CPC Type-model.
- CPC rated machine capacity in MSU.
- Adjusted MSU values for each LPAR. (Note that the zCAP algorithm has already been applied.)
- For each zCAP defining product with INPUTCSV DD data specified, the zCAP reduction applied per defining product, per partition.
- Total zCAP MSU reduction applied across all zCAP defining products.
- CPC Total MSU, which is the sum of the adjusted LPAR MSU values across all partitions for the adjusted product.

There are several additional values that can appear in section V5 of a sub-capacity report:

–

Indicates that there is no data for this interval.

dd(rr)

Indicates that the defined capacity limit for the LPAR was reached during this interval. The first number (*dd*) represents the defined capacity value used for this hour. The second number (*rr*) represents the 4-hour rolling average value prior to the capping.

G(rr)

Indicates that the sum of the partitions exceeded the group capacity limit; the group capacity was used for this hour for that set of LPARs. The *rr* value represents the 4-hour rolling average value prior to the capping.

P(rr)

Indicates that the sum of the partitions exceeded the processor capacity; the processor capacity was used. The *rr* value represents the 4-hour rolling average value prior to the capping.

(GSSP)

Indicates that a WebSphere program eligible for Getting Started Sub-Capacity Pricing received the GSSP value so there was no additional adjustment.

nnn(CPS)

Indicates that an LPAR was assigned to a Tailored Fit Pricing solution for this hour. In this case, the LPAR does not contribute to the product sub-capacity value. The *nnn* value represents what the LPAR sub-capacity value would have been if it were part of the standard (non-container) environment.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being

generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

The following descriptions of the IBM Z Collocated Application Pricing algorithm are based on the sample configuration in [“Sample configuration for IBM Z Collocated Application Pricing”](#) on page 237. Note that only a subset of the products running in the environment are evaluated.

Results for z/OS V2R1

After SCRT applied the appropriate algorithm hour by hour, LPAR by LPAR, and then summed the sub-capacity values for the same hour across all of the z/OS LPARs, SCRT reported 04 May 2024 at 14:00 as the earliest hour with the peak sub-capacity value for z/OS, as indicated in the sample Product Max Contributors section.

For 04 May 2024 at 14:00, SCRT applied the appropriate algorithms and calculated a sub-capacity value of 2649 MSUs for LPAR1. As shown in the **5635-A04 zCAP Partition** data column for LPAR1, this includes an 3 MSU zCAP reduction derived from the supplied INPUTCSV zCAP defining program general purpose processor CPU time in support of the new application workload, as specified as input in [“Sample configuration for IBM Z Collocated Application Pricing”](#) on page 237.

On LPAR1, the sub-capacity value also reflects a 49 MSU reduction for zCAP defining program 5655-Y04 (CICS TS for z/OS V5). This is derived from the usage data from the SMF type 89 records for CICS TS for z/OS V5 because the ONLYZCAPALL token was specified for LPAR1.

For LPAR2, SCRT applied the appropriate algorithms and calculated a sub-capacity value of 780 MSUs for 04 May 2015 at 14:00. As shown in the **5635-A04 zCAP Partition** data column for LPAR2, this includes a 42 MSU zCAP reduction derived from the usage data from the SMF type 89 records for IMS V13. The SMF type 89 records for IMS V13 are used on LPAR2 because the ONLYZCAP token was specified in the IMS V13 INPUTCSV data set, as specified in [“Sample configuration for IBM Z Collocated Application Pricing”](#) on page 237.

The total zCAP MSU reduction for both LPAR1 and LPAR2 on 04 May 2015 at 14:00 is 94 MSUs.

Figure 113 on page 247 illustrates these results.

=====V5=====									
MOBILE/ZCAP PRICING DETAILED DATA									
MLC Products									
5650-Z0S		z/OS V2							
Date Time of RMF Interval	Type	CPC MSU	Partition LPAR1	LPAR2	5635-A04 zCAP Partition LPAR1	LPAR2	5655-Y04 zCAP Partition LPAR1	Total zCAP MSU Reduction	CPC Total
04 May 2024 - 14:00	2964-796	9783	2649	780	3	42	49	94	3429
04 May 2024 - 15:00	2964-796	9783	2140	719	8	84	96	188	2859
04 May 2024 - 16:00	2964-796	9783	1709	628	13	128	161	302	2337
04 May 2024 - 17:00	2964-796	9783	1499	554	19	132	185	336	2053

Figure 113. Sample sub-capacity report for IBM Z Collocated Application Pricing: Mobile/zCap Pricing Detailed Data section – z/OS V2R1 results

Results for IBM MQ for z/OS V8

SCRT applied the process as described in [“Results for z/OS V2R1”](#) on page 247. Like z/OS, IBM MQ for z/OS V8 is based on the 4-hour rolling average MSU utilization of the LPAR.

Since IBM MQ for z/OS V8 ran on both z/OS LPARs during 04 May 2024 at 14:00, the application of the zCAP algorithm is similar to that previously described for z/OS. However, the allowed zCAP MSU reduction is greater than that for z/OS. As stated in the IBM Z Collocated Application Pricing terms and conditions,

up to 50% of the new application general purpose processor CPU time is removed from z/OS, but 100% of the new application general purpose processor CPU time is removed from other sub-capacity products.

A similar process applies for all sub-capacity programs that reported a peak MSU value of 3335 with 04 May 2024 at 14:00 as the earliest hour with the peak sub-capacity value.

Figure 114 on page 248 illustrates the results.

5655-W97		IBM MQ for z/OS V8				5635-A04 zCAP		5655-Y04 zCAP		Total	
Date	Time of RMF Interval	Type	CPC MSU	Partition LPAR1	LPAR2	Partition LPAR1	LPAR2	Partition LPAR1	LPAR2	zCAP MSU Reduction	CPC Total
04 May 2024	- 14:00	2964-796	9783	2597	738	6	84	98		188	3335
04 May 2024	- 15:00	2964-796	9783	2036	636	16	167	192		375	2672
04 May 2024	- 16:00	2964-796	9783	1535	500	27	256	321		604	2035
04 May 2024	- 17:00	2964-796	9783	1297	422	37	264	369		670	1719

Figure 114. Sample sub-capacity report for IBM Z Collocated Application Pricing: Mobile/zCap Pricing Detailed Data section — IBM MQ for z/OS V8 results

Results for CICS TS for z/OS V5

Since CICS TS for z/OS V5 is a zCAP defining program on LPAR1 in the sample configuration, SCRT applies a different algorithm on LPAR1 than on LPAR2.

For 04 May 2015 at 14:00, SCRT calculated a sub-capacity value of 98 MSUs for LPAR1. This is based on the usage data from the SMF type 89 records, not the LPAR utilization value, as specified in the IBM Addendum for IBM Z Collocated Application Pricing.

On LPAR2, CICS TS for z/OS V5 is not a zCAP defining program. Thus, its utilization is based on the 4-hour rolling average MSU utilization of the LPAR. SCRT calculated a peak sub-capacity value of 738 MSUs for LPAR2.

Since CICS TS for z/OS V5 is one of the zCAP defining programs, it does not receive the MSU reduction with the zCAP workload exploiting it. CICS TS for z/OS V5 does receive the IMS V13 (5635-A04) zCAP MSU reduction, resulting in a total 90 MSU zCAP reduction.

This results in a peak MSU value of 836 MSUs for 04 May 2015 at 14:00.

Figure 115 on page 248 illustrates these results.

5655-Y04		CICS TS for z/OS V5				5635-A04 zCAP		5655-Y04 zCAP		Total	
Date	Time of RMF Interval	Type	CPC MSU	Partition LPAR1	LPAR2	Partition LPAR1	LPAR2	Partition LPAR1	LPAR2	zCAP MSU Reduction	CPC Total
04 May 2024	- 14:00	2964-796	9783	98	738	6	84	0		90	836
04 May 2024	- 15:00	2964-796	9783	192	636	16	167	0		183	828
04 May 2024	- 16:00	2964-796	9783	321	500	27	256	0		283	821
04 May 2024	- 17:00	2964-796	9783	369	422	37	264	0		301	791

Figure 115. Sample sub-capacity report for IBM Z Collocated Application Pricing: Mobile/zCap Pricing Detailed Data section — CICS TS for z/OS V5 results

Results for IMS V13

Since IMS V13 is a zCAP defining program on both LPAR1 and LPAR2 in the sample configuration, SCRT applies the zCAP algorithm on both LPARs.

For 04 May 2015 at 14:00, SCRT calculated a sub-capacity value of 2603 MSUs for LPAR1. This is based on the LPAR utilization value, as specified in the IBM Addendum for IBM Z Collocated Application Pricing. The LPAR utilization value is used for LPAR1 because IMS V13 is running both zCAP and traditional sub-capacity workloads.

On LPAR2, IMS V13 is also a zCAP defining program. However, no traditional sub-capacity workloads are using IMS V13. Therefore, its utilization is based on the usage data from the SMF type 89 records, not the LPAR utilization value, as specified in the IBM Addendum for IBM Z Collocated Application Pricing. SCRT calculates a peak of 84 MSUs on LPAR2.

Since IMS V13 is one of the zCAP defining programs, it does not receive the MSU reduction associated with the zCAP workload exploiting it. IMS V13 does receive the CICS TS for z/OS V5 (5655-Y04) zCAP MSU reduction, resulting in a total 98 MSU zCAP reduction.

This results in a peak MSU value of 2867 MSUs for 04 May 2015 at 14:00.

Figure 116 on page 249 illustrates these results.

5635-A04		IMS V13				5635-A04 zCAP		5655-Y04 zCAP		Total zCAP MSU Reduction	CPC Total
Date	Time of RMF Interval	Type	CPC MSU	Partition LPAR1	LPAR2	Partition LPAR1	LPAR2	Partition LPAR1			
04 May 2024	- 14:00	2964-796	9783	2603	84	0	0	98		98	2687
04 May 2024	- 15:00	2964-796	9783	2052	167	0	0	192		192	2219
04 May 2024	- 16:00	2964-796	9783	1562	256	0	0	321		321	1818
04 May 2024	- 17:00	2964-796	9783	1334	264	0	0	369		369	1598

Figure 116. Sample sub-capacity report for IBM Z Collocated Application Pricing: Mobile/zCap Pricing Detailed Data section — IMS V13 results

Results for IBM Multi-site Workload Lifeline V2

SCRT applies the process as described in “Results for z/OS V2R1” on page 247 and “Results for IBM MQ for z/OS V8” on page 247. Like those products, IBM Multi-site Workload Lifeline V2 is based on the 4-hour rolling average MSU utilization of the LPAR.

However, IBM Multi-site Workload Lifeline V2 is only running in partition LPAR1. Thus, it is not eligible for any reduction applied on LPAR2. As a result, SCRT calculated an MSU value of 2597 for IBM Multi-site Workload Lifeline V2, as illustrated in Figure 117 on page 249.

5655-UM4		IBM Multi-site Workload Lifeline V2				5635-A04 zCAP		5655-Y04 zCAP		Total zCAP MSU Reduction	CPC Total
Date	Time of RMF Interval	Type	CPC MSU	Partition LPAR1		Partition LPAR1		Partition LPAR1			
04 May 2024	- 14:00	2964-796	9783	2597		6		98		104	2597
04 May 2024	- 15:00	2964-796	9783	2036		16		192		208	2036
04 May 2024	- 16:00	2964-796	9783	1535		27		321		348	1535
04 May 2024	- 17:00	2964-796	9783	1297		37		369		406	1297

Figure 117. Sample sub-capacity report for IBM Z Collocated Application Pricing: Mobile/zCap Pricing Detailed Data section — IBM Multi-site Workload Lifeline V2 results

Results for IBM WebSphere Application Server for z/OS V8

SCRT applied the Getting Started Sub-Capacity Pricing for z/OS IPLA Software process to IBM WebSphere Application Server for z/OS V8.

Since the terms and conditions of GSSP were met in the sample configuration, and the GSSP values were less than the traditional sub-capacity values for each hour of the reporting period, SCRT assigned IBM WebSphere Application Server for z/OS V8 the GSSP value for each hour in the reporting period.

When the GSSP value is used for a GSSP eligible product, the Mobile/zCAP Pricing Detailed Data section places the (GSSP) token in the zCAP defining program partition reduction column, as shown in [Figure 118 on page 250](#). In accordance with the IBM Addendum for IBM Z Collocated Application Pricing, the zCAP reduction is not applied when GSSP pricing is used instead of the traditional sub-capacity value.

In the sample configuration, the peak MSU value used for IBM WebSphere Application Server for z/OS V8 is 21, which is achieved during the 04 May 2024 17:00 hour.

For information about the GSSP process, see [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#).

5655-W65		IBM WebSphere Application Server for z/OS V8				5635-A04 zCAP	5655-Y04 zCAP	Total	
Date Time of RMF Interval	Type	CPC MSU	Partition LPAR1	Partition LPAR1	Partition LPAR1	Partition LPAR1	Partition LPAR1	zCAP MSU Reduction	CPC Total
04 May 2024 - 14:00	2964-796	9783	5	(GSSP)	(GSSP)	0	0	5	
04 May 2024 - 15:00	2964-796	9783	11	(GSSP)	(GSSP)	0	0	11	
04 May 2024 - 16:00	2964-796	9783	16	(GSSP)	(GSSP)	0	0	16	
04 May 2024 - 17:00	2964-796	9783	21	(GSSP)	(GSSP)	0	0	21	

Figure 118. Sample sub-capacity report for IBM Z Collocated Application Pricing: Mobile/zCap Pricing Detailed Data section — IBM WebSphere Application Server for z/OS V8 results

Chapter 12. Sub-capacity reporting for IBM Z Workload Pricing for Cloud (zWPC)

IBM Z Workload Pricing for Cloud (zWPC) is a sub-capacity pricing option under Advanced Workload License Charges (AWLC), Advanced Entry Workload License Charges (AEWLC), System z New Application License Charges (zNALC), or Country Multiplex License Charges (CMLC). It is intended to enable qualifying z/OS customers to mitigate the cost of growth for public cloud transactions processed by select zWPC defining programs, specifically in cases where the higher public cloud transaction volumes may cause a spike in machine utilization (as described in US Software Announcement Letter 216-319). This sub-capacity pricing option is effective for SMF data collected on 2 October 2016 and later, provided that all of the required zWPC conditions have been met.

Overview of IBM Z Workload Pricing for Cloud

IBM Z Workload Pricing for Cloud (zWPC) provides an enhanced way of reporting z/OS system utilization, which can improve price/performance for sub-capacity eligible programs when running in the same LPAR as new public cloud workloads processed by select zWPC defining programs. Public cloud workloads are defined as the transactions processed by a named public cloud offering, such as IBM Bluemix, Amazon Web Services, Microsoft Azure, and more.

Table 40 on page 251 lists the zWPC defining programs supported by SCRT V24.10.00 and later. These are the products whose sub-capacity pricing values will be reduced when the required set of zWPC conditions are met.

Note: For the most current list of zWPC defining programs, see [IBM Z Workload Pricing for Cloud \(zWPC\) Defining Programs \(www.ibm.com/downloads/documents/us-en/107a02e95b48f614\)](http://www.ibm.com/downloads/documents/us-en/107a02e95b48f614).

Table 40. zWPC defining programs

Program name	Program ID
CICS Transaction Server for z/OS, V4	5655-S97
CICS Transaction Server for z/OS, V5	5655-Y04
CICS Value Unit Edition (VUE) V5	5722-DFJ
IBM DB2 V9 for z/OS	5635-DB2
DB2 V9 for z/OS	5635-DB2
Db2 10 for z/OS	5605-DB2
Db2 11 for z/OS	5615-DB2
Db2 12 for z/OS	5650-DB2
DB2 VUE V9	5697-P12
Db2 10 VUE	5697-P31
Db2 11 VUE	5697-P43
Db2 12 VUE	5770-AF3
IMS V11	5635-A02
IMS V12	5635-A03
IMS V13	5635-A04
IMS V14	5635-A05
IMS DB VUE V12	5655-DSQ
IMS DB VUE V13	5655-DSM
IMS DB VUE V14	5655-DSE

Table 40. zWPC defining programs (continued)

Program name	Program ID
IMS TM VUE V12	5655-TM1
IMS TM VUE V12	5655-TM2
IMS TM VUE V12	5655-TM3
WebSphere MQ for z/OS, V6	5655-L82
WebSphere MQ for z/OS, V7	5655-R36
IBM MQ for z/OS, V8	5655-W97
IBM MQ for z/OS V9	5655-MQ9
WebSphere MQ VUE V7	5655-VUE
IBM MQ VUE V8	5655-VU8
IBM MQ VUE V9	5655-VU9
WebSphere Application Server for z/OS, V7	5655-N02
WebSphere Application Server for z/OS, V8	5655-W65
WebSphere Application Server for z/OS, V9	5655-WAS

IBM Z Workload Pricing for Cloud requires that you identify public cloud transactions. You can use two different mechanisms to accomplish this:

- zWPC reporting using an INPUTCSV file
- zWPC reporting using z/OS workload management (WLM) workload classification

Note: You can use both the input file method and the WLM workload classification method to pass public cloud transaction input data to SCRT in the same run. However, you must ensure that the data passed by each method represents data for a unique public cloud workload; you cannot pass data for the same public cloud workload using both input methods at the same time.

zWPC reporting using an INPUTCSV file

You must be able to track and report the general purpose processor time (CPU time) for public cloud transactions and report those values in a pre-defined format to IBM each month in order to take advantage of this pricing offering.

Whenever the required zWPC conditions have been met, SCRT reduces the sub-capacity value that is reported for all products by subtracting 60 percent of the 4-hour rolling average utilization of the reported public cloud transaction general purpose processor time consumed by a set of products referred to as *zWPC defining programs*, which are listed in [Table 40 on page 251](#).

For the case where a z/OS system runs as a guest under z/VM, each eligible z/OS guest system is treated as its own LPAR when applying any sub-capacity adjustments for zWPC.

SCRT calculates the 4-hour rolling average of the reported public cloud transaction general purpose processor time consumed by the zWPC defining programs and subtracts 60 percent of those values from the traditional sub-capacity MSUs for all sub-capacity eligible programs running in the same LPAR(s) as the public cloud workloads, on an hour-by-hour basis, per LPAR. The program values for the same hour are summed across all of the LPARs in which the program runs to create an adjusted sub-capacity value for the program, for the given CPC, for each hour. SCRT determines the billable MSU peak for a given program on a CPC using the adjusted MSU values.

zWPC reporting using WLM workload classification

The WLM workload classification support allows you to specify within the WLM service definition those transactions that are eligible for IBM Z Workload Pricing for Cloud. Workloads classified as zWPC using

the WLM workload classification support must exploit products in the set of products referred to as zWPC *defining programs*, which are listed in [Table 40](#) on page 251.

For more information about the WLM support for mobile workload classification, see [z/OS MVS Planning: Workload Management](#).

In order to use this offering, you must have the WLM support applied on your z/OS systems and be able to classify public cloud transactions. You must classify all of your public cloud transactions in one of either CATEGORYA or CATEGORYB. You must use the same classification field for all zWPC-eligible transactions in your environment.

Once implemented and activated on a z/OS image, WLM and RMF report the 4-hour rolling average of public cloud transaction CPU consumption for that image in either the SMF70LACA or SMF70LACB field of SMF type 70 records, depending on whether you selected CATEGORYA or CATEGORYB. Because the data is reported in the same SMF type 70 records that SCRT already requires, you need not perform any additional data tracking or collection.

You must use the **Assign zWPC** control statement on the SPECIAL DD to specify to SCRT which category assignment you have chosen for your environment. Whenever the required zWPC conditions are met, SCRT reduces the sub-capacity value that is reported for all products by subtracting 60 percent of the 4-hour rolling average utilization, as reported in the appropriate field of the SMF type 70 records (SMF70LACA or SMF70LACB).

For the case where a z/OS system runs as a guest under z/VM, each eligible z/OS guest system is treated as its own LPAR when applying any sub-capacity adjustments for zWPC.

SCRT subtracts 60 percent of the reported value from the traditional sub-capacity MSUs for all sub-capacity eligible programs running in the same LPARs as the public cloud workloads, on an hour-by-hour basis, per LPAR. The program values for the same hour are summed across all of the LPARs in which the program runs to create an adjusted sub-capacity value for the program, for the given CPC, for each hour.

SCRT uses the adjusted MSU values to determine the billable MSU peak for a given program on the CPC.

Qualifications for IBM Z Workload Pricing for Cloud

The following conditions must be met before IBM will accept and process sub-capacity reports where IBM Z Workload Pricing for Cloud is being reported via footnote (e). Some of these conditions are also described in more detail elsewhere in this chapter but are consolidated here for ease of reference.

1. The CPC from which the SMF data is collected must be licensed under AWLC, AEWLC, zNALC, or CMLC sub-capacity terms and conditions. The zWPC terms and conditions require that defining products be installed on an IBM z13, z13s, zEC12, or zBC12 CPC that has implemented sub-capacity pricing. Alternatively, you may qualify for zWPC by installing authorized zWPC defining programs on an IBM z196 or z114 CPC if a z13, z13s, zEC12, or zBC12 CPC is being operated in your enterprise. Note that the SMF data must be collected from one or more z/OS systems running on zEnterprise or later CPCs; SCRT only applies the zWPC algorithm when processing SMF data collected from zEnterprise or later CPCs.
2. The traditional sub-capacity values of programs are only adjusted when one or more zWPC defining programs (for which data was submitted via the INPUTCSV DD statement) were running in the same z/OS system (LPAR) during the same 4-hour window that was used to determine the traditional sub-capacity value of the zWPC defining program. Programs running in one z/OS system (LPAR) are not affected by zWPC defining programs running during the same hour in a different z/OS system (LPAR).
3. Only SMF data collected on 2 October 2016 and later is eligible for IBM Z Workload Pricing for Cloud, as set forth in the IBM Z Workload Pricing for Cloud announcement letter. Although SCRT will generate sub-capacity reports based on the zWPC algorithm for SMF data collected earlier than 2 October 2016, the IBM business process will reject those reports.
4. You must generate and submit to IBM a monthly sub-capacity report that contains zWPC adjustments during the normal sub-capacity submission window (business as usual).
5. You must use SCRT V24.10.00 or later to generate a sub-capacity report that is submitted for IBM Z Workload Pricing for Cloud.

6. When you intend to provide IBM Z Workload Pricing for Cloud data for a zWPC defining product to SCRT in CSV format, you must use the INPUTCSV DD statement to request that SCRT generate a sub-capacity report using the zWPC algorithms whenever all of the required zWPC conditions are met. The default SCRT processing does not use the zWPC algorithms. For details, see [“INPUTCSV DD statement for zWPC” on page 255](#).

When you provide IBM Z Workload Pricing for Cloud data to SCRT via WLM mobile workload classification, you must use the **Assign zWPC** control statement in the SPECIAL DD to instruct SCRT as to which categorization was used for cloud transaction data (CATEGORYA or CATEGORYB). When you specify Assign zWPC=CATEGORYA or Assign zWPC=CATEGORYB, SCRT uses the zWPC algorithms to generate a sub-capacity report whenever all of the required zWPC conditions are met. The default SCRT processing does not use the zWPC algorithms.

7. You must agree to the IBM Z Workload Pricing for Cloud Addendum prior to submitting a sub-capacity report where footnote (e) appears in the Product Information Summary section, since footnote (e) indicates that the sub-capacity values for one or more products were adjusted based on zWPC. Failure to formally agree to the zWPC Addendum prior to submitting a sub-capacity report with footnote (e) will cause the IBM business process to reject that report. The rejection will contain an explanation and request that you generate a replacement sub-capacity report without using the INPUTCSV DD statement to submit zWPC data, and submit it to IBM in order to continue to receive traditional sub-capacity pricing for the IBM products on that CPC.

Data collection requirements for IBM Z Workload Pricing for Cloud

The data collection requirements for IBM Z Workload Pricing for Cloud are:

- You must collect SMF type 70 and type 89 records for all LPARs on all CPCs each month, and you must retain these records for six months after the billing period for auditing purposes.

You must also collect data about public cloud transactions using the following mechanism:

- Manual public cloud transaction reporting via CSV file input:
 - You must collect the source data (CPU time) for the public cloud transactions that will be used in monthly reporting, and you must retain this data (CPU time and CSV files) for six months after the billing period for auditing purposes.
 - You must preprocess your public cloud transaction data into a predefined format to be loaded into SCRT for each sub-capacity reporting period. The data must consist of general purpose processor CPU seconds for public cloud transactions, summarized by hour, by LPAR, for all CPCs that process public cloud transactions.
- Public cloud transaction reporting via WLM mobile workload classification:
 - You must classify public cloud workloads as described in [z/OS MVS Planning: Workload Management](#).
 - The workload must be classified as either CATEGORYA or CATEGORYB classification in the WLM service class. You can choose either category, but the category must be the same for all CPCs and system images that will be processed by SCRT in a single execution.
 - Workloads that are classified as public cloud using either the CATEGORYA or CATEGORYB classification in the WLM service class do not require any manual aggregation or preprocessing of data.

zWPC monthly reporting process

When using IBM Z Workload Pricing for Cloud, you must complete the following requirements each month:

- For programs that have public cloud transactions tracked manually and provided to SCRT in CSV format:
 1. Track public cloud transactions, including CPU seconds, per program, on an hourly basis, per LPAR.

2. Produce a file in the comma-separated value (CSV) format that contains the mobile CPU consumption data for each program. The format of the data in this file is described in [“Format of the INPUTCSV input data for zWPC”](#) on page 256.
 3. Load the resulting data files along with the SMF type 70 and type 89 records into SCRT.
 4. Run SCRT and submit the resulting sub-capacity reports to IBM for each reporting period.
- For programs that are classified in WLM as a public cloud workload using either CATEGORYA or CATEGORYB classification:
 1. Load the SMF type 70 and type 89 records into SCRT.
 2. Indicate your chosen classification category by specifying one of the following control statements in the SPECIAL DD:
 - Assign ZWPC=CATEGORYA
 - Assign ZWPC=CATEGORYB
 3. Run SCRT and submit the resulting sub-capacity reports to IBM for each reporting period.

Notes:

- This monthly process requires an initial setup action in WLM to classify one or more workloads as public cloud workloads.

SCRT externals for IBM Z Workload Pricing for Cloud

The SCRT externals for IBM Z Workload Pricing for Cloud differ depending on whether you are supplying public cloud workload data using the INPUTCSV DD or via WLM workload classification.

SCRT externals for INPUTCSV-based IBM Z Workload Pricing for Cloud

SCRT provides the following externals for IBM Z Workload Pricing for Cloud using input data in CSV format. Where noted, these externals are specific to zWPC and are treated as optional when zWPC is not being requested or does not apply to the SMF data being processed by SCRT.

- SCRT uses the optional DD statement, INPUTCSV, which is used to provide input CPU consumption data in a predefined, CSV format for IBM Z Workload Pricing for Cloud.
- SCRT uses footnote (e) in the Product Summary Information section of the sub-capacity report to indicate that the **Tool MSUs** value for the corresponding product was adjusted by the zWPC algorithm.
- SCRT uses an optional **Generate_Detailed_Data** control statement to track the application of zWPC and other reductions per hour, per LPAR.

INPUTCSV DD statement for zWPC

The default SCRT processing does not support IBM Z Workload Pricing for Cloud. You must provide your public cloud transaction data as input on the INPUTCSV DD statement. The data in the data set pointed to by the INPUTCSV DD statement must be in a predefined format, as described in [“Format of the INPUTCSV input data for zWPC”](#) on page 256.

Note: IBM authorizes the use of the INPUTCSV DD statement for providing public cloud CPU consumption data to SCRT only when you have signed or otherwise agreed to the terms of the IBM Z Workload Pricing for Cloud Addendum.

For more information about this statement, see [“INPUTCSV DD statement”](#) on page 60.

zWPC workload footnote

For the case where the **Tool MSUs** value for one or more products listed in the Product Summary Information section have been adjusted for zWPC, SCRT applies the following changes to the Product Summary Information section:

- Generates an optional **Footnotes** column header for the MLC section.

- Posts footnote (e) in the **Footnotes** column for each zWPC defining product for which public cloud workload pricing data was supplied via the INPUTCSV DD statement.
- Generates a footnote subsection immediately after the last IPLA product is listed. The footnote subsection contains a footnote header row followed by a description of footnote (e).

Footnote (e) is optional. All three footnote-related items are only generated if one or more zWPC defining programs requires such a footnote.

Generate_Detailed_Data control statement

SCRT supports an optional **Generate_Detailed_Data** SPECIAL control statement. This is used to track the application of zWPC and other reductions per hour, per LPAR.

When specified, this control statement causes the Mobile/zCAP Pricing Detailed Data section (V5) to be generated in the Sub-capacity report.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

Mobile/zCAP Pricing Detailed Data section

The Mobile/zCAP Pricing Detailed Data section (V5) of the sub-capacity report shows, for each product, the hour-by-hour application of zWPC and other applicable reductions.

See [“Sample sub-capacity report for INPUTCSV-based IBM Z Workload Pricing for Cloud” on page 261](#) for a detailed description of the output in the Mobile/zCAP Pricing Detailed Data section.

Format of the INPUTCSV input data for zWPC

For IBM Z Workload Pricing for Cloud, the INPUTCSV DD statement must include the data sets that contain data formatted as described in this topic.

The data contained in each IBM Z Workload Pricing for Cloud input file must consist of general purpose processor CPU seconds for public cloud transactions. These are summarized by hour, by LPAR, for all CPCs that process public cloud transactions.

Each zWPC defining program product family (CICS, Db2, MQ, IMS, and WebSphere) must be contained in a separate CSV data set or data set member.

Note the following points about the INPUTCSV data set:

- Any completely blank row in the file is ignored.
- Any row that starts with a double quotation mark immediately followed by an asterisk ("*) is considered a comment and is ignored.
- The INPUTCSV data uses LPAR names, *not* system IDs (SYSIDs).

Figure 119 on page 256 shows an example of the format for the contents of the mobile transaction CSV file. The field separator character may be a comma (,) or a semicolon (;), but must be consistent throughout.

```
"Product-zWPC", "5655-S97", "5655-M15"
"Date", "2827-23456:SYSF", "2827-23456:SYSG", "2827-23456:SYSH"
"2014-03-02T00:00:00", 3026.23, 4103.88, 229.41
"2014-03-02T01:00:00", 1954.01, 3582.93, 257.33
"2014-03-02T02:00:00", 2444.22, 3876.60, 262.84
```

Figure 119. Example of the contents of a file in the INPUTCSV DD

Note: The double quotation marks shown in [Figure 119 on page 256](#) are part of the data and are required. You must use the keyboard double quotation mark character (EBCDIC X'7F' or ASCII X'22'), **not** typographical opening and closing double quotation marks.

The contents of the INPUTCSV data set consist of a product row (line), a header row, and one or more data rows.

Product row: The first row in the file is the product row.

- The product row must contain the "Product-zWPC" keyword as the first token.
- Following the "Product-zWPC" keyword, you must list the product ID (or IDs) for all versions of a particular product family (for instance, the product IDs associated with the CICS product family). If you are using multiple versions of a particular product, you must specify multiple product IDs. (A typical reason for this would be if you are migrating from one version of the product to another.) Each product ID must be enclosed in double quotation marks. The example in [Figure 119 on page 256](#) shows two products: "5655-S97" and "5655-M15".

Header row: The first non-blank, non-comment row after the product row is the header row.

- The first column of the header row must contain the date header information in one of the following formats:

- The "Date" token, which indicates that all times listed are local time, in ISO-8601 format.
- The "DateUTC" token, which indicates that all times listed are UTC time, in ISO-8601 format.

Note: The use of Hypervisor offsets (either LPAR Sysplex Timer offset or z/VM virtual clock, VTOD) requires that you enter the CSV time as though these values are already applied to the data. SCRT applies Hypervisor offset values as the SMF records are read and expects the CSV time values to reflect the Hypervisor offset. Normal z/OS time zone offsets do not require this adjustment.

For example, if the Hypervisor offset is used to run the system one month in the future, the CSV file should have that one month in the future date, and not the current day's date and times.

- The subsequent columns of the header row must contain the LPAR names in one of the following formats:

- The LPAR name (*not* the system ID), enclosed in double quotation marks

Example: "SYSF"

- Both the type-serial and LPAR name (*not* the system ID), separated by a colon, in the format *type-serial:LPARname*, enclosed in double quotation marks. If you have (or plan to have) duplicate LPAR names on different CPCs, you must use this format.

Example: "2827-23456:SYSF"

Data rows: Any non-blank, non-comment rows after the header row must contain the public cloud transaction data.

- The first column of a public cloud transaction data row must contain the date. The date must be in ISO-8601 format and enclosed in double quotation marks.

Example: A date value of "2014-03-21T03:00:00" indicates 3 a.m. on 21 March 2014.

- The subsequent cell values of the public cloud transaction data row must contain the CPU seconds for the hour specified in the first column of that row. The CPU second values must be placed in the column corresponding to the LPAR on which they were gathered. The CPU seconds are specified as a floating-point number and must not be enclosed in double quotation marks.

For more information about the INPUTCSV DD statement, see [“INPUTCSV DD statement” on page 60](#).

SCRT externals for WLM classification-based IBM Z Workload Pricing for Cloud

SCRT provides the following externals for IBM Z Workload Pricing for Cloud using WLM classification. Where noted, these externals are specific to zWPC and are treated as optional when zWPC does not apply to the SMF data being processed by SCRT.

- SCRT uses an optional **Assign ZWPC** control statement in the SPECIAL DD to allow you to indicate which WLM category or classification you are using to aggregate public cloud workload data.
- SCRT uses the `zWPC_Input_From_WLM` message in report section D5 to indicate when public cloud workload data from WLM has been requested and is reported. This message also specifies the WLM classification that was requested.
- SCRT uses footnote (e) in the Product Summary Information section of the sub-capacity report to indicate the zWPC defining program.
- SCRT uses an optional **Generate_Detailed_Data** control statement to track the application of public cloud, mobile, and z Systems Collocated Application Pricing (zCAP) reductions per hour, per LPAR.
- SCRT uses an optional **DIAGMSG=WLMZWPC** control statement to provide additional diagnostic messaging to help determine why an expected WLM zWPC adjustment does not appear on the sub-capacity report.

Assign ZWPC SPECIAL DD statement

SCRT supports an optional **Assign ZWPC** control statement in the SPECIAL DD. Use this control statement to indicate to SCRT which WLM workload classification category you are using to track public cloud workloads.

The format of this control statement is as follows:

```
Assign ZWPC={CATEGORYA|CATEGORYB}
```

Assign ZWPC=CATEGORYA

Indicates that the CATEGORYA classification (or reporting attribute) is being used as a source of public cloud transaction data.

Assign ZWPC=CATEGORYB

Indicates that the CATEGORYB classification (or reporting attribute) is being used as a source of public cloud transaction data.

Note: When selecting a WLM reporting attribute, use the same workload classification (or reporting attribute) for all zWPC-eligible public cloud workloads that are to be processed during a single execution of SCRT.

zWPC_Input_From_WLM message

For the case where the **Tool MSUs** value for one or more products listed in the Product Summary Information section have been adjusted for zWPC based on input from a WLM classification, SCRT generates an optional D5 report section. The `zWPC_Input_From_WLM` message appears in section D5, along with any other messages applicable to SCRT execution.

The `zWPC_Input_From_WLM` message also includes the WLM classification (or reporting attribute) category being used, as follows:

```
"zWPC_Input_From_WLM", "CATEGORYA"
```

or

```
"zWPC_Input_From_WLM", "CATEGORYB"
```


IBM Z Workload Pricing for Cloud footnote

For the case where the **Tool MSUs** value for one or more products listed in the Product Summary Information section have been adjusted for zWPC, SCRT applies the following changes to the Product Summary Information section:

- Generates an optional **Footnotes** column header for the MLC section.
- Posts footnote (e) in the **Footnotes** column for the highest z/OS operating system version.
- Generates a footnote subsection immediately after the last IPLA product is listed. The footnote subsection contains a footnote header row followed by a description of footnote (e).

Footnote (e) is optional. All three footnote-related items are only generated if the SMF type 70 data contains a valid, non-zero public cloud 4-hour rolling average value for at least one interval in the reporting period.

Generate_Detailed_Data control statement

SCRT supports an optional **Generate_Detailed_Data** SPECIAL control statement. This is used to track the application of zWPC and other reductions per hour, per LPAR.

When specified, this control statement causes the Mobile/zCAP Pricing Detailed Data section (V5) to be generated in the Sub-capacity report.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

DIAGMSG=WLMZWPC control statement

SCRT supports an optional **DIAGMSG=WLMZWPC** SPECIAL control statement. This control statement is used to provide additional diagnostic messaging to help determine why an expected WLM public cloud adjustment does not appear in the sub-capacity report.

When specified, this control statement causes SCRT to generate the SCRTTOOL612 or SCRTTOOL613 messages, when appropriate.

The SCRTTOOL612 message indicates that the specified LPAR does not appear to have the required z/OS maintenance level to support the new SMF70LACM field in the SMF type 70 record.

The SCRTTOOL613 message indicates that the specified LPAR did not report a non-zero value during the reporting period for the SMF field containing the public cloud workload 4-hour rolling average CPU consumption.

Mobile/zCAP Pricing Detailed Data section

The Mobile/zCAP Pricing Detailed Data section (V5) of the sub-capacity report shows, for each product, the hour-by-hour application of zWPC and other MSU adjustments.

See [“Sample sub-capacity report for WLM classification-based IBM Z Workload Pricing for Cloud”](#) on page 271 for a detailed description of the output in the Mobile/zCAP Pricing Detailed Data section.

Reporting sample for INPUTCSV-based IBM Z Workload Pricing for Cloud

This topic provides a detailed example of sub-capacity reporting for INPUTCSV-based IBM Z Workload Pricing for Cloud.

Sample configuration for INPUTCSV-based IBM Z Workload Pricing for Cloud

Table 41 on page 260 illustrates the sample hardware and software configuration for a CPC where an input CSV file was provided for one zWPC defining program that was running during the reporting period.

The sample configuration consists of a 2964 Model 7A1 CPC with a serial number of 02-F2345. The CPC is configured with one LPAR running z/OS natively. The z/OS system in LPAR1 has a z/OS SYSID of SYS1.

The z/OS system is running a combination of the following z/OS products:

- Db2 11 for z/OS
- CICS TS for z/OS V5
- IBM MQ for z/OS V8
- IMS V13
- IBM Enterprise COBOL for z/OS V4
- File Export for z/OS V1
- Migration Utility for z/OS V4

Table 41. Sample configuration for IBM Z Workload Pricing for Cloud	
CPC: 2964-7A1 (serial = 02-F2345)	
LPAR1	
z/OS V2.1 (SYSID = SYS1)	
Db2 11 for z/OS CICS TS for z/OS V5 IBM MQ for z/OS V8 IMS V13 IBM Enterprise COBOL for z/OS V4 File Export for z/OS V Migration Utility for z/OS V4	

Sample control statement input to SCRT for INPUTCSV-based IBM Z Workload Pricing for Cloud

SCRT supports additional externals for zWPC. The sample sub-capacity report discussed in [“Sample sub-capacity report for INPUTCSV-based IBM Z Workload Pricing for Cloud” on page 261](#) is based on the following SCRT control statement and JCL inputs to SCRT:

INPUTCSV DD statement

The following DD statement was used to specify the data set containing public cloud workload pricing data for the reporting period for the LPAR and SYSID described in [“Sample configuration for INPUTCSV-based IBM Z Workload Pricing for Cloud” on page 260](#):

```
//INPUTCSV DD DISP=SHR,DSN=MY.DSN.ID(CICSZWPC)
```

Specifying the INPUTCSV statement and supplying CPU consumption data for public cloud transactions in the appropriate format in the INPUTCSV data set is the method to request that SCRT use the zWPC

algorithm whenever all of the zWPC conditions have been met. See [“INPUTCSV DD statement for zWPC”](#) on page 255 for a detailed description and restrictions on its use.

In the sample configuration, the following public cloud consumption data is stored in MY.DSN.ID(CICSZWPC), which is the data set referenced on the sample INPUTCSV DD statement:

```
"Product-zWPC", "5655-Y04"
"Date", "2964-840F7:JL0"
"2015-02-08T12:00:00", 500.00
"2015-02-08T13:00:00", 500.00
"2015-02-08T14:00:00", 700.00
"2015-02-08T15:00:00", 800.00
"2015-02-08T16:00:00", 700.00
"2015-02-08T17:00:00", 600.00
```

For more information about the format of the INPUTCSV data set, see [“Format of the INPUTCSV input data for zWPC”](#) on page 256.

SPECIAL DD control statement

The following SCRT control statement can be specified in the SPECIAL DD to request that SCRT generate the Mobile/zCAP Detailed Data section of the sub-capacity report. This control statement has no effect on whether SCRT uses the zWPC algorithm to apply reductions based on the data supplied via the INPUTCSV DD statement.

```
Generate_Detailed_Data
```

Note that this control statement was not specified for this sample configuration.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

Sample sub-capacity report for INPUTCSV-based IBM Z Workload Pricing for Cloud

The sub-capacity report for IBM Z Workload Pricing for Cloud has the same basic format and report sections as the sub-capacity report for a base z/OS system. However, the sub-capacity report for zWPC always contains a footnote to indicate the case when public cloud CPU consumption data was provided for zWPC defining programs.

The following sample report sections describe the differences between a sub-capacity report for a base z/OS system and a sub-capacity report that uses IBM Z Workload Pricing for Cloud. For complete descriptions of all of the fields that appear in a base sub-capacity report see [“Sample sub-capacity report for single-version z/OS native systems”](#) on page 105.

Billing-related section

The billing-related section of the sub-capacity report displays the customer information, tool information, special conditions, product summary information, and detail data collection information, which are described in [“Billing-related section”](#) on page 105.

Figure 120 on page 262 shows an example of the Customer Information section and the Tool Information section.

```
==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====
```

```
Run Date/Time           03 Nov 2024 - 13:47
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter:  800-555-0123
```

```
Customer Name           ABC Corporation
Customer Number          9871234567
Machine Serial Number    02-F2345
Machine Type and Model    2964-7A1
Machine Rated Capacity (MSUs)      10171
Purchase Order Number      (optional)
Customer Comments (255 chars max)  (optional)
```

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

```
==C5=====
TOOL INFORMATION
```

```
Tool Release           30.1.0
Reporting Period        2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)
```

Figure 120. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Customer information and Tool Information sections

Customer Information section

There are no differences between the Customer Information section for a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC. For details about this section, see [“Customer information” on page 105](#).

Tool Information section

There are no differences between the Tool Information section for a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC. For details about this section, see [“Tool Information section” on page 107](#).

Special Conditions section

There are no differences between the Special Conditions section for a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC.

Product Summary Information section

An optional Footnotes column header appears in the MLC section whenever SCRT needs to post a footnote for a zWPC defining program listed in this section of the report. The **Footnotes** heading does not appear if there are no footnotes for this section.

There are no differences between the **MLC Product Name**, **MLC Product ID**, **Customer Comment**, **IPLA Product Name**, **IPLA Product ID**, and **Customer MSUs** fields on a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC. For descriptions of these fields, see [“Product Summary Information section” on page 110](#).

The sample Product Summary Information section appears in [Figure 121 on page 263](#). The following fields differ for zWPC reporting:

Tool MSUs

The values listed in the **Tool MSUs** column represent the peak sub-capacity values calculated by SCRT for each product using the appropriate sub-capacity rules for that product. SCRT uses the zWPC algorithm and the reported public cloud transaction CPU data to adjust the 4-hour rolling average sub-capacity MSU value for sub-capacity eligible programs on a given CPC.

See “Product Max Contributors section” on page 264 for a more detailed discussion of how the zWPC algorithm affects the sub-capacity programs that are listed in this section of the sub-capacity report. For additional general information about **Tool MSUs**, see “Product Summary Information section” on page 110.

Footnotes

SCRT posts footnote (e) for each zWPC defining program for which zWPC data was supplied via the INPUTCSV DD statement.

The **Tool MSUs** values are partially based on the zWPC algorithm whenever the posted value for a product is combined from multiple LPARs where the product's value is based on the traditional sub-capacity values from some LPARs and on the zWPC adjusted values from other LPARs.

Whenever SCRT posts footnote (e) in the **Footnotes** column, it also generates a row with a footnote header followed by a description for the footnote.

The Product Summary Information section in [Figure 121 on page 263](#) uses footnote (e) to indicate that data was specified on the INPUTCSV DD statement for zWPC defining program CICS TS for z/OS V5.

```

=====
PRODUCT SUMMARY INFORMATION

MLC Product Name          MLC Product ID    Tool MSUs    Customer MSUs    Customer Comments (255 chars max)    Footnotes
z/OS V2                   5650-Z05          1519
DB2 11 for z/OS           5615-DB2          1519
CICS TS for z/OS V5       5655-Y04          1519
IBM MQ for z/OS V8        5655-W97          1519
IMS V13                   5635-A04          1519
IBM Enterprise Cobol for z/OS V4  5655-S71          1519    (optional)    (conditional)

IPLA Product Name        IPLA Product ID    Tool MSUs    Customer MSUs    Customer Comments (255 chars max)    Footnotes
File Export for z/OS V1   5697-I12          1519    (optional)    (conditional)
Migration Utility for z/OS V4  5655-MGU          1519    (optional)    (conditional)

Footnotes:
(e) Workload Pricing for Cloud transaction data from this product was included from one or more LPARs

```

Figure 121. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Product Summary Information section

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the Product Summary Information section.

SMF / SCRT89 Input Data Statistics section

There are no differences between the SMF / SCRT89 Input Data Statistics section for a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC. For details about this section, see “SMF / SCRT89 Input Data Statistics section” on page 118.

The sample SMF / SCRT89 Input Data Statistics section appears in [Figure 122 on page 264](#).

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

          SYSID      Input Data Start      Input Data End
LPAR1      SYS1      08 Oct 2024 - 10:59  08 Oct 2024 - 19:59
CPC                08 Oct 2024 - 10:59  08 Oct 2024 - 19:59

```

Figure 122. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: SMF / SCRT89 Input Data Statistics section

Detail LPAR Data Section

There are no differences between the Detail LPAR Data Section for a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC. For details about this section, see [“Detail LPAR Data section” on page 119](#).

The sample Detail LPAR Data Section appears in [Figure 123 on page 264](#).

```

==N5=====
DETAIL LPAR DATA SECTION

          Highest Hour Count Date/Time          2nd Highest Hour Count Date/Time
LPAR1          1531         1 08 Oct 2024 - 17:00      1506         1 08 Oct 2024 - 16:00
CPC            1531         1 08 Oct 2024 - 17:00      1506         1 08 Oct 2024 - 16:00

```

Figure 123. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Detail LPAR Data section

Product Max Contributors section

The MSU values that are listed in the Product Max Contributors section are the same for a basic z/OS sub-capacity report and a IBM Z Workload Pricing for Cloud sub-capacity report, with one key difference. Whenever zWPC is requested and all conditions have been met, a Cloud MSU Reduction column header is generated. For each product that is eligible to be reduced, the calculated cloud reduction that was applied during the peak hour is reported.

This section provides an additional layer of detail behind the sub-capacity values that are reported for the products listed in the Product Summary Information section by reporting the date/time that the sub-capacity value occurred together with the product's sub-capacity contribution from each LPAR in which that product was running at that date and time. When the *same* peak sub-capacity value occurs for multiple hours, the Product Max Contributors section lists the data for the *earliest* hour. The impact of SCRT reporting a product's sub-capacity value based on the earliest peak hour whenever that peak value occurs for multiple hours is important enough to reiterate in the following explanations because the zWPC conditions for the other peak hours can differ from the zWPC conditions for the earliest peak hour.

Given the level of detail in this section, the sample Product Max Contributors section shown in [Figure 124 on page 265](#) is used to explain IBM Z Workload Pricing for Cloud in more detail. However, this explanation assumes you have a general understanding of the contents and use of each field in the Product Max Contributors section, as described in [“Product Max Contributors section” on page 121](#).

```

=====
==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR SYS1	Cloud MSU Reduction
z/OS V2	5650-Z0S	1519	08 Oct 2024 - 17:00	1519	12
DB2 11 for z/OS	5615-DB2	1519	08 Oct 2024 - 17:00	1519	12
CICS TS for z/OS V5	5655-Y04	1519	08 Oct 2024 - 17:00	1519	12
IBM MQ for z/OS V8	5655-W97	1519	08 Oct 2024 - 17:00	1519	12
IMS V13	5635-A04	1519	08 Oct 2024 - 17:00	1519	12
IBM Enterprise Cobol for z/OS V4	5655-S71	1519	08 Oct 2024 - 17:00	1519	12
File Export for z/OS V1	5697-I12	1519	08 Oct 2024 - 17:00	1519	12
Migration Utility for z/OS V4	5655-MGU	1519	08 Oct 2024 - 17:00	1519	12

Figure 124. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Product Max Contributors section

How SCRT implements IBM Z Workload Pricing for Cloud

It is important to understand that SCRT uses the appropriate algorithm to calculate the sub-capacity value for each product for each hour for each LPAR, independently of what is happening with that product on other LPARs or guest systems during that same hour. Then, SCRT sums each of the product's sub-capacity values across all of the LPARs for the same hour to calculate a sub-capacity value for the product for the CPC for that hour. After performing these sub-capacity calculations for the product across the CPC for all of the hours during the reporting period, SCRT selects the earliest hour across the CPC with the product's peak sub-capacity value and reports the date and time for that hour, along with the individual sub-capacity values for the product for each LPAR for that same hour, in the Product Max Contributors section.

For products like z/OS, whose sub-capacity is entirely based on the 4-hour rolling average MSU utilization of the LPAR, the 4-hour rolling average MSU utilization is always reported as the product's sub-capacity value. In the sample Product Max Contributors section, the peak z/OS sub-capacity value across the CPC was 1519 MSUs and that peak first occurred on 08 Oct 2024 at 17:00.

As mentioned earlier, a key difference between a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC is the Cloud MSU Reduction column in the Product Max Contributors section. This column is generated whenever public cloud pricing data was provided via the INPUTCSV DD statement, as was the case in the configuration (see [“Sample configuration for INPUTCSV-based IBM Z Workload Pricing for Cloud”](#) on page 260).

The Cloud MSU Reduction column shows the adjustment that SCRT applied to products' sub-capacity values using the zWPC algorithm. In the Product Max Contributors section for the sample configuration, the peak z/OS sub-capacity value and the value of all products based entirely on the 4-hour rolling average MSU utilization of the LPAR were adjusted by SCRT using the value shown in the Cloud MSU Reduction column.

In the Product Max Contributors section for the sample configuration, note that the MSU values shown for LPAR SYS1 already reflect the contribution of zWPC defining programs that ran in each LPAR during the peak hour. If multiple LPARs are present, the Cloud MSU Reduction column reflects the total reduction across all LPARs for that hour.

A note about MWP adjustments: The MSU values that are used for zWPC adjustments are derived from the public cloud transaction general purpose processor time values reported using the INPUTCSV DD statement. These adjustments are calculated internally by SCRT as specified in the IBM Addendum for z Systems Workload Pricing for Cloud. The Product Max Contributors section is only intended to show the contribution of each LPAR to the specified product's peak sub-capacity value for earliest peak hour.

Product Grid Snapshot section

There are no differences between the Product Grid Snapshot section for a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC. For details about this section, see [“Product Grid Snapshot section”](#) on page 123.

The sample Product Grid Snapshot section appears in [Figure 125 on page 266](#).

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                Product ID    LPAR1
z/OS V2                     5650-Z0S     1.10%
DB2 11 for z/OS             5615-DB2     1.10%
CICS TS for z/OS V5         5655-Y04     1.10%
IBM MQ for z/OS V8          5655-W97     1.10%
IMS V13                     5635-A04     1.10%

```

Figure 125. Sample sub-capacity report for Integrated Workload Pricing: Product Grid Snapshot section

Getting Started Sub-Capacity Pricing Assigned Hours section

There are no differences between the Getting Started Sub-Capacity Pricing Assigned Hours section for a basic z/OS sub-capacity report and for a sub-capacity report generated using zWPC. For details about this section, see [“Getting Started Sub-Capacity Pricing Assigned section” on page 179](#).

The sample Getting Started Sub-Capacity Pricing Assigned Hours section appears in [Figure 126 on page 266](#).

```

==S7=====
GETTING STARTED SUB-CAPACITY PRICING ASSIGNED - HOURS

Product Name                Product ID    LPAR    LPAR
                           SYS1          SYS2
IBM WebSphere Application Server for z/OS V8  5655-W65     4

```

Figure 126. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Getting Started Sub-Capacity Pricing Assigned Hours section

Mobile/zCAP Pricing Detailed Data section for zWPC

The optional Mobile/zCAP Pricing Detailed Data section appears when the optional **Generate_Detailed_Data** SPECIAL DD control statement is specified. This section shows the mobile MSU reduction applied by SCRT for each zWPC defining program, per LPAR. This information is shown for each sub-capacity product on an hour-by-hour basis.

As shown in [Figure 127 on page 268](#), the optional Mobile/zCAP Pricing Detailed Data section reports several columns for each product:

- Date and time of the adjusted interval.
- CPC Type-model.
- CPC rated machine capacity in MSU.
- Adjusted MSU values for each LPAR. (Note that the MWP algorithm has already been applied.)
- For each mobile defining product with INPUTCSV DD data specified, the MWP reduction applied per defining product, per partition.
- Total Mobile MSU reduction applied across all MWP defining products.
- CPC Total MSU, which is the sum of the adjusted LPAR MSU values across all partitions for the adjusted product.

There are several additional values that can appear in section V5 of a sub-capacity report:

–

Indicates that there is no data for this interval.

dd(rr)

Indicates that the defined capacity limit for the LPAR was reached during this interval. The first number (*dd*) represents the defined capacity value used for this hour. The second number (*rr*) represents the 4-hour rolling average value prior to the capping.

G(rr)

Indicates that the sum of the partitions exceeded the group capacity limit; the group capacity was used for this hour for that set of LPARs. The *rr* value represents the 4-hour rolling average value prior to the capping.

P(rr)

Indicates that the sum of the partitions exceeded the processor capacity; the processor capacity was used. The *rr* value represents the 4-hour rolling average value prior to the capping.

(GSSP)

Indicates that a WebSphere program eligible for Getting Started Sub-Capacity Pricing received the GSSP value so there was no additional adjustment.

nnn(CPS)

Indicates that an LPAR was assigned to a Tailored Fit Pricing solution for this hour. In this case, the LPAR does not contribute to the product sub-capacity value. The *nnn* value represents what the LPAR sub-capacity value would have been if it were part of the standard (non-container) environment.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

The following descriptions of the IBM Z Workload Pricing for Cloud algorithm are based on the sample configuration in [“Sample configuration for INPUTCSV-based IBM Z Workload Pricing for Cloud”](#) on page 260. Note that only a subset of the products running in the environment are evaluated.

Results for z/OS V2R1

After SCRT applied the appropriate algorithm hour by hour, LPAR by LPAR, and then summed the sub-capacity values for the same hour across all of the z/OS LPARs, SCRT reported 04 May 2024 at 14:00 as the earliest hour with the peak sub-capacity value for z/OS, as indicated in the sample Product Max Contributors section.

For 04 May 2024 at 14:00, SCRT applied the appropriate algorithms and calculated a sub-capacity value of 2693 MSUs for LPAR SYS1. As shown in the **CICS Mobile Partition** column for SYS1, this includes an 8 MSU zWPC reduction derived from the CICS mobile transaction general purpose processor time that was specified as input in the sample configuration in [“Sample configuration for INPUTCSV-based IBM Z Workload Pricing for Cloud”](#) on page 260.

For partition SYS2, SCRT calculated a sub-capacity value of 821 MSUs, which includes a 1 MSU zWPC reduction derived from the IMS mobile transaction general purpose processor time that was specified as input in the sample configuration.

[Figure 127 on page 268](#) illustrates these results.

```

==V5=====
MOBILE/ZCAP PRICING DETAILED DATA

MLC Products

5650-ZOS                z/OS V2

Date Time of RMF Interval  Type      CPC MSU  Partition  SYS2  CICS Mobile Partition  IMS Mobile Partition  Total Mobile MSU  CPC Total
                                SYS1      821      SYS1      SYS2      Reduction
04 May 2024 - 14:00      2964-796  9783    2693      821      8         1         9         3514
04 May 2024 - 15:00      2964-796  9783    2231      798     13         5         18        3029
04 May 2024 - 16:00      2964-796  9783    1864      749     19         7         26        2613
04 May 2024 - 17:00      2964-796  9783    1679      676     24        10        34        2355

```

Figure 127. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Mobile/zCap Pricing Detailed Data section — z/OS results

Results for IBM MQ for z/OS V8

SCRT applied the process as described in “Results for z/OS V2R1” on page 267. Like z/OS, IBM MQ for z/OS V8 is based on the 4-hour rolling average MSU utilization of the LPAR.

Since IBM MQ for z/OS V8 ran on both z/OS LPARs during 04 May 2024 at 14:00, the application of the zWPC algorithm mirrors that previously described for z/OS.

This same process applies for all sub-capacity programs that reported a peak MSU value of 3514 with 04 May 2024 at 14:00 as the earliest hour with the peak sub-capacity value.

Figure 128 on page 268 illustrates the results.

```

5655-W97                IBM MQ for z/OS V8

Date Time of RMF Interval  Type      CPC MSU  Partition  SYS2  CICS Mobile Partition  IMS Mobile Partition  Total Mobile MSU  CPC Total
                                SYS1      821      SYS1      SYS2      Reduction
04 May 2024 - 14:00      2964-796  9783    2693      821      8         1         9         3514
04 May 2024 - 15:00      2964-796  9783    2231      798     13         5         18        3029
04 May 2024 - 16:00      2964-796  9783    1864      749     19         7         26        2613
04 May 2024 - 17:00      2964-796  9783    1679      676     24        10        34        2355

```

Figure 128. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Mobile/zCap Pricing Detailed Data section — IBM MQ for z/OS results

Results for IBM Multi-site Workload Lifeline V2

SCRT applied the process as described in “Results for z/OS V2R1” on page 267 and “Results for IBM MQ for z/OS V8” on page 268. Like those products, IBM Multi-site Workload Lifeline V2 is based on the 4-hour rolling average MSU utilization of the LPAR.

However, IBM Multi-site Workload Lifeline V2 is only running in partition SYS1. As such, it is not eligible for the reduction derived from the IMS mobile defining program, as mobile transaction general purpose processor time was only specified for IMS on SYS2. As a result, SCRT calculated an MSU value of 2693 for IBM Multi-site Workload Lifeline V2, as illustrated in Figure 129 on page 269.

5655-UM4		IBM Multi-site Workload Lifeline V2				
Date Time of RMF Interval	Type	CPC MSU	Partition SYS1	CICS Mobile Partition SYS1	Total Mobile MSU Reduction	CPC Total
04 May 2024 - 14:00	2964-796	9783	2693	8	8	2693
04 May 2024 - 15:00	2964-796	9783	2231	13	13	2231
04 May 2024 - 16:00	2964-796	9783	1864	19	19	1864
04 May 2024 - 17:00	2964-796	9783	1679	24	24	1679

Figure 129. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Mobile/zCap Pricing Detailed Data section — IBM Multi-site Workload Lifeline results

Results for IBM WebSphere Application Server for z/OS V8

SCRT applied the Getting Started Sub-Capacity Pricing for z/OS IPLA Software process to IBM WebSphere Application Server for z/OS V8.

Since the terms and conditions for GSSP were met in the sample configuration, and the GSSP values were less than the traditional sub-capacity values for each hour of the reporting period, SCRT assigned IBM WebSphere Application Server for z/OS V8 the GSSP value for each hour in the reporting period.

When the GSSP value is used for a GSSP eligible product, the Mobile/zCAP Pricing Detailed Data section places the (GSSP) token in the zWPC defining program partition reduction column, as shown in Figure 130 on page 269. In accordance with the IBM Z Workload Pricing for Cloud addendum, the zWPC reduction is not applied when GSSP pricing is used instead of the traditional sub-capacity value.

In the sample configuration, the peak MSU value used for IBM WebSphere Application Server for z/OS V8 is 21, which is achieved during the 04 May 2024 17:00 hour.

For information about the GSSP process, see Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing (GSSP),” on page 165.

5655-W65			IBM WebSphere Application Server for z/OS V8					
Date Time of RMF Interval			Type	CPC MSU	Partition	CICS Mobile Partition	Total Mobile MSU Reduction	CPC Total
04 May 2024 - 14:00			2964-796	9783	5	(GSSP)	0	5
04 May 2024 - 15:00			2964-796	9783	11	(GSSP)	0	11
04 May 2024 - 16:00			2964-796	9783	16	(GSSP)	0	16
04 May 2024 - 17:00			2964-796	9783	21	(GSSP)	0	21

Figure 130. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Mobile/zCap Pricing Detailed Data section — IBM WebSphere Application Server for z/OS results

Reporting sample for WLM classification-based IBM Z Workload Pricing for Cloud

This topic provides a detailed example of sub-capacity reporting for WLM classification-based IBM Z Workload Pricing for Cloud.

Sample configuration for WLM classification-based IBM Z Workload Pricing for Cloud

Table 42 on page 270 illustrates the sample hardware and software configuration for a CPC where public cloud workloads are indicated using WLM classification during the reporting period. Although you must

define a given workload as public cloud using either the CATEGORYA or CATEGORYB classification in WLM, SCRT is not aware of the public cloud contributors under this approach.

The sample configuration consists of a 2817 Model 760 CPC with a serial number of 02-12345. The CPC is configured with one LPAR that runs z/OS natively. The z/OS system in LPAR1 has a z/OS SYSID of SYS1.

The z/OS systems are running a combination of the following z/OS products:

- Db2 11 for z/OS
- CICS TS for z/OS V5
- IBM MQ for z/OS V8
- IMS V13
- IBM Enterprise COBOL for z/OS V4
- IBM WebSphere Application Server for z/OS V8
- IBM Integration Bus for z/OS V9
- IBM Multi-site Workload Lifeline V2

<i>Table 42. Sample configuration for IBM Z Workload Pricing for Cloud</i>	
CPC: 2817-760 (serial = 02-12345)	
LPAR1	
z/OS V2.1 (SYSID = SYS1) Db2 11 for z/OS CICS TS for z/OS V5 IBM MQ for z/OS V8 IMS V13 IBM Enterprise COBOL for z/OS V4 IBM WebSphere Application Server for z/OS V8 IBM Integration Bus for z/OS V9 IBM Multi-site Workload Lifeline V2	

Sample control statement input to SCRT for WLM classification-based IBM Z Workload Pricing for Cloud

SCRT supports additional externals for WLM classification-based zWPC. The sample sub-capacity report discussed in [“Sample sub-capacity report for WLM classification-based IBM Z Workload Pricing for Cloud” on page 271](#) is based on the following SCRT control statement and JCL inputs to SCRT.

SPECIAL DD control statement

The following SCRT control statement was specified in the SPECIAL DD to indicate to SCRT which WLM classification was chosen for zWPC workloads. This specification is required for SCRT to process WLM classification-based IBM Z Workload Pricing for Cloud.

```
Assign ZWPC=CATEGORYA
```

In this sample environment, all zWPC eligible workloads are configured to have a reporting attribute (or classification) of CATEGORYA.

The following SCRT control statement was specified in the SPECIAL DD to request that SCRT generate the Mobile/zCAP Detailed Data section of the sub-capacity report. This control statement has no effect on whether SCRT uses the zWPC algorithm to apply reductions based on the WLM classification-based workloads found in the SMF data.

```
Generate_Detailed_Data
```

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

Sample sub-capacity report for WLM classification-based IBM Z Workload Pricing for Cloud

The sub-capacity report for IBM Z Workload Pricing for Cloud has the same basic format and report sections as the sub-capacity report for a base z/OS system. However, the sub-capacity report for zWPC always contains a Special Condition message when public cloud data was provided via WLM workload classification. Additionally, a zWPC footnote is always present on the highest z/OS operating system level in use on the CPC to indicate the case when public cloud CPU consumption data was provided via WLM.

The following sample report sections describe the differences between a sub-capacity report for a base z/OS system and a sub-capacity report that uses WLM classification-based IBM Z Workload Pricing for Cloud. For a complete description of all of the fields appearing in a base sub-capacity report, see [“Sample sub-capacity report for single-version z/OS native systems” on page 105](#).

Billing-related section

The billing-related section of the sub-capacity report displays the customer information, tool information, special conditions, product summary information, and detail data collection information, which are described in [“Billing-related section” on page 105](#).

Figure 131 on page 271 shows an example of the Customer Information section and the Tool Information section.

```

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time                03 Nov 2024 - 11:09
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter:  800-555-0123

Customer Name                ABC Corp
Customer Number              US-S000000000
Machine Serial Number        02-12345
Machine Type and Model       2817-760
Machine Rated Capacity (MSUs) 5001
Purchase Order Number        (optional)
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust
the billable MSUs in inventory for all MLC Products listed under the MLC Product Name
column on this report. In accordance with our agreement, IBM will treat a change in product
licensed capacity as an order. If the MSUs have changed since the last report, software billing
based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to
bill those IPLA products listed under the IPLA Product Name column in this report which
exceed your entitled capacity. In accordance with our agreement, IBM will treat the use
of a product in excess of its entitled capacity as an order and you will be billed for
the amount in excess of your entitlement.

==C5=====
TOOL INFORMATION

Tool Release                30.1.0
Reporting Period            2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

```

Figure 131. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Customer information and Tool Information sections

Customer Information section

There are no differences between the Customer Information section for a basic z/OS sub-capacity report and a sub-capacity report that is generated using zWPC. For details about this section, see [“Customer information” on page 105](#).

Tool Information section

There are no differences between the Tool Information section for a basic z/OS sub-capacity report and a sub-capacity report that is generated using zWPC. For details about this section, see [“Tool Information section” on page 107](#).

Special Conditions section

The key difference between the Special Conditions section for a basic z/OS sub-capacity report and a sub-capacity report for WLM classification-based zWPC is the presence of the following message:

```
zWPC_Input_From_WLM      CATEGORYA
```

This message is generated when the **Assign ZWPC** control statement is specified in the SPECIAL DD and SCRT detects and applies valid WLM public cloud workload data in the SMF input data. This message reflects the WLM category that was specified on the **Assign ZWPC** control statement.

The sample Special Conditions section appears in [Figure 132 on page 272](#).

```
==D5=====
SPECIAL CONDITIONS
zWPC_Input_From_WLM      CATEGORYA
```

Figure 132. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Special Conditions section

Product Summary Information section

An optional Footnotes column header appears in the MLC section whenever SCRT needs to post a footnote for a zWPC defining program listed in this section of the report. The **Footnotes** heading does not appear whenever there are no footnotes for this section.

There are no differences between the **MLC Product Name**, **MLC Product ID**, **Customer Comment**, **IPLA Product Name**, **IPLA Product ID**, and **Customer MSUs** fields on a basic z/OS sub-capacity report and on a sub-capacity report generated using zWPC. For a description of these fields see [“Product Summary Information section” on page 110](#).

The sample Product Summary Information section appears in [Figure 133 on page 273](#). The following fields differ for zWPC reporting:

Tool MSUs

The values listed in the **Tool MSUs** column represent the peak sub-capacity values calculated by SCRT for each product using the appropriate sub-capacity rules for that product. SCRT uses the zWPC algorithm and the reported public cloud transaction CPU data to adjust the 4-hour rolling average sub-capacity MSUs for sub-capacity eligible programs on a given CPC.

See [“Product Max Contributors section” on page 264](#) for a more detailed discussion of how the zWPC algorithm affects the sub-capacity programs that are listed in this section of the sub-capacity report. For additional general information about **Tool MSUs**, see [“Product Summary Information section” on page 110](#).

Footnotes

When public cloud workload data is provided via WLM classification, SCRT posts footnote (e) for the first z/OS operating system version that appears in the sub-capacity report.

The **Tool MSUs** values are partially based on the zWPC algorithm whenever the posted value for the product is combined from multiple LPARs where the product's value is based on the traditional sub-capacity values from some LPARs and on the zWPC adjusted values from other LPARs.

Whenever SCRT posts footnote (e) in the **Footnotes** column, it also generates a row with a footnote header followed by a description of the footnote.

The Product Summary Information section in [Figure 133 on page 273](#) uses footnote (e) to indicate that public cloud data is based on WLM zWPC workload classification.

SCRT also posts footnote (iv) whenever the **Tool MSUs** value posted for a product is based totally or partially on having been adjusted by the Getting Started Sub-Capacity Pricing algorithm. See [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165](#) for a detailed description of Getting Started Sub-Capacity Pricing.

```

=====
PRODUCT SUMMARY INFORMATION
=====

```

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)	Footnotes
z/OS V2	5650-Z0S	1284			(e)
DB2 11 for z/OS	5615-DB2	1284			
CICS TS for z/OS V5	5655-Y04	1284			
IBM MQ for z/OS V8	5655-W97	1284			
IMS V13	5635-A04	1284			
IBM Enterprise Cobol for z/OS V4	5655-S71	1284	(optional)	(conditional)	

IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)	Footnotes
File Export for z/OS V1	5697-I12	1284	(optional)	(conditional)	
IBM WebSphere Application Server for z/OS V8	5655-W65	3			(iv)
IBM Integration Bus for z/OS V9	5655-IBB	28			(iv)
Migration Utility for z/OS V4	5655-MGU	1284	(optional)	(conditional)	
IBM Multi-site Workload Lifeline V2	5655-UM4	1284			

Footnotes:
 (iv) Product's Tool MSUs are based on Getting Started Sub-capacity Pricing from one or more LPARs
 (e) Workload Pricing for Cloud transaction data from this product was included from one or more LPARs

Figure 133. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Product Summary Information section

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the Product Summary Information section.

SMF / SCRT89 Input Data Statistics section

There are no differences between the SMF / SCRT89 Input Data Statistics section for a basic z/OS sub-capacity report and a sub-capacity report that is generated using zWPC. For details about this section, see [“SMF / SCRT89 Input Data Statistics section” on page 118](#).

The sample SMF / SCRT89 Input Data Statistics section appears in [Figure 134 on page 273](#).

```

=====
L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
=====
M5=====
SMF/SCRT89 INPUT DATA STATISTICS
=====

```

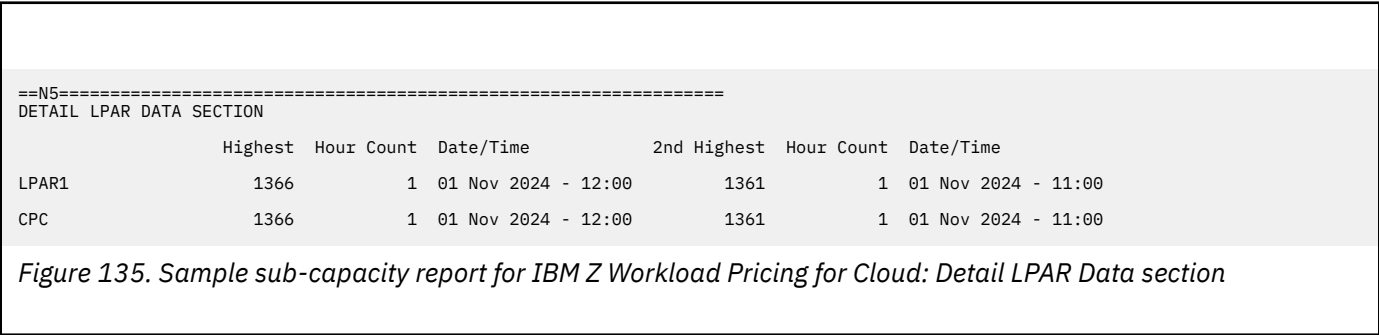
	SYSID	Input Data Start	Input Data End
LPAR1	SYS1	30 Oct 2024 - 22:59	01 Nov 2024 - 23:59
CPC		30 Oct 2024 - 22:59	01 Nov 2024 - 23:59

Figure 134. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: SMF / SCRT89 Input Data Statistics section

Detail LPAR Data Section

There are no differences between the Detail LPAR Data Section for a basic z/OS sub-capacity report and a sub-capacity report that is generated using zWPC. For details about this section, see “Detail LPAR Data section” on page 119.

The sample Detail LPAR Data Section appears in [Figure 135 on page 274](#).



Product Max Contributors section

The MSU values that are listed in the Product Max Contributors section are the same for a basic z/OS sub-capacity report and a Getting Started Sub-Capacity Pricing sub-capacity report, with one key difference. Whenever zWPC is requested and all conditions have been met, a **Cloud MSU Reduction** column header is generated. For each product that is eligible to be reduced, the **Cloud MSU Reduction** column reports the calculated cloud reduction that was applied during the peak hour.

This section provides an additional layer of detail behind the sub-capacity values reported for the products listed in the Product Summary Information section by reporting the date/time that the sub-capacity value occurred together with the product's sub-capacity contribution from each LPAR in which that product was running at that date and time. When the *same* peak sub-capacity value occurs for multiple hours, the Product Max Contributors section lists the data for the *earliest* hour. The impact of SCRT reporting a product's sub-capacity value based on the earliest peak hour whenever that peak value occurs for multiple hours is important enough to reiterate in the following explanations because the zWPC conditions for the other peak hours can differ from the zWPC conditions for the earliest peak hour.

In the sample configuration, note that no **Cloud MSU Reduction** value is listed for IBM WebSphere Application Server for z/OS V8 or IBM Integration Bus for z/OS V9 because the MSU values listed for these products are based on Getting Started Sub-Capacity Pricing, not on the standard sub-capacity 4-hour rolling average value.

Given the level of detail in this section, the sample Product Max Contributors section shown in [Figure 136 on page 275](#) is used to explain IBM Z Workload Pricing for Cloud in more detail. However, this explanation assumes you have a general understanding of the contents and use of each field in the Product Max Contributors section, as described in “Product Max Contributors section” on page 121.


```

=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	Cloud MSU Reduction
z/OS V2	5650-Z0S	1284	01 Nov 2024 - 12:00	1284	82
DB2 11 for z/OS	5615-DB2	1284	01 Nov 2024 - 12:00	1284	82
CICS TS for z/OS V5	5655-Y04	1284	01 Nov 2024 - 12:00	1284	82
IBM MQ for z/OS V8	5655-W97	1284	01 Nov 2024 - 12:00	1284	82
IMS V13	5635-A04	1284	01 Nov 2024 - 12:00	1284	82
IBM Enterprise Cobol for z/OS V4	5655-S71	1284	01 Nov 2024 - 12:00	1284	82
File Export for z/OS V1	5697-I12	1284	01 Nov 2024 - 12:00	1284	82
IBM WebSphere Application Server for z/OS V8	5655-W65	3	01 Nov 2024 - 03:00	3	0
IBM Integration Bus for z/OS V9	5655-IBB	28	01 Nov 2024 - 12:00	28	0
Migration Utility for z/OS V4	5655-MGU	1284	01 Nov 2024 - 12:00	1284	82
IBM Multi-site Workload Lifeline V2	5655-UM4	1284	01 Nov 2024 - 12:00	1284	82

Figure 136. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Product Max Contributors section

How SCRT implements IBM Z Workload Pricing for Cloud

It is important to understand that SCRT uses the appropriate algorithm to calculate the sub-capacity value for each product for each hour for each LPAR, independently of what is happening with that product on other LPARs or guest systems during that same hour. Then, SCRT sums each of the product's sub-capacity values across all of the LPARs for the same hour to calculate a sub-capacity value for the product for the CPC for that hour. After performing these sub-capacity calculations for the product across the CPC for all of the hours during the reporting period, SCRT selects the earliest hour across the CPC with the product's peak sub-capacity value and reports the date and time for that hour, along with the individual sub-capacity values for the product for each LPAR for that same hour, in the Product Max Contributors section.

For products like z/OS, whose sub-capacity is entirely based on the 4-hour rolling average MSU utilization of the LPAR, the 4-hour rolling average MSU utilization is always reported as the product's sub-capacity value. In the sample Product Max Contributors section, the peak z/OS sub-capacity value across the CPC was 1284 MSUs and that peak first occurred on 01 Nov 2024 at 12:00.

As mentioned earlier, a key difference between a basic z/OS sub-capacity report and a sub-capacity report generated using zWPC is the Cloud MSU Reduction column in the Product Max Contributors section. This column is generated whenever public cloud pricing data is provided via the **Assign ZWPC** control statement in the SPECIAL DD, as was the case in the configuration (see [“Sample configuration for INPUTCSV-based IBM Z Workload Pricing for Cloud”](#) on page 260).

The Cloud MSU Reduction column shows the adjustment that SCRT applied to products' sub-capacity values using the zWPC algorithm. In the Product Max Contributors section for the sample configuration, the peak z/OS sub-capacity value and the value of all products based entirely on the 4-hour rolling average MSU utilization of the LPAR were adjusted by SCRT using the value shown in the Cloud MSU Reduction column.

For products that are not based on the 4-hour rolling average MSU utilization of the LPAR, the product is not adjusted by zWPC. In the sample configuration, an example of such a product is IBM WebSphere Application Server for z/OS V8, which is based on the Getting Started Sub-Capacity Pricing algorithm.

In the Product Max Contributors section for the sample configuration, note that the MSU values shown for LPAR SYS1 already reflect the contribution of zWPC WLM public cloud workloads that ran in each LPAR during the peak hour. If multiple LPARs are present, the Cloud MSU Reduction column reflects the total reduction across all LPARs for that hour.

Product Grid Snapshot section

There are no differences between the Product Grid Snapshot section for a basic z/OS sub-capacity report and a sub-capacity report that is generated using zWPC. For details about this section, see [“Product Grid Snapshot section”](#) on page 123.

The sample Product Grid Snapshot section appears in [Figure 137 on page 276](#).

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                                Product ID    LPAR1
z/OS V2                                    5650-Z0S      3.30%
DB2 11 for z/OS                           5615-DB2      3.30%
CICS TS for z/OS V5                       5655-Y04      3.40%
IBM MQ for z/OS V8                        5655-W97      3.40%
IMS V13                                   5635-A04      3.40%
IBM WebSphere Application Server for z/OS V8 5655-W65      3.30%
IBM Integration Bus for z/OS V9            5655-IBB      3.40%
IBM Multi-site Workload Lifeline V2        5655-UM4      3.30%

```

Figure 137. Sample sub-capacity report for Integrated Workload Pricing: Product Grid Snapshot section

Getting Started Sub-Capacity Pricing Assigned Hours section

There are no differences between the Getting Started Sub-Capacity Pricing Assigned Hours section for a basic z/OS sub-capacity report and a sub-capacity report that is generated using zWPC. For details about this section, see [“Getting Started Sub-Capacity Pricing Assigned section” on page 179](#).

The sample Getting Started Sub-Capacity Pricing Assigned Hours section appears in [Figure 138 on page 276](#).

```

==S7=====
GETTING STARTED SUB-CAPACITY PRICING ASSIGNED - HOURS

Product Name                                Product ID    LPAR
                                                LPAR1
IBM WebSphere Application Server for z/OS V8 5655-W65      24
IBM Integration Bus for z/OS V9             5655-IBB      24

```

Figure 138. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Getting Started Sub-Capacity Pricing Assigned Hours section

Mobile/zCAP Pricing Detailed Data section for zWPC

The optional Mobile/zCAP Pricing Detailed Data section appears when the optional **Generate_Detailed_Data** control statement is specified in the SPECIAL DD. This section shows the public cloud MSU reduction that SCRT applied for WLM-based IBM Z Workload Pricing for Cloud, per LPAR. This information is shown for each sub-capacity product on an hour-by-hour basis.

As shown in [Figure 139 on page 278](#), the optional Mobile/zCAP Pricing Detailed Data section reports several columns for each product:

- Date and time of the adjusted interval.
- CPC Type-model.
- CPC rated machine capacity, in MSU.
- Adjusted MSU values for each LPAR. (Note that the zWPC algorithm has already been applied.)
- The WLM-based zWPC reduction that was applied per defining product, per partition.
- Total cloud MSU reduction applied across all WLM-based zWPC defining products and any INPUTCSV-based zWPC defining products.
- CPC Total MSU, which is the sum of the adjusted LPAR MSU values across all partitions for the adjusted product.

There are several additional values that can appear in section V5 of a sub-capacity report:

-

Indicates that there is no data for this interval.

dd(rr)

Indicates that the defined capacity limit for the LPAR was reached during this interval. The first number (*dd*) represents the defined capacity value used for this hour. The second number (*rr*) represents the 4-hour rolling average value prior to the capping.

G(rr)

Indicates that the sum of the partitions exceeded the group capacity limit; the group capacity was used for this hour for that set of LPARs. The *rr* value represents the 4-hour rolling average value prior to the capping.

P(rr)

Indicates that the sum of the partitions exceeded the processor capacity; the processor capacity was used. The *rr* value represents the 4-hour rolling average value prior to the capping.

(GSSP)

Indicates that a WebSphere program eligible for Getting Started Sub-Capacity Pricing received the GSSP value so there was no additional adjustment.

nnn(CPS)

Indicates that an LPAR was assigned to a Tailored Fit Pricing solution for this hour. In this case, the LPAR does not contribute to the product sub-capacity value. The *nnn* value represents what the LPAR sub-capacity value would have been if it were part of the standard (non-container) environment.

Guideline: The information in the Mobile/zCAP Detailed Data section of the sub-capacity report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating sub-capacity reports for submission to IBM. In the event that IBM requests that you submit a sub-capacity report that includes section V5, you can use the LMS web application to submit the report.

The following descriptions of the IBM Z Workload Pricing for Cloud algorithm are based on the sample configuration in [“Sample configuration for INPUTCSV-based IBM Z Workload Pricing for Cloud”](#) on page 260.

Note: Only a subset of the products running in the sample environment are included here.

Results for z/OS V2R1

After SCRT applied the appropriate algorithm hour by hour, SCRT reported 01 Nov 2024 at 12:00 as the earliest hour with the peak sub-capacity value for z/OS, as indicated in the sample Product Max Contributors section.

For 01 Nov 2024 at 12:00, SCRT applied the appropriate algorithms and calculated a sub-capacity value of 1284 MSUs for LPAR1. As shown in the WLM Cloud Partition column for LPAR1, this includes a 82-MSU zWPC reduction derived from the WLM classification 4-hour rolling average data found in the SMF type 70 records.

[Figure 139 on page 278](#) illustrates these results.

```
==V5=====
```

MOBILE/ZCAP PRICING DETAILED DATA

MLC Products

5650-ZOS

z/OS V2

Date Time of RMF Interval	Type	CPC MSU	Partition LPAR1	WLM Cloud Partition LPAR1	Total Cloud MSU Reduction	CPC Total
30 Oct 2024 - 23:00	2817-760	5001	-	-	0	0
01 Nov 2024 - 00:00	2817-760	5001	778	49	49	778
01 Nov 2024 - 01:00	2817-760	5001	781	49	49	781
01 Nov 2024 - 02:00	2817-760	5001	829	52	52	829
01 Nov 2024 - 03:00	2817-760	5001	809	52	52	809
01 Nov 2024 - 04:00	2817-760	5001	793	50	50	793
01 Nov 2024 - 05:00	2817-760	5001	795	50	50	795
01 Nov 2024 - 06:00	2817-760	5001	803	51	51	803
01 Nov 2024 - 07:00	2817-760	5001	848	54	54	848
01 Nov 2024 - 08:00	2817-760	5001	947	60	60	947
01 Nov 2024 - 09:00	2817-760	5001	1072	68	68	1072
01 Nov 2024 - 10:00	2817-760	5001	1176	75	75	1176
01 Nov 2024 - 11:00	2817-760	5001	1280	81	81	1280
01 Nov 2024 - 12:00	2817-760	5001	1284	82	82	1284
01 Nov 2024 - 13:00	2817-760	5001	1274	81	81	1274
01 Nov 2024 - 14:00	2817-760	5001	1270	80	80	1270
01 Nov 2024 - 15:00	2817-760	5001	586	37	37	586
01 Nov 2024 - 16:00	2817-760	5001	169	10	10	169
01 Nov 2024 - 17:00	2817-760	5001	340	22	22	340
01 Nov 2024 - 18:00	2817-760	5001	467	29	29	467
01 Nov 2024 - 19:00	2817-760	5001	589	37	37	589
01 Nov 2024 - 20:00	2817-760	5001	565	35	35	565
01 Nov 2024 - 21:00	2817-760	5001	497	31	31	497
01 Nov 2024 - 22:00	2817-760	5001	480	30	30	480
01 Nov 2024 - 23:00	2817-760	5001	488	31	31	488

Figure 139. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Mobile/zCap Pricing Detailed Data section — z/OS results

Results for Db2 11 for z/OS

SCRT applied the process as described in “Results for z/OS V2R1” on [page 277](#). Like z/OS, Db2 11 for z/OS is based on the 4-hour rolling average MSU utilization of the LPAR.

This same process applies for all sub-capacity programs that reported a peak MSU value of 1284 with 01 Nov 2024 at 17:00 as the earliest hour with the peak sub-capacity value.

[Figure 140 on page 279](#) illustrates the results.

5615-DB2		DB2 11 for z/OS					
Date	Time of RMF Interval	Type	CPC MSU	Partition LPAR1	WLM Cloud Partition LPAR1	Total Cloud MSU Reduction	CPC Total
30 Oct 2024	- 23:00	2817-760	5001	-	-	0	0
01 Nov 2024	- 00:00	2817-760	5001	778	49	49	778
01 Nov 2024	- 01:00	2817-760	5001	781	49	49	781
01 Nov 2024	- 02:00	2817-760	5001	829	52	52	829
01 Nov 2024	- 03:00	2817-760	5001	809	52	52	809
01 Nov 2024	- 04:00	2817-760	5001	793	50	50	793
01 Nov 2024	- 05:00	2817-760	5001	795	50	50	795
01 Nov 2024	- 06:00	2817-760	5001	803	51	51	803
01 Nov 2024	- 07:00	2817-760	5001	848	54	54	848
01 Nov 2024	- 08:00	2817-760	5001	947	60	60	947
01 Nov 2024	- 09:00	2817-760	5001	1072	68	68	1072
01 Nov 2024	- 10:00	2817-760	5001	1176	75	75	1176
01 Nov 2024	- 11:00	2817-760	5001	1280	81	81	1280
01 Nov 2024	- 12:00	2817-760	5001	1284	82	82	1284
01 Nov 2024	- 13:00	2817-760	5001	1274	81	81	1274
01 Nov 2024	- 14:00	2817-760	5001	1270	80	80	1270
01 Nov 2024	- 15:00	2817-760	5001	586	37	37	586
01 Nov 2024	- 16:00	2817-760	5001	169	10	10	169
01 Nov 2024	- 17:00	2817-760	5001	340	22	22	340
01 Nov 2024	- 18:00	2817-760	5001	467	29	29	467
01 Nov 2024	- 19:00	2817-760	5001	589	37	37	589
01 Nov 2024	- 20:00	2817-760	5001	565	35	35	565
01 Nov 2024	- 21:00	2817-760	5001	497	31	31	497
01 Nov 2024	- 22:00	2817-760	5001	480	30	30	480
01 Nov 2024	- 23:00	2817-760	5001	488	31	31	488

Figure 140. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Mobile/zCap Pricing Detailed Data section — Db2 11 for z/OS results

Results for IBM WebSphere Application Server for z/OS V8

SCRT applied the Getting Started Sub-Capacity Pricing for z/OS IPLA Software process to IBM WebSphere Application Server for z/OS V8. Since the terms and conditions of GSSP were met in the sample configuration, and the GSSP values were less than the traditional sub-capacity values for each hour of the reporting period, SCRT assigned IBM WebSphere Application Server for z/OS V8 the GSSP value for each hour in the reporting period.

When the GSSP value is used for a GSSP eligible product, the Mobile/zCAP Pricing Detailed Data section places the (GSSP) token in the **WLM Cloud Partition** column, as shown in Figure 141 on page 280. In accordance with the IBM Z Workload Pricing for Cloud addendum, the zWPC reduction is not applied when GSSP pricing is used instead of the traditional sub-capacity value.

In the sample configuration, the peak MSU value used for IBM WebSphere Application Server for z/OS V8 is 3, which is achieved during the 01 Nov 2024 - 03:00 hour.

For information about the GSSP process, see [Chapter 8, “Sub-capacity reporting for Getting Started Sub-Capacity Pricing \(GSSP\),” on page 165.](#)

5655-W65		IBM WebSphere Application Server for z/OS V8				Total		CPC Total
Date	Time of RMF Interval	Type	CPC MSU	Partition LPAR1	WLM Cloud Partition LPAR1	Cloud MSU Reduction	MSU	
30 Oct 2024	- 23:00	2817-760	5001	-	-	0	0	0
01 Nov 2024	- 00:00	2817-760	5001	1	(GSSP)	0	0	1
01 Nov 2024	- 01:00	2817-760	5001	1	(GSSP)	0	0	1
01 Nov 2024	- 02:00	2817-760	5001	2	(GSSP)	0	0	2
01 Nov 2024	- 03:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 04:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 05:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 06:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 07:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 08:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 09:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 10:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 11:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 12:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 13:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 14:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 15:00	2817-760	5001	3	(GSSP)	0	0	3
01 Nov 2024	- 16:00	2817-760	5001	2	(GSSP)	0	0	2
01 Nov 2024	- 17:00	2817-760	5001	2	(GSSP)	0	0	2
01 Nov 2024	- 18:00	2817-760	5001	2	(GSSP)	0	0	2
01 Nov 2024	- 19:00	2817-760	5001	2	(GSSP)	0	0	2
01 Nov 2024	- 20:00	2817-760	5001	2	(GSSP)	0	0	2
01 Nov 2024	- 21:00	2817-760	5001	2	(GSSP)	0	0	2
01 Nov 2024	- 22:00	2817-760	5001	2	(GSSP)	0	0	2
01 Nov 2024	- 23:00	2817-760	5001	3	(GSSP)	0	0	3

Figure 141. Sample sub-capacity report for IBM Z Workload Pricing for Cloud: Mobile/zCap Pricing Detailed Data section — IBM WebSphere Application Server for z/OS results

Chapter 13. Sub-capacity reporting for z/TPF systems

With release 1.1 of z/Transaction Processing Facility Enterprise Edition (z/TPF) (5748-T15), z/Architecture CPCs that contain both z/OS systems and z/TPF systems can report sub-capacity MSUs for z/TPF and z/TPFDF. For this initial method of sub-capacity reporting, z/OS must be running natively on the LPAR (not as a guest under z/VM) and must be running concurrently with z/TPF on the CPC. That z/OS system must also collect both SMF70 and SMF89 records; otherwise, SCRT rejects all of the SMF data collected for that z/OS system.

Starting with z/TPF 1.1 PUT 02, you have the option to allow z/TPF to produce SCRT89 accounting records to report sub-capacity MSUs on CPCs that are running z/TPF and z/TPFDF. If you elect to have z/TPF, which must be running natively on the LPAR (not as a guest under z/VM), produce SCRT89 records, this method allows you to have CPCs that run standalone z/TPF and you only need a minimum of one system or LPAR (or one VM guest) running z/OS in your enterprise. This minimum of one z/OS system is needed in order to use SCRT to process the SCRT89 records.

For both of these methods (using or not using SCRT89 records), your system must follow the LPAR naming convention described in [“Requirements for sub-capacity pricing for z/TPF” on page 281](#).

If you do not use SCRT89 records, sub-capacity pricing is based on the capacity of the LPARs in which z/TPF or z/TPFDF run:

- For the LPARs with dedicated engines, the pricing is based on their maximum capacity.
- For LPARs with shared engines, the capacity is based on the 4-hour rolling average utilization of the LPAR.

If you collect and use SCRT89 records for native z/TPF systems or for z/TPF guest systems, sub-capacity pricing is based on the 4-hour rolling average utilization for those systems. Such sub-capacity pricing may be less than the sub-capacity pricing based on LPARs for which no SCRT89 records are collected because the SCRT89 records exclude the time used by the z/TPF Recoup Utility. Thus, IBM recommends collecting SCRT89 records from the z/TPF systems that run the Recoup Utility. For details about how SCRT deals with CPC configurations where at least one z/TPF system collects SCRT89 records and at least one does not, see [Table 44 on page 289](#).

This chapter describes how to use both of these methods to obtain sub-capacity pricing for sub-capacity eligible software running on z/TPF systems. The chapter also shows sample sub-capacity reports that include z/TPF LPARs and describes the z/TPF-specific fields of the report. For a description of other fields in the reports, see [Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103](#).

Note: z/TPF APARs PJ38004 (z/TPF PUT 07) and PJ38530 must be applied before running z/TPF on a zEC12, z196, zBC12, or z114 CPC in order for z/TPF to correctly calculate the sub-capacity MSU values for these CPCs.

For CPCs that support power save mode (where the CPC can reduce its power requirements during periods where the full capacity of the CPC is not required, but at a decrease in CPC performance), see [“Static power save mode” on page 506](#) for a description of power save mode and additional details about the z/TPF support for power save mode.

Requirements for sub-capacity pricing for z/TPF

In addition to the requirements listed in [“Preparing your hardware and software” on page 11](#), IBM requires all z/TPF customers to use a naming convention for any LPAR running z/TPF. Failure to use the naming convention can result in an incorrect sub-capacity report and billing at full capacity. SCRT must be able to recognize a z/TPF LPAR and produce the necessary report for VWLC billing. A z/TPF LPAR name must be in the format TPF α nnnn, where α is either the character P for production or the character T for test, and where nnnn is any alphanumeric character string that you choose.

For z/TPF guest systems, you must either provide the SCRT89 records for all of the z/TPF guests running within an LPAR, or you must not provide any SCRT89 records for any of the z/TPF guests running in that LPAR. This requirement applies on an LPAR-by-LPAR basis. If the CPC with z/TPF systems also has LPARs that run other operating systems, avoid mixing z/TPF guest systems together with other operating system guests in the same LPAR unless you collect SCRT89 records from all of the z/TPF guest systems in that LPAR. Do not use the format required for z/TPF LPAR names (TPF $xnnnn$) to name the LPARs that are only used by other operating systems. If your existing LPAR names already use this LPAR naming convention for non-z/TPF operating systems, see the guideline in [“Unusual situations that affect utilization”](#) on page 17.

Identities reported for z/TPF systems

SCRT uses several formats to report the identities of z/TPF systems based on the configuration in which those z/TPF systems run, whether or not SCRT89 records are being collected, and, sometimes, whether or not APAR PJ38004 has been applied. Table 43 on page 282 summarizes the formats that SCRT uses to identify z/TPF systems that are reported in the Detail Data Collection section, SMF / SCRT89 Input Data Statistics section, Detail LPAR Data Section, Product Max Contributors section, and Product Grid Snapshot section.

Table 43. Summary of z/TPF system identities reported by SCRT

z/TPF configuration	z/TPF APAR PJ38004 applied?	Collecting SCRT89 records?	Format of z/TPF system identity
z/TPF native system with no other sub-capacity operating systems sharing the LPAR	–	yes	<i>LPAR_name</i> Example: TPFP1234
z/TPF native system sharing the LPAR with z/OS at different times	–	yes	<i>LPAR_name (z/TPF)</i> Example: TPFP1234(z/TPF)
z/TPF native system or z/TPF guest system	–	no	<i>LPAR_name*</i> Example: TPFP1234*
z/TPF guest system(z/VM level-1 guest)	yes	yes	<i>LPAR_name (VM_userid)</i> Example: TPFP1234(USER1)
	no	yes	<i>*lpar*(VM_userid)</i> Example: *lpar*(USER1)
z/TPF guest system(z/VM level-2 or higher guest)	yes	yes	<i>LPAR_name (VM_userid-VM_level)</i> Example: TPFP1234(USER1-2)
	no	yes	<i>*lpar*(VM_userid-VM_level)</i> Example: *lpar*(USER1-2)

Sub-capacity reporting without SCRT89 records

In order to simplify the process of collecting and managing the SCRT89 record files for large numbers of z/TPF guest systems, IBM recommends (but does not require) that you not collect SCRT89 records from z/TPF guest systems running in the same LPAR, whenever *all* of the following conditions are met. Note that this recommendation works equally well when only a small number of z/TPF systems run in the same LPAR.

1. There is at least one native z/OS system running on the same CPC whenever any z/TPF guest system is running in the z/TPF guest LPAR, and that native z/OS system is collecting SMF70 and SMF89 records. This requirement can also be satisfied by a native z/TPF system running on the same CPC whenever any z/TPF guest system is running in the z/TPF guest LPAR, and that native z/TPF system is collecting SCRT89 records.

2. None of the z/TPF guest systems running in the z/TPF guest LPAR are running the Recoup utility.
3. All of the guest systems are z/TPF systems. There are no non-z/TPF guests running at any time in the z/TPF guest LPAR.

For every hour that SCRT does not receive SCRT89 records from a busy z/TPF LPAR, it automatically calculates the 4-hour rolling average for the z/TPF systems running in that LPAR based on the busy time of the entire z/TPF LPAR. When the systems running in the z/TPF guest LPAR meet all of the previously listed conditions, SCRT will calculate an MSU value that would be the equivalent of the sum of the MSU values reported on the SCRT89 records for each of those z/TPF systems. See [Table 44 on page 289](#) for more details about how SCRT reports on native and guest z/TPF systems that do not collect SCRT89 records.

The methodology of not collecting SCRT89 records from such z/TPF guest systems has the added advantage of having SCRT calculate a more accurate MSU value for the z/TPF work running in the LPAR whenever idle z/TPF guests are running. The SCRT calculation is more accurate because z/TPF assigns a minimum of one MSU per SCRT89 record for every z/TPF system, while SCRT calculates the z/TPF MSUs based on the actual busy time of the z/TPF systems running in the LPAR, without regard for the number of z/TPF systems.

Whenever one or more of the z/TPF guests use the Recoup utility or whenever non-z/TPF guests run in the z/TPF guest LPAR, IBM recommends collecting SCRT89 records from all z/TPF guests running in that LPAR; otherwise, SCRT automatically uses the busy time for the entire LPAR to calculate the z/TPF MSUs for that LPAR. In this case, the busy time for that LPAR includes the time used by the Recoup utility and the busy time used by all of the non-z/TPF guests running in that LPAR, which means that SCRT will calculate a higher MSU value for such z/TPF LPARs than if SCRT were to process the SCRT89 records for all of the z/TPF systems that ran in that LPAR.

Sub-capacity reporting using SCRT89 records

z/TPF produces SCRT89 accounting records for use in sub-capacity reporting. SCRT processes the data from the SCRT89 records to generate a sub-capacity report based on 4-hour rolling averages of z/TPF usage. If you have z/TPF 1.1 PUT 02 or higher, you can use the following steps to record and process SCRT89 data for z/TPF.

Note: z/TPF assigns a minimum of one MSU to every z/TPF system that collects SCRT89 records. The result is that an LPAR which supports a large number of z/TPF guest systems, each of which is collecting SCRT89 records, may cause SCRT to report an MSU value for all of the z/TPF guests logged on under that LPAR which is larger than the busy time used by those z/TPF guests. Consider using the technique of not collecting SCRT89 records from such guest LPARs (as discussed in [“Sub-capacity reporting without SCRT89 records” on page 282](#)) when the limitations of that technique do not apply to the LPAR—for example, when the Recoup utility is not run and all of the guests in the LPAR are z/TPF guests.

Running z/TPF and non-z/TPF guests in the same LPAR

Although IBM recommends not running both z/TPF and non-z/TPF guests in the same z/TPF guest LPAR, your development or test processes may require you to do so. In order for SCRT to report on the capacity used exclusively by the z/TPF systems running in a such z/TPF guest LPAR, your process must collect the SCRT89 records from all of the z/TPF guest systems running in the LPAR and have those records processed by SCRT. However, it is possible that the non-z/TPF guests might be running during hours when no z/TPF guest is running. For those hours, SCRT will automatically use the busy time of non-z/TPF systems running in that LPAR to calculate the MSUs to be assigned to z/TPF, which could result in SCRT reporting a higher MSU value for z/TPF than would be the case if there were no non-z/TPF guests.

You can take the following actions to prevent SCRT from reporting the busy time of non-z/TPF systems, as long as all of the z/TPF guests in the LPAR are running with z/TPF APAR PJ38004 applied and all z/TPF guests are collecting SCRT89 records. When both of these conditions are met, you may use the SCRT **Exclude** control statement, to exclude all of the busy time from the LPAR in which those z/TPF guests ran. SCRT will then ignore the busy time for all non-z/TPF systems when the required conditions are met. This methodology excludes the busy time used by the Recoup utility when calculating the z/TPF MSUs.

1. Set the **Image_Id** value on the **Exclude** control statement to the name of the z/TPF guest LPAR. When multiple z/TPF LPARs meet the required conditions for using an **Exclude** control statement, as described here, each z/TPF LPAR requires its own **Exclude** control statement.
2. Omit the use of the **Product_Id** parameter on the **Exclude** control statement because you are excluding all of the data from this LPAR and because z/TPF and z/TPFDF are considered to be operating system products and **Exclude** does not support the specific exclusion of products that are operating systems.
3. Set the values for the **Start** and **Resume** parameters to include the entire reporting period.

If the z/TPF guest LPAR always meets the required conditions, then the **Start** value could be chosen to be one day before the start of the first reporting period where this **Exclude** control statement is used, and the **Resume** value should be set to 2099/12/31/23 (with the UTC suffix added, if necessary). If the dates are set this way, this same **Exclude** control statement can be used in future SCRT jobs as long as the required conditions continue to be met.
4. In the Special Conditions section of the sub-capacity report, replace the (required) placeholder in the **Customer Comments** column with the following text for each **Exclude** control statement listed that was used to handle this condition:

This LPAR only runs z/TPF guest systems which collect SCRT89 records or non-z/TPF guest systems.

See [“Sub-capacity report for a CPC running z/TPF systems and non-z/TPF systems”](#) on page 290 for an example of an **Exclude** control statement that was used for a z/TPF guest LPAR where both z/TPF and non-z/TPF guests were running.

IBM will reject sub-capacity reports that use the **Exclude** control statements without meeting all of the specified conditions.

Multiple sub-capacity operating systems

When z/TPF shares a CPC with z/OS and/or z/VSE, the SCRT89 records for z/TPF and z/VSE and the SMF records for z/OS running on the same CPC must be input to the same execution of SCRT so that the z/TPF, z/VSE, and z/OS products appear on the same report. Because z/TPF will be creating a single file of SCRT89 records for all of the z/TPF systems running in the same z/TPF complex, you must input the SCRT89 records for z/TPF and z/VSE and the SMF records for z/OS for all of the CPCs that are in the same z/TPF complex in the same execution of SCRT.

You can input the SCRT89 records for all of the z/TPF and z/VSE systems in the data center together with all of the SMF records from all of the z/OS systems in the data center into a single execution of SCRT, if desired.

Then, produce a sub-capacity report by having SCRT process all of the SCRT89 records from the z/TPF and z/VSE systems, together with the SMF records from all of the z/OS systems which share the same CPC as the z/TPF systems. You can also consolidate all of your SMF and SCRT89 records for a CPC into a single DASD or tape file before running SCRT.

Steps for z/TPF data collection and processing

The following steps use z/TPF's **ZSUBC** command. For details about the **ZSUBC** command, including syntax and parameter descriptions, see *z/TPF Operations*. For command output, see *z/TPF Messages*. For online help on z/TPF, enter: ZSUBC HELP

1. **Initialize SCRT89 data recording.** To start the recording of SCRT89 records on z/TPF, after IPL enter the **ZSUBC INITIALIZE** command. This command must be entered once on any logical z/TPF processor in the loosely coupled complex to create sub-capacity records. The command initializes a control structure on DASD that is shared by all logical processors within the loosely coupled complex and creates the /IBMSubcapacity directory in the z/TPF file system. You can specify an optional target subsystem where the files will be created, specifying the four-character subsystem name or BSS

for the basic subsystem. For example, the following command specifies that the files be created on the basic subsystem:

```
ZSUBC Initialize Target-BSS
```

2. **Specify a recording period.** To specify the number of days to keep sub-capacity records in the predefined directory, issue the command:

```
ZSUBC KEEP nnn
```

where *nnn* specifies the number of days to keep sub-capacity files in the directory before they are deleted automatically. Specify a number from 0 to 999. If the parameter is set to 0 (the default), no sub-capacity files are created. For example, to keep records for 365 days, enter the command:

```
ZSUBC KEEP 365
```

3. **Identify the SCRT89 records.** Each z/TPF LPAR in the complex creates a new file of SCRT89 records for every day that the SCRT89 recording has been active. The SCRT89 file names take the following format:

```
MachineType_PlantOfManufacture_Serialnumber_LPARname_date.scr
```

In the file name:

- The *date* can be in any of the following formats, depending upon the user's system setting: *mmddyyyy*, *ddmmyyyy*, *yyyymmdd* as specified by the **DATEFMT** parameter of the SIP **CONFIG** macro.
- The *LPARname* can be in one of two formats:
 - For z/TPF systems that are running natively on an LPAR, *LPARname* is the actual name of the LPAR.
 - When the z/TPF system is a VM guest system, *LPARname* is the user ID assigned to the virtual machine under which the z/TPF guest system runs.

For example, the following file name is for SCRT89 records from 4 March 2006, for an LPAR named TFPF7 on a model 2084 CPC with a serial number of d153f:

```
2084_02_d153f_TFPF7_03042006.scr
```

4. **Collect the SCRT89 data.** The reporting period for SCRT is from the second day of the month to the first day of the following month. However, creating a file of SCRT89 records which starts exactly on the first day of the reporting period and ends exactly on the last day of the reporting period could result in SCRT reporting higher MSUs for z/TPF than it should.

This potential reporting problem can occur when the sub-capacity report uses UTC dates and times. When that happens, SCRT converts local dates/times on the SMF and SCRT89 records for every z/OS, z/TPF, and z/VSE system on that CPC to UTC dates/times. Based on that conversion, the earliest UTC hour on the SCRT89 records from some z/TPF systems could be shifted several hours later than the earliest UTC hour for other z/OS or z/TPF systems. The result is that SCRT uses the LPAR busy time from the SMF or SCRT89 records to calculate the z/TPF MSUs for the z/TPF systems whose SCRT89 records were provided in the SCRT89 record file and whose earliest UTC hour is now one or more hours after the start of the UTC reporting period.

The same condition can also occur at the end of the UTC reporting period, except that the last UTC hour being reported on SCRT89 records for some z/TPF systems are earlier than the last hour in the UTC reporting period.

Since z/TPF excludes the busy time from certain functions when it calculates and reports the z/TPF MSUs on SCRT89 records, the MSUs reported on those SCRT89 records might be less than the MSUs calculated by SCRT based on the LPAR's busy time. For more information, see [“Sub-capacity reporting without SCRT89 records”](#) on page 282.

Guideline: When creating a file with the SCRT89 records for a reporting period, always include the SCRT89 records from one day before and one day after the reporting period. Creating a file with those additional SCRT89 records will provide SCRT with the best opportunity for using the z/TPF MSUs reported on an SCRT89 record, instead of calculating the potentially higher z/TPF MSUs based on the z/TPF LPAR busy time. The remainder of the instructions for this step are based on this guideline.

On the third day of the month, use the **ZSUBC COPY** command to copy the daily SCRT89 files into a larger file for the preceding month. Specify the beginning and ending dates and a DD name for the file to contain the records. For example, to copy the files from 1 February 2007 through 2 March 2007 into a file with DD name FEBFILE, use the following syntax:

```
ZSUBC Copy Begin-02012007 End-03022007 DD-FEBFILE
```

The DD name must be defined to the z/TPF system using the **ZDSMG DEFINE** command. Enter this command on the subsystem specified on the **Target** parameter (the default is the basic subsystem: BSS). The begin date and end date must follow the format specified by the **DATEFMT** parameter of the **SIP CONFIG** macro.

5. **Transfer and process the SCRT89 records.** You must transfer the SCRT89 records from your z/TPF system to the target system where those SCRT89 records will be processed by SCRT. The target system can be a z/OS system or a Linux or Windows system.

- **For a z/OS target system:**

You can transfer the SCRT89 records either by using blocked tape files or by using FTP, as described in the following scenarios. Choose whichever medium (tape or files for FTP) you prefer to use for archiving your SCRT89 data for the required six-month period.

- **Scenario 1 – Using tape:** After collecting the SCRT89 records, transfer all of your SCRT89 records for the reporting period to one or more blocked tape files using z/TPF commands. For example, use commands such as:

```
ZTLBL SUB C 0 BLK COMP F$ZTPF.SUBCAPACITY$
ZTLBL SUB U
ZDSMG DEF SUBTAPE TAPE-SUB
ZTMNT SUB cuu AO           where cuu is the tape address
ZSUBC COPY BEGIN-02012007 END-03022007 DD-SUBTAPE
```

The **ZTLBL** and **ZDSMG** commands only need to be entered once.

Note: In this scenario the SCRT89 record utility is not required and is not used.

- **Scenario 2 – Using FTP:** After collecting the SCRT89 records, use the **ZSUBC COPY** command to transfer all of your SCRT89 records for the reporting period to a single z/TPF file. For example:

```
ZDSMG DEF FEBFILE PATH- '/IBMSubcapacity/ComplexA_February2006_Subcapacity.scr'
ZSUBC COPY BEGIN-02012007 END-03022007 DD-FEBFILE
```

Then, FTP that file to a z/OS UNIX file system. You can then use the SCRT89 record utility to process the records. If you already have procedures in place to retain, archive, and purge old SMF data, it would probably be simplest to use the SCRT89 record utility to convert these SCRT89 records to conventional z/OS VBS formatted (or perhaps just VBS formatted) records and archive them with your other SMF records.

Using the SCRT89 record utility: You can use the SCRT89 record utility to convert the SCRT89 records in the z/OS UNIX file system into a standard z/OS VBS file. Then, produce a sub-capacity report by concatenating the converted SCRT89 file with the SMF files from all of the z/OS systems as the input to the existing SMF DD. You can also consolidate all of your SMF and SCRT89 records into a single DASD or tape file before running SCRT. For more information, see [“Installing and using the SCRT89 record utility”](#) on page 287.

- **For a Linux or Windows target system:**

You must use FTP to transfer the SCRT89 records from your z/TPF system to the Linux or Windows system where those records will be processed by SCRT.

After collecting the SCRT89 records, use the **ZSUBC COPY** command to transfer all of the SCRT89 records for the reporting period to a single z/TPF file. For example:

```
ZDSMG DEF FEBFILE PATH- '/IBMsubcapacity/ComplexA_February2006_Subcapacity.scr'
ZSUBC COPY BEGIN-02012007 END-03022007 DD-FEBFILE
```

Then, FTP that file in binary mode to the desired Linux or Windows system.

6. **Run SCRT.** If z/OS, z/TPF, and z/VSE run on the same CPC, you must process the SMF records for all of the z/OS systems and the SCRT89 records for all of the z/TPF and z/VSE systems in the same execution of SCRT.
 - If you are running SCRT on a z/OS system, see [Chapter 3, “Using SCRT on z/OS systems,”](#) on page 21 for details about how to specify the required DD statements and report parameters.
 - If you are running SCRT on a Linux or Windows system, see [Chapter 4, “Using SCRT on Windows and Linux systems,”](#) on page 69 for details about how to specify report parameters.
7. **Retain the SCRT data.** IBM's sub-capacity contract requires you to save the SCRT89 records for at least six months in order to be able to resolve questions and problems. If you choose to have SCRT process the SCRT89 records directly from the UNIX file system, you must retain the SCRT89 records for six months either on the z/TPF system or in the UNIX file system.

Installing and using the SCRT89 record utility

Note: The SCRT89 record utility is only required if you are processing SCRT89 records on a z/OS system. If you are using the Linux or Windows version of SCRT, the SCRT89 record utility is not required.

Download the SCRT89 record utility (SCRT89RU) from [SCRT89 Record Utility \(ibm.biz/SCRT89RU\)](http://ibm.biz/SCRT89RU).

The SCRT89 record utility is contained in a self-extracting .zip file. After downloading the SCRT89 utility perform the following steps:

1. Run the `scrt89ru.exe` file. The extracted, uncompressed file is called `scrt89ru.bin`.
2. Create a data set on the host system:
 - Use ISPF panel **3.2** to allocate a new 80-byte, fixed data set called `userid.T00L.JCL`.
 - Use these values for the data set:

Attribute	Value
Record format	FB (fixed block)
Record length	80
Directory blocks	0
Block size	0 (system-determined block size)

3. Upload the `scrt89ru.bin` file to the host system from your workstation:
 - Open a command prompt on your workstation.
 - Change to the directory that contains the `scrt89ru.bin` file.
 - Begin an FTP session:

```
ftp hostname
```

- When prompted, enter your user ID and password on that host.
- Change the FTP transfer mode to binary:

```
bin
```

- Move the `scrt89ru.bin` file to the host, overwriting the `userid.TOOL.JCL` data set that you just created:

```
put scrt89ru.bin tool.jcl
```

4. On the host system, customize the job. The job has two steps. The first step deletes the output data set. The second step reads the data from the z/OS UNIX file system and reformats it into an SMF compatible data set.

- Customize the job card to fit your environment:

```
//SCRT89RU JOB (, ),MSGLEVEL=1,MSGCLASS=0,NOTIFY=NAME
```

- Change the SMFUSSFS DD to point to your data in the z/OS UNIX file system:

```
//SMFUSSFS DD PATH='/u/HLQ/ComplexA_February2006_Subcapacity.scr'
```

- Change the output dataset name to fit your environment on both steps. The names should match.

```
//DD1      DD DSN=DSN=HLQ.COMPLEXA_FEB06_SCRT89,UNIT=SYSDA,
//          SPACE=(TRK,(1,0)),DISP=(MOD,DELETE)

//SMFVB     DD DISP=(,CATLG),DSN=DSN=HLQ.COMPLEXA_FEB06_SCRT89,
//          UNIT=SYSDA,SPACE=(CYL,(2,2))
```

Note: The SCRT89 record utility is only intended for use with SCRT89 records; it does not support the processing of any other records.

Using the SCRT89 record utility

Use the SCRT89 record utility to convert SCRT89 records into VBS format records that can be processed by the SCRT. The SCRT89 records must first be FTPed to a z/OS UNIX file as described in step “5” on page 286. JCL to run the SCRT89 record utility is shown in Figure 142 on page 288. Update the JCL as follows:

1. Specify the UNIX path of the input file name on the SMFUSSFS DD statement.
2. Specify the desired output file name in the DD1 DD statement and in the SMFVB DD statement.

```
//SCRT89RU JOB (????,????),MSGCLASS=0,NOTIFY=HLQ
//*
//* SCRT89 RECORD UTILITY
//*
//* THIS IS A PROGRAM TO CONVERT SCRT89 RECORDS PRODUCED BY
//* TPF AND TRANSPORTED TO A UNIX FILE SYSTEM ON A Z/OS SYSTEM
//* TO CONVENTIONAL SMF VBS RECORDS WHICH CAN BE PROCESSED BY
//* TRADITIONAL SMF UTILITIES.
//*
//* OUTPUT FILE IS DELETED IN THE DELETE STEP AND CREATED
//* IN THE FINAL STEP. SPECIFY THE NAME IN BOTH PLACES.
//*
//* INPUT FILE IS AT BOTTOM WITH DD NAME SMFUSSFS.
//* SPECIFY THE UNIX PATH HERE.
//
// EXEC PGM=IEFBR14
//DD1      DD DSN=HLQ.SCRT89RU.SMF,UNIT=SYSDA,
//          SPACE=(TRK,(1,0)),DISP=(MOD,DELETE)
// EXEC PGM=LOADER
//SMFUSSFS DD PATH='/u/hlq/scrt/may06.scr'
//SMFVB     DD DISP=(,CATLG),DSN=HLQ.SCRT89RU.SMF,
//          UNIT=SYSDA,SPACE=(CYL,(2,2))
//SYSPRINT DD SYSOUT=*
//SYSLIN   DD *
```

Figure 142. SCRT89 record utility JCL

When you run the JCL, the data in the output file is converted to VBS format. For explanations of any messages from the SCRT89 record utility, see “SCRT89 Record Utility messages” on page 495.

Summary of sub-capacity reporting for sub-capacity eligible products running under z/TPF

Table 44 on page 289 shows how the number of MSUs associated with sub-capacity eligible products running under z/TPF is calculated in different configurations of z/TPF and z/OS.

Note: IBM recommends that the z/TPF Recoup Utility only be run on z/TPF systems that collect and supply SCRT89 records to SCRT because the busy times for the Recoup Utility have been removed from the times reported in the SCRT89 records. However, the Recoup Utility time will be counted if no SCRT89 records are provided and SCRT must use the busy time for the LPAR to report on z/TPF activity in that LPAR.

Table 44. How product MSUs are calculated for z/TPF-based programs using SCRT89 records

If the z/TPF-based program ran in...		And z/OS ran concurrently at all times with those z/TPF LPAR(s) in z/OS native LPAR(s)...	Then the number of MSUs associated with the z/TPF-based program is based on the program's peak hour utilization across all LPARs during the reporting period, with the program utilization for each hour being the highest combined...
z/TPF native LPAR(s)...	And z/TPF guest LPAR(s)...	LPAR(s)...	
Yes with SCRT89 records from at least one LPAR	No	No	1. The 4-hour rolling average utilization across all of the z/TPF native LPARs that supply SCRT89 records, <i>plus...</i>
Yes with or without SCRT89 records	No	Yes	2. The 4-hour rolling average LPAR utilization across all of the z/TPF native LPARs that have shared engines and do not supply SCRT89 records, <i>plus...</i> 3. The maximum capacity of all of the z/TPF native LPARs that have dedicated engines and do not supply SCRT89 records.
Yes with SCRT89 records from at least one LPAR	Yes ¹	No	1. The 4-hour rolling average utilization across all of the z/TPF native LPARs and z/TPF guest systems that supply SCRT89 records, <i>plus...</i>
Yes with or without SCRT89 records	Yes ¹	Yes	2. The 4-hour rolling average LPAR utilization across all of the z/TPF native and guest LPARs that have shared engines and do not supply SCRT89 records, <i>plus...</i> 3. The maximum capacity of all of the z/TPF native and guest LPARs that have dedicated engines and do not supply SCRT89 records.
No	Yes ²	No	The 4-hour rolling average utilization across all of the z/TPF guest LPARs
No	Yes ¹	Yes	1. The 4-hour rolling average utilization across all of the z/TPF guest systems that supply SCRT89 records, <i>plus...</i> 2. The 4-hour rolling average LPAR utilization across all of the z/TPF guest LPARs that have shared engines and do not supply SCRT89 records, <i>plus...</i> 3. The maximum capacity of all of the z/TPF guest LPARs that have dedicated engines and do not supply SCRT89 records.

Table 44. How product MSUs are calculated for z/TPF-based programs using SCRT89 records (continued)

If the z/TPF-based program ran in...		And z/OS ran concurrently at all times with those z/TPF LPAR(s) in z/OS native LPAR(s)...	Then the number of MSUs associated with the z/TPF-based program is based on the program's peak hour utilization across all LPARs during the reporting period, with the program utilization for each hour being the highest combined...
z/TPF native LPAR(s)...	And z/TPF guest LPAR(s)...	LPAR(s)...	
No	Yes, without SCRT89 records from every guest system	No	Full machine capacity
Yes without SCRT89 records	No	No	
Yes without SCRT89 records	Yes, without SCRT89 records from every guest system	No	

Footnotes:

1. You must either provide the SCRT89 records for all of the z/TPF guests within an LPAR, or you must not provide any SCRT89 records for any of the z/TPF guests in the LPAR. This requirement applies on an LPAR by LPAR basis.
2. You must provide the SCRT89 records for all of the z/TPF guests running on the machine.

For guest systems, the following rule applies: Either all of the guest systems within the same LPAR must provide SCRT89 records or none of the guest systems in the same LPAR should provide SCRT89 records.

Sub-capacity report for a CPC running z/TPF systems and non-z/TPF systems

Table 45 on page 291 shows the sample hardware and software configuration for a CPC where a combination of z/TPF native and guest systems ran together with non-z/TPF guests, including Linux guests.

- A native z/TPF production system ran in LPAR TPF220 for the entire reporting period and collected SCRT89 records. The Recoup utility ran on this system several times during the reporting period.
- Thirty-nine z/TPF development guests ran at various times in LPAR TPFTDEV during the reporting period. APAR PJ38004 (z/TPF PUT 07) is installed on some, but not all, of the z/TPF guests. No SCRT89 records were collected from any of the guests. None of those guests used the Recoup utility. No non-z/TPF guests ran in this LPAR during the reporting period.
- Three z/TPF test guests (z/VM level-1 guests) ran at various times in LPAR TPFTEST during the reporting period. APAR PJ38004 is installed on all three z/TPF guests prior to the reporting period. Each of these guests ran the Recoup utility and each collected SCRT89 records. Non-z/TPF guests were also run in this LPAR to prepare for the testing to be done in the z/TPF guest systems. These non-z/TPF systems did not support SCRT89 records.
- Various Linux for z Systems guest systems ran in LPAR LINUX during the reporting period. No z/TPF guests ran in this LPAR.

The TPFTEST LPAR supported both z/TPF and non-z/TPF systems, so it is possible that one or more of the non-z/TPF systems might have been running when none of the z/TPF systems were running, which would have caused SCRT to calculate the z/TPF busy time for that hour based on the busy time for the z/TPF LPAR (the non-z/TPF systems) because SCRT received no SCRT89 records for that hour. However, since the z/TPF guests running in that LPAR met all of the conditions listed in [“Running z/TPF and non-z/TPF](#)

guests in the same LPAR” on page 283, the following **Exclude** control statement was used to prevent SCRT from counting any of the non-z/TPF guest busy time in that LPAR:

```
* Used to prevent SCRT from counting non-z/TPF busy time in TPFTEST
Exclude CPC=2097-12345,Image_Id=TPFTEST,
Start=2024/10/01/00,Resume=2099/12/31/23
```

Table 45. Sample configuration with z/TPF native and guest systems and non-z/TPF systems

CPC: 2097-705 (serial = 02-12345)			
TPFP220	TPFTDEV	TPFTEST	LINUX
z/TPF (A001) z/TPFDF (Recoup) (SCRT89 records collected)	z/VM z/TPF guest IDs: DEVU10 (B002) : DEVU39 (B002) Non-z/TPF guest IDs: (none)	z/VM z/TPF guest IDs: TESTU1 (C004) (Recoup) TESTU2 (C004) (Recoup) TESTU3 (C004) (Recoup) Non-z/TPF guest IDs: PREPU1 PREPU2 PREPU3 (SCRT89 records collected)	z/VM Linux guests

The following figures show sections of the sample sub-capacity report that were generated based on the SCRT89 records collected from the configuration shown in [Table 45 on page 291](#) and the **Exclude** control statement shown earlier.

```
==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====
```

```
Run Date/Time           03 Nov 2024 - 09:17
Name of Person Submitting Report: John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter: 444-999-9999
```

```
Customer Name           ABC Corporation
Customer Number          8971234567
Machine Serial Number    02-12345
Machine Type and Model   2097-705
Machine Rated Capacity (MSUs) 488
Purchase Order Number    (optional)
Customer Comments (255 chars max) (optional)
```

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

```
==C5=====
TOOL INFORMATION
```

```
Tool Release           30.1.0
Reporting Period       2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)
```

Figure 143. Sample sub-capacity report for z/TPF: Customer information and Tool Information sections

```

==D5=====
SPECIAL CONDITIONS

SMF/SCRT89 Exclusion Control Statement   Image ID   Product ID   Start       Resume      Customer Comments (255 chars max)
Exclude Control Statement               TPFTTEST   *ALL        01 Oct 2024 - 00:00   31 Dec 2099 - 23:00   (required)

SMF/SCRT89 Exclusion Results             Image ID   Product ID   Start       Resume      Hours
SMF/SCRT89 records excluded            TPFTTEST   *ALL        01 Oct 2024 - 00:00   31 Dec 2099 - 23:00   338

```

Figure 144. Sample sub-capacity report for z/TPF: Special Conditions section with Exclude

The Special Conditions section shown in Figure 144 on page 292 appears as a result of the **Exclude** control statement that was input to SCRT to prevent it from counting the busy time for the non-z/TPF guest systems running in the TPFTTEST LPAR. See “Running z/TPF and non-z/TPF guests in the same LPAR” on page 283 for a more detailed explanation of the situation being addressed and for a complete list of conditions that must be satisfied by the z/TPF systems running in that LPAR before such use of the **Exclude** control statement can be allowed. Note that if you created such a sub-capacity report, you would be required to replace the (required) placeholder in the **Customer Comments** column with the following text:

This LPAR only runs z/TPF guest systems which collect SCRT89 records or non-z/TPF guest systems.

```

==D5=====
SPECIAL CONDITIONS

z/TPF data corruption detected

```

Figure 145. Sample sub-capacity report for z/TPF: Special Conditions section for data corruption detected

Note: In cases where internal z/TPF data corruption occurs during the reporting period, the Special Conditions section will appear with the message, z/TPF data corruption detected, as shown in Figure 145 on page 292. Whenever you submit a sub-capacity report containing this message, IBM will contact you to inquire about the conditions that caused this message to be generated.

```

==E5=====
PRODUCT SUMMARY INFORMATION

MLC Product Name                MLC Product ID   Tool MSUs   Customer MSUs   Customer Comments (255 chars max)
z/TPF V1                        5748-T15         455         455
z/TPFDF V1                      5748-F15         455

```

Figure 146. Sample sub-capacity report for z/TPF: Product Summary Information section

SCRT uses several different formats for reporting the identities of z/TPF systems based on the conditions under which each z/TPF system was running, as described in “Identities reported for z/TPF systems” on page 282. Once SCRT assigns an identity to a z/TPF system, SCRT uses that same identity in the Detail Data Collection section, SMF / SCRT89 Input Data Statistics section, Product Max Contributors section, and Product Grid Snapshot section of the sub-capacity report.

The remainder of the sub-capacity report sections are for the scenario, as originally described, where no SCRT89 records were collected for LPAR TPFTDEV.

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

          SYSID      Input Data Start      Input Data End
TPFP220      A001      01 Oct 2024 - 23:00  01 Nov 2024 - 23:30
TPFTDEV*      02 Oct 2024 - 00:00  01 Nov 2024 - 23:30
TPFTTEST(TESTU1)  C004      03 Oct 2024 - 09:00  29 Oct 2024 - 20:00
TPFTTEST(TESTU2)  C004      12 Oct 2024 - 05:00  24 Oct 2024 - 12:00
TPFTTEST(TESTU3)  C004      06 Oct 2024 - 08:00  27 Oct 2024 - 14:00

CPC              01 Oct 2024 - 23:00  01 Nov 2024 - 23:30

```

```

==N5=====
DETAIL LPAR DATA SECTION

          Highest Hour Count Date/Time          2nd Highest Hour Count Date/Time
TPFP220      251          1 31 Oct 2024 - 22:00      249          1 01 Nov 2024 - 05:00
TPFTDEV*      157          1 23 Oct 2024 - 18:00      141          3 12 Oct 2024 - 16:00
TPFTTEST(TESTU1)  51          1 06 Oct 2024 - 11:00      241          2 06 Oct 2024 - 10:00
TPFTTEST(TESTU2)  18          2 11 Oct 2024 - 19:00       51          4 11 Oct 2024 - 18:00
TPFTTEST(TESTU3)  23          2 04 Oct 2024 - 17:00       51          3 04 Oct 2024 - 16:00

CPC              454          1 06 Oct 2024 - 10:00      450          1 06 Oct 2024 - 11:00

* LPAR data based on PR/SM provided utilizations

```

```

==P5=====
PRODUCT MAX CONTRIBUTORS

Product Name      Product ID  Highest Date/Time      LPAR      LPAR      LPAR      LPAR      LPAR
                   5748-T15    455 11 Oct 2024 - 20:00  216      169      49      0      21
z/TPF V1          5748-F15    455 11 Oct 2024 - 20:00  216      169      49      0      21
z/TPFDF V1

```

Figure 147. Sample sub-capacity report for z/TPF: SMF / SCRT89 Input Data Statistics, Detail LPAR Data, and Product Max Contributors sections

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name      Product ID  TPFP220 TPFTDEV* TPFTTEST(TESTU1) TPFTTEST(TESTU2) TPFTTEST(TESTU3)
z/TPF V1          5748-T15    91.5%   92.5%      23.6%      11.7%      19.1%
z/TPFDF V1        5748-F15    91.5%   92.5%      23.6%      11.7%      19.1%

```

```

==W5=====
This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.

```

Figure 148. Sample sub-capacity report for z/TPF: Product Grid Snapshot section

The Sub-Capacity Capping Enforced by SCRT section and the Group Capacity LPARs section will not appear on a sub-capacity report generated for a CPC which is only running z/TPF systems because z/TPF does not support defined capacity or group capacity.

Interpreting the detail data section of the sub-capacity report for z/TPF

Only the report sections and fields that are different than those for a sub-capacity report for z/OS data without z/TPF are described in this section. You can find information about all other fields in [Chapter 5, “Sub-capacity reporting for z/OS native systems,”](#) on page 103.

z/TPF LPARs are listed in all of the Detail Data sections of the sub-capacity report. The following unique z/TPF conditions are shown in the preceding sample report:

- In order to be eligible for sub-capacity pricing, z/TPF is required to run in LPARs whose names are of the form TPFPxxxx or TPFTxxxx.

- SCRT uses PR/SM data to calculate the MSUs for z/TPF LPARs, which do not provide SCRT89 records, and identifies such LPARs by placing an asterisk (*) after those LPAR names. In the sample report, the data referenced by the name TPFTDEV, is based on PR/SM data collected from the z/TPF systems that ran in TPFP220.

For a description of the remaining fields in this report, see [“Sample sub-capacity report for single-version z/OS native systems” on page 105.](#)

Sub-capacity report for a CPC running z/OS and z/TPF

This section shows a sub-capacity report with LPARs that are running z/OS and z/TPF, including an LPAR that ran both z/OS and z/TPF at different times during the reporting period. [Table 46 on page 294](#) shows this sample configuration.

<i>Table 46. Sample configuration with z/OS and z/TPF systems</i>				
CPC: 2097-705 (serial = 02-12345)				
LPAR1	TPFP3		TPFTDEV	LPARVM
2 October - 1 November	2 October - 2 October	3 October - 1 November	2 October - 1 November	2 October - 1 November
z/OS V1 (SYSID = SYSH) DB2 UDB V8 CICS TS V2 CICS TS MQSeries V5 NetView V2 WebSphere Portal V4 WebSphere MQ Integrator Broker V2	z/OS V1 (SYSID = SYSK) DB2 UDB V8 CICS TS V2 CICS TS MQSeries V5 NetView V2 WebSphere Portal V4 WebSphere MQ Integrator Broker V2	z/TPF (A004) z/TPFDF (SCRT89 records collected)	z/TPF (B006) z/TPFDF (SCRT89 records collected)	z/VM CMS

The sub-capacity report based on the configuration shown in [Table 46 on page 294](#) starts with [Figure 149 on page 295.](#)

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 03 Nov 2024 - 09:17
 Name of Person Submitting Report: John Customer
 E-mail Address of Report Submitter: customer@abc.com
 Phone Number of Report Submitter: 444-999-9999

Customer Name ABC Corporation
 Customer Number 8971234567
 Machine Serial Number 02-12345
 Machine Type and Model 2097-705
 Machine Rated Capacity (MSUs) 488
 Purchase Order Number (optional)
 Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

==C5=====

TOOL INFORMATION

Tool Release 30.1.0
 Reporting Period 2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

Figure 149. Sample sub-capacity report for z/OS and z/TPF: Customer information and Tool Information sections

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS V1	5694-A01	471		
DB2 UDB for OS/390 V7	5675-DB2	471		
CICS TS for OS/390 V2	5697-E93	421		
CICS TS for OS/390	5655-147	471		
MQSeries for OS/390 V5	5655-F10	471		
NetView Perf Monitor V2	5655-043	471	(optional)	(conditional)
Tivoli NetView for z/OS	5697-ENV	471	(optional)	(conditional)
z/TPF V1	5748-T15	68		
z/TPFDF V1	5748-F15	68		
IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
IBM WebSphere Portal for z/OS and OS/390 V4	5655-K12	471		
IBM WebSphere MQ Integrator Broker for z/OS V2	5655-I58	471		

Figure 150. Sample sub-capacity report for z/OS and z/TPF: Product Summary Information section

```

=====
L5
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
=====
M5
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
LPAR1	SYSH	01 Oct 2024 - 23:00	02 Nov 2024 - 00:00
TPFP3	SYSK	01 Oct 2024 - 23:00	02 Oct 2024 - 00:00
TPFP3(z/TPF)	A004	03 Oct 2024 - 00:00	02 Nov 2024 - 23:00
TPFTDEV*		02 Oct 2024 - 00:00	01 Nov 2024 - 23:37
LPARVM	-	-	-
CPC		01 Oct 2024 - 23:00	02 Nov 2024 - 00:00

```

=====
N5
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
LPAR1	475	2	06 Oct 2024 - 10:00	473	8	06 Oct 2024 - 09:00
TPFP3	40	1	02 Oct 2024 - 20:00	39	1	02 Oct 2024 - 19:00
TPFP3(z/TPF)	56	3	11 Oct 2024 - 20:00	51	1	07 Oct 2024 - 19:00
TPFTDEV*	17	3	19 Oct 2024 - 11:00	16	8	19 Oct 2024 - 10:00
CPC	471	1	06 Oct 2024 - 10:00	465	1	06 Oct 2024 - 11:00

* LPAR data based on PR/SM provided utilizations

```

=====
P5
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR TPFP3	LPAR TPFP3(z/TPF)	LPAR TPFTDEV*
z/OS V1	5694-A01	471	06 Oct 2024 - 10:00	471	0	0	0
DB2 UDB for OS/390 V7	5675-DB2	471	06 Oct 2024 - 10:00	471	0	0	0
CICS TS for OS/390 V2	5697-E93	421	31 Oct 2024 - 22:00	471	0	0	0
CICS TS for OS/390	5655-147	471	06 Oct 2024 - 10:00	421	0	0	0
MQSeries for OS/390 V5	5655-F10	471	06 Oct 2024 - 10:00	471	0	0	0
NetView Perf Monitor V2	5655-043	471	06 Oct 2024 - 10:00	471	0	0	0
Tivoli NetView for z/OS	5697-ENV	471	06 Oct 2024 - 10:00	471	0	0	0
IBM WebSphere Portal for z/OS and OS/390 V4	5655-K12	471	06 Oct 2024 - 10:00	471	0	0	0
IBM WebSphere MQ Integrator Broker for z/OS V2	5655-I58	471	06 Oct 2024 - 10:00	471	0	0	0
z/TPF V1	5748-T15	68	11 Oct 2024 - 20:00	0	0	56	12
z/TPFDF V1	5748-F15	68	11 Oct 2024 - 20:00	0	0	56	12

Figure 151. Sample sub-capacity report for z/OS and z/TPF: SMF / SCRT89 Input Data Statistics, Detail LPAR Data Section, and Product Max Contributors sections

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                Product ID    LPAR1    TFPF3    TFPF3(z/TPF)    TPFTDEV*
z/OS V1                     5694-A01     99.1%    0.0%
DB2 UDB for OS/390 V7       5675-DB2     99.1%    0.0%
CICS TS for OS/390 V2       5697-E93     99.1%    0.0%
CICS TS for OS/390         5655-147     18.9%    0.0%
MQSeries for OS/390 V5      5655-F10     99.1%    0.0%

z/TPF V1                     5748-T15
z/TPFDF V1                   5748-F15
                                92.4%      90.3%
                                92.4%      90.3%

```

```

==S5=====
SUB-CAPACITY CAPPING ENFORCED BY SCRT - HOURS

Product Name                Product ID    LPAR    LPAR    LPAR    LPAR
                        LPAR1    LPAR1    TFPF3    TFPF3(z/TPF)    TPFTDEV*
z/OS V1                     5694-A01         2
DB2 UDB for OS/390 V7       5675-DB2         2
CICS TS for OS/390 V2       5697-E93         2
CICS TS for OS/390         5655-147
MQSeries for OS/390 V5      5655-F10         2
NetView Perf Monitor V2     5655-043         2
Tivoli NetView for z/OS     5697-ENV         2
IBM WebSphere Portal for OS/390 and z/OS V4 5655-K12         2
IBM WebSphere MQ Integrator Broker for z/OS V2 5655-I58         2

```

```

==W5=====
This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.

```

Figure 152. Sample sub-capacity report for z/OS and z/TPF: Product Grid Snapshot and Sub-Capacity Capping Enforced by SCRT sections

Interpreting the detail data sections of the sub-capacity report for z/OS and z/TPF

Only the report sections and fields that are different than those for a sub-capacity report for z/OS data without z/TPF are described in this section. You can find information about all other fields in [Chapter 5, “Sub-capacity reporting for z/OS native systems,”](#) on page 103. z/TPF LPARs are listed in the Detail Data Collection section and in all of the Detail Data sections of the sub-capacity report. The following unique z/TPF conditions are shown in the preceding sample report:

- In order to be eligible for sub-capacity pricing, z/TPF is required to run in LPARs whose names are of the form TFPFxxxx or TPFTxxxx.
- Whenever both z/OS and z/TPF run in the same LPAR at different times during the same reporting period, SCRT uses the LPAR name alone to indicate the statistics for the z/OS system running in that LPAR and uses the *LPARname* (z/TPF) designation to indicate the statistics for the z/TPF system that ran in that LPAR. In the sample report, the TFPF3 LPAR was used by a z/OS system for all 24 hours on 2 October. That same LPAR was used by a z/TPF system for the remainder of the reporting period. All of the statistics for the z/OS system running in that LPAR are listed under TFPF3. All of the statistics for the z/TPF system running in that LPAR are listed under TFPF3(z/TPF).
- SCRT uses PR/SM data to calculate the MSUs for z/TPF LPARs that do not provide SCRT89 records and identifies such LPARs by placing an asterisk (*) after those LPAR names. In the sample report, the data referenced by the name TPFTDEV, is based on PR/SM data collected from either the z/OS system, which ran in LPAR1, or from the z/TPF system that ran in TFPF3.

For a description of the remaining fields in this report, see [“Sample sub-capacity report for single-version z/OS native systems”](#) on page 105.

Note that whenever multiple sub-capacity eligible operating system families run on the same CPC during the same reporting period (for example, z/OS and z/TPF), the SMF/SCRT89 data from all of

those operating system families must be processed by the same execution of SCRT, so that all of those operating system families appear on the same sub-capacity report. IBM will only accept a single sub-capacity report per CPC per reporting period.

Chapter 14. Reporting for the z/TPF Transformation Engines Offering

With the introduction of the optional z/TPF Transformation Engines Offering, SCRT reports on an optional (hardware) processor-related offering, in addition to reporting on the existing software sub-capacity offerings available for z/TPF. The SCRT support for the z/TPF Transformation Engines Offering described in this chapter is separate from the sub-capacity pricing support for the z/TPF (5748-T15) and z/TPFDF (5748-F15) products that is described in [Chapter 13, “Sub-capacity reporting for z/TPF systems,”](#) on page 281.

You must follow all of the processes and meet all of the requirements set forth in [Chapter 13, “Sub-capacity reporting for z/TPF systems,”](#) on page 281 in order to generate a valid sub-capacity report which will include reporting on the z/TPF Transformation Engines Offering whenever you also follow the processes described in this chapter. You should continue to use your current submission process for sending sub-capacity reports to IBM, including when those reports also report on the z/TPF Transformation Engines Offering. In the latter case, IBM will use the same sub-capacity report for both sub-capacity reporting purposes and for z/TPF Transformation Engines Offering reporting purposes.

Under the z/TPF Transformation Engines Offering, z/TPF sub-capacity customers can purchase CP engines (transformation engines) at a special price and use them to process a special category of z/TPF work which is outside of the z/TPF mainline function. This category of work includes work performed by z/TPF supported utilities, such as OpenLDAP and MySQL, and z/TPF support for XML, SOAP, SSL, HTTP client, and HTTP Server. The SCRT support for the z/TPF Transformation Engines Offering consists of measuring and reporting the number of MSUs being used to perform this special category of work at the same hour that the z/TPF product (5748-T15) peak occurs for sub-capacity purposes. Go to [IBM z/Transaction Processing Facility \(www.ibm.com/products/z-transaction-processing-facility\)](http://www.ibm.com/products/z-transaction-processing-facility) for additional information about this offering, including the offering’s benefits and eligibility requirements.

For z/TPF systems that generate SCRT89 records, SCRT assigns the MSUs for the special category of z/TPF work to the z/TPF Transformation Engines Offering entry for every hour that has an SCRT89 record. SCRT does this for both native z/TPF systems and z/TPF guest systems running under z/VM. For z/TPF systems that do not provide SCRT89 records, SCRT assigns a value of zero MSUs to the z/TPF Transformation Engines Offering for that hour, as there is no basis to report on the special category of work without the SCRT89 records. SCRT makes those MSU assignments to the z/TPF Transformation Engines Offering on an hour-by-hour and z/TPF system-by-system basis.

When reporting the peak MSU value for the z/TPF Transformation Engines Offering, SCRT determines the peak date/hour for the z/TPF product (5748-T15) across all z/TPF systems running on that CPC, then reports the sum of the MSUs assigned to the z/TPF Transformation Engines Offering across all of those systems for the same date/hour. If the same MSU peak occurs for multiple dates/hours for the z/TPF product (5748-T15), SCRT reports the MSU value for the z/TPF Transformation Engines Offering corresponding to the peak z/TPF product date/hour when the Offering has the highest MSU value for the CPC.

Qualifications for the z/TPF Transformation Engines Offering

All of the following conditions must be met before you can receive any benefit from the z/TPF Transformation Engines Offering:

1. The CPC must be an IBM zEnterprise EC12 (zEC12) or zEnterprise BC12 (zBC12) or a later processor family and must be under sub-capacity pricing for z/TPF. SCRT will report on the z/TPF Transformation Engines Offering for all processor families, including those that are not eligible for the offering, so that you can use SCRT to assess the size of the benefit, if any, that you might derive from the offering if you were running the same z/TPF workloads on an eligible processor family. The IBM billing process will accept sub-capacity reports created for ineligible processor families, but will not apply any z/TPF Transformation Engines Offering benefits for those reports.

2. All of the z/TPF systems running on the CPC must be running at the V1.1 level or later, as that is the earliest z/TPF level that supports sub-capacity reporting. Running an earlier z/TPF level on the same CPC makes that CPC ineligible for the sub-capacity pricing and, therefore, ineligible for the z/TPF Transformation Engines Offering.
3. One or more of the z/TPF systems running on the CPC must be at the V1.1 PUT 02 level or later and must be collecting SCRT89 records for all of the hours during which other z/TPF systems are running on that CPC and not collecting SCRT89 records. z/TPF V1.1 systems running with PUT 11 gain additional benefits for the z/TPF Transformation Engines Offering, as PUT 11 includes additional functions in the special category of work reported under the Offering. The V1.1 PUT 02 level is the earliest level of z/TPF that generates SCRT89 records. At least one z/TPF V1.1 system at the PUT 02 level must be collecting SCRT89 records to provide the data necessary to determine the MSUs reported for the z/TPF Transformation Engines Offering. Without SCRT89 records, SCRT has no data to report on the z/TPF Transformation Engines Offering.
4. You must purchase one or more transformation engines for the CPC for use in accordance with the terms and conditions of the z/TPF Transformation Engines Offering.
5. You must use the **z/TPF_Transformation_Engines_Offering** control statement when generating your sub-capacity report in order for SCRT to report the MSU value for the z/TPF Transformation Engines Offering.

SCRT restriction on the z/TPF Transformation Engines Offering

SCRT does not support using the **Exclude** control statement to specifically exclude the z/TPF Transformation Engines Offering use, as the Offering does not have a real product ID. Any attempt to exclude the product ID of 0000-000 which appears on the sub-capacity report for the Offering will result in SCRT issuing the SCRTTOOL123 message and terminating processing. However, like the z/TPF and z/TPFDF products, the z/TPF Transformation Engines Offering will be excluded for the dates and times for which the entire z/TPF system is excluded—that is, by not using the **Product_Id** parameter on the **Exclude** control statement.

SCRT externals for the z/TPF Transformation Engines Offering

SCRT provides the following externals for the optional z/TPF Transformation Engines Offering. All of these externals are specific to the Offering and are ignored when the Offering is not being used.

- SCRT supports the optional **z/TPF_Transformation_Engines_Offering** control statement, which you must use to request that SCRT report on the z/TPF Transformation Engines Offering.
- SCRT reports the use of the **z/TPF_Transformation_Engines_Offering** control statement in the Special Conditions section of the sub-capacity report.
- SCRT uses a footnote, (ix), in the Product Summary Information section of the sub-capacity report to indicate that the matching offering entry was reported using the special algorithm for the z/TPF Transformation Engines Offering.

In addition, when SCRT reports on the z/TPF Transformation Engines Offering, the following sub-capacity report sections are affected:

- Customer Information section
- Special Conditions section
- Product Summary Information section
- Product Max Contributors section
- Product Grid Snapshot section

The changes to these report sections are described later in this chapter, along with sample sub-capacity report sections to illustrate the changes.

About the sample sub-capacity report sections in this chapter:

- The sample sub-capacity report sections show how those sections are changed when SCRT reports on the z/TPF Transformation Engines Offering. The z/TPF Transformation Engines Offering does not affect any other report sections; unaffected report sections are not shown.
- The sample report sections are based on the z/TPF configuration shown in [Table 45 on page 291](#).
- The z/TPF configuration was deliberately chosen to illustrate several different conditions that are reported by SCRT and may not reflect an actual customer configuration.

z/TPF_Transformation_Engines_Offering control statement

The default SCRT processing does not support the z/TPF Transformation Engines Offering. You must specify the **z/TPF_Transformation_Engines_Offering** control statement to request that SCRT report on the Offering when generating sub-capacity reports. For more information about this control statement, see “z/TPF_TRANSFORMATION_ENGINES_OFFERING control statement” on page 60.

When you specify the **z/TPF_Transformation_Engines_Offering** control statement, SCRT reports on the Offering for each CPC on which one or more z/TPF systems have created one or more SCRT89 records during the reporting period. SCRT does not report on the Offering for CPCs that do not meet these conditions.

Addition to the Customer Information section

Sub-capacity reports that report on the z/TPF Transformation Engines Offering contain an additional paragraph in the Customer Information section (B5) of the sub-capacity report that describes the IBM business intent for using the values reported for the z/TPF Transformation Engines Offering.

[Figure 153 on page 301](#) shows an example of the Customer Information section for a sub-capacity report for the z/TPF Transformation Engines Offering.

```

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time                03 Nov 2024 - 09:17
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter:  444-999-9999

Customer Name                ABC Corporation
Customer Number              8971234567
Machine Serial Number        02-12345
Machine Type and Model        2097-705
Machine Rated Capacity (MSUs) 488
Purchase Order Number         (optional)
Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust
the billable MSUs in inventory for all MLC Products listed under the MLC Product Name
column on this report. In accordance with our agreement, IBM will treat a change in product
licensed capacity as an order. If the MSUs have changed since the last report, software billing
based on inventory MSUs will increase or decrease accordingly.

For the product entry(s) identified as offering(s) by footnotes, the data supplied in this
report will be used in the fulfillment of each offering in accordance with the offering's
terms and conditions.

```

Figure 153. Sample sub-capacity report for z/TPF Transformation Engines Offering: Customer information section

Addition to the Special Conditions section

Whenever SCRT encounters the **z/TPF_Transformation_Engines_Offering** control statement, it displays a header line followed by the control statement itself in the Special Conditions section of every sub-capacity report generated by that execution of SCRT. SCRT does this even if it never reported on the z/TPF Transformation Engines Offering on that sub-capacity report due to all of the Offering's requirements not being met.

[Figure 154 on page 302](#) shows an example of the Special Conditions section for a sub-capacity report for the z/TPF Transformation Engines Offering for the z/TPF configuration shown in [Table 45 on page 291](#).

```

==D5=====
SPECIAL CONDITIONS

z/TPF Transformation Engines Offering Control Statement
z/TPF_Transformation_Engines_Offering

SMF/SCRT89 Exclusion Control Statement   Image ID   Product ID   Start      Resume      Customer Comments (255 chars max)
Exclude Control Statement               TPFTEST    *ALL         01 Oct 2024 - 00:00  31 Dec 2099 - 23:00  (required)

SMF/SCRT89 Exclusion Results             Image ID   Product ID   Start      Resume      Hours
SMF/SCRT89 records excluded            TPFTEST    *ALL         01 Oct 2024 - 00:00  31 Dec 2099 - 23:00  338

```

Figure 154. Sample sub-capacity report for z/TPF Transformation Engines Offering: Special Conditions section

Changes to the Product Summary Information section

Whenever SCRT processes a **z/TPF_Transformation_Engines_Offering** control statement along with one or more SCRT89 records within the reporting period, the following changes appear in the Product Summary Information section of the sub-capacity report:

- A **Footnotes** column header appears in the MLC section.
- The z/TPF Transformation Engines Offering entry and its MSU value are listed in the row following the z/TPFDF product. Note that an MSU value of zero is considered a valid value for the Offering, as valid conditions can cause a value of zero to occur.
- A value of 0000-000 is listed in the MLC Product ID column for the Offering. The z/TPF Transformation Engines Offering is an offering, not an actual product, so it does not have a software product ID.
- The (ix) footnote is listed in the **Footnotes** column for the z/TPF Transformation Engines Offering entry to indicate that the value being reported is different from the values being reported for the other entries in this report section.
- A description of footnote (ix) appears in the Footnotes subsection, which follows the last sub-capacity product listed in this section.

Figure 155 on page 302 shows an example of the Product Summary Information section for a sub-capacity report for the z/TPF Transformation Engines Offering for the z/TPF configuration shown in Table 45 on page 291.

```

==E5=====
PRODUCT SUMMARY INFORMATION

MLC Product Name                MLC Product ID   Tool MSUs   Customer MSUs   Customer Comments (255 chars max)   Footnotes
z/TPF V1                        5748-T15         455
z/TPFDF V1                      5748-F15         455
z/TPF Transformation Engines Offering 0000-000         102                                (ix)

```

Figure 155. Sample sub-capacity report for z/TPF Transformation Engines Offering: Product Summary Information section

Changes to the Product Max Contributors section

Whenever SCRT reports on the z/TPF Transformation Engines Offering in the Product Summary Information section, the following changes also appear in the Product Max Contributors section of the sub-capacity report:

- The z/TPF Transformation Engines Offering entry with a product ID of 0000-000 is listed in the row following the z/TPFDF product.
- The date/time value is listed for the hour whose MSU value is being posted for the offering. This date/time corresponds to one of the peak hours for the z/TPF product (5748-T15).

If the same MSU peak occurs for multiple dates/hours for the z/TPF product (5748-T15), SCRT reports the MSU value for the z/TPF Transformation Engines Offering corresponding to the peak z/TPF product

date/hour when the Offering has the highest MSU value. In this case, the date/time listed for the Offering will differ from the date/time listed for the z/TPF product.

Figure 156 on page 303 shows an example of the Product Max Contributors section for a sub-capacity report for the z/TPF Transformation Engines Offering for the z/TPF configuration shown in Table 45 on page 291.

```

=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR TPFP220	LPAR TPFTDEV*	LPAR TPFTEST(TESTU1)	LPAR TPFTEST(TESTU2)	LPAR TPFTEST(TESTU3)
z/TPF V1	5748-T15	455	11 Oct 2024 - 20:00	216	169	49	0	21
z/TPFDF V1	5748-F15	455	11 Oct 2024 - 20:00	216	169	49	0	21
z/TPF Transformation Engines Offering	0000-000	102	25 Oct 2024 - 22:00	95	0	5	0	2

Figure 156. Sample sub-capacity report for z/TPF Transformation Engines Offering: Product Max Contributors section

Changes to the Product Grid Snapshot section

Whenever SCRT reports on the z/TPF Transformation Engines Offering in the Product Summary Information section, the following changes also appear in the Product Grid Snapshot section of the sub-capacity report:

- The z/TPF Transformation Engines Offering entry with a product ID of 0000-000 are listed in the row following the z/TPFDF product.
- The percentage of hours in the reporting period for which SCRT assigned a value for the z/TPF Transformation Engines Offering is listed for each z/TPF system. Since the same SCRT89 records report on the z/TPF product, the z/TPFDF product, and the z/TPF Transformation Engines Offering, the same percentage value will be listed for all three entries.

Figure 157 on page 303 shows an example of the Product Grid Snapshot section for a sub-capacity report for the z/TPF Transformation Engines Offering for the z/TPF configuration shown in Table 45 on page 291.

```

=====
PRODUCT GRID SNAPSHOT

```

Product Name	Product ID	TPFP220	TPFTDEV*	TPFTEST(TESTU1)	TPFTEST(TESTU2)	TPFTEST(TESTU3)
z/TPF V1	5748-T15	91.5%		23.6%	11.7%	19.1%
z/TPFDF V1	5748-F15	91.5%		23.6%	11.7%	19.1%
z/TPF Transformation Engines Offering	0000-000	91.5%		23.6%	11.7%	19.1%

Figure 157. Sample sub-capacity report for z/TPF Transformation Engines Offering: Product Grid Snapshot section

Chapter 15. Sub-capacity reporting for z/VSE systems

This chapter tells the system programmer how to use SCRT to generate sub-capacity reports based on SCRT89 records collected from z/VSE systems. It describes the process for generating sub-capacity reports for z/VSE systems by running SCRT on a z/OS system and by running the SCRT in a Windows or Linux environment (including Linux on IBM Z).

Introduction to sub-capacity pricing for z/VSE systems

IBM offers sub-capacity pricing for z/VSE and selected VSE middleware products. Before you begin, check [Sub-Capacity for z/VSE](http://www.ibm.com/downloads/documents/us-en/10c31775c6d400a2) (www.ibm.com/downloads/documents/us-en/10c31775c6d400a2) and [Midrange Workload License Charges \(MWLC\)](http://www.ibm.com/downloads/documents/us-en/107a02e95ac8f5d0) (www.ibm.com/downloads/documents/us-en/107a02e95ac8f5d0) for the latest sub-capacity pricing requirements and list of z/VSE products that are sub-capacity eligible.

Starting with z/VSE V4.1, you have the option to allow z/VSE to produce SCRT89 accounting records. The SCRT89 records report sub-capacity MSUs on CPCs that are running z/VSE V4.1 or later. z/VSE uses the Capacity Measurement Tool (CMT) to produce SCRT89 accounting records. Sub-capacity pricing is based on the 4-hour rolling average utilization of the z/VSE LPARs or z/VM guests (or a combination of the two) observed within a one month reporting period. The reporting periods run from midnight (00:00) at the very beginning of the second day of the month, up to but not including midnight (00:00) at the very beginning of second day of the next month.

In order for the CMT to collect the required SCRT89 records from an IBM z17, z16, z15, z14, z13, zEC12, z196, z13s, zBC12, or z114 CPC, the appropriate APARs listed in [Table 47 on page 306](#) must be applied prior to using the CMT to collect SCRT89 records on those CPCs.

Given that many z/VSE native and guest systems often use capacities that are much less than one MSU, the z/VSE CMT process has been enhanced so that SCRT89 records generated by z/VSE V4.1 and higher systems with the required APAR level will record capacities less than one MSU for each hour during which such low capacities occur. Previously, the z/VSE CMT process recorded a minimum of one MSU for every z/VSE system, native or guest.

SCRT recognizes and uses z/VSE system capacities of less than one MSU when calculating the overall capacities for z/VSE and z/VSE middleware products across all of the z/VSE systems on the same CPC. However, SCRT will report a minimum of one MSU for every z/VSE product that has been used during the reporting period. This enhancement only applies to z/VSE systems that have collected SCRT89 records with the required APARs applied; it does not apply to z/OS or z/TPF systems. SCRT89 records collected from z/VSE systems without the required APARs applied will continue to report a minimum of one MSU for every z/VSE system. Versions of SCRT prior to V18.1.0 will also report a minimum of one MSU for every z/VSE system, even when the SCRT89 records were collected with the required APARs applied.

With the introduction of z/VSE V6, up to three different versions of z/VSE may be running on the same CPC during the same SCRT reporting period. This causes SCRT to report additional z/VSE information on the sub-capacity report for the CPC. For more information, see [“Sub-capacity reporting for multiple z/VSE versions” on page 308](#).

For CPCs that support power save mode (where the CPC can reduce its power requirements during periods where the full capacity of the CPC is not required, but at a decrease in CPC performance), see [“Static power save mode” on page 506](#) for a description of power save mode and additional details about the z/VSE support for power save mode.

If z/OS runs on the same CPC as z/VSE and you also have sub-capacity pricing for z/OS, use SCRT on z/OS to process the SCRT89 records from your z/VSE systems together with the SMF type 70 and type 89 records from your z/OS systems on this CPC to generate one sub-capacity report.

If z/OS (and z/TPF) do not run on the same CPC as z/VSE, or if z/OS (and z/TPF) on that CPC do not have sub-capacity pricing, use SCRT in a Windows or Linux environment (including Linux on IBM Z) to process the SCRT89 records collected from this CPC to generate the sub-capacity report for this CPC. However, if there are other CPCs in your data center where z/OS and z/TPF run on the same CPC as z/VSE, then it might be easier to generate all of your sub-capacity reports on one of your z/OS systems.

Requirements for sub-capacity pricing for z/VSE systems

In addition to the requirements listed in “Preparing your hardware and software” on page 11, there are additional requirements for sub-capacity pricing on z/VSE:

- The participating CPC must be an IBM z17 (z17), IBM z16 (z16), IBM z15 (z15), IBM z14 (z14), IBM z13 (z13), z13s, zEnterprise EC12 (zEC12), zEnterprise BC12 (zBC12), zEnterprise 196 (z196), zEnterprise 114 (z114), System z10 EC, System z10 BC, System z9 EC, or System z9 BC, but cannot be a z10 BC Model A01 nor a z9 BC Model A01.
- All of the z/VSE systems running on a participating CPC must be z/VSE V4.1 or higher.
- The z/VSE system can run in an LPAR or as a z/VM guest using z/VM V5.2 or higher.
- All of the z/VSE systems running on a participating CPC must run CMT. [Table 47 on page 306](#) shows the recommended APAR levels for z/VSE CMT.

Table 47. Recommended APAR levels for z/VSE CMT

z/VSE level	APAR level	APAR description
z/VSE V6.1	DY47647	Support to extract SCRT89 records for the current month using the MONTH=CURRENT option of IJBCMTRP
z/VSE V5.2	DY47646	Support to extract SCRT89 records for the current month using the MONTH=CURRENT option of IJBCMTRP
z/VSE V5.1	DY47645	Support to extract SCRT89 records for the current month using the MONTH=CURRENT option of IJBCMTRP
z/VSE V4.3	DY47197	Support to extract SCRT89 records with DATE or TIME offsets using IJBCMTRP
z/VSE V4.2	DY47029	Support for reporting z/VSE system image capacities of less than one MSU and support for saving the LPAR name used by guest systems on SCRT89 records
	DY47152	Support for processing CMT data files using IJBCMTRP that exceed 32767 tracks on ECKD disks
	DY47111	Support for z13, z13s, zEC12, zBC12, z196, and z114 CPCs, machine types 2964, 2965, 2827, 2828, 2817 and 2818, respectively (includes DY47029 and DY47152) ¹
	DY47194	Support to extract SCRT89 records with DATE or TIME offsets using IJBCMTRP

Table 47. Recommended APAR levels for z/VSE CMT (continued)

z/VSE level	APAR level	APAR description
z/VSE V4.1	DY47027	Support for reporting z/VSE system image capacities of less than one MSU and support for saving the LPAR name used by guest systems on SCRT89 records
	DY47152	Support for processing CMT data files using IJBCMTRP that exceed 32767 tracks on ECKD disks
	DY47110	Support for z13, z13s, zEC12, zBC12, z196, and z114 CPCs, machine types 2964, 2965, 2827, 2828, 2817 and 2818, respectively (includes DY47027 and DY47152) ¹
	DY47196	Support to extract SCRT89 records with DATE or TIME offsets using IJBCMTRP

Footnotes:

1. This APAR is *required* in order for CMT to be able to collect SCRT89 records on z16, z15, z14, z13, z13s, zEC12, zBC12, z196, and z114 CPCs.

Reporting of IBM z/VSE products on 21st Century Software VSEⁿ

IBM z/VSE products running on 21st Century Software VSEⁿ on a CPC that participates in sub-capacity pricing must be included in the sub-capacity report. For that purpose, SCRT89 accounting records must be collected on every 21st Century Software VSEⁿ system that is running IBM z/VSE products on the CPC that uses sub-capacity pricing. These SCRT89 accounting records from the 21st Century Software VSEⁿ systems then must be provided, in addition to the SCRT89 records from any sub-capacity eligible IBM operating system products (such as IBM z/VSE systems), as input to SCRT when producing the sub-capacity report for that CPC. Any IBM NO89 products running on 21st Century Software VSEⁿ systems on that CPC must be declared in the same manner as they would when running on IBM z/VSE.

Starting with SCRT 29.1.1, SCRT includes limited support for reporting IBM z/VSE products running on 21st Century Software VSEⁿ 6.3. Note that SCRT does not report on the 21st Century Software VSEⁿ operating system nor on any 21st Century Software VSEⁿ products. Neither does SCRT reflect any 21st Century Software VSEⁿ operating system nor any 21st Century Software VSEⁿ products in the optional (A11) entries used for Multi-Version Measurement (MVM). You must use the appropriate level of SCRT to report your IBM z/VSE products.

The following special condition will appear in the Special Conditions section of the sub-capacity report when SCRT89 accounting records from a 21st Century Software VSEⁿ system were processed:

```
==D5=====
SPECIAL CONDITIONS
SCRT89 records from a non-IBM VSE operating system product were processed
```

Note that IBM z/VSE products that are under a sub-capacity pricing capable metric running on 21st Century Software VSEⁿ can be billed as sub-capacity only while the associated IBM z/VSE operating system license on that CPC has not been canceled.

The following IBM z/VSE products are not reported when running on 21st Century Software VSEⁿ and are not eligible for sub-capacity pricing on 21st Century Software VSEⁿ:

- IBM ACF/VTAM® for VSE/ESA V4 (5686-065)
- IBM CICS Transaction Server for VSE/ESA (5648-054)
- IBM CICS Transaction Server for z/VSE V2 (5655-VSE)

Identities reported for z/VSE systems

SCRT uses several formats to report the identities of individual z/VSE systems based on whether those systems are native systems or guest systems and, if guest systems, based on the level of service that has been applied. Table 48 on page 308 summarizes the formats that SCRT uses for z/VSE systems that are reported in the Detail Data Collection section, SMF / SCRT89 Input Data Statistics section, Detail LPAR Data Section, Product Max Contributors section and Product Grid Snapshot section of the sub-capacity report.

Table 48. Summary of z/VSE system identities reported by SCRT

z/VSE configuration	z/VSE APAR applied? (V4.1: DY47027; V4.2: DY47029; V4.3 and V5: included)	Format of z/VSE system identity
z/VSE native system	–	<i>LPAR_name</i> Example: VSEPROD
z/VSE guest system (z/VM level-1 guest)	Yes	<i>LPAR_name(VM_userid)</i> Example: VSEVM(VGDEV1)
	No	<i>*lpar*(VM_userid)</i> Example: *lpar*(VGDEV1)
z/VSE guest system (z/VM level-2 or higher guest)	Yes	<i>LPAR_name(VM_userid-VM_level)</i> Example: VSEVM(VGDEV1-2)
	No	<i>*lpar*(VM_userid-VM_level)</i> Example: *lpar*(VGDEV1-2)

Sub-capacity reporting for multiple z/VSE versions

The format of the sub-capacity report will be extended whenever any combination of two or more of the following z/VSE products are running on the same CPC during the same reporting period:

- z/VSE V6
- z/VSE V5 (VSE Central Functions V9)
- z/VSE V4 (VSE Central Functions V8)

Each version of z/VSE is considered a separate sub-capacity product, and SCRT reports on each of these products independently of one another. However, in addition to reporting separate values for individual versions of z/VSE products, SCRT also reports an additional sub-capacity value which represents the concurrent sub-capacity peak value for all versions of the z/VSE products, as if they were a single product. The sub-capacity report entry that lists this additional value is assigned a product name of z/VSE without any version designation and is assigned a product ID value of (A11). This entry appears prior to any entries for individual z/VSE versions in the following sub-capacity report sections, but only when multiple z/VSE versions appear on the same sub-capacity report.

- Product Summary Information
- Product Max Contributors
- Product Grid Snapshot

The additional z/VSE (A11) entry does not appear when only one version of z/VSE is listed on the sub-capacity report.

IBM uses the sub-capacity value reported for the new z/VSE (A11) entry to determine the billing for the z/VSE products whenever multiple versions of z/VSE are reported on the same sub-capacity report. Under Multi-Version Measurement (MVM), IBM waives the charges for the older version(s) and charges for the latest version of z/VSE based on the sub-capacity value reported for the z/VSE (A11) entry, or based on

the sub-capacity value reported for the single version of z/VSE whenever the z/VSE (A11) entry is not reported. For more information about MVM, see [Chapter 17, “Sub-capacity reporting and Multi-Version Measurement \(MVM\),”](#) on page 367.

Whenever a sub-capacity value of less than three MSUs is reported for a product, that product will be billed for the system minimum of three MSUs.

The sub-capacity values reported for the IBM middleware running on any version of z/VSE continue to be based on the 4-hour rolling average utilization of the z/VSE native system or guest system on which those products run, regardless of the version of that z/VSE system. See [“Sample sub-capacity report for z/VSE V4 and V5”](#) on page 327 and [“Sample sub-capacity report for z/VSE V6, z/VSE V5, and z/VSE V4”](#) on page 330 for examples of sub-capacity reports where multiple versions of z/VSE are reported.

Overview of the CMT

z/VSE uses the Capacity Measurement Tool (CMT) to produce SCRT89 accounting records. To produce SCRT89 records, the CMT requires three files:

- A single control file: IJSYSCC
- Two data files: IJSYSC1 and IJSYSC2

The IJSYSC1 and IJSYSC2 data files are also referred to as SCRT89 files because they contain the SCRT89 records. To define these files, see [“Overview of skeletons for CMT”](#) on page 310.

The length of an SCRT89 record depends on the CPC or z/VM guest configuration (for example, the number of LPARs or the number of processors) and can change during the measurement period if the configuration changes. The minimum size required for the IJSYSC1 and IJSYSC2 data files guarantees that the SCRT89 records for a given month do not overwrite the SCRT89 records for the previous month.

The IJSYSCC, IJSYSC1, and IJSYSC2 files have the following characteristics:

- Must be sequential (BAM) disk files
- Can reside on disks of either type 3390 or FBA (mixed disk types are not allowed)
- Can only have one extent for each file
- Must have an expiration date of 99/366, so that they cannot expire
- Must not reside on the DOSRES or SYSWK1 volumes to avoid conflicts with the IBM supplied disk layout
- Must not be shared among multiple systems. Each system requires its own set of files
- Can reside on a shared disk as long as disk sharing is set up correctly and the files are not shared among multiple systems

Attention: When the files reside on a shared disk and disk sharing is set up incorrectly, data will be lost. Therefore, IBM recommends against placing the files on a shared disk.

- Must not be relocated (neither moved on the disk nor moved to a new disk), as this will cause data to be lost.

The files are created during initialization of the CMT. After that, they must not be relocated. If you need to move the files, create a new set of files and initialize them. For the report, you can then extract the data from both sets of files.

Before starting the CMT, you must initialize all of the CMT files. When the CMT is active, an SCRT89 record is written every 30 minutes to one of the data files. The CMT starts writing to the IJSYSC1 data file. When the IJSYSC1 data file is full, the CMT automatically switches to the IJSYSC2 data file. When the IJSYSC2 data file is full, the CMT switches back to the IJSYSC1 data file. That means that at least one month's SCRT89 records are found in both data files. The IJSYSCC control file contains information about the usage of the data files.

At the end of a reporting period, you must prepare the input to SCRT:

- **SCRT89 records:** The SCRT89 records for a reporting period must be extracted from the IJSYSC1 and IJSYSC2 CMT data files. CMT provides a program named IJBCMTRP to extract the SCRT89 records for the reporting period into a VSAM ESDS cluster.

- **PARMS control statements:** The PARMs control statements contain the SCRT parameters for your data. You must modify the sample PARMs control statements to meet your needs. For details, see [“Customizing the PARMs control statements” on page 320](#).
- **NO89 control statements:** The NO89 control statements identify all the sub-capacity-eligible z/VSE products that were active in the reporting period, but were not reported in the SCRT89 records. VSE Central Functions, ACF/VTAM and CICS TS for VSE/ESA are not included in the NO89 control statements as these products are reported in the SCRT89 records. You must modify the sample NO89 control statements to match your system configuration. For details, see [“Customizing the NO89 control statements” on page 320](#).
- The **SPECIAL control statements:** The SPECIAL control statements are optional and should only be used to deal with very unusual reporting situations where SCRT89 data needs to be excluded from SCRT processing. Such conditions might involve using the SPECIAL control statements under the guidance of the SCRT support team. If such use becomes necessary, see [“SPECIAL DD statement” on page 40](#) for an explanation of the EXCLUDE control statement.

Configuring the CMT files

The section covers the following topics:

- [“DLBL and EXTENT statements” on page 310](#)
- [“System GETVIS 24-bit requirements” on page 310](#)
- [“Overview of skeletons for CMT” on page 310](#)

DLBL and EXTENT statements

You must add **DLBL** and **EXTENT** statements to the system standard label group for the IJSYSCC, IJSYSC1 and IJSYSC2 files. Do not specify a logical unit in the **EXTENT** statement; CMT dynamically assigns the logical unit. Sample **DLBL** and **EXTENT** statements appear under the CMTLAB procedure in [“Overview of skeletons for CMT” on page 310](#).

System GETVIS 24-bit requirements

The CMT requires approximately 20 KB of system GETVIS 24-bit storage. The value might increase when you increase the number of LPARs, or processors, or both.

Overview of skeletons for CMT

The VSE/ICCF library 59 provides the following skeletons to help you to prepare your system for CMT:

- [“SKCMT skeleton” on page 311](#)
- [“SKCMTINI skeleton” on page 313](#)
- [“SKCMTREP skeleton” on page 313](#)
- [“SKUSERBG skeleton” on page 313](#)

You can use one of the following methods to access the job skeletons shipped with z/VSE in ICCF library 59:

- Use the Interactive User Interface (IUI) to copy the skeletons to a library of your choice. Select **Program Development**, then **Program Development Library**. On the **Program Development Library** panel, enter your user library number as the primary library number and 59 as the secondary library number, specify SKCMT as the member prefix, and select option **2** to display the secondary library, as shown in [Figure 158 on page 311](#).

```

IESLIBI                                PROGRAM DEVELOPMENT LIBRARY

ENTER the numbers of the Libraries you want to access

    PRIMARY..... 10                You may read and write in this library
    SECONDARY.... 59__             You may only read from this library

SPECIFY optionally

    PREFIX..... SKCMT__            To list only members whose names begin
                                    with the specified prefix

SELECT which library you want to display

    OPTION..... 2                  1 = Primary   Library
                                    2 = Secondary Library

PF1=HELP      2=REDISPLAY  3=END      4=RETURN

```

Figure 158. Interactive User Interface (IUI): Program Development Library panel

On the Secondary Library panel, specify option **4** next to the SKCMT, SKCMTINI, and SKCMTREP members, as shown in [Figure 159 on page 311](#). You can then edit the copies in your chosen destination library.

```

IESLIBC                                SECONDARY LIBRARY                                Page 1 of 1
SECONDARY (READ ONLY):  59              PRIMARY:  10              PREFIX:  SKCMT
OPTIONS: 1 = DISPLAY  3 = PRINT  4 = COPY TO PRIMARY  7 = SUBMIT  8 = COMPILE

OPT  MEMBER NAME      NEW NAME              OWNER      PASSW  PRIVATE
 4   SKCMT             -----            05/27/2011  AAAA
 4   SKCMTINI          -----            05/27/2011  AAAA
 4   SKCMTREP          -----            05/27/2011  AAAA

PF1=HELP      2=REFRESH  3=END      4=RETURN
                9=SORT.DATE 10=SORT.NAME 11=SORT.SIZE 12=LIST QUEUE
LOCATE MEMBER/LIST QUEUE PREFIX ==> _____ MEMBER PREFIX (PF2) ==> SKCMT

```

Figure 159. Interactive User Interface (IUI): Secondary Library panel

- If you do not have access to the Interactive User Interface and as long as you did not delete the DTS file, you can:
 - Use the DTSUTIL utility to PRINT or PUNCH the skeletons. For details about DTSUTIL, see *VSE/ICCF Administration and Operation*, SC33-8329.
 - Use the VSE Navigator to access the ICCF library contents. This requires setting up the VSE Connector Server on the z/VSE system and the VSE Navigator on your workstation. This will actually run DTRUTIL in the background.

SKCMT skeleton

The SKCMT skeleton contains samples for the CMTLAB and CMTSTART procedures.

CMTLAB

When started, procedure CMTLAB adds labels for the IJSYSCC, IJSYSC1, and IJSYSC2 files to the system standard label group.

The samples that follow use these variables:

aa

Begin track/block

bb

Number of tracks/blocks

valid

Volume serial number of the disk (do not use DOSRES or SYSWK1)

Figure 160 on page 312 shows sample **DLBL** and **EXTENT** statements for the IJSYSCC file. The value for *bb* is:

- 1 track for an ECKD disk
- 64 blocks for an FBA disk

```
// DLBL IJSYSCC,'CMT.CONTROL.FILE',99/366,SD
// EXTENT ,valid,1,0,aa,bb
```

Figure 160. Sample DLBL and EXTENT for file IJSYSCC

Figure 161 on page 312 shows sample **DLBL** and **EXTENT** statements for the IJSYSC1 and IJSYSC2 files. The minimum value for *bb* is:

- 1584 tracks for an ECKD disk
- 175830 blocks for an FBA disk

These required minimum sizes for the IJSYSC1 and IJSYSC2 data files guarantee that the SCRT89 records for a given month do not overwrite the SCRT89 records for the previous month. IBM strongly recommends that you use these minimum values for ECKD or FBA disks.

Note: If you choose to override the recommended defaults and you define data files that exceed 32 767 tracks on an ECKD disk, you must install the fix for APAR DY47152 in order to process such large files with IJBCMTRP.

```
// DLBL IJSYSC1,'CMT.SCRT89.DATA.FILE1',99/366,SD
// EXTENT ,valid,1,0,aa,bb

// DLBL IJSYSC2,'CMT.SCRT89.DATA.FILE2',99/366,SD
// EXTENT ,valid,1,0,aa,bb
```

Figure 161. Sample DLBL and EXTENT statements for files IJSYSC1 and IJSYSC2

CMTSTART

When started, the CMTSTART procedure performs these steps:

- Runs the CMTLAB procedure
- Loads the CMT phase IJBCMT0 into the SVA
- Starts CMT using the following **EXEC** statement:

```
// EXEC IJBCMT,PARM='START ID=xxxx'
```

where xxxx is a four-character CMT ID, or SYSID, that uniquely identifies your z/VSE system within the CPC. The xxxx value can consist of any combination of the characters A-Z and 0-9 (for example, VSE1).

Rule: Do not use a CMT ID that is already in use for another z/VSE system within the CPC (either LPAR or VM guest) or a z/OS system within the CPC (**SYSID** parameter).

SKCMTINI skeleton

The SKCMTINI skeleton contains a sample job, CMTINIT, to initialize the IJSYSCC, IJSYSC1, and IJSYSC2 files. The CMTINIT job runs the CMTLAB procedure and initializes the files using this **EXEC** statement:

```
// EXEC IJBCMT,PARM='INITDS'
```

SKCMTREP skeleton

The SKCMTREP skeleton contains a sample job, CMTREPO, to extract the SCRT89 records for a given month into the CMTREPO file.

SKUSERBG skeleton

The SKUSERBG skeleton provides a skeleton to automatically start CMT during IPL using this **EXEC** statement:

```
// EXEC PROC=CMTSTART
```

Steps for preparing z/VSE to use CMT

Perform the following steps to prepare your z/VSE system to use CMT:

1. Customize the skeletons for your z/VSE system:

- Customize the SKCMT skeleton and catalog the CMTLAB and CMTSTART procedures.
- Customize the SKCMTINI skeleton to prepare the CMTINIT job.
- Customize the SKUSERBG skeleton to start CMT automatically during IPL through the CMTSTART procedure.

For details, see [“Overview of skeletons for CMT” on page 310](#)

2. Initialize the IJSYSCC, IJSYSC1, and IJSYSC2 files.

You can use the CMTINIT job that you customized in step 1. Initialize the files once, before you start CMT the very first time.

Rule: After CMT has written the SCRT89 records, you must not run the CMTINIT job again. Otherwise, all of your SCRT89 records will be lost because the CMTINIT job erases any data in the IJSYSCC, IJSYSC1, and IJSYSC2 files.

3. When you customize your system following the procedure in step 1, CMT starts automatically when you IPL your z/VSE system.

IBM recommends that CMT be automatically started during IPL. However, if you did not place the CMTSTART procedure in your USERBG.PROC, ensure that CMT always starts when you IPL your z/VSE system.

Note: Failure to report capacity for any LPAR or z/VM guest will result in all of the MWLC products on the machine being charged at the full machine capacity.

If you have customized your system, initialized the IJSYSCC, IJSYSC1, and IJSYSC2 data files and now want to start CMT without an IPL of your z/VSE system, you can:

- Load phase IJBCMT0 into the SVA through **SET SDL**.
- Start CMT using this **EXEC** statement:

```
// EXEC IJBCMT,PARM='START ID=xxxx'
```

Note: If you run CMTSTART.PROC after you run CMTINIT without an IPL of your z/VSE system, you might get the error messages 1L30D and 1S12D. This is because the labels for the data files were

already added to the standard label area. You can accept each error; the CMTSTART . PROC will start CMT.

4. You know the CMT task is active when the **AR SIR** command displays:

```
CAPACITY MEASUREMENT ACTIVE
```

If you must stop CMT, use:

```
// EXEC IJBCMT,PARM='STOP'
```

After stopping CMT, no more SCRT89 records are written.

Sub-capacity reporting using SCRT89 records

z/VSE produces accounting records called SCRT89 records for use in sub-capacity pricing. SCRT processes the data from the SCRT89 records to generate a sub-capacity report based on the 4-hour rolling average utilization of the z/VSE systems.

Whenever the SCRT89 records are generated by a z/VSE V4.1 system with APAR DY47027 or a z/VSE V4.2 system with APAR DY47029, or later, SCRT will recognize and use z/VSE system capacities of less than one MSU when calculating the overall z/VSE capacity across all of the z/VSE systems on the CPC. This enhancement is further discussed where applicable under [“Sample sub-capacity reports for z/VSE”](#) on page 322.

If you run both z/VSE and z/OS on your participating CPC and have sub-capacity pricing for both z/VSE and z/OS, you must produce the sub-capacity report on the z/OS system. In this case, use the SCRT89 records from your z/VSE systems and the SMF type 70 and type 89 records from your z/OS systems as input to the same execution of SCRT, so that they appear on the same sub-capacity report. Follow the procedure in [“Processing SCRT89 data on a z/OS system”](#) on page 314.

If you only have sub-capacity pricing for z/VSE or do not have z/OS in your environment, you can produce the sub-capacity report by running SCRT on a Windows or Linux system. Use the SCRT89 records from all of your z/VSE systems as input to the same execution of SCRT. SCRT produces a separate sub-capacity report for each CPC that it encounters in the SCRT89 records. Follow the procedure in [Chapter 4, “Using SCRT on Windows and Linux systems,”](#) on page 69.

Processing SCRT89 data on a z/OS system

Before you begin processing SCRT89 data by running SCRT on a z/OS system, follow the procedure in [“Preparing to process SCRT89 data on a z/OS system”](#) on page 316. Then, once each month, follow the procedure in [“Steps for processing SCRT89 data on a z/OS system”](#) on page 317 to process and transfer the SCRT89 records for each z/VSE system on the participating CPC.

Overview of processing SCRT89 data on a z/OS system

Figure 162 on [page 315](#) shows an overview of processing SCRT89 data from z/VSE systems and running SCRT on a z/OS system.




Figure 162. Overview of processing SCRT89 data from z/VSE systems and running SCRT on a z/OS system

The numbered, high-level steps shown in [Figure 162 on page 315](#) are:

- 1** CMT continuously writes SCRT89 records into its BAM data sets.
Then, once each month, the following steps must be done:
- 2** EXTRACT: On each z/VSE system, run the CMTREPO job to extract the SCRT89 data for the reporting period into a VSAM ESDS file.
- 3** TRANSFER: Transfer the VSAM ESDS files (containing the SCRT89 records) from the z/VSE system(s) to the z/OS system where SCRT will be run.
- 4** RUN SCRT: Run SCRT on the z/OS system. The input for SCRT consists of all of the SCRT89 records from all of the z/VSE systems. SCRT processes the SCRT89 records and writes the sub-capacity report to the REPORT file.

5

Transfer the sub-capacity report to your workstation. Verify the report. Then, submit the report to IBM.

[“Steps for processing SCRT89 data on a z/OS system” on page 317](#) describes these steps in detail.

Preparing to process SCRT89 data on a z/OS system

Do the following to prepare for processing your SCRT89 data with SCRT running on a z/OS system:

1. **Download the SCRT deliverable** that runs on a z/OS system from the SCRT home page. For instructions, see [“Obtaining and installing SCRT on a z/OS system” on page 21](#).
2. **Define a file as the target for the extracted SCRT89 records on your z/VSE systems.** The output file for the CMTREPO job must be a VSAM entry-sequenced data set (ESDS) with the following attributes:

Attribute	Value
Control interval size	32768
Average record size	1024
Maximum record size	32756

Define the output file as Work File (REUSE). Choose any name for the file and, later, use this name in the SKCMTREP skeleton for the VSAM label, CMTREPO. There should be enough space available for the file to keep at least 2000 records. The following examples use the name SCRT89 for the VSAM file.

The following example job shows how to define the SCRT89 record file:

```
* $$ JOB JNM=SCRT89,CLASS=0,DISP=D,NTFY=YES
// JOB SCRT89 - DEFINE SCRT89 RECORD FILE
// EXEC IDCAMS,SIZE=AUTO
  DEFINE CLUSTER ( -
    NAME (SCRT89) -
    RECORDS (2000 2000) -
    RECORDSIZE (1024 32756) -
    REUSE -
    SHAREOPTIONS (1) -
    VOLUMES (MYVOL1 MYVOL2 MYVOL3) -
    NONINDEXED /* ESDS */ -
    TO (99366) -
  ) -
  DATA ( -
    NAME (SCRT89.DATA) -
    CONTROLINTERVALSIZE (32768) -
  ) -
  CATALOG (MYCAT)
  IF LASTCC NE 0 THEN CANCEL JOB
/*
/&
* $$ E0J
```

3. Customize the TCP/IP stack on your z/VSE systems.

To use FTP to transfer the CMTREPO file to your workstation or to a z/OS system, you must define the file in the TCP/IP partition. The following definitions are valid for the TCP/IP for VSE/ESA product. If you use another TCP/IP product, adjust the definitions accordingly.

- a. Define the CMTREPO file in the TCP/IP partition by editing the TCP/IP startup job and adding the following definition. Adjust for your system and use the file name that you chose in step [“2” on page 316](#).

```
// DLBL CMTREPO, 'SCRT89', ,VSAM,CAT=MYCAT
```

- b. Make the files accessible in TCP/IP by editing the TCP/IP initialization member (for example, IPINIT00.L) and adding the following definition:

```
MODIFY FILE,PUBLIC=CMTREPO,TYPE=ESDS,RECFM=V,LRECL=32756
```

Tip: Depending on your system configuration, you might need to use the **DEFINE FILE** command instead of the **MODIFY FILE** command.

- c. Ensure there is a FTP daemon defined in your TCP/IP (DEFINE FTPD).
- d. Restart the TCP/IP partition to activate the changes.

Steps for processing SCRT89 data on a z/OS system

Before you begin: Review the illustration and high-level steps in “Overview of processing SCRT89 data on a z/OS system” on page 314 and make sure you have completed the steps in “Preparing to process SCRT89 data on a z/OS system” on page 316.

Do the following to process your SCRT89 data with SCRT running on a z/OS system:

1. Extract the SCRT89 records on your z/VSE systems.

The reporting period for SCRT is from the second day of the month to the first day of the following month. For example, a report for June would include the SCRT89 records from 2 June through 1 July. On the second day of the month, run the CMTREPO job to extract the SCRT89 records from the BAM data sets into the CMTREPO VSAM file.

Customize a CMTREPO job using the SKCMTREP skeleton. Replace the BAM label called CMTREPO with a VSAM label of the same name.

Example: The following is a VSAM label with a VSAM file name of SCRT89:

```
// DLBL CMTREPO, 'SCRT89', ,VSAM,CAT=MYCAT
```

If you specified MONTH=LAST, the job extracts the SCRT89 records for the previous month from the data files IJSYSC1 and IJSYSC2 and writes them to the CMTREPO VSAM file.

Example: If you run the job on 2 March and specify MONTH=LAST, the job extracts SCRT89 records from 2 February through 1 March.

If you specify MONTH=*month/year*, the job extracts the records from the second day of *month* through the first day of the following month.

On z/VSE V5.1 systems with APAR DY47645, z/VSE V5.2 systems with APAR DY47646, or z/VSE V6.1 systems with APAR DY47647, if you specify MONTH=CURRENT, the job extracts the SCRT89 records for the current month.

Example: If you run the job on 6 March and specify MONTH=CURRENT, the job extracts the SCRT89 records from 2 March through 1 April.

The **OUTPUTDS** parameter specifies the target data set for the extracted records. The value of the **OUTPUTDS** parameter must be a valid Language Environment file name, hence it is preceded by DD:.

On z/VSE V4.1 systems with APAR DY47027 or z/VSE V4.2 systems with APAR DY47029, you can use the **OUTPUTMODE** parameter to indicate whether the CMTREPO output file is to be overwritten or appended. The default is OUTPUTMODE=REUSE and indicates that the output file is to be overwritten each time the CMTREPO job is run, which is the same behavior as before the application of the specified APARs. The OUTPUTMODE=APPEND setting indicates that records are to be appended to the end of the output file.

2. Transfer the extracted SCRT89 records (CMTREPO file) to your workstation.

Use FTP to transfer the extracted SCRT89 records (CMTREPO file) from your z/VSE system to your workstation. See step 3 under “Preparing to process SCRT89 data on a z/OS system” on page 316. On your workstation, ensure that you choose a unique file name for the CMTREPO file of each z/VSE system.

The FTP transfer must be in binary format (ftp> bin) and in block mode (ftp> quote mode b).

Example: Figure 163 on page 318 shows a sample FTP session to transfer the extracted SCRT89 records (CMTREPO) from z/VSE to the workstation.

```

C:\cmtdata> ftp zvsehost
Connected to zvsehost.
220-TCP/IP for VSE Internal FTPDAEMN 01.05 E 20061010 11.26
    Copyright (c) 1995,2006 Connectivity Systems Incorporated
220 Ready for new user
User (zvsehost:(none)): zvseuser
331 User name okay, need password
Password: zvsepass
230 User logged in, proceed
ftp> bin
200 Command okay
ftp> quote mode b
200 Command okay
ftp> get CMTREPO scrt89-vse1.bin
200 Command okay
150-About to open data connection
    File:CMTREPO
    Type:Binary Recfm:V Lrecl: 32756
    CC=ON UNIX=OFF RECLF=OFF TRCC=OFF CRLF=ON NAT=NO CONT=OFF
    MODE=Block STRU=File
150 File status okay; about to open data connection
226-Bytes sent: 36,215
    Transfer Seconds:      .20 ( 177K per second)
    File I/O Seconds:     .02 ( 1768K per second)
226 Closing data connection
ftp: 36215 bytes received in 0,80Seconds 45,21Kbytes/sec.
ftp> bye
221 FTPDaemn closing control connection

```

Figure 163. Sample FTP transfer of SCRT89 records (CMTREPO) from z/VSE to the workstation

Alternative method: Transfer the SCRT89 records directly to the z/OS system

You can also transfer the SCRT89 records directly from your z/VSE systems to the z/OS system on which you will run SCRT. This method eliminates the need to first transfer the SCRT89 records to the workstation, as described earlier in this step, and then transfer them from the workstation to the z/OS system, as described in step “3” on page 318.

Figure 164 on page 318 shows a sample z/VSE job to transfer a file directly to a z/OS system.

```

* $$ JOB JNM=SENDDATA,CLASS=0,DISP=D
// JOB SENDDATA
// EXEC FTP,SIZE=FTP,PARM='ID=00'
* This FTP job transfers the z/VSE file
* VSEDATA to the specified target file
* on a z/OS system.
NOUPRPT
OPEN zoshost
USER zosuser
PASS zospass
LOPEN
LUSER zvseuser
LPASS zvsepass
BINARY
MODE B
QUOTE MODE B
PUT VSEDATA 'zosuser.SCRTT00L.VSE.SCRT89.DATA'
BYE
LCLOSE
/*
/&
* $$ E0J

```

Figure 164. Sample z/VSE job to transfer a file to a z/OS system

3. Transfer the workstation files containing your z/VSE systems' SCRT89 records into z/OS data sets, as follows:

- a. Allocate one data set on z/OS for each z/VSE system on the participating CPC. Each z/OS data set will contain the SCRT89 records of one z/VSE system. The files you transfer from your workstation

to z/OS are the (CMTREPO) files you downloaded from your z/VSE systems in step 2 (for example, `scrt89-vse1.bin` file). Use the following attributes to allocate the z/OS data sets:

Attribute	Value
Directory blocks	0
Record format	VB
Record length	32756
Block size	32760

Adjust the size of the z/OS data sets according to the size of the SCRT89 records (CMTREPO) file you downloaded from z/VSE in step 2.

- b. Use FTP to upload the SCRT89 files of all your z/VSE systems from your workstation into the new z/OS data sets.

Example: Figure 165 on page 319 shows a sample FTP session to upload the `scrt89-vse1.bin` file (containing a month of SCRT89 records of one z/VSE system) into the z/OS data set named `SCRT89.RECORDS.VSE1`.

```
C:\Temp> ftp zoshost
Connected to zoshost.
220-FTP, 08:28:03 on 2006-12-27.
220 Connection will close if idle for more than 60 minutes.
User (zoshost:(none)): zosuser
331 Send password please.
Password: zospass
230 ZOSUSER is logged on. Working directory is "ZOSUSER.".
ftp> bin
200 Representation type is Image
ftp> quote mode b
200 Data transfer mode is Block
ftp> put scrt89-vse1.bin SCRT89.RECORDS.VSE1
200 Port request OK.
125-Waiting for recall of data set ZOSUSER.SCRT89.RECORDS.VSE1
125 Storing data set ZOSUSER.SCRT89.RECORDS.VSE1
250 Transfer completed successfully.
ftp: 8619 bytes sent in 0,00Seconds 8619000,00Kbytes/sec.
ftp> bye
221 Quit command received. Goodbye.
```

Figure 165. Sample FTP transfer of z/VSE SCRT89 records from the workstation to z/OS

You know you are done when you receive the message: Transfer completed successfully.

4. Customize the SCRT job and run SCRT on a z/OS system, as follows:

- a. Update the JCL to include all the z/OS data sets with z/VSE SCRT89 records for processing (for example, `SCRT89.RECORDS.VSE1`). See the instructions in [Chapter 3, "Using SCRT on z/OS systems,"](#) on page 21.
- b. Customize the PARMS control statements according to the instructions in ["Customizing the PARMS control statements"](#) on page 320.
- c. Customize the NO89 control statements according to the instructions in ["Customizing the NO89 control statements"](#) on page 320.
- d. (Optional) Customize the SPECIAL control statements, if needed, according to the instructions in ["Customizing the SPECIAL control statements"](#) on page 321.
- e. Run SCRT to generate a sub-capacity report for each CPC participating in sub-capacity pricing.
- f. Download the sub-capacity report (.csv) files that were generated by SCRT to your workstation and submit them to IBM. For details, see [Chapter 26, "Submitting sub-capacity reports to IBM,"](#) on page 427.

Processing SCRT89 data on a Windows or Linux system

If you want to process SCRT89 data from one or more z/VSE systems by running SCRT on a Windows or Linux system, see [Chapter 4, “Using SCRT on Windows and Linux systems,” on page 69.](#)

Customizing the SCRT control statements

You must customize the PARMS, NO89, and (optionally) SPECIAL control statements, as described in the following topics.

Customizing the PARMS control statements

Customize the SCRT PARMS control statements, either in the SCRT JCL that runs on z/OS or in the SCRT for Windows and Linux user interface, so that they contain the requested information. For details about the PARMS control statements, see [“PARMS DD statement” on page 34.](#)

SCRT will not generate a sub-capacity report if the values for any of the required keywords are left blank.

Customizing the NO89 control statements

The process to customize the NO89 control statements differs depending on whether you are running SCRT on a z/OS system or using SCRT for Windows and Linux.

Process when running SCRT on a z/OS system

The SCRT deliverable that runs on z/OS contains two sets of NO89 control statements: one set for z/OS products and one set for z/VSE products. Both sets of control statements are preceded by comments indicating the operating system belonging to the corresponding set of statements. Only use both sets of NO89 control statements when the input to SCRT contains both SMF type 89 records from z/OS systems and SCRT89 records from z/VSE systems. If you are only running z/VSE systems, delete the set of z/OS NO89 control statements. For example, you would delete the z/OS statements that follow this section:

```
* ----- Z/OS NO89S -----
*
*   DELETE THE FOLLOWING STMTS IF Z/OS IS *NOT* RUNNING ON THE CPC
*
```

Customize the values on each of the NO89 control statements as described in [“Customizing NO89 values” on page 320.](#)

Process when running SCRT for Windows and Linux

See [“Using the SCRT graphical user interface” on page 73](#) or [“Using the SCRT command line interface” on page 90](#) for details about how to specify NO89 products when using SCRT for Windows and Linux.

Customizing NO89 values

You must enter a valid value for each product in the list of z/VSE NO89 control statements. For additional information, see [“NO89 DD statement” on page 61.](#)

- If you do not have a product installed or are not using that product on the CPC(s) being processed by that execution of SCRT, enter *NONE next to the product number—do not delete products from the list of z/VSE NO89 control statements.
- If a product in the list runs in all LPARs on all of the CPCs being processed by that execution of SCRT, enter *ALL.
- If a product runs in one or more LPARs, enter a list of the LPAR names in which it runs.
- If the z/VSE system is running as a z/VM guest, list the virtual machine's user ID for that z/VSE system, instead of the LPAR name.

Guidelines for selected NO89 products: The following NO89 products might require special attention when completing the NO89 control statements, depending on how you ordered and used these products:

- DITTO/ESA FOR VSE (5648-099)
 - If you ran DITTO under z/VSE, use the NO89 control statement to assign 5648-099 to every z/VSE native system (LPAR name) and every z/VSE guest system (z/VM guest ID) where DITTO was used.
 - If you ran DITTO under z/VM in an LPAR without any z/VSE guests, do not assign 5648-099 to that LPAR because you are already being billed separately for DITTO under z/VM.
- Db2 SERVER FOR VSE & VM (5697-F42)
 - If you ran Db2 under z/VSE, use the NO89 control statement to assign 5697-F42 to every z/VSE native system (LPAR name) and every z/VSE guest system (z/VM guest ID) where any of the DB2 features were used, with the exception of the Data Restore for VM/VSE feature. When the Data Restore feature was the only Db2 feature being used on a z/VSE system, then do not assign 5697-F42 to that z/VSE native or guest system because you are already being billed separately for this feature.
 - If you ran Db2 under z/VM in an LPAR without any z/VSE guests, do not assign 5697-F42 to that LPAR because you are already being billed separately for Db2 under z/VM.
- HIGH LEVEL ASSEMBLER VSE & VM (5696-234)
 - If you ran High Level Assembler (HLASM) only under z/VSE and not under z/VM, use the NO89 control statement to assign 5696-234 to every z/VSE native system (LPAR name) and every z/VSE guest system (z/VM guest ID) where HLASM was used.
 - If you ran HLASM under both z/VM and z/VSE, do not assign 5696-234 to any LPAR because you are already being billed separately for HLASM under z/VM.

Figure 166 on page 321 shows a sample of the NO89 control statements and comments.

```
*
* DB2 SERVER FOR VSE & VM
5697-F42=
* DITTO/ESA FOR VSE
5648-099=
* DL/1 DOS/VS
5746-XX1=
* HIGH LEVEL ASSEMBLER VSE & VM
5696-234=
* IBM C FOR VSE/ESA V1
5686-A01=
* IBM COBOL VSE/ESA
5686-068=
* IBM DFSORT/VSE V3
5746-SM3=
* IBM PL/1 FOR VSE/ESA
5686-069=
* MQSERIES FOR VSE/ESA V2
5686-A06=
* WEBSPHERE MQ FOR VSE/ESA V3
5655-U97=
* TCP/IP FOR VSE
5686-A04=
* (NEW 23.5.0) TCP/IP FOR Z/VSE V2
5686-CS1
* IBM RATIONAL COBOL RUNTIME FOR Z/VSE V7
5648-F66=
* IPv6/VSE V1
5686-BS1
```

Figure 166. Sample NO89 control statements for z/VSE

Customizing the SPECIAL control statements

The SPECIAL control statements are optional and should only be used to deal with unusual reporting situations where SCRT89 records need to be excluded from SCRT processing. Such conditions might involve using the SPECIAL control statements under the guidance of the SCRT support team.

For explanations of the SPECIAL control statements, see [“SPECIAL DD statement”](#) on page 40.

To use the SPECIAL control statements, either modify the JCL for your SCRT job on z/OS, or, for SCRT on Windows and Linux, see [“Using the SCRT graphical user interface”](#) on page 73 or [“Using the SCRT command line interface”](#) on page 90.

Backup considerations

It is your responsibility to create a back up copy of the extracted SCRT89 records, together with the PARMS and NO89 statements. You can create a copy of the SCRT89 records on z/VSE using a tape or a VSAM file (with this method you have to back up the PARMS and NO89 statements separately). IBM recommends that you create a new file for the SCRT89 records of each report period. This allows you to selectively delete the SCRT89 records of old report periods.

Rule: IBM requires that customers retain their SCRT89 records for at least six months.

If you want to create a copy of the SCRT89 records using a tape unit as target for the written records, create the **OUTPUTDS** parameter as either: `DD:lu-tlbl` or `DD:lu`. The *lu* specifies a logical unit in the form `SYSnnn`. You must assign a tape to this logical unit. The *tbl* (tape label), concatenated by a dash, specifies to write to a labeled tape file. If *tbl* is not specified, the data is written to an unlabeled tape file. The specified *tbl* must be present in the system standard, class, or partition label area. Multivolume files and multiple files on a volume are supported. For example:

```
OUTPUTDS = DD:SYS004-TAPEFIL
```

To back up the data records to a VSAM file:

1. Create a VSAM ESDS file, with a maximum record size of 32756 bytes and variable record length.
2. Add the **DLBL** for the file to the CMTREPO job.
3. Specify the **DLBL** name in the **OUTPUTDS** parameter.

For example:

```
OUTPUTDS = DD:CMT0612
```

Sample sub-capacity reports for z/VSE

With few exceptions, the sub-capacity reports for z/VSE systems have the same format and same fields as the sub-capacity reports for z/OS systems. You can find the information about these fields in [Chapter 5, “Sub-capacity reporting for z/OS native systems,”](#) on page 103.

The only differences in sub-capacity reporting between z/VSE and z/OS systems are:

- The products that are reported
- The reporting of z/VSE systems whose capacities are less than one MSU

See [“Introduction to sub-capacity pricing for z/VSE systems”](#) on page 305 for a general description of this support. Reporting differences due to z/VSE system capacities less than one MSU are explained where applicable for the sample sub-capacity reports that follow.

- The guest naming conventions used by SCRT for z/OS guest systems versus z/VSE guest systems
z/OS guest systems are either shown by the LPAR name in which the z/OS guest runs (for example, LPAR1) or by the format `sysid(guest)` where `sysid` is the z/OS SYSID for the guest system (for example, V2A(guest)). SCRT uses the format `LPAR_name(VM_userid)` to show z/VSE guest systems (for example, VSEVM(VGDEV1), VSEVM(VGDEV2) in the sample reports that follow). See [Table 48 on page 308](#) for a more detailed description for all z/VSE systems.
- The additional z/VSE information that appears on a sub-capacity report when multiple z/VSE versions run on the same CPC during the same reporting period

[“Sample sub-capacity report for z/VSE V4 only”](#) on page 323 illustrates a sample sub-capacity report that only contains z/VSE V4 systems, some of whose capacities are less than one MSU.

“Sample sub-capacity report for z/VSE V4 and V5” on page 327 illustrates a sample sub-capacity report that includes a mix of z/VSE V4 and V5 native and guest systems. All of the z/VSE systems used capacities of more than one MSU.

“Sample sub-capacity report for z/OS and z/VSE V4” on page 333 illustrates a sample sub-capacity report that includes z/VSE systems, some of whose capacities are less than one MSU, and one z/OS system.

“Sample sub-capacity report for z/VSE V6, z/VSE V5, and z/VSE V4” on page 330 illustrates a sample sub-capacity report that includes three different z/VSE versions.

Sample sub-capacity report for z/VSE V4 only

The following sample sub-capacity report is for a CPC with three LPARS named VSEPROD, VSEVM and LINUX:

- VSEPROD runs a native z/VSE V4 (Central Functions V8) system, VPRD (CMT ID = VPRD).
- VSEVM runs a z/VM system that supports two z/VSE V4 guest systems, whose virtual machines have been assigned VM user IDs of VGDEV1 and VGDEV2 (CMT ID = DEV1 and DEV2).
- LINUX runs a z/VM system that supports a number of LINUX guest systems and uses approximately 60 percent of the capacity of the CPC.

Table 49 on page 323 shows the sample configuration.

CPC: 2097-405 (serial = 02-123AB)			
VSEPROD	VSEVM (z/VM)		LINUX (z/VM)
z/VSE V4 (Central Func V8) (CMT ID = VPRD)	z/VSE V4 (Central Func V8) (CMT ID = DEV1) (VM user ID = VGDEV1)	z/VSE V4 (Central Func V8) (CMT ID = DEV2) (VM user ID = VGDEV2)	Linux (Uses about 60% of the CPC)
ACF/VTAM CICS TS Db2 Server DITTO/ESA DL/1 HL Assembler IBM C IBM DFSort IBM PL/1 MQSeries TCP/IP	ACF/VTAM CICS TS Db2 Server HL Assembler IBM C IBM COBOL IBM DFSort IBM PL/1 TCP/IP	ACF/VTAM HL Assembler IBM C IBM COBOL IBM DFSort IBM PL/1 MQSeries TCP/IP	

Figure 167 on page 324 through Figure 171 on page 326 show a sample sub-capacity report based on the configuration given in Table 49 on page 323.

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 03 Nov 2024 - 09:17
 Name of Person Submitting Report: John Customer
 E-mail Address of Report Submitter: customer@abc.com
 Phone Number of Report Submitter: 444-999-9999

Customer Name ABC Corporation
 Customer Number 8971234567
 Machine Serial Number 02-123AB
 Machine Type and Model 2097-405
 Machine Rated Capacity (MSUs) 118
 Purchase Order Number (optional)
 Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

==C5=====

TOOL INFORMATION

Tool Release 30.1.0
 Reporting Period 2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

Figure 167. Sample sub-capacity report for z/VSE V4 only: Customer information and Tool Information sections

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/VSE V4 (Central Functions V8)	5686-CF8	37		
ACF/VTAM V4 VSE/ESA	5686-065	37		
CICS TS for VSE/ESA	5648-054	37		
DB2 Server for VSE & VM	5697-F42	37	(optional)	(conditional)
DITTO/ESA for VSE	5648-099	37	(optional)	(conditional)
DL/1 DOS/VSE	5746-XX1	37	(optional)	(conditional)
High Level Assembler VSE & VM	5696-234	37	(optional)	(conditional)
IBM C for VSE/ESA V1	5686-A01	37	(optional)	(conditional)
IBM COBOL VSE/ESA	5686-068	1	(optional)	(conditional)
IBM DFSORT/VSE V3	5746-SM3	37	(optional)	(conditional)
IBM PL/1 for VSE/ESA	5686-069	37	(optional)	(conditional)
MQSeries VSE/ESA	5686-A06	37	(optional)	(conditional)
TCP/IP for VSE	5686-A04	37	(optional)	(conditional)

Figure 168. Sample sub-capacity report for z/VSE V4 only: Product Summary Information section

Even though SCRT processed SCRT89 records from z/VSE guest systems VGDEV1 and VGDEV2, both of which had capacities of less than one MSU, all of the **Tool MSUs** values reported in the Product Summary Information section (Figure 168 on page 324) have been converted to integers (whole number values). In this sample configuration, the COBOL product only ran on the VGDEV1 and VGDEV2 systems, whose peak summed capacities were slightly higher than one MSU, so SCRT reported the capacity of the COBOL product as one MSU. However, if the peak summed capacities of those systems had been less than one MSU, SCRT would still have reported a minimum of one MSU for that product, since one MSU is the lowest value that SCRT will report for a product that ran during the reporting period.

You can find more information about the fields in this section and other sections of the sub-capacity report in [Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103.](#)

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

          SYSID      Input Data Start      Input Data End
VSEPROD      VPRD      01 Oct 2024 - 00:00  02 Nov 2024 - 22:00
VSEVM(VGDEV1) DEV1      01 Oct 2024 - 06:05  02 Nov 2024 - 23:45
VSEVM(VGDEV2) DEV2      01 Oct 2024 - 00:00  30 Oct 2024 - 22:30
LINUX        -         -                 -
CPC           -         01 Oct 2024 - 00:00  02 Nov 2024 - 23:45

==N5=====
DETAIL LPAR DATA SECTION

          Highest Hour Count Date/Time      2nd Highest Hour Count Date/Time
VSEPROD      36          1 03 Oct 2024 - 07:00      35          5 03 Oct 2024 - 06:00
VSEVM(VGDEV1) 1          1 02 Oct 2024 - 11:00      (v)         688 03 Oct 2024 - 10:00
VSEVM(VGDEV2) (v)        680 02 Oct 2024 - 00:00
CPC           37          5 03 Oct 2024 - 12:00      36          14 03 Oct 2024 - 11:00

Footnotes:
(v) Value is less than 1 MSU

```

Figure 169. Sample sub-capacity report for z/VSE V4 only: SMF / SCRT89 Input Data Statistics and Detail LPAR Data sections

The fact that z/VSE guest systems VGDEV1 and VGDEV2 both had capacities of less than one MSU during the reporting period, as either their highest value or their second highest value, triggers a change in the format of the Detail LPAR Data section (N5) in the sample sub-capacity report shown in [Figure 169](#) on [page 325](#). Whenever SCRT would have reported a value of less than one MSU, SCRT uses the footnote (v), instead of the MSU value, and adds a description of the footnote at the end of that report section.

In section N5 of the sample sub-capacity report, VGDEV1 had a single hour where its capacity was one MSU and had a value of less than one MSU for all of the remaining hours in the reporting period. SCRT reports this by using footnote (v) as the capacity value in the **2nd Highest** column and reports all hours that had a value less than one MSU as the corresponding **Hour Count** value. The highest capacity for VGDEV2 was less than one MSU, so SCRT reports that by using footnote (v) as the capacity value in the **Highest** column for VGDEV2. SCRT reports all of the VGDEV2 hours as the **Hour Count** value, and reports no **2nd Highest** value for VGDEV2.

Do not interpret the lack of footnote (v) in section N5 to mean that the z/VSE systems listed in this section did not have any hours when their capacity was less than one MSU. The lack of that footnote simply means that the **Highest** and **2nd Highest** values for that z/VSE system were higher than one MSU. It is still possible that such a system could have had other hours when its capacity was less than one MSU and so could contribute less than one MSU to the earliest peak hour for a product in the Product Max Contributors section (P5), where SCRT would then report footnote (v) for that z/VSE system.

You can find more information about the fields in this section and other sections of the sub-capacity report in [Chapter 5, “Sub-capacity reporting for z/OS native systems,”](#) on [page 103](#).

[Figure 170](#) on [page 326](#) shows the Product Max Contributors section of the sample sub-capacity report for z/VSE.

Whenever the contribution of a z/VSE system to the **Highest** product value reported in the Product Max Contributors section is less than one MSU, SCRT uses footnote (v) to report that system's contribution to the product's MSU value. Whenever SCRT reports footnote (v) in the Product Max Contributors section, SCRT adds a description of that footnote at the end of that report section.

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR VSEPROD	LPAR VSEVM (VGDEV1)	LPAR VSEVM (VGDEV2)
z/VSE V4 (Central Functions V8)	5686-CF8	37	03 Oct 2024 - 07:00	36	(v)	(v)
ACF/VTAM V4 VSE/ESA	5686-065	37	03 Oct 2024 - 07:00	36	(v)	(v)
CICS TS for VSE/ESA	5648-054	37	03 Oct 2024 - 07:00	36	(v)	0
DB2 Server for VSE & VM	5697-F42	37	03 Oct 2024 - 07:00	36	(v)	(v)
DITTO/ESA for VSE	5648-099	37	03 Oct 2024 - 07:00	36	(v)	(v)
DL/1 DOS/VS	5746-XX1	37	03 Oct 2024 - 07:00	36	(v)	(v)
High Level Assembler VSE & VM	5696-234	37	03 Oct 2024 - 07:00	36	(v)	(v)
IBM C for VSE/ESA V1	5686-A01	37	03 Oct 2024 - 07:00	36	(v)	(v)
IBM COBOL VSE/ESA	5686-068	1	03 Oct 2024 - 11:00	0	1	(v)
IBM DFSort/VSE V3	5746-SM3	37	03 Oct 2024 - 07:00	36	(v)	(v)
IBM PL/1 for VSE/ESA	5686-069	37	03 Oct 2024 - 07:00	36	(v)	(v)
MQSeries VSE/ESA	5686-A06	37	03 Oct 2024 - 07:00	36	(v)	(v)
TCP/IP for VSE	5686-A04	37	03 Oct 2024 - 07:00	36	(v)	(v)

Footnotes:
(v) Value is less than 1 MSU

Figure 170. Sample sub-capacity report for z/VSE V4 only: Product Max Contributors section

In the sample Product Max Contributors section shown in [Figure 170 on page 326](#), the contribution to the **Highest** values from VGDEV1 and VGDEV2 is less than one MSU for all z/VSE products except for the IBM COBOL VSE/ESA product. SCRT reports that condition by using footnote (v) for all products except COBOL. However, when added together, VGDEV1 and VGDEV2 contributed one MSU, which is added to the VSEPROD contribution of 36 MSUs, for a total product value of 37 MSUs for all products except COBOL, which did not run on the VSEPROD system. In this sample report, the sum of the VGDEV1 and VGDEV2 systems totalled slightly more than one MSU for the COBOL product, which SCRT reported as one MSU. The peak value for the VGDEV1 and VGDEV2 systems was higher for the COBOL product than for the other products because COBOL did not run on VSEPROD and, thus, peaked at a different time than the other products. However, if the VGDEV1 and VGDEV2 systems had summed to less than one MSU, SCRT still would have reported a minimum of one MSU for COBOL because one MSU is the product minimum, but SCRT would have reported a total of 36 MSUs for all of the other products.

Although the peak capacity of the VGDEV1 system was one MSU for one hour according to the Detail LPAR Data section, that hour occurred later than the earliest product peak hour for all of the products except COBOL. Because the Product Max Contributors section only reports on the earliest peak hour, the one MSU rating for the VGDEV1 system was not posted for any of the products listed in that section except for COBOL.

You can find more information about the fields in this section and other sections of the sub-capacity report in [Chapter 5, "Sub-capacity reporting for z/OS native systems," on page 103](#).

```

==Q5=====
PRODUCT GRID SNAPSHOT

```

Product Name	Product ID	VSEPROD	VSEVM (VGDEV1)	VSEVM (VGDEV2)
z/VSE V4 (Central Functions V8)	5686-CF8	91.4%	92.7%	91.5%
ACF/VTAM V4 VSE/ESA	5686-065	91.4%	92.7%	91.5%
CICS TS for VSE/ESA	5648-054	91.4%	92.7%	

```

==W5=====

```

This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such Customer is solely responsible for the completeness and accuracy of information and data used to create this report. Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report. Any questions regarding the contents of this report should be directed to the Customer.

Figure 171. Sample sub-capacity report for z/VSE V4 only: Product Grid Snapshot section

The Sub-Capacity Capping Enforced by SCRT section and the Group Capacity LPARs section will not appear on an SCRT generated for a CPC which is only running z/VSE systems because z/VSE does not support defined capacity or group capacity.

Sample sub-capacity report for z/VSE V4 and V5

The following sample sub-capacity report is for a CPC with three LPARs named VSEPROD4, VSEPROD5, and VSEVM2:

- VSEPROD4 runs a native z/VSE V4 (Central Functions V8) system, PRD4 (CMT ID = PRD4).
- VSEPROD5 runs a native z/VSE V5 (Central Functions V9) system, PRD5 (CMT ID = PRD5).
- VSEVM2 runs a z/VM system that supports two z/VSE guest systems:
 - A z/VSE V4 system, DEV4 (CMT ID = DEV4), that runs under a z/VM user ID of VGDEV4
 - A z/VSE V5 system, DEV5 (CMT ID = DEV5), that runs under a z/VM user ID of VGDEV5

[Table 50 on page 327](#) shows the sample configuration.

CPC: 2097-403 (serial = 02-234CD)			
VSEPROD4	VSEPROD5	VSEVM2	
z/VSE V4 (Central Func V8) (CMT ID = PRD4)	z/VSE V5 (Central Func V9) (CMT ID = PRD5)	z/VSE V4 (Central Func V8) (CMT ID = DEV4) (VM user ID = VGDEV4)	z/VSE V5 (Central Func V9) (CMT ID = DEV5) (VM user ID = VGDEV5)
ACF/VTAM CICS TS Db2 Server DITTO/ESA DL/1 HL Assembler IBM C	ACF/VTAM Db2 Server DITTO/ESA DL/1 HL Assembler IBM C	ACF/VTAM CICS TS Db2 Server HL Assembler IBM C IBM COBOL IBM DFSort IBM PL/1 MQSeries TCP/IP	ACF/VTAM Db2 Server HL Assembler IBM C IBM COBOL IBM PL/1 TCP/IP

Except for the differences described in “Sample sub-capacity report for z/VSE V4 only” on page 323 and the differences described in this sample report, the sub-capacity reports have the same format and same data fields as those for z/OS systems. You can find information about these fields in [Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103](#).

The sample configuration in this example represents two permanent, native production z/VSE systems that support different workloads, and two permanent z/VSE development/test guest systems, each of which supports one of the native production systems. In this configuration, the customer has migrated one of the native production systems (CMT ID = PRD5) and its corresponding development guest system (CMT ID = DEV5) to z/VSE V5. The migration to z/VSE V5 was complete prior to the reporting period shown in the sample sub-capacity report.

[Figure 172 on page 328](#) through [Figure 176 on page 330](#) show the sample sub-capacity report based on the configuration shown in [Table 50 on page 327](#).

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 03 Nov 2024 - 09:17
 Name of Person Submitting Report: John Customer
 E-mail Address of Report Submitter: customer@abc.com
 Phone Number of Report Submitter: 444-999-9999

Customer Name ABC Corporation
 Customer Number 8971234567
 Machine Serial Number 02-234CD
 Machine Type and Model 2097-403
 Machine Rated Capacity (MSUs) 75
 Purchase Order Number (optional)
 Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

==C5=====

TOOL INFORMATION

Tool Release 30.1.0
 Reporting Period 2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

Figure 172. Sample sub-capacity report for z/VSE V4 and V5: Customer information and Tool Information sections

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/VSE	(All)	67		
z/VSE V5 (Central Functions V9)	5686-CF9	36		
z/VSE V4 (Central Functions V8)	5686-CF8	38		
ACF/VTAM V4 VSE/ESA	5686-065	67		
CICS TS for VSE/ESA	5648-054	38		
DB2 Server for VSE & VM	5697-F42	67	(optional)	(conditional)
DITTO/ESA for VSE	5648-099	49	(optional)	(conditional)
DL/1 DOS/VS	5746-XX1	49	(optional)	(conditional)
High Level Assembler VSE & VM	5696-234	67	(optional)	(conditional)
IBM C for VSE/ESA V1	5686-A01	67	(optional)	(conditional)
IBM COBOL VSE/ESA	5686-068	18	(optional)	(conditional)
IBM DFSort/VSE V3	5746-SM3	38	(optional)	(conditional)
IBM PL/1 for VSE/ESA	5686-069	67	(optional)	(conditional)
MQSeries VSE/ESA	5686-A06	38	(optional)	(conditional)
TCP/IP for VSE	5686-A04	67	(optional)	(conditional)

Figure 173. Sample sub-capacity report for z/VSE V4 and V5: Product Summary Information section

Since both z/VSE V4 (5686-CF8) and z/VSE V5 (5686-CF9) are reported in the Product Summary Information section, SCRT also reports an additional z/VSE entry which represents the peak concurrent sub-capacity value for both z/VSE V4 and z/VSE V5, as if they were treated as a single product. Under Multi-Version Measurement (MVM) rules, IBM will waive the charges for the z/VSE V4 product and will charge for the z/VSE V5 product based on the sub-capacity value reported for the z/VSE entry, which is 67 MSUs in this example. The IBM middleware products continue to be charged based on their reported peak sub-capacity values, which are independent of the z/VSE version on which those middleware products run. See [“Sub-capacity reporting for multiple z/VSE versions” on page 308](#) for more information.

Note that the additional z/VSE entry is reported whenever both z/VSE V4 and z/VSE V5 are running on the same CPC at any time during the reporting period.

You can find more information about the fields in this section and other sections of the sub-capacity report in [Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103](#).

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
VSEPROD4	PRD4	02 Oct 2024 - 00:00	01 Nov 2024 - 22:00
VSEPROD5	PRD5	02 Oct 2024 - 00:00	01 Nov 2024 - 22:00
VSEVM2(VGDEV4)	DEV4	02 Oct 2024 - 06:05	01 Nov 2024 - 23:45
VSEVM2(VGDEV5)	DEV5	02 Oct 2024 - 00:00	30 Oct 2024 - 22:30
CPC		02 Oct 2024 - 00:00	01 Nov 2024 - 23:45

```

==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
VSEPROD4	30	3	03 Oct 2024 - 07:00	29	5	03 Oct 2024 - 06:00
VSEPROD5	27	1	02 Oct 2024 - 10:00	26	7	02 Oct 2024 - 09:00
VSEVM2(VGDEV4)	12	6	03 Oct 2024 - 11:00	11	11	03 Oct 2024 - 10:00
VSEVM2(VGDEV5)	12	1	04 Oct 2024 - 12:00	11	4	04 Oct 2024 - 11:00
CPC	67	5	03 Oct 2024 - 10:00	66	14	03 Oct 2024 - 09:00

Figure 174. Sample sub-capacity report for z/VSE V4 and V5: SMF / SCRT89 Input Data Statistics and Detail LPAR Data sections

The sections in Figure 174 on page 329 are not affected when both z/VSE V4 and z/VSE V5 run on the same CPC during the same reporting period because these sections do not report any product-specific information.

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR VSEPROD4	LPAR VSEPROD5	LPAR VSEVM2(VGDEV4)	LPAR VSEVM2(VGDEV5)
z/VSE	(All)	67	03 Oct 2024 - 10:00	26	24	8	9
z/VSE V5 (Central Functions V9)	5686-CF9	36	05 Oct 2024 - 11:00	0	29	0	7
z/VSE V4 (Central Functions V8)	5686-CF8	38	04 Oct 2024 - 14:00	28	0	10	0
ACF/VTAM V4 VSE/ESA	5686-065	67	03 Oct 2024 - 10:00	26	24	8	9
CICS TS for VSE/ESA	5648-054	38	04 Oct 2024 - 14:00	28	0	10	0
DB2 Server for VSE & VM	5697-F42	67	03 Oct 2024 - 10:00	26	24	8	9
DITTO/ESA for VSE	5648-099	49	03 Oct 2024 - 09:00	25	24	0	0
DL/1 DOS/V5	5746-XX1	49	03 Oct 2024 - 09:00	25	24	0	0
High Level Assembler VSE & VM	5696-234	67	03 Oct 2024 - 10:00	26	24	8	9
IBM C for VSE/ESA V1	5686-A01	67	03 Oct 2024 - 10:00	26	24	8	9
IBM COBOL VSE/ESA	5686-068	18	05 Oct 2024 - 14:00	0	0	11	9
IBM DFSort/VSE V3	5746-SM3	38	04 Oct 2024 - 14:00	28	0	10	0
IBM PL/1 for VSE/ESA	5686-069	67	03 Oct 2024 - 10:00	26	24	8	9
MQSeries VSE/ESA	5686-A06	38	04 Oct 2024 - 14:00	28	0	10	0
TCP/IP for VSE	5686-A04	67	03 Oct 2024 - 10:00	26	24	8	9

Figure 175. Sample sub-capacity report for z/VSE V4 and V5: Product Max Contributors section

As shown in Figure 175 on page 329, the Product Max Contributors section will also report the z/VSE entry whenever both z/VSE V4 and z/VSE V5 run on the same CPC anytime during the reporting period. The z/VSE entry reports the MSU contribution from z/VSE V4 and z/VSE V5 for the earliest concurrent peak value across both versions of z/VSE.

You can find more information about the fields in this section and other sections of the sub-capacity report in [Chapter 5, “Sub-capacity reporting for z/OS native systems,”](#) on page 103.

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                                Product ID    VSEPROD4    VSEPROD5    VSEVM2(VGDEV4)  VSEVM2(VGDEV5)
z/VSE                                       (All)         91.4%       92.7%       92.3%           91.5%
z/VSE V5 (Central Functions V9)           5686-CF9      91.4%       92.7%       92.3%           91.5%
z/VSE V4 (Central Functions V8)           5686-CF8      91.4%       92.7%       92.3%           91.5%
ACF/VTAM V4 VSE/ESA                       5686-065      91.4%       92.7%       92.3%           91.5%
CICS TS for VSE/ESA                       5648-054      91.4%       92.7%       92.3%           91.5%

==W5=====
This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.

```

Figure 176. Sample sub-capacity report for z/VSE V4 and V5: Product Grid Snapshot section

As shown in [Figure 176 on page 330](#), the Product Grid Snapshot section will also report the z/VSE entry whenever both z/VSE V4 and z/VSE V5 run on the same CPC anytime during the reporting period.

You can find more information about the fields in this section and other sections of the sub-capacity report in [Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103](#).

Sample sub-capacity report for z/VSE V6, z/VSE V5, and z/VSE V4

The following sample sub-capacity report is for a CPC with two LPARs named VSEPROD1 and VSEVM1:

- VSEPROD1 runs a native z/VSE V5 (Central Functions V9) system, PRD1 (CMT ID = PRD1).
- VSEVM1 runs a z/VM system that supports two z/VSE guest systems:
 - A z/VSE V6 system, TST1 (CMT ID = TST1), that runs under a z/VM user ID of VMTST1
 - A z/VSE V4 (Central Functions V8) system, TST2 (CMT ID = TST2), that runs under a z/VM user ID of VMTST2

[Table 51 on page 330](#) shows the sample configuration.

CPC: 2817-722 (serial = 02-34567)		
VSEPROD1	VSEVM1	
z/VSE V5 (Central Functions V9) (CMT ID = PRD1)	z/VSE V6 (CMT ID = TST1) (VM user ID = VMTST1)	z/VSE V4 (Central Functions V8) (CMT ID = TST2) (VM user ID = VMTST2)
ACF/VTAM CICS TS Db2 Server DITTO/ESA	Db2 Server DITTO/ESA	Db2 Server DITTO/ESA

Except for the differences described in [“Sample sub-capacity report for z/VSE V4 only” on page 323](#) and the differences described in this sample report, the sub-capacity reports have the same format and same data fields as those for z/OS systems. You can find information about these fields in [Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103](#).

[Figure 177 on page 331](#) through [Figure 180 on page 332](#) show the sample sub-capacity report based on the configuration shown in [Table 51 on page 330](#).

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 26 Oct 2024 - 14:05
 Name of Person Submitting Report: John Customer
 E-mail Address of Report Submitter: customer@abc.com
 Phone Number of Report Submitter: 444-999-9999

Customer Name ABC Corporation
 Customer Number 8971234567
 Machine Serial Number 02-34567
 Machine Type and Model 2817-722
 Machine Rated Capacity (MSUs) 2224
 Purchase Order Number (optional)
 Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

==C5=====

TOOL INFORMATION

Tool Release 30.1.0
 Reporting Period 2 May, 2024 - 1 Jun, 2024 inclusive (31 days)

Figure 177. Sample sub-capacity report for z/VSE V6, V5, and V4: Customer information and Tool Information sections

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/VSE	(A11)	476		
z/VSE V6	5686-VS6	201		
z/VSE V5 (Central Functions V9)	5686-CF9	101		
z/VSE V4 (Central Functions V8)	5686-CF8	201		
ACF/VTAM V4 VSE/ESA	5686-065	101		
CICS TS for VSE/ESA	5648-054	101		
DB2 Server for VSE & VM	5697-F42	476	(optional)	(conditional)
DITTO/ESA for VSE	5648-099	476	(optional)	(conditional)

Figure 178. Sample sub-capacity report for z/VSE V6, V5, and V4: Product Summary Information section

Since z/VSE V6 (5686-VS6), z/VSE V5 (5686-CF9), and z/VSE V4 (5686-CF8) are reported in the Product Summary Information section, SCRT also reports an additional z/VSE (A11) entry which represents the peak concurrent sub-capacity value for z/VSE V6, z/VSE V5, and z/VSE V4, as if they were a single product.

If the z/VSE V4 and V5 products are under Migration Pricing Option (MPO) on this CPC, IBM waives the charges for the z/VSE V4 and V5 products and charges for the z/VSE V6 product based on the sub-capacity value reported for the z/VSE (A11) entry, which is 476 MSUs in this example. If the MPO period has expired, IBM charges for the z/VSE V4 or V5 products based on their reported sub-capacity values.

Note that when multiple versions of z/VSE are in use, any combination of z/VSE V6, z/VSE V5, or z/VSE V4 on the same CPC results in the z/VSE (A11) line being reported. The additional line represents the peak concurrent sub-capacity value for the listed z/VSE operating systems.

You can find more information about the fields in this section and other sections of the sub-capacity report in [Chapter 5, “Sub-capacity reporting for z/OS native systems,”](#) on page 103.

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

```

	SYSID	Input Data Start	Input Data End
VSEPROD1	PROD1	20 May 2024 - 16:32	25 May 2024 - 15:00
VSEVM1(VMTST1)	TST1	20 May 2024 - 16:30	25 May 2024 - 15:00
VSEVM1(VMTST2)	TST2	20 May 2024 - 16:30	25 May 2024 - 15:00
CPC		20 May 2024 - 16:30	25 May 2024 - 15:00

```

==N5=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
VSEPROD1	101	14	20 May 2024 - 21:00	99	1	20 May 2024 - 20:00
VSEVM1(VMTST1)	201	11	22 May 2024 - 13:00	199	2	22 May 2024 - 12:00
VSEVM1(VMTST2)	201	11	22 May 2024 - 13:00	199	2	22 May 2024 - 12:00
CPC	476	1	22 May 2024 - 12:00	470	1	22 May 2024 - 11:00

Figure 179. Sample sub-capacity report for z/VSE V6, V5, and V4: SMF / SCRT89 Input Data Statistics and Detail LPAR Data sections

The sections in Figure 179 on page 332 are not affected when z/VSE V6, z/VSE V5, and z/VSE V4 run on the same CPC during the same reporting period because these sections do not report any product-specific information.

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR VSEPROD1	LPAR VSEVM1(VMTST1)	LPAR VSEVM1(VMTST2)
z/VSE	(All)	476	22 May 2024 - 12:00	78	199	199
z/VSE V6	5686-VS6	201	22 May 2024 - 13:00	0	201	0
z/VSE V5 (Central Functions V9)	5686-CF9	101	20 May 2024 - 21:00	101	0	0
z/VSE V4 (Central Functions V8)	5686-CF8	201	22 May 2024 - 13:00	0	0	201
ACF/VTAM V4 VSE/ESA	5686-065	101	20 May 2024 - 21:00	101	0	0
CICS TS for VSE/ESA	5648-054	101	20 May 2024 - 21:00	101	0	0
DB2 Server for VSE & VM	5697-F42	476	22 May 2024 - 12:00	78	199	199
DITTO/ESA for VSE	5648-099	476	22 May 2024 - 12:00	78	199	199

```

==Q5=====
PRODUCT GRID SNAPSHOT

```

Product Name	Product ID	VSEPROD1	VSEVM1(VMTST1)	VSEVM1(VMTST2)
z/VSE	(All)	92.7%	15.9%	15.9%
z/VSE V6	5686-VS6		15.9%	
VSE Central Functions V9	5686-CF9	92.7%		
VSE Central Functions V8	5686-CF8			15.9%
ACF/VTAM V4 VSE/ESA	5686-065	92.7%		
CICS TS for VSE/ESA	5648-054	92.7%		

```

==W5=====
This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such
Customer is solely responsible for the completeness and accuracy of information and data used to create this report.
Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report.
Any questions regarding the contents of this report should be directed to the Customer.

```

Figure 180. Sample sub-capacity report for z/VSE V6, V5, and V4: Product Max Contributors section and Product Grid Snapshot section

As shown in Figure 180 on page 332, the Product Max Contributors and Product Grid Snapshot sections also report the z/VSE (All) entry whenever two or more versions of z/VSE run on the same CPC anytime during the reporting period. The z/VSE (All) entry reports the MSU contribution from z/VSE V6, z/VSE V5, and z/VSE V4 for the earliest concurrent peak value across all versions of z/VSE.

You can find more information about the fields in this section and other sections of the sub-capacity report in Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103.

Sample sub-capacity report for z/OS and z/VSE V4

The following sample sub-capacity report is for a CPC with four LPARS named VSEPROD, VSEVM, ZOSPROD, and LINUX:

- VSEPROD runs a native z/VSE V4 (Central Functions V8) system, VPRD (CMT ID = VPRD).
- VSEVM runs a z/VM system which supports two z/VSE V4 (Central Functions V8) guest systems, DEV1 and DEV2 (CMT IDs = DEV1 and DEV2).
- ZOSPROD runs a native z/OS system, SYSP (SID = SYSP).
- LINUX runs a z/VM system that supports a number of Linux guest systems and uses approximately 45 percent of the capacity of the CPC.

Table 52 on page 333 shows this configuration.

Table 52. Sample configuration with z/OS and z/VSE systems				
CPC: 2097-405 (serial = 02-123AB)				
VSEPROD	VSEVM (z/VM)		ZOSPROD	LINUX (z/VM)
z/VSE V4 (Central Func V8) (CMT ID = VPRD)	z/VSE V4 (Central Func V8) (CMT ID = DEV1) (VM user ID = VGDEV1)	z/VSE V4 (Central Func V8) (CMT ID = DEV2) (VM user ID = VGDEV2)	z/OS V1 (SYSID = SYSP)	Linux (Uses about 45 percent of the CPC)
ACF/VTAM CICS TS Db2 Server	ACF/VTAM CICS TS Db2 Server	ACF/VTAM	CICS TS V3 DB2 UDB V8 IMS V8 MQSeries V6 File Manager V6	
MQSeries DITTO/ESA DL/1 HL Assembler IBM C	HL Assembler IBM C IBM COBOL IBM DFSort IBM PL/1 TCP/IP	MQSeries HL Assembler IBM C IBM COBOL IBM PL/1 TCP/IP	Cobol V3	
IBM DFSort IBM PL/1 TCP/IP			WebSphere App. Server V6 WebSphere MQ Workflow V3	

The sub-capacity reports for z/VSE systems have the same format and same fields as the sub-capacity reports for z/OS systems. You can find the information about these fields in Chapter 5, “Sub-capacity reporting for z/OS native systems,” on page 103. For examples of the differences between the sub-capacity reports for z/VSE systems and z/OS systems, see the description in “Sample sub-capacity reports for z/VSE” on page 322.

Figure 181 on page 334 through Figure 187 on page 336 show a sample sub-capacity report based on the configuration given in Table 52 on page 333.

Note that, although the z/VSE hardware and software configuration for the following sample sub-capacity report represents a mixed z/VSE and z/OS environment, the z/VSE portion of the configuration is identical to the z/VSE configuration shown in “Sample sub-capacity report for z/VSE V4 only” on page 323. The net result is that the SCRT reporting for the z/VSE systems is identical to the sample sub-capacity report shown in “Sample sub-capacity report for z/VSE V4 only” on page 323. See the descriptions of the corresponding report sections in “Sample sub-capacity report for z/VSE V4 only” on page 323 for explanations of the impacts of z/VSE systems whose capacities are less than one MSU.

==B5===== SCRT SUB-CAPACITY REPORT - IBM Corp =====

Run Date/Time 03 Nov 2024 - 09:17
 Name of Person Submitting Report: John Customer
 E-mail Address of Report Submitter: customer@abc.com
 Phone Number of Report Submitter: 444-999-9999

Customer Name ABC Corporation
 Customer Number 8971234567
 Machine Serial Number 02-123AB
 Machine Type and Model 2097-405
 Machine Rated Capacity (MSUs) 118
 Purchase Order Number (optional)
 Customer Comments (255 chars max) (optional)

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

Figure 181. Sample sub-capacity report for z/OS and z/VSE: Customer information section

==C5=====

TOOL INFORMATION

Tool Release 30.1.0
 Reporting Period 2 Oct, 2024 - 1 Nov, 2024 inclusive (31 days)

Figure 182. Sample sub-capacity report for z/OS and z/VSE: Tool Information section

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
z/OS V1	5694-A01	31		
DB2 UDB for z/OS V8	5625-DB2	31		
CICS TS for z/OS V3	5655-M15	31		
MQSeries for z/OS V6	5655-L82	31		
IMS V8	5655-C56	31		
IBM Enterprise Cobol for z/OS and OS/390 V3	5655-G53	31	(optional)	(conditional)
z/VSE V4 (Central Functions V8)	5686-CF8	37		
ACF/VTAM V4 VSE/ESA	5686-065	37		
CICS TS for VSE/ESA	5648-054	37		
DB2 Server for VSE & VM	5697-F42	37	(optional)	(conditional)
DITTO/ESA for VSE	5648-099	37	(optional)	(conditional)
DL/1 DOS/VS	5746-XX1	37	(optional)	(conditional)
High Level Assembler VSE & VM	5696-234	37	(optional)	(conditional)
IBM C for VSE/ESA V1	5686-A01	37	(optional)	(conditional)
IBM COBOL VSE/ESA	5686-068	1	(optional)	(conditional)
IBM DFSort/VSE V3	5746-SM3	37	(optional)	(conditional)
IBM PL/1 for VSE/ESA	5686-069	37	(optional)	(conditional)
MQSeries VSE/ESA	5686-A06	37	(optional)	(conditional)
TCP/IP for VSE	5686-A04	37	(optional)	(conditional)
IPLA Product Name	IPLA Product ID	Tool MSUs	Customer MSUs	Customer Comments (255 chars max)
File Manager for z/OS V6	5655-P17	31	(optional)	(conditional)
IBM WebSphere Application Server for z/OS V6	5655-N01	31		
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	31		

Figure 183. Sample sub-capacity report for z/OS and z/VSE: Product Summary Information section

```

==L5=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS

          SYSID      Input Data Start      Input Data End
VSEPROD      VPRD      01 Oct 2024 - 00:00  02 Nov 2024 - 22:00
VSEVM(VGDEV1) DEV1      01 Oct 2024 - 06:05  02 Nov 2024 - 23:45
VSEVM(VGDEV2) DEV2      01 Oct 2024 - 00:00  30 Oct 2024 - 22:30
ZOSPROD      SYS1      01 Oct 2024 - 00:00  02 Nov 2024 - 23:45
LINUX        -         -         -
CPC          -         01 Oct 2024 - 00:00  02 Nov 2024 - 23:45

```

Figure 184. Sample sub-capacity report for z/OS and z/VSE: SMF / SCRT89 Input Data Statistics section

```

==N5=====
DETAIL LPAR DATA SECTION

          Highest Hour Count Date/Time          2nd Highest Hour Count Date/Time
VSEPROD      36          1 03 Oct 2024 - 07:00      35          5 03 Oct 2024 - 06:00
VSEVM(VGDEV1) 1          1 02 Oct 2024 - 11:00      (v)         688 03 Oct 2024 - 10:00
VSEVM(VGDEV2) (v)        680 02 Oct 2024 - 00:00
ZOSPROD      33          3 05 Oct 2024 - 09:00      32          9 05 Oct 2024 - 08:00
CPC          70          5 03 Oct 2024 - 12:00      69          14 03 Oct 2024 - 11:00

Footnotes:
(v) Value is less than 1 MSU

```

Figure 185. Sample sub-capacity report for z/OS and z/VSE: Detail LPAR Data Section

```

==P5=====
PRODUCT MAX CONTRIBUTORS

Product Name          Product ID Highest Date/Time          LPAR LPAR LPAR LPAR
                        VSEPROD VSEVM(VGDEV1) VSEVM(VGDEV2) ZOSPROD
z/OS V1                5694-A01 31 05 Oct 2024 - 09:00      0      0      0      31
DB2 UDB for z/OS V8     5625-DB2 31 05 Oct 2024 - 09:00      0      0      0      31
CICS TS for z/OS V3     5655-M15 31 05 Oct 2024 - 09:00      0      0      0      31
MQSeries for z/OS V6    5655-L82 31 05 Oct 2024 - 09:00      0      0      0      31
IMS V8                  5655-C56 31 05 Oct 2024 - 09:00      0      0      0      31
IBM Enterprise Cobol for z/OS and OS/390 V3 5655-G53 29 03 Oct 2024 - 12:00      0      0      0      31
File Manager for z/OS V6 5655-P17 31 05 Oct 2024 - 09:00      0      0      0      31
IBM WebSphere Application Server for z/OS V6 5655-N01 31 05 Oct 2024 - 09:00      0      0      0      31
IBM WebSphere MQ Workflow for z/OS V3      5655-BPM 31 05 Oct 2024 - 09:00      0      0      0      31

z/VSE V4 (Central Functions V8) 5686-CF8 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
ACF/VTAM V4 VSE/ESA 5686-065 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
CICS TS for VSE/ESA 5648-054 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
DB2 Server for VSE & VM 5697-F42 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
DITTO/ESA for VSE 5648-099 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
DL/I DOS/VS 5746-XX1 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
High Level Assembler VSE & VM 5696-234 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
IBM C for VSE/ESA V1 5686-A01 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
IBM COBOL VSE/ESA 5686-068 1 03 Oct 2024 - 11:00      0      1      (v)      0
IBM DFSort/VSE V3 5746-SM3 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
IBM PL/1 for VSE/ESA 5686-069 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
MQSeries VSE/ESA 5686-A06 37 03 Oct 2024 - 07:00      36      (v)      (v)      0
TCP/IP for VSE 5686-A04 37 03 Oct 2024 - 07:00      36      (v)      (v)      0

Footnotes:
(v) Value is less than 1 MSU

```

Figure 186. Sample sub-capacity report for z/OS and z/VSE: Product Max Contributors section

```
==Q5=====
PRODUCT GRID SNAPSHOT
```

Product Name	Product ID	VSEPROD	VSEVM(VGDEV1)	VSEVM(VGDEV2)	ZOSPROD
z/OS V1	5694-A01				97.2%
DB2 UDB for z/OS V8	5625-DB2				97.2%
CICS TS for z/OS V3	5655-M15				97.2%
MQSeries for z/OS V6	5655-L82				97.2%
IMS V8	5655-C56				97.2%
IBM WebSphere Application Server for z/OS V6	5655-N01				97.2%
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM				97.2%
z/VSE V4 (Central Functions V8)	5686-CF8	91.4%	92.7%	91.5%	
ACF/VTAM V4 VSE/ESA	5686-065	91.4%	92.7%	91.5%	
CICS TS for VSE/ESA	5648-054	91.4%	92.7%		

```
==W5=====
```

This report is prepared by the IBM System z customer identified above ("Customer") or its authorized designee, and such Customer is solely responsible for the completeness and accuracy of information and data used to create this report. Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report. Any questions regarding the contents of this report should be directed to the Customer.

Figure 187. Sample sub-capacity report for z/OS and z/VSE: Product Grid Snapshot section

Sub-capacity reporting for z/VSE systems in a country multiplex report

IBM Country Multiplex Pricing (CMP) is a sub-capacity pricing offering for z/OS and z/TPF systems.

For eligible environments, a Country Multiplex Pricing report (or simply a *multiplex report*) treats all CPCs within a single country as one machine for the purpose of sub-capacity reporting. For details about CMP, see Chapter 16, “Sub-capacity reporting for Country Multiplex Pricing (CMP),” on page 343.

SCRT supports standard sub-capacity reporting of z/VSE systems in a multiplex report. For z/OS or z/TPF customers that have a Country Multiplex Pricing agreement, this removes the need to generate a separate sub-capacity report for z/VSE systems.

Note the following points about z/VSE systems in a multiplex report:

- The z/VSE products do not appear in the Multiplex Summary section.
- The z/VSE product peaks that SCRT reports in the Product Summary section for each CPC are the traditional sub-capacity peaks that occurred on that machine.
- The z/OS and z/TPF peaks that SCRT reports in the Product Summary section for each CPC are the sub-capacity values for that machine during the multiplex peak hour.

SCRT supports z/VSE systems and products in a multiplex report for the following configurations:

- A CPC with only z/VSE systems as part of a country multiplex
- A CPC with a mix of z/VSE and non-z/VSE systems as part of a country multiplex

Sample Country Multiplex Pricing report for z/OS and z/VSE systems

The sub-capacity report for CMP (that is, the multiplex report) for z/OS and z/VSE systems has the same basic format and report sections as a multiplex report without a z/VSE system.

When z/VSE data is present in the input to SCRT, the following differences appear in the multiplex report:

- An indicator in the machine input parameters section of the multiplex summary to highlight the cases when z/VSE is present on a CPC
- The products that are reported
- Other differences as described in “Sample sub-capacity reports for z/VSE” on page 322

The following multiplex report is for three CPCs with the following configurations:

- CPC 1 is a 2965-Z06 (serial = 00-0FF23) with three LPARS:
 - LPARs VSE1 and VSE2 run native z/VSE V5 systems.
 - LPAR VSE3 runs a native z/VSE V6 system.
- CPC 2 is a 2827-710 (serial = 02-0A123) with two LPARs:
 - LPARs LPARa and LPARB run native z/OS V2 systems.
- CPC 3 is a 2827-708 (serial = 02-0B123) with two LPARS:
 - LPARs LPAR1 and LPAR2 run native z/OS V2 systems.

Multiplex summary section and machine summary section

The multiplex summary section for z/VSE systems have the same format and the same fields as the multiplex report for z/OS systems. For details about these fields, see [“Multiplex summary section” on page 347](#).

```

==B5===== SCRT MULTIPLEX REPORT - IBM Corp =====
SCRT Tool Release                23.10.00
Name of Person Submitting Report: Joe Customer
E-Mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter: 845-555-1212

Customer Name                    ABC Corp
Run Date/Time                    3 Jun 2024 - 14:13
Report Period                    2 May, 2024 - 1 Jun, 2024 inclusive (31 days)
Number of processors in Multiplex 3
  
```

Figure 188. Sample multiplex report for Country Multiplex Pricing for z/VSE: Multiplex summary section

A major difference between a multiplex report for z/OS systems and a multiplex report that includes z/VSE data is the addition of a **z/VSE Data** row in the machine summary section, as shown in [Figure 189 on page 337](#). A value of Y under a CPC column indicates that the CPC has at least one LPAR with z/VSE data.

Machine identifier	M1C1	M2C2	M3C2
Customer number	2000000000	8971234567	8971234567
Machine Serial Number	00-0FF23	02-0A123	02-0B123
Machine Type and Model	2965-Z06	2827-710	2827-708
Machine Rated Capacity (MSUs)	884	1473	1224
Machine Model Changed	N	N	N
Exclude Data	Y	Y	Y
Missing LPAR Data	Y	N	N
Missing CPC Data	Y	N	N
z/VSE Data	Y	N	N

MLC Product Name	Number	MSU	Time		
z/OS V2	5650-Z0S	970	04 May 2024 - 19:00 UTC	598	372
DB2 10 for z/OS	5605-DB2	970	04 May 2024 - 19:00 UTC	598	372
CICS TS for z/OS V4	5655-S97	718	01 Jun 2024 - 05:00 UTC	718	
IBM MQ for z/OS V8	5655-W97	718	01 Jun 2024 - 05:00 UTC	718	
IMS V13	5635-A04	718	01 Jun 2024 - 05:00 UTC	718	0

IPLA Product Name	Number	MSU	Time		
Debug Tool Utilities Adv Function	5655-S16	464	01 Jun 2024 - 04:00	464	

Figure 189. Sample multiplex report for Country Multiplex Pricing for z/VSE: Machine summary section

In the example configuration, CPC 1 (machine identifier M1C1) has only z/VSE systems. There are no z/OS systems, so M1C1 has no contribution to the multiplex peaks for any of the listed MLC or IPLA products.

Because z/VSE products are still reported under standard sub-capacity rules when CMP is requested, no z/VSE products appear in the machine summary section.

Machine summary section for CPC M1C1

A Machine Summary section is generated for each CPC found in the multiplex. The header for each Machine summary section contains the machine identifier within the header itself. For instance, the Machine Summary section for CPC M1C1 contains the M1C1 value within the header, as shown in [Figure 190 on page 338](#).

Machine Information section for CPC M1C1

Figure 190 on [page 338](#) shows the Machine Information section for z/VSE system M1C1, which contains z/VSE LPARs. There is no difference between a machine information section for a CPC with z/VSE systems and one without z/VSE systems.

```
==M1C1=====
```

Machine identifier	M1C1
Customer number	US-S000000000
Machine serial number	00-0FF23
Machine Type and Model	2965-Z06
Machine Rated Capacity (MSUs)	884

Figure 190. Sample multiplex report for Country Multiplex Pricing for z/VSE: Machine Summary section for CPC M1C1 - Machine information

The Machine Information section contains the machine identifier, customer number, serial number, machine type and model, and rated capacity. For more information about these fields, see [“Machine summary section” on page 348](#).

Product Summary Information section for CPC M1C1

Figure 191 on [page 338](#) shows the Product Summary Information section for system M1C1 in the sample configuration.

```
==E5=====
```

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs
z/VSE	(All)	453
z/VSE V6	5686-VS6	370
z/VSE V5 (Central Functions V9)	5686-CF9	302
ACF/VTAM V4 VSE/ESA	5686-065	453
CICS TS for z/VSE	(All)	453
CICS TS for z/VSE V2	5655-VSE	370
CICS TS for VSE/ESA	5648-054	302
TCP/IP for VSE/ESA	5686-A04	371

Figure 191. Sample multiplex report for Country Multiplex Pricing for z/VSE: Machine Summary section for CPC M1C1 - Product Summary Information

The Product Summary Information section for z/VSE products in a multiplex report differs from that for z/OS or z/TPF products. For z/OS and z/TPF products in a multiplex, the **Tool MSUs** values correspond to the contribution of the CPC to the multiplex peak value, instead of any individual peak that may have occurred on the CPC. However, for the z/VSE products shown in [Figure 191 on page 338](#), the **Tool MSUs**

values correspond to the products' individual peaks on machine M1C1; they are the same values as in a standard sub-capacity report.

Missing LPAR Data Details section and Missing CPC Data Details section for CPC M1C1

The Missing LPAR Data Details section is generated to notify you that SCRT detected missing LPAR data for the CPC.

The Missing CPC Data Details section is generated to notify you that SCRT detected missing CPC data in the reporting period.

When z/VSE systems are processed as part of a multiplex report, the CMP data reporting rules must be met.

Figure 192 on page 339 shows an example of the these sections for machine M1C1.

```

==H4=====
MISSING LPAR DATA DETAILS

Active LPARs with no SMF/SCRT89 records      Detected              Resolved (missing hours)      Est Peak MSU      Justification for low data
S97BLP04                                     15 May 2024 - 16:00          16 May 2024 - 14:00 (23 hours)      10                collection (255 chars max)
Total Missing Hours                          23
OS Product Name      OS Product ID      Multiplex Peak MSU      Est Peak MSU      Date/Time      VSE1      VSE2      VSE3      S97BLP04
z/VSE                (All)              900                    460               15 May 2024 - 16:00      150       150       150       10

==H6=====
MISSING CPC DATA DETAILS

CPC with no SMF/SCRT89 data from any LPAR      Detected              Resolved (missing hours)      Justification for low data
M1C1                                           02 May 2024 - 00:00 UTC      15 May 2024 - 15:00 UTC (326 hours)      collection (255 chars max)
M1C1                                           16 May 2024 - 15:00 UTC      16 May 2024 - 23:00 UTC (7 hours)          (required)
M1C1                                           17 May 2024 - 00:00 UTC      02 Jun 2024 - 00:00 UTC (383 hours)          (required)
Total Missing Hours                          716

```

Figure 192. Sample multiplex report for Country Multiplex Pricing for z/VSE: Machine Summary section - Missing LPAR Data Details section and Missing CPC Data Details section

For details about the CMP data reporting requirements, see “Missing LPAR Data Details section” on page 353 and “Missing CPC Data Details section for CPC M1C1” on page 354.

Machine summary section for CPC M2C2 and CPC M3C2

CPC M2C2 and M3C2 have no z/VSE LPARs or products. For these CPCs, the presence of z/VSE products does not impact the machine summary section.

Customer verification section for CPC M1C1

The customer verification section of the multiplex report contains information that supports and expands on the information in the Product Summary Information section.

Other sections that are not displayed for the sample configuration may be displayed under different conditions. For information about the possible sections, see “Customer verification section ” on page 118.

The detailed data sections for customer verification follow the Detail Data Sections header in the multiplex report, as shown in Figure 193 on page 340.

```
==L7=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
```

Figure 193. Sample multiplex report for Country Multiplex Pricing for z/VSE: Customer verification section - Detail Data Sections header

Machine information section for CPC M1C1

Figure 194 on page 340 shows the Machine Information section for CPC 2965 Model Z06 (M1C1) in the sample configuration.

```
==M7=====

Machine identifier          M1C1
Customer number            US-S000000000
Machine serial number      00-0FF23
Machine Type and Model     2965-Z06
Machine Rated Capacity (MSUs) 884
```

Figure 194. Sample multiplex report for Country Multiplex Pricing for z/VSE: Machine Summary section for CPC M1C1 - Machine information

The Machine Information section contains the machine identifier, customer number, serial number, machine type and model, and rated capacity. For more information about these fields, see [“Machine summary section”](#) on page 348.

Detail LPAR Data section for CPC M1C1

There are no differences in the format of the Detail LPAR Data section in a standard sub-capacity report and a multiplex report. There is, however, a key difference in the values that are represented for z/VSE systems and non-z/VSE systems:

- For z/OS and z/TPF systems, the values and times shown in a multiplex report correspond to the contribution of the LPAR to the multiplex peak concurrent MSU value.
- For z/VSE, the values and times shown represent the standard sub-capacity peak values for the CPC. This is equivalent to the z/VSE system representation in a standard sub-capacity report.

Figure 195 on page 340 shows an example of the Detail LPAR Data Section for CPC M1C1 for the sample configuration.

```
==N5=====
DETAIL LPAR DATA SECTION
```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
VSE1	151	1	16 May 2024 - 13:00 UTC	132	2	16 May 2024 - 12:00 UTC
VSE2	185	2	15 May 2024 - 21:00 UTC	151	2	16 May 2024 - 12:00 UTC
VSE3	370	2	16 May 2024 - 02:00 UTC	324	1	16 May 2024 - 01:00 UTC
CPC	453	1	16 May 2024 - 13:00 UTC	434	1	16 May 2024 - 12:00 UTC

Figure 195. Sample multiplex report for Country Multiplex Pricing for z/VSE: Detail LPAR Data section for CPC M1C1

For more information about the Detail LPAR Data section, see [“Detail LPAR Data section”](#) on page 119.

Product Max Contributors section for CPC M1C1

There are no differences in the format of the Product Max Contributors section in a standard sub-capacity report and a multiplex report. There is, however, a key difference in the values that are represented. The values and time shown in a multiplex report correspond to the contribution of the LPAR to the multiplex peak concurrent MSU value.

There is a key difference in the values represented for z/VSE products in the Product Max Contributors section and non-z/VSE systems:

- For z/OS and z/TPF products, the values and times shown in a multiplex report correspond to the contribution of the LPAR to the multiplex peak concurrent MSU value.
- For z/VSE products, the values and times shown represent the standard sub-capacity peak values for the CPC. This is equivalent to the z/VSE product representation in a standard sub-capacity report.

Figure 196 on page 341 shows the Product Max Contributors section for CPC M1C1 in the sample configuration.

```

==P5=====
PRODUCT MAX CONTRIBUTORS

Product Name                Product ID      Highest  Date/Time      LPAR   LPAR   LPAR
z/VSE                       (All)          453     16 May 2024 - 13:00 UTC  151    151    151
z/VSE V6                    5686-VS6       370     16 May 2024 - 02:00 UTC   0       0     370
z/VSE V5 (Central Functions V9) 5686-CF9       302     16 May 2024 - 13:00 UTC  151    151    0
ACF/VTAM V4 VSE/ESA         5686-065       453     16 May 2024 - 13:00 UTC  151    151    151
CICS TS for z/VSE           (All)          453     16 May 2024 - 13:00 UTC  151    151    151
CICS TS for z/VSE V2        5655-VSE       370     16 May 2024 - 02:00 UTC   0       0     370
CICS TS for VSE/ESA         5648-054       302     16 May 2024 - 13:00 UTC  151    151    0
TCP/IP for VSE/ESA          5686-A04       371     16 May 2024 - 02:00 UTC   1       0     370

```

Figure 196. Sample multiplex report for Country Multiplex Pricing: Product Max Contributors section for CPC M1C1

For more information about the Product Max Contributors section, see [“Product Max Contributors section”](#) on page 121.

Product Grid Snapshot section for CPC M1C1

There are no differences in the format or values of the Product Grid Snapshot section in a standard sub-capacity report and a multiplex report.

Figure 197 on page 341 shows the Product Grid Snapshot section for CPC M1C1 in the sample configuration.

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                Product ID      VSE1    VSE2    VSE3
z/VSE                       (All)          2.90%   2.80%   2.80%
z/VSE V6                    5686-VS6       2.90%   2.80%   2.80%
z/VSE V5 (Central Functions V9) 5686-CF9       2.90%   2.80%   2.80%
ACF/VTAM V4 VSE/ESA         5686-065       2.90%   2.80%   2.80%
CICS TS for z/VSE           (All)          2.90%   2.80%   2.80%
CICS TS for z/VSE V2        5655-VSE       2.90%   2.80%   2.80%
CICS TS for VSE/ESA         5648-054       2.90%   2.80%   2.80%

```

Figure 197. Sample multiplex report for Country Multiplex Pricing for z/VSE: Product Grid Snapshot section for CPC M1C1

For more information about the Product Grid Snapshot section, see [“Product Grid Snapshot section”](#) on page 123.

Customer verification sections for CPCs M2C2 and M3C2

CPCs M2C2 and M3C2 have no z/VSE LPARs or products. For these CPCs, the presence of z/VSE products does not impact the customer verification section.

Chapter 16. Sub-capacity reporting for Country Multiplex Pricing (CMP)

IBM Country Multiplex Pricing (CMP) is a sub-capacity pricing offering that allows you to use your IBM Z capacity within a given country without the constraints of sysplex aggregation rules and many of the limitations of previous reporting methodologies. CMP is a new way of measuring and pricing sub-capacity MSUs that is designed to allow more flexible deployment and movement of workloads across all CPCs within a country.

CMP also simplifies the process for software migrations with a shift from Single Version Charging (SVC) and limited migration windows to Multi-Version Measurement (MVM) reporting, where there are no time constraints. In addition, the cost of growth per MSU for a given sub-capacity program is based on the total reported multiplex MSUs, which means that any given sub-capacity program has a single price point for growth in the country regardless of where the workload runs.

The report that is generated when you request CMP is referred to as a *multiplex report* to differentiate it from a standard sub-capacity report where the CMP algorithm has not been requested.

Country Multiplex Pricing is effected for SMF data collected on 2 October 2015 and later, provided that all of the required CMP conditions have been met.

Overview of Country Multiplex Pricing

Country Multiplex Pricing simplifies many aspects of earlier sub-capacity pricing structures with the objective of creating a flexible, country-wide IBM Z platform where you can configure and balance workloads more easily with fewer constraints.

A *multiplex* is the collection of all eligible IBM Z CPCs or sysplexes or both within a single country, measured as one machine for purposes of sub-capacity reporting. You can request that a multiplex report be generated. The multiplex report provides sub-capacity program utilization peaks across all CPCs in the multiplex simultaneously, rather than separate peaks by CPC.

When eligibility requirements are met, as described in US Software Announcement letter 215-230, you can transition to this offering. While all existing pricing metrics for z/OS and z/TPF workloads are eligible to transition to CMP, a new pricing metric called Country Multiplex License Charges (CMLC) will be used for all multiplexes. In addition, a new Multiplex zNALC (MzNALC) metric replaces the original zNALC metric for use within a multiplex. All sub-capacity eligible programs within a multiplex must be priced according to CMLC or zNALC sub-capacity terms. CPCs that are full capacity (such as PSLC) prior to CMP are eligible for inclusion in a multiplex but must transition to CMP sub-capacity terms.

Note that if you are still using the original NALC metric, you must adopt zNALC terms and conditions prior to the transition to CMP. Only the CMLC and MzNALC metrics are supported on multiplex CPCs for sub-capacity eligible z/OS and z/TPF workloads.

The topics in this chapter describe the requirements, initial transition process, and changes to prior rules and guidelines, as well as how MSU peak calculations and monthly billing work under CMP.

Qualifications for Country Multiplex Pricing

The following conditions must be met before IBM will accept and process a multiplex report for IBM Country Multiplex Pricing.

1. CPCs that are part of a multiplex must be licensed under CMLC or MzNALC sub-capacity terms and conditions. Because these terms and conditions require that all products be licensed to IBM zEnterprise System servers (IBM z13, z13s, zEC12, zBC12, z196, or z114 servers, or later), SMF and SCRT89 data must be collected from z/OS or z/TPF systems running on zEnterprise CPCs.
2. When a multiplex report is requested, the input for SCRT is one or more sequential data sets containing SMF type 70 subtype 1 and SMF type 89 subtypes 1 and 2 records, together with the

SCRT89 records, for all of the LPARs on all CPCs that are members of the multiplex for the entire reporting period. Failure to include a CPC that is part of the multiplex will cause IBM to reject the affected multiplex report.

3. The SMF data must be collected from z/OS V1 (5694-A01), z/OS V2 (5650-ZOS), or z/TPF (5748-T15) operating systems. Note that other relevant sub-capacity reporting options, such as Mobile Workload Pricing (MWP) or Integrated Workload Pricing (IWP), can be specified, where applicable, as they were when generating sub-capacity reports. In cases where these options are used, any requirements for those options must still be met.
4. For z/OS guest systems running under z/VM, all z/OS guest systems must be at z/OS V1.7 or later. Note that native z/OS systems, including z/OS V1.6, are not affected by the SCRT rules for z/OS V1.6 guests. Products running on native z/OS V1.10 and later systems are eligible for IBM Z Collocated Application Pricing (zCAP) when all of the other zCAP conditions are met, even when z/OS V1.6 or earlier guest systems are also running on the same CPC. The guest restriction only applies to z/OS guest systems and has no effect on native z/OS systems. When using z/OS guest systems, either z/VM V5.4 with APAR VM64798 or z/VM V6.1 with APAR VM64879 are required.
5. Only SMF data collected on 2 October 2015 and later is eligible for Country Multiplex Pricing, as set forth in the Country Multiplex Pricing announcement letter. Although SCRT will generate multiplex reports for SMF data collected earlier than 2 October 2015, the IBM business process will reject those reports.
6. You must configure all CPCs in the multiplex with Global Performance Data Control Authority, as documented in the appropriate *PR/SM Planning Guide* for that CPC.
7. You must generate and submit a monthly multiplex report to IBM during the normal sub-capacity report submission window (business as usual).
8. You must use SCRT V23.10.0 or later to generate the multiplex report.
9. You must use the CMP control statement to request that SCRT generate a multiplex report using the CMP algorithms whenever all of the required CMP conditions are met. The default SCRT processing does not use the CMP algorithms. For details, see [“Country_Multiplex_Pricing control statement” on page 344](#).
10. Prior to submitting a multiplex report, you must meet any eligibility requirements and agree to the relevant Country Multiplex Pricing documents, as defined in IBM US Software Announcement 215-230. Failure to formally agree to these documents prior to submitting a multiplex report will cause the IBM business process to reject that report.

SCRT externals for Country Multiplex Pricing

SCRT provides the following externals for Country Multiplex Pricing.

1. SCRT supports an optional **Country_Multiplex_Pricing** control statement to request that SCRT apply the CMP algorithm to generate a multiplex report. Note that the format of the multiplex report is significantly different from the standard sub-capacity report output.
2. SCRT uses an optional **Generate_Detailed_Data** control statement to track the application of mobile and CMP reductions per hour, per LPAR.

SCRT continues to report externals associated with other requested or relevant reporting options.

Country_Multiplex_Pricing control statement

The default SCRT processing does not support Country Multiplex Pricing. You must specify the **Country_Multiplex_Pricing** SPECIAL control statement to request that SCRT generate a multiplex report.

Note: IBM authorizes the use of the CMP control statement only when you have signed or otherwise agreed to the terms of the attachment for IBM Country Multiplex Pricing. IBM will reject the multiplex report if the CMP contractual terms and conditions are not in place.

For the format of the CMP control statement and other details, see [“SPECIAL DD control statements” on page 346](#).

For details about the format of the multiplex report, see [“Sample sub-capacity report for Country Multiplex Pricing” on page 347](#).

Generate_Detailed_Data control statement

SCRT supports an optional **Generate_Detailed_Data** SPECIAL control statement.

When a CMP multiplex report is requested, this control statement is used to display detailed data about the contribution of each LPAR on each CPC towards the multiplex total. Data is displayed for every hour in the reporting period for which SCRT received SMF data.

When specified, this control statement causes the Detailed Interval Data section (W3) to be generated in the multiplex report.

Guideline: The information in the Detailed Interval Data section of the multiplex report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating multiplex reports for submission to IBM. In the event that IBM requests that you submit a multiplex report that includes section W3, you can use the LMS Web application to submit the report.

Note that when MWP or zCAP data is provided using the INPUTCSV DD statement and a multiplex report has been requested, the **Generate_Detailed_Data** control statement will not generate a Mobile/zCAP Pricing Detailed Data section (V5). Instead, MWP and zCAP application information will be displayed in section W3.

Sample configuration for Country Multiplex Pricing

Table 53 on page 346 illustrates the sample hardware and software configuration for a two-CPC multiplex. While the mix of software products and versions used in this sample configuration might not be realistic, it serves to illustrate some of the key points of the CMP reporting.

The sample configuration consists of a two CPCs, as follows:

- A 2817 Model 760 CPC with serial number 02-12345. The CPC is configured with one LPAR running a native z/OS V2 image. This CPC is known as M1C1.
- A 2964 Model 796 CPC with serial number 02-56789. The CPC is configured with one LPAR running a native z/OS V2 image. This CPC is known as M2C1.

The CPCs are running a combination of the following z/OS products across the respective LPARs defined to them:

- z/OS V2
- Db2 11 for z/OS
- CICS TS for z/OS V4
- IBM MQ for z/OS V8
- IMS V12
- IMS V13
- IBM WebSphere Application Server for z/OS V8
- IBM Integration Bus for z/OS V9
- IBM Multi-site Workload Lifeline V2

Table 53. Sample configuration for Country Multiplex Pricing

Multiplex	
CPC 1 (M1C1): 2817-760 (serial = 02-12345)	CPC 2 (M2C1): 2964-796 (serial = 02-56789)
LPAR name = M1_LPAR	LPAR name = M2_LPAR
z/OS V2 Db2 11 for z/OS CICS TS for z/OS V4 IBM MQ for z/OS V8 IMS V12 IBM WebSphere Application Server for z/OS V8 IBM Integration Bus for z/OS V9 IBM Multi-site Workload Lifeline V2	z/OS V2 Db2 11 for z/OS CICS TS for z/OS V4 IBM MQ for z/OS V8 IMS V13

Sample control statement input to SCRT for Country Multiplex Pricing

SCRT supports additional externals for Country Multiplex Pricing. The sample multiplex report discussed in [“Sample sub-capacity report for Country Multiplex Pricing”](#) on page 347 is based on the following SCRT control statement and JCL inputs to SCRT:

SPECIAL DD control statements

The following SCRT control statement was specified in the SPECIAL DD to request that SCRT generate a multiplex report using the SMF or SCRT89 input data provided on the SMF DD statement. See [“Country_Multiplex_Pricing control statement”](#) on page 344 for a detailed description and restrictions on its use.

```
Country_Multiplex_Pricing
```

The following SCRT control statement was specified in the SPECIAL DD to request that SCRT generate the Detailed Interval Data section (W3) of the multiplex report. This control statement has no effect on whether SCRT uses the CMP algorithm to generate a multiplex report.

```
Generate_Detailed_Data
```

Guideline: The information in the Detailed Interval Data section of the multiplex report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating multiplex reports for submission to IBM. In the event that IBM requests that you submit a multiplex report that includes section W3, you can use the LMS Web application to submit the report.

INPUTCSV DD statement

MWP data or zCAP data or both can be provided as input when a multiplex report is requested via the **Country_Multiplex_Pricing** control statement. The requirements for MWP or zCAP or both must still be met.

In the sample environment, MWP data is provided using the INPUTCSV DD statement. The MWP data is for transactions occurring on MWP defining program IMS V13 for z/OS on LPAR M2_LPAR, which is running on the 2964-796 CPC in the sample configuration.

For more information about the INPUTCSV DD statement, see [“INPUTCSV DD statement”](#) on page 60.

Sample sub-capacity report for Country Multiplex Pricing

The sub-capacity report for Country Multiplex Pricing, includes many significant changes from the standard sub-capacity report. Due to these differences, when the **Country_Multiplex_Pricing** control statement is specified, the report that is generated is referred to as a *multiplex report*.

A multiplex report consists of a summary section, which contains billing and a summary of the CPCs and products that make up the multiplex environment. Following the summary section, there are additional report sections that detail information per CPC in the multiplex; some of these sections mirror those in a standard sub-capacity report.

The following topics describe the format of a multiplex report and, where appropriate, the differences between a sub-capacity report and a multiplex report.

Multiplex summary section

The multiplex summary section contains information about the version of SCRT used to generate the multiplex report, customer information, a summary of the machine configuration that makes up the multiplex, and the multiplex MSU values for the reporting period.

Figure 198 on page 347 shows the multiplex summary section for the sample configuration.

```

==B5===== SCRT MULTIPLEX REPORT - IBM Corp =====
SCRT Tool Release                23.10.00
Name of Person Submitting Report: John Customer
E-Mail Address of Report Submitter: customer@abc.com
Phone Number of Report Submitter: 800-555-0123

Customer Name                    ABC Corp
Run Date/Time                    10 Aug 2024 - 10:27
Report Period                    2 May, 2024 - 1 Jun, 2024 inclusive (31 days)
Number of processors in Multiplex 2

```

Figure 198. Sample multiplex report for Country Multiplex Pricing: Multiplex summary section

The fields shown in [Figure 198 on page 347](#) are:

SCRT Tool Release

The **SCRT Tool Release** field indicates the SCRT version, release, and modification level that was used to produce this multiplex report.

Customer Information section

The Customer Information section for a multiplex report contains many of the same fields that are shown on a standard sub-capacity report. For information about these fields, see [“Customer information” on page 105](#).

Run Date/Time

The **Run Date/Time** field shows the date and time at which the multiplex report was generated by SCRT.

Report Period

The **Report Period** field shows the start and end date of the reporting period, which includes the second day of one month through the first day of the next month.

Integrated Workload Pricing

When the **Integrated_Workload_Pricing** control statement is specified, an optional field is generated to indicate that IWP is active for the reporting period. In the sample configuration, IWP was not requested, so this optional field does not appear in the sample multiplex report.

Number of processors in Multiplex

The **Number of processors in Multiplex** field shows the number of distinct CPCs found in the SMF or SCRT89 data that was submitted on the SMF DD statement. SCRT applies the CMP algorithm based on the data found for these CPCs.

There are two CPCs in the sample configuration, as shown in [Figure 198 on page 347](#).

Note that the Country Multiplex Pricing terms and conditions require that SMF and SCRT89 data be collected for SCRT usage for all CPCs that are part of the multiplex.

Machine summary section

A major difference between a standard sub-capacity report and a multiplex report is that a multiplex report analyzes and reports on product usage across all of the CPCs that make up the multiplex.

SCRT provides a summary of the CPCs that make up the multiplex. The summary contains basic information about the CPCs that were found as well as the parameters that were provided as input to SCRT.

[Figure 199 on page 348](#) shows an example of the machine summary section of the multiplex report for the sample configuration.

Machine identifier				M1C1	M2C1
Customer number				US-S000000000	US-S000000000
Machine Serial Number				02-12345	02-56789
Machine Type and Model				2817-760	2964-796
Machine Rated Capacity (MSUs)				5001	9783
Machine Peak Utilization				1396	852
Machine MSU Consumed				3038	7130
Machine Model Changed				N	N
Exclude Data				N	N
Missing LPAR Data				Y	Y
Missing CPC Data				Y	Y
MWP Data				N	Y
MLC Product Name	Number	MSU	Time		
z/OS V2	5650-Z05	2244	04 May 2024 - 12:00	1396	848
DB2 11 for z/OS	5615-DB2	2244	04 May 2024 - 12:00	1396	848
CICS TS for z/OS V5	5655-Y04	2244	04 May 2024 - 12:00	1396	848
IBM MQ for z/OS V8	5655-W97	2244	04 May 2024 - 12:00	1396	848
IMS	(A11)	2244	04 May 2024 - 12:00	1396	848
IMS V13	5635-A04	848	04 May 2024 - 12:00		848
IMS V12	5635-A03	1396	04 May 2024 - 12:00	1396	
IPLA Product Name	Number	MSU	Time		
IBM WebSphere Application Server for z/OS V8	5655-W65	4	04 May 2024 - 12:00	4	
IBM Integration Bus for z/OS V9	5655-IBB	24	04 May 2024 - 12:00	24	
IBM Multi-site Workload Lifeline V2	5655-UM4	1396	04 May 2024 - 12:00	1396	
IPLA IMS-Based	(A11)	2244	04 May 2024 - 12:00	1396	848
z/OS Unreduced	(A11)	2248	04 May 2024 - 12:00	1396	852
IPLA z/OS-Based	(A11)	2244	04 May 2024 - 12:00	1396	848

Figure 199. Sample multiplex report for Country Multiplex Pricing: Machine summary section

The following topics describe the fields and sections shown in [Figure 199 on page 348](#).

Machine identifier

The machine identifier is a token that SCRT assigns to each CPC in the multiplex environment in order to differentiate between CPCs in the multiplex.

The format of the machine identifier is $MxCy$, where:

x

A number assigned to the CPC. It is unique within a set of CPCs assigned to a given customer number.

y

A number used to differentiate between the customer numbers assigned to CPCs in the multiplex.

In the sample configuration, the token M1C1 is assigned to the 2817 Model 760 CPC, and M2C1 is assigned to the 2964 Model 796 CPC. The number following the C is the same for both machine identifiers (that is, the y value is 1) because both CPCs are assigned to the same customer number.

For each CPC in the multiplex, the machine identifier tokens appear in separate columns following (to the right of) the **Machine identifier** row header. Additional information about each CPC appears in the same

columns (beneath the respective machine identifier values) on subsequent rows in the machine summary section.

Machine identifier assignments are not retained across multiple SCRT executions; thus, the machine identifiers may change on a subsequent run.

Customer number

The **Customer number** field in the multiplex report follows the same guidelines as in a standard sub-capacity report. For details about the customer number, see [“Customer information” on page 105](#).

SCRT supports the use of the **Assign Customer_Number** parameter on the PARMS DD statement. This parameter allows SCRT to assign additional customer numbers to specific CPCs within a multiplex. For more information about this parameter, see [“PARMS parameters that apply to specific CPCs” on page 37](#).

Machine Serial Number, Type and Model, and Rated Capacity (MSUs)

The **Machine Serial Number**, **Machine Type and Model**, and **Machine Rated Capacity (MSUs)** fields follow the same guidelines as in a standard sub-capacity report. For details, see [“Customer information” on page 105](#).

In a multiplex report, the values for these fields appear in the same column as the machine identifier for the CPC to which they pertain.

Machine Peak Utilization

The highest observed 4-hour rolling average MSU value for the CPC based on all systems reported by SCRT. This value corresponds to the CPC value reported in the Detail LPAR Data section (N5).

Machine MSU Consumed

The total MSU consumption for the CPC based on all systems reported by SCRT. This value corresponds to the CPC value reported in the Detail LPAR Usage Data section (N7).

Machine Model Changed

For each CPC in the multiplex, a value of Y (yes) or N (no) appears in the column under the respective machine identifier. A value of Y indicates that a machine model change was observed for that CPC one or more times during the reporting period. In such a case, the displayed **Machine Type and Model** value will be the highest observed value during the reporting period. For details, see [“Machine software model changes” on page 441](#).

Machine input parameters section

The rows that represent the machine input parameters display information that indicates the status of control statements and parameters for this execution of SCRT.

For these rows, a value of Y or N appears in the column under the respective machine identifier to indicate whether the condition identified by the row header was or was not met.

SCRT always displays Y or N values for the following rows:

Exclude Data

A value of Y under a given CPC indicates that an **Exclude** control statement was supplied for the CPC or for an LPAR on the CPC.

Missing LPAR Data

A value of Y for a given CPC indicates that SMF or SCRT89 data for an active LPAR was missing from the input data. For more information, see [Figure 204 on page 354](#).

Missing CPC Data

A value of Y for a given CPC indicates that SMF or SCRT89 data for the CPC was missing for the reporting period. For more information, see [Figure 205 on page 354](#).

In the sample configuration, SCRT determined that LPAR data was missing from both CPCs. SCRT also determined that CPC data was missing for periods of time for both CPCs.

SCRT may also display the following optional rows. These rows are displayed if, for any CPC, SCRT determines that the option was specified. The optional rows are:

MWP Data

A value of Y indicates that Mobile Workload Pricing data was provided to SCRT using the INPUTCSV DD statement.

zCAP Data

A value of Y indicates that IBM Z Collocated Application Pricing data was provided to SCRT using the INPUTCSV DD statement.

In the sample configuration, MWP data was supplied for the CPC identified by M2C1 (the 2964 Model 796 CPC), so the **MWP Data** row was generated in the sample multiplex report.

Product Summary Information section

The Product Summary Information section of the machine summary contains the data for the current reporting period, and contains many of the same fields as a standard sub-capacity report. However, since the CMP algorithm analyzes data across a multiplex (not just a single CPC), the reported data has a different scope. This topic highlights the differences between the Product Summary Information section generated using the CMP algorithms and that of a standard sub-capacity report.

The Product Summary Information section for a multiplex report contains the following fields:

MLC Product Name

IPLA Product Name

The **MLC Product Name** and **IPLA Product Name** fields in a multiplex report contain similar information as in a standard sub-capacity report. For more information about this field, see [“Product Summary Information section” on page 110](#).

A key difference when Country Multiplex Pricing has been requested is the application of Multi-Version Measurement (MVM). When multiple versions of a given sub-capacity eligible program (such as Db2 for z/OS) are in use in the multiplex, all MSUs for multiple versions within a given program family of program IDs (PIDs) are reported on a concurrent basis. This is similar to how SCRT measures z/OS when multiple z/OS versions are in use. MVM replaces Single Version Charging (SVC) when MVM is in effect. See IBM US Software Announcement 215-230 for details.

When multiple versions of a supported sub-capacity eligible program are in use, the **MLC Product Name** field contains the program family name. Subsequent rows contain the individual family product versions that are in use across the multiplex.

In [Figure 199 on page 348](#), multiple versions of IMS are in use in the multiplex. The program family row for IMS (as indicated by the (A11) PID value in the **Number** column) precedes the rows for the individual IMS versions.

Number

The Number field lists the program IDs of all supported MLC and IPLA products that were active on CPCs in the multiplex. A value of (A11) indicates that the row represents an MVM program family entry.

MSU

The **MSU** field contains the concurrent sub-capacity peak utilization, in MSUs, across all CPCs in the multiplex. This differs from the value in a non-multiplex environment, where separate peaks are calculated per CPC. The peak concurrent sub-capacity MSU value for the multiplex is calculated for each sub-capacity eligible product listed on the report.

When multiple versions of supported sub-capacity eligible programs are in use, the MSU value for the MVM program family entry contains the concurrent sub-capacity peak utilization calculated by adding summing the LPAR values where any version of the product is running. For each version in the MVM program family, the individual peak utilization for that version is also calculated and shown on the row for that program.

In Figure 199 on page 348, the MVM program family peak for IMS occurred on 4 May 2024 at 12:00 with a value of 2244 MSUs. The individual peaks for IMS V13 and IMS V12 also occurred during this time period in the sample configuration. Note that the individual version peaks can occur at different hours in the reporting period than the MVM program family peak.

For more information about sub-capacity MSU calculations, see **Tool MSUs** in “[Product Summary Information section](#)” on page 110.

Time

The Time field shows the date and time when the highest sub-capacity utilization peak in the multiplex was first reached.

Machine identifier MSU fields

For each CPC in the multiplex, these fields show the contribution to the multiplex concurrent peak, in MSUs. These values are shown for each individual program and, where applicable, for each MVM program family.

For example, in the sample configuration, the CPC identified by M1C1 contributed 1396 MSUs to the multiplex peak for IMS V12, while M2C1 contributed 0 MSUs. However, M2C1 contributed 848 MSUs to the peak for IMS V13, while M1C1 contributed 0 MSUs. These version peaks occurred during the same hour; thus, M1C1 contributed 1396 MSUs and M2C1 contributed 848 MSUs to the multiplex peak of 2244 MSUs for the IMS MVM program family.

Machine Summary section for CPC M1C1

A Machine Summary section is generated for each CPC found in the multiplex. The header for each Machine Summary section contains the machine identifier within the header itself. For instance, the Machine Summary section for CPC M1C1 contains the M1C1 value within the header, as shown in [Figure 200 on page 351](#).

Machine Information section for CPC M1C1

[Figure 200 on page 351](#) shows the Machine Information section for CPC 2817 Model 760 (M1C1) in the sample configuration.

```
==M1C1=====
Machine identifier      M1C1
Customer number        US-S000000000
Machine serial number   02-12345
Machine Type and Model  2817-760
Machine Rated Capacity (MSUs) 5001
Machine Peak Utilization 1396
Machine MSU Consumed    13038
```

Figure 200. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M1C1 - Machine information

The Machine Information section contains the machine identifier, customer number, serial number, machine type and model, rated capacity, machine peak utilization, and total MSU consumption. For more information about these fields, see “[Machine summary section](#)” on page 348.

Special Conditions section

The difference between the optional Special Conditions section for a multiplex report and a standard sub-capacity report is the addition of the following possible condition:

Global Performance Data Control Authority was not enabled during one or more intervals.

This condition is displayed whenever SCRT detects that Global Performance Data Control Authority is not enabled. Global Performance Data Control Authority is a requirement for CPCs that are part of a multiplex. For details, see the appropriate *PR/SM Planning Guide* for your CPC type.

For more information, see [“Special Conditions section” on page 107](#).

No special conditions were present in the sample configuration, so no Special Conditions section appears in the sample multiplex report.

Product Summary Information section for CPC M1C1

The Product Summary Information section for a multiplex report differs from a standard sub-capacity report in that the **Tool MSUs** values correspond to the contribution of the CPC to the multiplex peak value, instead of any individual peak that may have occurred on the CPC.

Figure 201 on page 352 shows the Product Summary Information section for the sample configuration.

```

=====
PRODUCT SUMMARY INFORMATION

MLC Product Name                MLC Product ID    Tool MSUs

z/OS V2                        5650-Z05          1396
DB2 11 for z/OS                5615-DB2          1396
CICS TS for z/OS V5            5655-Y04          1396
IBM MQ for z/OS V8             5655-W97          1396
IMS                            (A11)             1396
IMS V12                        5635-A03          1396

IPLA Product Name              IPLA Product ID    Tool MSUs

IBM WebSphere Application Server for z/OS V8  5655-W65           4   (iv)
IBM Integration Bus for z/OS V9              5655-IBB           24  (iv)
IBM Multi-site Workload Lifeline V2          5655-UM4          1396
IPLA IMS-Based                             (A11)             1396
z/OS Unreduced                             (A11)             1396
IPLA z/OS-Based                             (A11)             1396

Footnotes:
(iv) Product's Tool MSUs are based on Getting Started Sub-capacity Pricing from one or more LPARs

```

Figure 201. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M1C1 - Product Summary Information

In the sample configuration, CPC M1C1 contributed 1396 MSUs to each of the products shown in [Figure 201 on page 352](#). When multiple versions of a program in a program family are in use in the multiplex, the MVM program family name is displayed in the product list, regardless of where multiple versions are in use on the LPARs in a single CPC.

For more information about the fields, see [“Product Summary Information section” on page 110](#).

As with standard sub-capacity reports, the optional footnotes column is generated if it is necessary to apply a footnote to any product. For more information about the footnotes, see [“Product Summary Information section” on page 110](#).

Product Features section

There are no differences in the format of the Product Features section in a standard sub-capacity report and a multiplex report. [Figure 202 on page 353](#) shows the Product Features section for CPC M1C1 in the sample configuration.

```

==F5=====
PRODUCT FEATURES

Product Name                Product ID      Entitlement ID   Feature Name
z/OS V2                     5650-Z0S       S017291         z/OS V2 DFSMS dsshsm
z/OS V2                     5650-Z0S       S017292         z/OS V2 DFSMS imm
z/OS V2                     5650-Z0S       S017293         z/OS V2 DFSMS tvs
z/OS V2                     5650-Z0S       S017294         z/OS V2 DFSORT
z/OS V2                     5650-Z0S       S01729C        z/OS V2 RMF
z/OS V2                     5650-Z0S       S01729D        z/OS V2 SDSF
z/OS V2                     5650-Z0S       S01729F        z/OS V2 Security Server

OTHER REPORTABLE FEATURES

Product Name                Product ID      Feature Name
z/OS V2                     5650-Z0S       C/C++/DEBUG
z/OS V2                     5650-Z0S       JES2
z/OS V2                     5650-Z0S       RACF
z/OS V2                     5650-Z0S       TCP/IP BASE
z/OS V2                     5650-Z0S       TCP/IP CICS

```

Figure 202. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M1C1 - Product Features

System Recovery Boost Activity section

There are no differences in the format of the System Recovery Boost Activity section in a standard sub-capacity report and a multiplex report. [Figure 203 on page 353](#) shows the System Recovery Boost Activity section for CPC M1C1 in the sample configuration.

```

==G8=====
HOURS WITH BOOST ACTIVITY

NO BOOST ACTIVITY

```

Figure 203. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M1C1 - Hours with Boost Activity

Missing LPAR Data Details section

[Figure 204 on page 354](#) shows an example of the Missing LPAR Data Details section.

==H4=====

MISSING LPAR DATA DETAILS

LPARs with unmatched SMF70 and SMF89 records				Detected	Resolved (missing hours)	Est Peak MSU	Justification for low data collection (255 chars max)				
M1_LPAR				04 May 2024 - 08:00	04 May 2024 - 09:00 (1 hours)	0	1 hour missing data (no SMF70 data collected)				
Active LPARs with no SMF/SCRT89 records				Detected	Resolved (missing hours)	Est Peak MSU	Justification for low data collection (255 chars max)				
M1_MISS1				04 May 2024 - 09:00	04 May 2024 - 19:00 (10 hours)	5	(required)				
M1_MISS2				04 May 2024 - 09:00	04 May 2024 - 19:00 (10 hours)	1794	(required)				
M1_MISS3				04 May 2024 - 09:00	04 May 2024 - 19:00 (10 hours)	149	(required)				
M1_MISS4				04 May 2024 - 09:00	04 May 2024 - 19:00 (10 hours)	73	(required)				
M1_MISS5				04 May 2024 - 09:00	04 May 2024 - 19:00 (10 hours)	431	(required)				
M1_MISS6				04 May 2024 - 09:00	04 May 2024 - 19:00 (10 hours)	1	(required)				
Total Missing Hours				61							

		Multiplex	Est Peak							
OS Product Name	OS Product ID	Peak MSU	MSU	Date/Time	M1_LPAR	M1_MISS1	M1_MISS2	M1_MISS3	M1_MISS4	M1_MISS5
z/OS V2	5650-Z0S	10463	3582	04 May 2024 - 12:00	1396	5	1548	144	69	419

Figure 204. Sample multiplex report for Country Multiplex Pricing: Machine Summary section - Missing LPAR Data Details section

In the sample configuration, LPAR data details are missing for M1_LPAR, as shown in Figure 204 on page 354. SCRT displays the notification message because the SMF data is incomplete for a given interval.

Missing CPC Data Details section for CPC M1C1

The Missing CPC Data Details section is generated to notify you that SCRT detected missing CPC data in the reporting period. This section is always generated, regardless of whether any data is missing for the CPC.

When SCRT detects that no data was provided for any LPAR on the CPC, it generates the following notification message in the Missing CPC Data Details section:

CPC with no SMF/SCRT89 data from any LPAR

For each incident, you must explain why the condition occurred. You must either provide your explanation by editing the multiplex report using a spreadsheet application or by entering your explanation in the LMS web application. The (required) placeholder appears in the field under the **Justification for low data collection** column when an explanation is necessary. Your explanation should replace the (required) placeholder. No explanation is necessary when the (required) placeholder does not appear.

Figure 205 on page 354 shows an example of the Missing CPC Data Details section.

==H6=====

MISSING CPC DATA DETAILS

CPC with no SMF/SCRT89 data from any LPAR		Detected	Resolved	Justification for low data collection (255 chars max)
M1C1		02 May 2024 - 00:00	04 May 2024 - 08:00 (56 hours)	(required)
M1C1		04 May 2024 - 19:00	02 Jun 2024 - 00:00 (677 hours)	(required)
Total Missing Hours		733		

Figure 205. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M1C1 - Missing CPC Data Details section

In the sample configuration, there are two intervals where data was missing from CPC M1C1, as shown in Figure 205 on page 354.

Machine summary section for CPC M2C1

“Machine Summary section for CPC M1C1” on page 351 provides an overview of the multiplex report sections that are generated on a per-CPC basis for each CPC in the multiplex.

The similar report sections for CPC M2C1 are shown here, but only key points and differences are discussed. For full descriptions of each report section, see the corresponding section in “Machine Summary section for CPC M1C1” on page 351.

Machine Information section for CPC M2C1

Figure 206 on page 355 shows the Machine Information section for CPC 2964 Model 796 (M2C1) in the sample configuration.

```
==M2C1=====

Machine identifier          M2C1
Customer number            US-S000000000
Machine serial number      02-56789
Machine Type and Model     2964-796
Machine Rated Capacity (MSUs) 9783
Machine Peak Utilization   852
Machine MSU Consumed       7130
```

Figure 206. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M2C1 - Machine information

Product Summary Information section for CPC M2C1

A key difference in the Product Summary Information section between CPCs M1C1 and M2C1 is the presence of the optional **Footnotes** column. The Footnotes column is generated because Mobile Workload Pricing data was provided via the INPUTCSV DD statement. The MWP footnote (d) appears next to the **Tool MSUs** value for the IMS V13 product, and the **Footnotes** header and a description row (per footnote) appears after the last product in section E5. (For more information about the MWP footnote, see “Mobile Workload Pricing footnote” on page 203.)

Figure 207 on page 355 shows the Product Summary Information section for the sample configuration.

```
==E5=====
PRODUCT SUMMARY INFORMATION

MLC Product Name          MLC Product ID    Tool MSUs  Footnotes

z/OS V2                   5650-Z0S         848
DB2 11 for z/OS           5615-DB2         848
CICS TS for z/OS V5       5655-Y04         848
IBM MQ for z/OS V8        5655-W97         848
IMS                        (A11)            848
IMS V13                   5635-A04         848      (d)

IPLA Product Name        IPLA Product ID    Tool MSUs

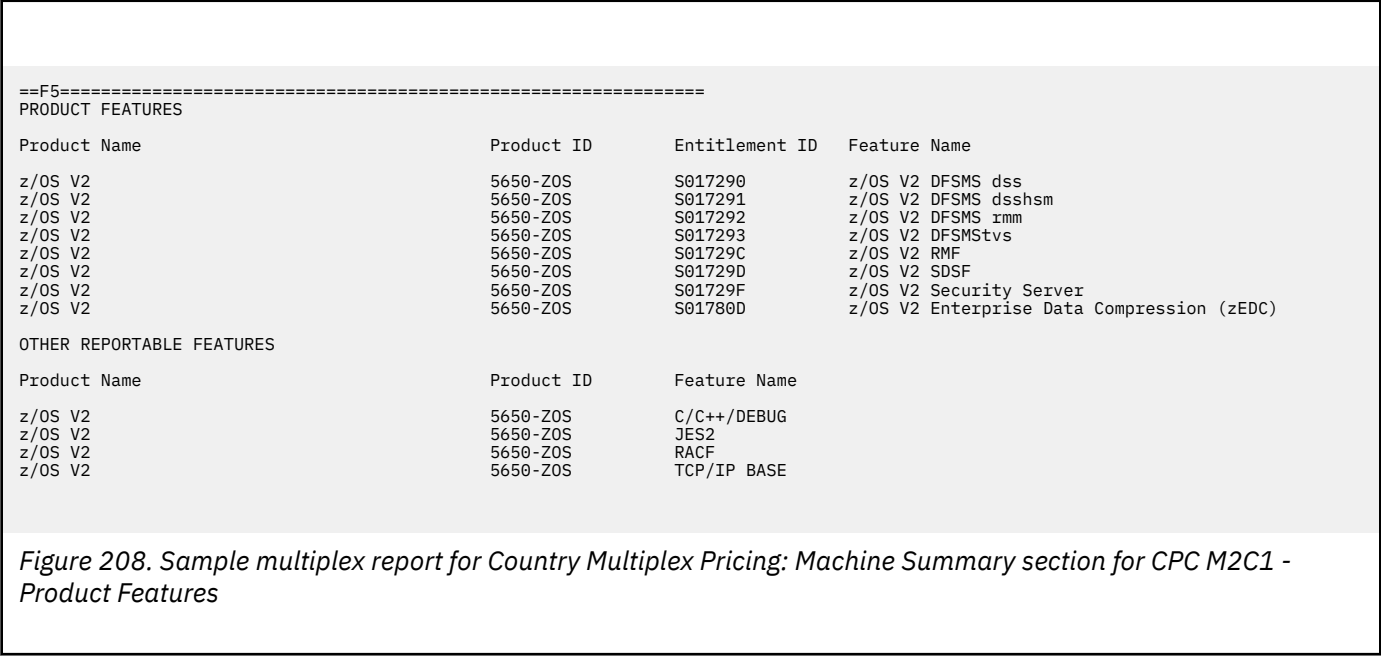
IPLA IMS-Based            (A11)             848
z/OS Unreduced            (A11)             852
IPLA z/OS-Based           (A11)             848

Footnotes:
(d) Mobile Workload Pricing transaction data from this product was included from one or more LPARs
```

Figure 207. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M2C1 - Product Summary Information

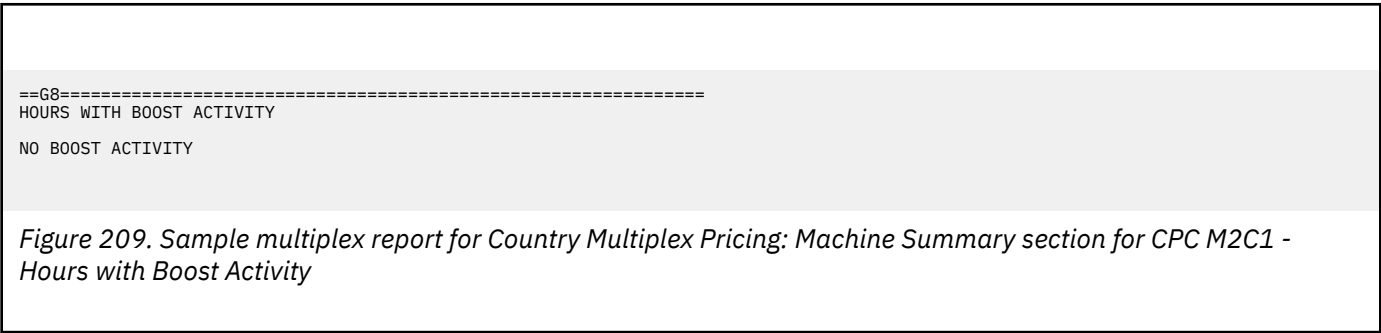
Product Features section

There are no differences in the format of the Product Features section in a standard sub-capacity report and a multiplex report. [Figure 208 on page 356](#) shows the Product Features section for CPC M2C1 in the sample configuration.



System Recovery Boost Activity section

There are no differences in the format of the System Recovery Boost Activity section in a standard sub-capacity report and a multiplex report. [Figure 209 on page 356](#) shows the System Recovery Boost Activity section for CPC M2C1 in the sample configuration.



Missing LPAR Data Details section

[Figure 210 on page 357](#) shows an example of the Missing LPAR Data Details section of a multiplex report.

```

==H4=====
MISSING LPAR DATA DETAILS

LPARs with unmatched
SMF70 and SMF89 records   Detected          Resolved (missing hours)      Est Peak      Justification for low data
M1_LPAR                   04 May 2024 - 08:00          04 May 2024 - 09:00 (1 hours)  MSU           collection (255 chars max)
collected)              0                               0                               0             1 hour missing data (no SMF70 data)

Active LPARs with no
SMF/SCRT89 records       Detected          Resolved (missing hours)      Est Peak      Justification for low data
M2_MISS1                 04 May 2024 - 09:00          04 May 2024 - 19:00 (10 hours) 3434          (required)
M2_MISS2                 04 May 2024 - 09:00          04 May 2024 - 19:00 (10 hours) 812           (required)
M2_MISS3                 04 May 2024 - 09:00          04 May 2024 - 19:00 (10 hours) 345           (required)
M2_MISS4                 04 May 2024 - 09:00          04 May 2024 - 19:00 (10 hours) 3            (required)
M2_MISS5                 04 May 2024 - 09:00          04 May 2024 - 19:00 (10 hours) 2            (required)
M2_MISS6                 04 May 2024 - 09:00          04 May 2024 - 19:00 (10 hours) 449          (required)
M2_MISS7                 04 May 2024 - 09:00          04 May 2024 - 19:00 (10 hours) 518          (required)
M2_MISS8                 04 May 2024 - 09:00          04 May 2024 - 19:00 (10 hours) 733          (required)

Total Missing Hours              81

OS Product Name      OS Product ID      Multiplex      Est Peak
M2_MISS5  M2_MISS6  M2_MISS7  M2_MISS8      Peak MSU      MSU      Date/Time      M2_LPAR  M2_MISS1  M2_MISS2  M2_MISS3  M2_MISS4
z/OS V2          3          2          444          432          572      10463      6881      04 May 2024 - 12:00      848      3434      812      334

```

Figure 210. Sample multiplex report for Country Multiplex Pricing: Machine Summary section - Missing LPAR Data Details section

Missing CPC Data Details section for CPC M2C1

Figure 211 on page 357 shows an example of the Missing CPC Data Details section of the multiplex report.

```

==H6=====
MISSING CPC DATA DETAILS

CPC with no SMF/SCRT89
data from any LPAR       Detected          Resolved          Justification for low data
M2C1                    02 May 2020 - 00:00          04 May 2020 - 08:00 (56 hours) collection (255 chars max)
M2C1                    04 May 2020 - 19:00          02 Jun 2020 - 00:00 (677 hours) (required)
                                           (required)

Total Missing Hours              733

```

Figure 211. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M2C1 - Missing CPC Data Details section (M2C1)

Customer verification section

The customer verification section of the multiplex report contains information that supports and expands on the information in the Product Summary Information section.

Other sections that are not displayed for the sample configuration may be displayed under different conditions. For information about the possible sections, see [“Customer verification section”](#) on page 118.

The detailed data sections for customer verification follow the Detail Data Sections header in the multiplex report, as shown in [Figure 212 on page 358](#).

```
==L7=====
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
```

Figure 212. Sample multiplex report for Country Multiplex Pricing: Customer verification section - Detail Data Sections header

Detail data sections for CPC M1C1

The detail data sections for CPC M1C1 include the Machine Information section, Detailed LPAR Data section, the Product Max Contributors section, and the Product Grid Snapshot section.

Machine information section for CPC M1C1

Figure 213 on page 358 shows the Machine Information section for CPC 2817 Model 760 (M1C1) in the sample configuration.

```
==M7=====

Machine identifier      M1C1
Customer number        US-S000000000
Machine serial number   02-12345
Machine Type and Model  2817-760
Machine Rated Capacity (MSUs) 5001
```

Figure 213. Sample multiplex report for Country Multiplex Pricing: Machine Summary section for CPC M1C1 - Machine information

The Machine Information section contains the machine identifier, customer number, serial number, machine type and model, and rated capacity. For more information about these fields, see “Machine summary section” on page 348.

Detail LPAR Data section for CPC M1C1

There are no differences in the format of the Detail LPAR Data section in a standard sub-capacity report and a multiplex report. There is, however, a key difference in the values that are represented. The values and time shown in a multiplex report correspond to the contribution of the LPAR to the multiplex peak concurrent MSU value.

Figure 214 on page 358 shows the Detail LPAR Data section for CPC M1C1 in the sample configuration.

```
==N5=====
DETAIL LPAR DATA SECTION
```

	Highest	Hour	Count	Date/Time	2nd Highest	Hour	Count	Date/Time
M1_LPAR	1396		1	04 May 2024 - 12:00	1394		1	04 May 2024 - 13:00
CPC	1396		1	04 May 2024 - 12:00	1394		1	04 May 2024 - 13:00

Figure 214. Sample multiplex report for Country Multiplex Pricing: Detail LPAR Data section for CPC M1C1

For more information about the Detail LPAR Data section, see “Detail LPAR Data section” on page 119.

Detail LPAR Usage Data section for CPC M1C1

There are no differences in the format of the Detail LPAR Usage Data section in a standard sub-capacity report and a multiplex report.

Figure 215 on page 359 shows the Detail LPAR Usage Data section for CPC M1C1 in the sample configuration.

```

==N7=====
DETAIL LPAR USAGE DATA SECTION

M1_LPAR      Total MSU Consumed   Peak Hour Consumption   Date/Time   OS
              13038              1483   04 May 2024 - 12:00   z/OS
CPC              13038              1483   04 May 2024 - 12:00

```

Figure 215. Sample multiplex report for Country Multiplex Pricing: Detail LPAR Usage Data section for CPC M1C1

For more information about the Detail LPAR Data section, see [“Detail LPAR Usage Data section”](#) on page 120.

Product Max Contributors section for CPC M1C1

There are no differences in the format of the Product Max Contributors section in a standard sub-capacity report and a multiplex report. There is, however, a key difference in the values that are represented. The values and time shown in a multiplex report correspond to the contribution of the LPAR to the multiplex peak concurrent MSU value.

Figure 216 on page 359 shows the Product Max Contributors section for CPC M1C1 in the sample configuration.

```

==P5=====
PRODUCT MAX CONTRIBUTORS

Product Name                Product ID   Highest   Date/Time   LPAR
                                M1_LPAR
z/OS V2                     5650-ZOS    1396    04 May 2024 - 12:00    1396
DB2 11 for z/OS             5615-DB2    1396    04 May 2024 - 12:00    1396
CICS TS for z/OS V5         5655-Y04    1396    04 May 2024 - 12:00    1396
IBM MQ for z/OS V8          5655-W97    1396    04 May 2024 - 12:00    1396
IMS                          (All)       1396    04 May 2024 - 12:00    1396
IMS V12                     5635-A03    1396    04 May 2024 - 12:00    1396
IBM WebSphere Application Server for z/OS V8  5655-W65     4      04 May 2024 - 12:00     4
IBM Integration Bus for z/OS V9  5655-IBB    24      04 May 2024 - 12:00    24
IBM Multi-site Workload Lifeline V2  5655-UM4    1396    04 May 2024 - 12:00    1396
IPLA IMS-Based              (All)       1396    04 May 2024 - 12:00    1396
IPLA z/OS-Based             (All)       1396    04 May 2024 - 12:00    1396

```

Figure 216. Sample multiplex report for Country Multiplex Pricing: Product Max Contributors section for CPC M1C1

For more information about the Product Max Contributors section, see [“Product Max Contributors section”](#) on page 121.

Product Grid Snapshot section for CPC M1C1

There are no differences in the format of the Product Grid Snapshot section in a standard sub-capacity report and a multiplex report.

Figure 217 on page 360 shows the Product Grid Snapshot section for CPC M1C1 in the sample configuration.

```

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                                Product ID                                M1_LPAR
z/OS V2                                    5650-Z05                                    1.3%
DB2 11 for z/OS                            5615-DB2                                    1.3%
CICS TS for z/OS V5                        5655-Y04                                    1.3%
IBM MQ for z/OS V8                         5655-W97                                    1.3%
IMS                                         (All)                                       1.3%
IMS V12                                    5635-A03                                    1.3%
IBM WebSphere Application Server for z/OS V8 5655-W65                                    1.3%
IBM Integration Bus for z/OS V9            5655-IBB                                    1.3%
IBM Multi-site Workload Lifeline V2        5655-UM4                                    1.3%
IPLA IMA-Based                             (All)                                       1.3%
IPLA z/OS-Based                             (All)                                       1.3%

```

Figure 217. Sample multiplex report for Country Multiplex Pricing: Product Grid Snapshot section for CPC M1C1

For more information about the Product Grid Snapshot section, see [“Product Grid Snapshot section”](#) on page 123.

Detail data sections for CPC M2C1

The detail data sections of the multiplex report for CPC M2C1 include the same sections as for CPC M1C1.

Machine information section for CPC M2C1

Figure 218 on page 360 shows the Machine Information section for CPC 2964 Model 796 (M2C1) in the sample configuration.

```

==M7=====

Machine identifier          M2C1
Customer number            US-S000000000
Machine serial number      02-56789
Machine Type and Model     2964-796
Machine Rated Capacity (MSUs) 9783

```

Figure 218. Sample multiplex report for Country Multiplex Pricing: Machine information section for CPC M2C1

The Machine Information section contains the machine identifier, customer number, serial number, machine type and model, and rated capacity. For more information about these fields, see [“Machine summary section”](#) on page 348.

Detail LPAR Data section for CPC M2C1

There are no differences in the format of the Detail LPAR Data section in a standard sub-capacity report and a multiplex report. There is, however, a key difference in the values that are represented. The values and time shown in a multiplex report correspond to the contribution of the LPAR to the multiplex peak concurrent MSU value.

Figure 219 on page 361 shows the Detail LPAR Data section for CPC M2C1 in the sample configuration.

```

==N5=====
DETAIL LPAR DATA SECTION

M2_LPAR      Highest      Hour Count      Date/Time      2nd Highest      Hour Count      Date/Time
              852          1              04 May 2015 - 12:00      836          1              04 May 2015 - 13:00
CPC              852          1              04 May 2015 - 12:00      836          1              04 May 2015 - 13:00

```

Figure 219. Sample multiplex report for Country Multiplex Pricing: Detail LPAR Data section for CPC M2C1

For more information about the Detail LPAR Data section, see [“Detail LPAR Data section” on page 119](#).

Detail LPAR Usage Data section for CPC M2C1

There are no differences in the format of the Detail LPAR Usage Data section in a standard sub-capacity report and a multiplex report. There is, however, a key difference in the values that are represented. The values and time shown in a multiplex report correspond to the contribution of the LPAR to the multiplex peak concurrent MSU value.

Figure 220 on [page 361](#) shows the Detail LPAR Usage Data section for CPC M2C1 in the sample configuration.

```

==N7=====
DETAIL LPAR USAGE DATA SECTION

M1_LPAR      Total MSU Consumed      Peak Hour Consumption      Date/Time      OS
              7130              970              04 May 2024 - 12:00      z/OS
CPC              7130              970              04 May 2024 - 12:00

```

Figure 220. Sample multiplex report for Country Multiplex Pricing: Detail LPAR Usage Data section for CPC M2C1

For more information about the Detail LPAR Data section, see [“Detail LPAR Usage Data section” on page 120](#).

Product Max Contributors section for CPC M2C1

There are no differences in the format of the Product Max Contributors section in a standard sub-capacity report and a multiplex report. There is, however, a key difference in the values that are represented. The values and time shown in a multiplex report correspond to the contribution of the LPAR to the multiplex peak concurrent MSU value.

Figure 221 on [page 362](#) shows the Product Max Contributors section for CPC M2C1 in the sample configuration. Note that, for CPC M2C1, the Mobile MSU reduction is shown.

```

=====
==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR M2_LPAR	Mobile MSU Reduction
z/OS V2	5650-Z0S	848	04 May 2024 - 12:00	848	4
DB2 11 for z/OS	5615-DB2	848	04 May 2024 - 12:00	848	4
CICS TS for z/OS V5	5655-Y04	848	04 May 2024 - 12:00	848	4
IBM MQ for z/OS V8	5655-W97	848	04 May 2024 - 12:00	848	4
IMS	(A11)	848	04 May 2024 - 12:00	848	4
IMS V13	5635-A04	848	04 May 2024 - 12:00	848	4
IPLA IMS-Based	(A11)	848	04 May 2024 - 12:00	848	4
IPLA z/OS-Based	(A11)	848	04 May 2024 - 12:00	848	4

Figure 221. Sample multiplex report for Country Multiplex Pricing: Product Max Contributors section for CPC M2C1

For more information about the Product Max Contributors section, see [“Product Max Contributors section” on page 121.](#)

Product Grid Snapshot section for CPC M2C1

There are no differences in the format of the Product Grid Snapshot section in a standard sub-capacity report and a multiplex report.

Figure 222 on page 362 shows the Product Grid Snapshot section for CPC M2C1 in the sample configuration.

```

=====
==Q5=====
PRODUCT GRID SNAPSHOT

```

Product Name	Product ID	M2_LPAR
z/OS V2	5650-Z0S	1.3%
DB2 11 for z/OS	5615-DB2	1.3%
CICS TS for z/OS V5	5655-Y04	1.3%
IBM MQ for z/OS V8	5655-W97	1.3%
IMS	(A11)	1.3%
IMS V13	5635-A04	1.3%
IPLA IMS-Based	(A11)	1.3%
IPLA z/OS-Based	(A11)	1.3%

Figure 222. Sample multiplex report for Country Multiplex Pricing: Product Grid Snapshot section for CPC M2C1

For more information about the Product Grid Snapshot section, see [“Product Grid Snapshot section” on page 123.](#)

Detailed Interval Data section

The Detailed Interval Data section (W3) is generated when you specify the **Generate_Detailed_Data** SPECIAL DD control statement. This optional section displays information that gives insight into the factors contributing to the multiplex MSU value on a per-hour basis for each sub-capacity product. This information includes the application of any MWP or zCAP reduction.

The products are listed in this section in the same order as they appear in other report sections. MLC products are listed first, followed by IPLA products.

Guideline: The information in the Detailed Interval Data section of the multiplex report is intended for customer informational purposes only and may result in a large amount of report data being generated. Therefore, you should typically not specify the **Generate_Detailed_Data** control statement when generating multiplex reports for submission to IBM. In the event that IBM requests that you submit a multiplex report that includes section W3, you can use the LMS web application to submit the report.

Figure 223 on page 363 shows the beginning of the Detailed Interval Data section for the sample configuration. Due to its size, the entire Detailed Interval Data section is not shown here; instead, portions of this section are shown to illustrate its use for a subset of products.

Figure 223 on page 363 shows the Detailed Interval Data section header and the z/OS V2 product entry, including the 12:00 multiplex concurrent peak. Notice that this interval includes a 4-MSU MWP reduction, based on mobile transaction data that was provided for the IMS V13 MWP defining program. In accordance with the terms and conditions for Mobile Workload Pricing, the mobile reduction is only applied to LPAR M2_LPAR on CPC M2C1.

```

=====
Detailed Interval Data
MLC Products
5650-ZOS
Processor Type Serial      z/OS V2
Date Time of RMF Interval  Type      CPC MSU  Partition  CPC Total  2964-56789
                                     M1_LPAR                                     Type      CPC MSU  Partition
04 May 2024 - 09:00      2817-760 5001    1220      1220      2964-796 9783    643
04 May 2024 - 10:00      2817-760 5001    1284      1284      2964-796 9783    741
04 May 2024 - 11:00      2817-760 5001    1358      1358      2964-796 9783    827
04 May 2024 - 12:00      2817-760 5001    1396      1396      2964-796 9783    848
04 May 2024 - 13:00      2817-760 5001    1394      1394      2964-796 9783    830
04 May 2024 - 14:00      2817-760 5001    1389      1389      2964-796 9783    814
04 May 2024 - 15:00      2817-760 5001    1371      1371      2964-796 9783    790
04 May 2024 - 16:00      2817-760 5001    1351      1351      2964-796 9783    742
04 May 2024 - 17:00      2817-760 5001    1293      1293      2964-796 9783    670
04 May 2024 - 18:00      2817-760 5001    1208      1208      2964-796 9783    575

IMS Mobile Partition  Total Mobile MSU  CPC Total  MultiPlex
M2_LPAR      Reduction                                     Total
0             0             643        1863
0             0             741        2025
0             0             827        2185
4             4             848        2244
6             6             830        2224
8             8             814        2203
13            13            790        2161
14            14            742        2093
16            16            670        1963
13            13            575        1783

```

Figure 223. Sample multiplex report for Country Multiplex Pricing: Detailed Interval Data section — MLC products

The Detailed Interval Data section displays the following columns for each product:

Date Time of RMF interval

The date and time of the adjusted interval.

Type

For each CPC, the machine type and model of the CPC.

CPC MSU

For each CPC, the machine rated capacity, in MSUs.

Partition *partitionName*

The adjusted MSU values for each LPAR (*partitionName*) on the CPC. Note that any application of the MWP or zCAP algorithm has already occurred.

product Mobile Partition partitionName

product zCAP Partition partitionName

For each mobile or zCAP defining product with data specified on the INPUTCSV DD statement, the MWP or zCAP reduction applied for that defining product, per partition.

Total Mobile MSU Reduction

Total zCAP MSU Reduction

Total Mobile or zCAP MSU reduction applied across all MWP or zCAP defining products.

CPC Total

For each CPC, the sum of the adjusted LPAR MSU values across all partitions for the adjusted product.

Multiplex Total

The sum of the CPC (and, therefore, LPAR) MSU values in the multiplex.

Figure 224 on page 364 shows the Detailed Interval Data sections for Db2 11 for z/OS, CICS TS for z/OS V5, and IBM MQ for z/OS V8.

5615-DB2													
Processor Type Serial				DB2 11 for z/OS									
Date	Time	of RMF	Interval	Type	CPC	MSU	Partition	CPC	Total	Type	CPC	MSU	Partition
				2817-12345			M1_LPAR			2964-56789			M2_LPAR

Figure 226 on page 365 shows the Detailed Interval Data sections for IPLA products for IPLA IMS-Based, z/OS Unreduced, and IPLA z/OS-Based products.

(All)															
Processor	Type	Serial	IPLA IMS-Based			2964-56789			IMS Mobile	Total		MultiPlex			
Date	Time	of RMF Interval	Type	CPC	MSU	Partition	CPC	Total	Type	CPC	MSU	Total			
2817-12345															
M1_LPAR															
04 May 2024	- 09:00		2817-760	5001		1220		1220	2964-796	9783	643	0	643	1863	
04 May 2024	- 10:00		2817-760	5001		1284		1284	2964-796	9783	741	0	741	2025	
04 May 2024	- 11:00		2817-760	5001		1358		1358	2964-796	9783	827	0	827	2185	
04 May 2024	- 12:00		2817-760	5001		1396		1396	2964-796	9783	848	4	4	848	2244
04 May 2024	- 13:00		2817-760	5001		1394		1394	2964-796	9783	830	6	6	830	2224
04 May 2024	- 14:00		2817-760	5001		1389		1389	2964-796	9783	814	8	8	814	2203
04 May 2024	- 15:00		2817-760	5001		1371		1371	2964-796	9783	790	13	13	790	2161
04 May 2024	- 16:00		2817-760	5001		1351		1351	2964-796	9783	742	14	14	742	2093
04 May 2024	- 17:00		2817-760	5001		1293		1293	2964-796	9783	670	16	16	670	1963
04 May 2024	- 18:00		2817-760	5001		1208		1208	2964-796	9783	575	13	13	575	1783
z/OS Unreduced															
Processor	Type	Serial	2817-12345			2964-56789			IMS Mobile	Total		MultiPlex			
Date	Time	of RMF Interval	Type	CPC	MSU	Partition	CPC	Total	Type	CPC	MSU	Total			
M1_LPAR															
04 May 2024	- 09:00		2817-760	5001		1220		1220	2964-796	9783	643	0	0	643	1863
04 May 2024	- 10:00		2817-760	5001		1284		1284	2964-796	9783	741	0	0	741	2025
04 May 2024	- 11:00		2817-760	5001		1358		1358	2964-796	9783	827	0	0	827	2185
04 May 2024	- 12:00		2817-760	5001		1396		1396	2964-796	9783	852	0	0	852	2248
04 May 2024	- 13:00		2817-760	5001		1394		1394	2964-796	9783	836	0	0	836	2230
04 May 2024	- 14:00		2817-760	5001		1389		1389	2964-796	9783	822	0	0	822	2211
04 May 2024	- 15:00		2817-760	5001		1371		1371	2964-796	9783	803	0	0	803	2174
04 May 2024	- 16:00		2817-760	5001		1351		1351	2964-796	9783	756	0	0	756	2107
04 May 2024	- 17:00		2817-760	5001		1293		1293	2964-796	9783	686	0	0	686	1979
04 May 2024	- 18:00		2817-760	5001		1208		1208	2964-796	9783	588	0	0	588	1796
IPLA z/OS-Based															
Processor	Type	Serial	2817-12345			2964-56789			IMS Mobile	Total		MultiPlex			
Date	Time	of RMF Interval	Type	CPC	MSU	Partition	CPC	Total	Type	CPC	MSU	Total			
M1_LPAR															
04 May 2024	- 09:00		2817-760	5001		1220		1220	2964-796	9783	643	0	0	643	1863
04 May 2024	- 10:00		2817-760	5001		1284		1284	2964-796	9783	741	0	0	741	2025
04 May 2024	- 11:00		2817-760	5001		1358		1358	2964-796	9783	827	0	0	827	2185
04 May 2024	- 12:00		2817-760	5001		1396		1396	2964-796	9783	848	4	4	848	2244
04 May 2024	- 13:00		2817-760	5001		1394		1394	2964-796	9783	830	6	6	830	2224
04 May 2024	- 14:00		2817-760	5001		1389		1389	2964-796	9783	814	8	8	814	2203
04 May 2024	- 15:00		2817-760	5001		1371		1371	2964-796	9783	790	13	13	790	2161
04 May 2024	- 16:00		2817-760	5001		1351		1351	2964-796	9783	742	14	14	742	2093
04 May 2024	- 17:00		2817-760	5001		1293		1293	2964-796	9783	670	16	16	670	1963
04 May 2024	- 18:00		2817-760	5001		1208		1208	2964-796	9783	575	13	13	575	1783

Figure 226. Sample multiplex report for Country Multiplex Pricing: Detailed Interval Data section — IPLA products (continued)

Chapter 17. Sub-capacity reporting and Multi-Version Measurement (MVM)

Multi-Version Measurement (MVM) for z/OS and z/VSE removes time limits for running multiple eligible versions of a software program.

With MVM, you can run different versions of a program simultaneously for an unlimited duration during a program version upgrade. You can also choose to run multiple different versions of a program simultaneously for an unlimited duration in a production environment. MVM allows you to selectively deploy new software version and provides more flexible control over their program upgrade cycles.

MVM replaces Single Version Charging (SVC), the Migration Pricing Option (MPO), and the IPLA migration grace period.

Overview of Multi-Version Measurement

Whenever multiple versions of a sub-capacity program are running, such as Db2 for z/OS, SCRT reports all MSUs for the multiple versions with that program family on a concurrent basis, similar to how it measures z/OS today. SCRT calculates the combined concurrent peak (indicated by (A11) in the report) by adding up the LPAR values where any version of the program is running. For billing purposes, all MSUs within a given family are reported on a concurrent peak basis and are priced at the cost of the latest version of the program within the family. Essentially, you are paying for all versions and releases of a program, such as Db2, as though it was a single version.

Multi-Version Measurement is described in U.S. Announcement letter 217-093. MVM is effective for SMF data collected on or after 2 April 2017.

For the current list of MLC MVM program families, see [Sub-Capacity Eligible and MVM Eligible MLC Programs \(www.ibm.com/downloads/documents/us-en/10a99803f2afd789\)](http://www.ibm.com/downloads/documents/us-en/10a99803f2afd789).

For the current list of IPLA MVM program families, see [IPLA Execution-Based Programs \(www.ibm.com/downloads/documents/us-en/10a99803f52fd884\)](http://www.ibm.com/downloads/documents/us-en/10a99803f52fd884).

Prerequisites for Multi-Version Measurement

The prerequisites for Multi-Version Measurement depend on whether you use sub-capacity pricing or Country Multiplex Pricing.

- For Country Multiplex Pricing, there are no additional requirements.
- For sub-capacity pricing, provide SCRT with SMF data collected on or after 2 April 2017.

Run SCRT for each reporting period and submit the resulting report to IBM on a monthly basis, as usual.

SCRT externals for Multi-Version Measurement

SCRT V24.11.0, and later, accepts the following optional control statements in the SPECIAL DD input.

MLC_ALL

Indicates that SCRT is to report an MLC family line for MVM eligible products even when only one product in the program family is in use.

IPLA_ALL

Indicates that SCRT is to report an IPLA family line for MVM eligible products even when only one product in the program family is in use.

Sub-capacity report differences for Multi-Version Measurement

Multi-Version Measurement introduces the following differences in the content of a standard sub-capacity report:

- The section B5 header is changed to indicate that MVM reporting is in use.
- In section E5, for each program family in which multiple product versions are running, an additional program family line is generated. The program family line, identified by (A11), indicates the simultaneous peak of all product versions within that program family.

If you specified the **MLC_ALL** or **IPLA_ALL** control statements in the SPECIAL DD, the program family line is generated even if only one product within a program family is in use.

Sample sub-capacity report for Multi-Version Measurement

A sub-capacity report where MVM is in use has the same basic format and report sections as a standard sub-capacity report with only a few differences.

Product Summary Information section

The Product Summary Information section is updated to include a line for each program family for which multiple versions of a product were in use. This includes changes in the contents of the following fields:

MLC Product Name

IPLA Product Name

When multiple versions of a supported program are in use, the product name field contains the program family name. Subsequent rows contain the individual family product versions that are in use.

MLC Product ID

IPLA Product ID

In the program family row, the product ID value is (A11) and indicates that the row represents an MVM program family entry.

Tool MSUs

The **Tool MSUs** field for a program family contains the peak concurrent MSU value for all products within the program family, calculated by summing the LPAR MSU values for each hour during which a product in the program family ran and taking the peak value.

The tool MSU value for a program family may represent a different peak hour than any of the individual product versions.

Figure 227 on [page 369](#) shows an excerpt from the Product Summary Information section of a sample MVM report.

```

=====
PRODUCT SUMMARY INFORMATION

MLC Product Name                MLC Product ID    Tool MSUs    Customer MSUs
z/OS V2                        5650-Z0S          658
DB2 for z/OS                  (A11)            617
DB2 11 for z/OS                5615-DB2          488
DB2 10 for z/OS                5605-DB2          597
IMS V12                        5635-A03          576
COBOL                        (A11)            451
Enterprise COBOL for z/OS V5    5655-W32          284
IBM Enterprise Cobol for z/OS V4 5655-S71          257    (optional)

IPLA Product Name              IPLA Product ID    Tool MSUs    Customer MSUs
CICS TS VUE V5                 5722-DFJ          509
Fault Analyzer              (A11)            91
Fault Analyzer for z/OS V7      5655-R46          42    (optional)
Fault Analyzer for z/OS V6      5655-P16          64    (optional)
File Export for z/OS V1         5697-I12          657    (optional)
IBM Integration Bus          (A11)            35
IBM Integration Bus for z/OS V9 5655-IBB          7

```

Figure 227. Sample sub-capacity report for Multi-Version Measurement: Product Summary Information section

The report shows MVM program family lines for MLC products Db2 and COBOL, and for IPLA products Fault Analyzer and IBM Integration Bus.

For instance, the report shows a DB2 for z/OS line with an **MLC Product ID** value of (A11). The product lines for the individual products, DB2 10 for z/OS and DB2 11 for z/OS, are reported as they were prior to MVM. The **Tool MSUs** value for the program family line represents the peak concurrent MSU value for all versions of Db2.

Note: SCRT may report both a line for the program family and a line for a program suite that contains several programs. SCRT can only report on the programs that it observes in a customer's input data; however, it is unaware of how any individual program is licensed for that customer.

Chapter 18. Sub-capacity reporting for products that report a non-MSU metric

SCRT supports sub-capacity reporting for products that have a required license capacity value that is not based on MSU metrics (non-MSU metrics).

The SCRT support for products that report a non-MSU metric differs from the support for products that report a standard sub-capacity value for MLC or IPLA products. For products whose usage is not based on MSUs, SCRT reports a metric that is based on the characteristics of the specific product and the associated value unit exhibit.

The following topics describe the reporting considerations for these unique metrics and the differences in a sub-capacity report that contains a non-MSU-based product.

Non-MSU metric types

This topic describes each non-MSU metric type and how SCRT determines the peak value to report for the products that use each metric type.

To understand the specific pricing details for a product, review the announcement letter for that product.

Simultaneous instances

For products that report on the *simultaneous instances* metric, SCRT determines the required license capacity by examining, for each hour in the reporting period, the number of instances reported by LPAR. SCRT then cross references the product execution by LPAR to determine the maximum concurrent LPAR instance count—the highest combined instance count of LPARs in which each product executes during the reporting period.

To determine the simultaneous instance count for a given product, SCRT evaluates the instance count as reported in the SMF89T2NumInstances field of the SMF type 89 subtype 2 records.

For more information about what an instance represents for a specific product, see the announcement letter for that product.

Number of CPCs

For products that are licensed at the CPC level, SCRT determines the required license capacity by identifying whether or not the product was active on the CPC at any point during the reporting period. For a standard sub-capacity report, these products will report 1 unit to indicate that the product was active on the CPC. For Country Multiplex Pricing reports or Enterprise Tailored Fit reports, these products will report the number of CPCs on which the product was active.

In most cases, SCRT determines a product to be active based on the instance count as reported in the SMF89T2NumInstances field of the SMF type 89 subtype 2 records or by the total combined TCB and SRB time as reported in the SMF89UCT and SMF89USR fields of the SMF type 89 subtype 2 records. However, some products are treated as active if they are reported as enabled in the IFAPRDxx member of parmlib as reported in the SMF89T2FLAGS field of the SMF type 89 subtype 2 records.

SCRT externals for reporting of non-MSU metrics

SCRT automatically applies non-MSU product-specific rules to all of the non-MSU-based products that it finds in the SMF records. In addition to reporting the appropriate value (the non-MSU metric) for each eligible product, SCRT also generates the following information in the sub-capacity report:

- A non-MSU product header in report section E5

- Footnotes that indicate the product's non-MSU metric type

IPLA non-MSU-based product header

When SCRT encounters any non-MSU-based IPLA product, the sub-capacity report includes a new product subsection with the following header within the Product Summary Information section:

IPLA Product Name - non-MSU based

The subsection contains the product name, IPLA product ID, tool units, and a footnotes column. For details about this subsection, see [“Sample sub-capacity report for non-MSU-based products” on page 372](#).

Note: The inclusion of "IPLA" in the product header is historical. Not all products that might be reported under this header are IPLA products. Review the announcement materials for each product to understand the terms specific to each product.

Non-MSU-based product footnote

When SCRT reports a non-MSU-based product under the IPLA non-MSU-based product header in report section E5, SCRT also posts a metric-specific footnote in the **Footnotes** column.

SCRT also generates a footnote subsection immediately after the last IPLA non-MSU-based product. The footnote subsection contains a footnote header row followed by a description of the footnote.

SCRT reports a different footnote for each non-MSU metric type, as listed in [Table 54 on page 372](#).

Table 54. Non-MSU metric types and footnotes

Metric type	Footnote
Simultaneous instances	(n1)
Product is enabled on this CPC	(n2)

Sample sub-capacity report for non-MSU-based products

The sub-capacity report for non-MSU-based products has the same basic format and report sections as the sub-capacity report for a base z/OS system, with the addition of an E5 subsection.

This topic describes the differences between a sub-capacity report for a base z/OS system and a sub-capacity report that includes non-MSU-based products. For complete descriptions of all of the fields that appear in a standard sub-capacity report, see [“Sample sub-capacity report for single-version z/OS native systems” on page 105](#).

Note that only those sections that differ from a standard sub-capacity report are described here.

Billing-related section

The billing-related section of the sub-capacity report displays the customer information, tool information, special conditions, product summary information, and detail data collection information, which are all described in [“Billing-related section” on page 105](#).

There is no difference in the billing-related section when non-MSU-based products are reported.

Product Summary Information section

The IPLA non-MSU-based product header appears after the standard IPLA section whenever SCRT reports on a non-MSU-based product.

There are no differences between the **MLC Product Name**, **MLC Product ID**, **Customer Comment**, **IPLA Product Name**, and **IPLA Product ID** fields on a basic z/OS sub-capacity report and a sub-capacity report that includes non-MSU-based products.

The reporting of non-MSU-based products impacts the following fields:

Tool Units

SCRT reports the non-MSU-based product metric in the **Tool Units** column. The value in this column represents the peak product value for the reporting period, as defined in [“Non-MSU metric types”](#) on page 371.

Footnotes

SCRT posts a footnote in this column to indicate the non-MSU metric for each product listed in the IPLA non-MSU-based product subsection. For a list of footnotes and their associated non-MSU metrics, see [“Non-MSU-based product footnote”](#) on page 372.

Figure 228 on page 373 shows a sample Product Summary Information section that includes a non-MSU-based product.

```

==E5=====
PRODUCT SUMMARY INFORMATION

MLC Product Name                MLC Product ID    Tool MSUs    Footnotes
z/OS V1                        5694-A01          500
DB2 11 for z/OS                5615-DB2          500
CICS TS for z/OS V5            5655-Y04          372    (viii)
MQSeries for z/OS V6           5655-L82          500
IMS V13                        5635-A04          372    (viii)

IPLA Product Name              IPLA Product ID    Tool MSUs    Footnotes
Fault Analyzer for z/OS V12     5655-W69          500
File Manager for z/OS V12       5655-W68          500
WebSphere Service Registry and  5655-WBS          1    (iv)
Portal Enable for z/OS V8       5655-Y16          1    (iv)
IBM WebSphere Application Server 5655-N02         160    (iv)
for z/OS V7                     5655-Y02          1    (iv)
IBM Business Process Manager for 5655-W96          1    (iv)
z/OS V8

IPLA Product Name - non-MSU based  IPLA Product ID    Tool Units    Footnotes
IBM z/OS Connect EE V2          5655-CEE          4    (n1)

Footnotes:
(iv) Product's Tool MSUs are based on Getting Started Sub-capacity Pricing from one or more LPARs
(viii) Product's Tool MSUs are based on Integrated Workload Pricing from one or more LPARs
(n1) Product Unit is Simultaneous Instance

```

Figure 228. Sample sub-capacity report with non-MSU-based products: Product Summary Information section

In the sample report, IBM z/OS Connect EE V2 is used as the non-MSU-based product, and it appears in the new IPLA non-MSU-based product subsection. In this example, the peak number of simultaneous instances (the metric type indicated by footnote (n1)) for IBM z/OS Connect EE V2 is 4, as listed in the **Tool Units** column.

Customer verification section

The customer verification section of the sub-capacity report contains information that supports and expands on the information in the Product Summary section.

There are no changes to the customer verification section of the sub-capacity report when non-MSU-based products are reported.

The Product Max Contributors section reports the non-MSU-based product. There are no changes to the section header or column headers, but the value that is reported for the product is the product unit (based on the metric type), as reported in the Product Summary Information section. The LPAR columns reflect the product unit on each LPAR during the peak hour.

Figure 229 on page 374 shows a sample Product Max Contributors section that includes a non-MSU-based product.

```
==P5=====
```

PRODUCT MAX CONTRIBUTORS

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1
z/OS V1	5694-A01	500	22 May 2024 - 13:00	500
DB2 11 for z/OS	5615-DB2	500	22 May 2024 - 13:00	500
CICS TS for z/OS V5	5655-Y04	372	22 May 2024 - 16:00	372
MQSeries for z/OS V6	5655-L82	500	22 May 2024 - 13:00	500
IMS V13	5635-A04	372	22 May 2024 - 16:00	372
Fault Analyzer for z/OS V12	5655-W69	500	22 May 2024 - 13:00	500
File Manager for z/OS V12	5655-W68	500	22 May 2024 - 13:00	500
WebSphere Service Registry and Repository for z/OS V8	5655-WBS	1	22 May 2024 - 14:00	1
WebSphere Portal Enable for z/OS V8	5655-Y16	1	22 May 2024 - 13:00	1
IBM WebSphere Application Server for z/OS V7	5655-N02	160	22 May 2024 - 16:00	160
IBM Business Process Manager for z/OS V8	5655-Y02	1	22 May 2024 - 13:00	1
IBM z/OS Connect EE V2	5655-CEE	4	22 May 2024 - 13:00	4

Figure 229. Sample sub-capacity report with non-MSU-based products: Product Max Contributors section

In Figure 229 on page 374, the peak simultaneous instance count for IBM z/OS Connect EE V2 is 4 instances. LPAR1 is the only LPAR running in the example environment and, thus, contributed 4 instances to the product peak.

There are no differences in the Product Grid Snapshot section, the Sub-Capacity Capping Enforced by SCRT section, the Group Capacity LPARs section, or any of the other remaining report sections.

Chapter 19. Reporting for Tailored Fit Pricing for IBM Z

Tailored Fit Pricing for IBM Z (Tailored Fit Pricing) provides simplified software pricing for qualified solutions and combines flexible deployment options with competitive economics that are directly relevant to those solutions.

Tailored Fit Pricing can scale from collocated solutions within existing LPARs, through separate LPARs, up to multi-LPAR solutions, without directly impacting the cost of unrelated workloads.

Tailored Fit Pricing is a framework that IBM solutions can use; it is not linked to a single, announced offering or solution. Each container-eligible solution has its own set of terms and conditions that apply to it.

Overview of Tailored Fit Pricing for IBM Z

Tailored Fit Pricing for IBM Z provides a mechanism for a qualified solution to be isolated from directly impacting the cost of traditional workloads in a customer environment, whether collocated on an existing LPAR or located on separate LPARs.

The Tailored Fit Pricing framework allows you to deploy a solution where you want it based on the best technical fit.

SCRT provides isolated reporting for each Tailored Fit Pricing solution that it finds on a CPC or in a multiplex. Each container reports a metric appropriate to the solution type associated with the container, such as peak 4-hour rolling average utilization or total MSU consumption. There is no concept of product sub-capacity tracking within a container.

For a new, collocated solution, z/OS provides enhanced tracking capability using tenant resource groups (TRGs) in z/OS Workload Management (WLM) to meter and report on a specific solution. This process requires a one-time setup step to define and classify the solution in WLM.

After this one-time setup step, z/OS meters and reports on the workload monthly. There is no requirement to generate INPUTCSV files or to otherwise modify your monthly reporting process. Solution-specific data is captured in enhanced SMF type 70 subtype 1 and type 89 subtype 1 and 2 records. SCRT analyzes the SMF data and removes the impact of any container TRGs from the LPARs on which they ran.

For a separate LPAR solution, a SPECIAL DD command associates an LPAR with a Tailored Fit Pricing solution. During SCRT processing, the named LPARs are removed from the traditional non-container environment. These LPARs do not contribute to product sub-capacity values.

For both collocated and separate LPAR solutions, Tailored Fit Pricing solutions that are found to be running in your environment are reported in a separate section of the SCRT report, isolated from the traditional sub-capacity environment.

When used in conjunction with Country Multiplex Pricing, SCRT provides aggregation of a container across multiple CPCs.

Note: Tailored Fit Pricing for IBM Z is only supported for z/OS systems.

For more information about Tailored Fit Pricing for IBM Z and announced solutions, see "Tailored fit pricing options" at [Pricing for IBM Z \(www.ibm.com/z/pricing\)](http://www.ibm.com/z/pricing).

Requirements for Tailored Fit Pricing for IBM Z

The following conditions must be met before IBM can accept and process sub-capacity reports in which Tailored Fit Pricing for IBM Z solutions are reported. Some of these conditions are further described in more detail in other topics but are consolidated here for easy reference. These requirements are in addition to any specific requirements that are stated in a Tailored Fit Pricing solution announcement.

- An IBM z17, IBM z16, IBM z15, IBM z14, IBM z13s®, or IBM z13 CPC that has implemented sub-capacity pricing.
- z/OS V2.2, or later, with the PTFs applied for the following APARs:
 - OA52312 (WLM)
 - OA52694 (RMF)
 - OA53033 (SMF)
 - PI82528 (SDSF)

These and any other related PTFs are associated with the **IBM.Function.PricingInfrastructure** fix category.

- Sub-Capacity Reporting Tool (SCRT) V25.2, or later.
- SCRT is run for each sub-capacity reporting period and the resulting sub-capacity report is submitted to IBM on a monthly basis.
- A unique solution ID is used to identify the Tailored Fit Pricing solution. IBM will provide you with the solution ID through the License Management Support (LMS) website upon approval of the solution. (See [License Management Support \(www.ibm.com/software/lms\)](http://www.ibm.com/software/lms).)
 - For solutions deployed on an existing sub-capacity LPAR (collocated solutions), you must associate the solution with the TRGs that define the qualified solution. Only work units (that is, address spaces and enclaves) that are associated with the qualified solution are to be classified as part of the TRG.
 - For solutions deployed on separate LPARs, you must use either of the following methods to associate the solution ID with the LPAR or system:
 - Use the **CONTAINER** command to specify only those LPARs that are part of the qualified solution. The **CONTAINER** command associates the LPARs with the solution ID.
 - Specify the **SOLUT** system parameter when IPLing a z/OS system as part of the qualified solution. The **SOLUT** system parameter associates the system with the solution ID. This option is available with the appropriate level of z/OS service.
- A solution name is used to identify the Tailored Fit Pricing solution. IBM will provide you with the solution ID through the LMS website upon approval of the solution.
 - The solution name that you use for a given container must be the one that is associated with the unique solution ID in LMS.
 - You can update the solution name in LMS. However, at the time of report submission, the solution name must match the name as defined in LMS for the container (with the associated solution ID).

Format of a Tailored Fit Pricing solution ID

The solution ID is a 64-character string (included embedded hyphens) and is intended to be an opaque token, but it does, in most circumstances, encode certain information that can be useful in understanding the nature of the corresponding solution.

The following character positions within the solution ID indicate information about the associated solution:

Character position	Description
1	If the character in this position is D , the solution ID is a sample and is not valid for submission to IBM.
29	The character in this position indicates the associated metric type: <ul style="list-style-type: none"> S: The associated solution is based on a consumption metric. T: The associated solution is based on a traditional 4HRA metric.

Character position	Description
30 - 31	The characters in these positions indicate the associated Tailored Fit Pricing for IBM Z offering type: DT: Application Development and Test Solution EC: Enterprise Consumption Solution NA: New Application Solution ZZ: Other or undefined

Example: Consider the following solution ID:

D911111-N33B4FD-449E85ACF3-NSECZZZZ-D127-4ADE-BEF9-E16E53-5E52D3

The values in the previously described character positions indicate that this is a sample solution ID for an Enterprise Consumption Solution which is based on a consumption metric.

Interactions with traditional sub-capacity offerings

Tailored Fit Pricing for IBM Z provides isolation for qualified solutions. While each solution has specific terms and conditions, there are a set of basic interactions and rules that dictate the impact of a container on your traditional sub-capacity environment. In general, within a container, IBM pricing offerings such as Mobile Workload Pricing, do not apply.

Collocated solutions

SMF type 70 subtype 1 records provide SCRT with a summary view of all TRGs in your environment. Only TRGs with a valid solution ID in the SMF records are evaluated for the Tailored Fit Pricing algorithm.

For each LPAR in the environment, on a per-hour basis, SCRT aggregates the weighted 4-hour rolling average of each TRG that has a valid solution ID. When determining the 4-hour rolling average for a TRG for a given hour, the duration of the LPAR is used for the hour. One hundred percent of this aggregate 4-hour rolling average is removed from the standard sub-capacity (SMF70LAC-based) value for the LPAR on a per-hour basis.

Additionally, SCRT uses enhanced TRG tracking via SMF type 89 subtype 1 and 2 records to adjust product usage and instance values, respectively. This adjustment occurs only for products that are executing as part of a TRG that can be associated with a Tailored Fit Pricing solution.

For Getting Started Sub-Capacity Pricing (GSSP) eligible products and IBM Z Collocated Application Pricing (zCAP) net-new defining programs, on a per-hour basis, SCRT aggregates the general-purpose processor time for each TRG that has a valid solution ID, and removes it from the product-level LPAR CPU time values. This has the effect of lowering the CPU time that is used to calculate the product's GSSP or zCAP usage value.

For zCAP, Mobile Workload Pricing (MWP), and IBM Z Workload Pricing for Cloud (zWPC) INPUTCSV processing, SCRT similarly adjusts the general-purpose processor LPAR CPU time. This serves to decrease the maximum allowed INPUTCSV time for the defining program, as appropriate.

For instance based products, such as z/OS Connect EE, on a per-hour basis, SCRT aggregates the instances of the product running in TRGs with a valid solution ID and removes them from the product-level LPAR instance values.

For all products, both SMF type 89 subtype 1 CPU time and subtype 2 instance counts are evaluated on a per-hour basis to determine whether products executing in a container TRG are also active in the traditional, non-container environment. In cases where a product did not accrue any CPU time or instances outside of a container TRG, SCRT considers that product to be inactive for the hour.

Separate LPAR solutions

Tailored Fit Pricing for IBM Z allows for an entire LPAR to be selected as part of a solution. The solution ID and applicable systems are provided to SCRT via one of more of the following options:

- If a solution ID is provided on the **SOLUT** system parameter when a z/OS system is IPLed, the solution ID is included in the SMF type 89 records from that system. SCRT automatically associates the system with the corresponding solution. Solution IDs assigned in this manner are applied on an hour-by-hour basis. A system may move between containers (or out of a container) during the reporting period.
- When a **CONTAINER** control statement is specified, SCRT associates the specified LPARs with the specified solution. A **CONTAINER** control statement applies to the entire reporting period. If a **CONTAINER** statement is supplied for a system that also has a solution ID specified via the **SOLUT** system parameter, the solution ID specified on the **CONTAINER** statement takes precedence and the system parameter value is discarded.

When an LPAR or system is associated with a solution ID, SCRT removes the named LPAR from the standard sub-capacity configuration. These systems do not contribute to any traditional product sub-capacity values for the reporting period and, therefore, are removed from aggregation.

SCRT supports the execution of co-located Tailored Fit Pricing for IBM Z solutions on LPARs that are part of a separate LPAR solution (that is, a container within a container). When a co-located solution runs on an LPAR that is associated with a different solution, SCRT removes the impact of the co-located solution from the solution associated with the LPAR.

SCRT externals for Tailored Fit Pricing for IBM Z

SCRT provides the following externals for Tailored Fit Pricing for IBM Z:

- The optional **CONTAINER** command (control statement) lets you to specify the LPARs that are part of a qualified container solution.
- The **UPDATE CONTAINER** command (control statement) lets you to specify the required solution name for each container.
- The optional **UPDATE SOLUTION ID** command (control statement) lets you to substitute a new or corrected solution ID found in the SMF data.
- The **IGNORE CONTAINER** command (control statement) lets you specify that container processing is to be ignored for a specific container for the indicated time interval.
- The optional **CONTAINER_INTERVAL_CONSUMPTION** command (control statement) lets you see MSU consumption on a per-container basis with 5-minute granularity.
- The **CPSTRACE DD** statement lets you specify an output data set or file to contain detailed information about container solution execution.

CONTAINER control statement

In cases where an LPAR is dedicated to a qualified Tailored Fit Pricing solution, use the **CONTAINER** control statement (or the appropriate SCRT GUI panel) to associate the LPAR with a container solution ID.

IBM authorizes the use of the **CONTAINER** control statement only in cases where you have signed or otherwise agreed to the terms and conditions of a Tailored Fit Pricing solution. Any LPARs that are part of a Tailored Fit Pricing solution must be dedicated to the solution or workload indicated as part of the agreement. The solution ID that you specify on the **CONTAINER** statement must match the solution ID provided on the IBM License Management Support (LMS) website for the qualified workload.

For the format of the **CONTAINER** command, see [“CONTAINER control statement” on page 41](#).

UPDATE CONTAINER control statement

Use the **UPDATE CONTAINER** control statement to associate the LMS-provided solution name with a Tailored Fit Pricing solution based on the associated solution ID and to designate NO89 products that run in the container. The **UPDATE CONTAINER** statement is required for Tailored Fit Pricing solutions.

Each Tailored Fit Pricing solution, whether found in SMF data or specified using the **CONTAINER** control statement, must have a solution name associated with it via an **UPDATE CONTAINER** statement. The name specified on the **UPDATE CONTAINER** statement must match the name specified for the container (as identified by the unique solution ID) in the LMS web application.

For the format of the **UPDATE CONTAINER** command, see [“UPDATE CONTAINER control statement” on page 59.](#)

UPDATE SOLUTION_ID control statement

Use the optional **UPDATE SOLUTION_ID** control statement to replace a solution ID associated with a container TRG with a new or corrected solution ID. The replacement solution ID must correspond to the qualified workload that ran in the container.

The **UPDATE SOLUTION_ID** control statement is only intended for use under rare circumstances, such as when a TRG is defined with a solution ID that is incorrect or not applicable.

For the format of the **UPDATE SOLUTION_ID** command, see [“UPDATE SOLUTION_ID control statement” on page 59.](#)

IGNORE CONTAINER control statement

Use the optional **IGNORE CONTAINER** control statement to request that SCRT ignore a Tailored Fit Pricing solution for a specified time interval.

The **IGNORE CONTAINER** control statement is only intended for use under rare circumstances and as directed by IBM.

For the interval specified on an **IGNORE CONTAINER** control statement, SCRT does not apply the Tailored Fit Pricing algorithm and considers any Tailored Fit Pricing tenant resource groups (TRGs) and dedicated LPARs that are associated with the specified solution ID to be part of the standard sub-capacity environment. Specifically:

- Tailored Fit Pricing LPARs are aggregated into standard sub-capacity product reporting.
- The 4-hour rolling average contribution and product data for the Tailored Fit Pricing TRG are not removed from the product values for the LPAR.
- Container peak calculations ignore the specified time range.

For the format of the **IGNORE CONTAINER** command, see [“IGNORE CONTAINER control statement” on page 53.](#)

CONTAINER_INTERVAL_CONSUMPTION control statement

Use the **CONTAINER_INTERVAL_CONSUMPTION** control statement to see MSU consumption on a per-container basis with 5-minute granularity. This control statement is optional and should be used only at the direction and authorization of IBM. An SMF interval length of 5 minutes is required for all systems that are part of the container.

CPSTRACE DD statement

You can use the optional CPSTRACE DD statement to specify a data set or file that is to contain detailed, hour-by-hour tracking for Tailored Fit Pricing reporting. The trace output is organized as comma-separated fields to allow the file to be read by a spreadsheet application.

The CPSTRACE output contains a section for each container. Each container section begins with a header record, followed by an hour-by-hour view of the TRGs and dedicated LPARs that contributed to the container's 4-hour rolling average utilization for each hour.

- When using the default SCRT processing and multiple CPCs are found in the input SMF data, the CPSTRACE output displays a Tailored Fit Pricing solution that spans multiple CPCs on a per CPC basis. The default SCRT processing does not support aggregation across CPCs.

- When using Country Multiplex Pricing, the CPSTRACE output for a single container spans each CPC on which the container is found.

The CPSTRACE output is divided into two sub-reports labeled CPSTRACE 1 and CPSTRACE 2. Each sub-report contains a section for each container.

- The CPSTRACE 1 sub-report output contains a section for each container. Each container section begins with a header record followed by an hour-by-hour view of the TRGs and dedicated LPARs that contributed to the container's 4-hour rolling average utilization for each hour.
- The CPSTRACE 2 sub-report output contains a section for each container. Each container section begins with a header record followed by an hour-by-hour view of the TRGs and dedicated LPARs that contributed to the container's MSU consumption for each hour.

Figure 230 on page 380 shows an example of the format of the CPSTRACE output.

```

==CPSTRACE=====
Container Pricing Detailed Data

==CPSTRACE 1=====
Four Hour Rolling Average

CPS1  DevTest Solution  Z111111-N31BB29-8FC80FDC07-NSDTZZZZ-91CC-465B-98CA-0565FF-E42B4F
CPS1  Date Time        Processor Partition - LPAR1 LPAR2 CPC Container
CPS1  (tttt-sssss) TRG - (lpar) TRG1 Total Total
CPS1  02 Aug 2018 - 00:00 3906-12345 1450 49 1449 1449
CPS1  02 Aug 2018 - 01:00 3906-12345 2283 251 2534 2534
CPS1  02 Aug 2018 - 02:00 3906-12345 2638 275 2913 2913
CPS1  02 Aug 2018 - 03:00 3906-12345 2795 531 3326 3326
CPS1  02 Aug 2018 - 04:00 3906-12345 1976 841 2817 2817
:

==CPSTRACE 2=====
MSU Consumption

CPS1  DevTest Solution  Z111111-N31BB29-8FC80FDC07-NSDTZZZZ-91CC-465B-98CA-0565FF-E42B4F
CPS1  Date Time        Processor Partition - LPAR1 LPAR2 CPC Container
CPS1  (tttt-sssss) TRG - (lpar) TRG1 Total Total
CPS1  02 Aug 2018 - 00:00 3906-12345 5803 197 6000 6000
CPS1  02 Aug 2018 - 01:00 3906-12345 3330 810 4140 4140
CPS1  02 Aug 2018 - 02:00 3906-12345 1423 93 1516 1516
CPS1  02 Aug 2018 - 03:00 3906-12345 627 1024 1651 1651
CPS1  02 Aug 2018 - 04:00 3906-12345 2529 1437 3966 3966
:

```

Figure 230. Example of CPSTRACE output (formatted)

Each row for a specific Tailored Fit Pricing solution (CPS) begins with a solution token assigned by SCRT. Each CPS token begins with CPS followed by a number (for instance, CPS1), and corresponds to the CPS token for that container (as indicated by the unique solution ID) in the SCRT report for a CPC.

The first CPS_n row contains the unique IBM-assigned solution ID for the Tailored Fit Pricing solution.

The following fields appear in the header record:

Date Time

The day, month, year, and time corresponding to the data presented on the row of the CPSTRACE output.

Processor

The machine type (tttt) and serial number (sssss) of the CPC on which the Tailored Fit Pricing solution ran.

Partition

The LPARs on which the Tailored Fit Pricing solution ran, either in one or more TRGs or in one or more dedicated LPARs or both.

TRG

The name assigned to the tenant resource group in z/OS WLM, or (lpar) for a dedicated LPAR (assigned using the CONTAINER command). If a dedicated LPAR has co-located workloads for other solutions, the hourly values reflect the net contribution of the LPAR to the container, already reduced by co-located solutions.

CPC Total

The total 4-hour rolling average utilization or MSU consumption for the container on the CPC for the hour.

Container Total

The total 4-hour rolling average utilization or MSU consumption for the container for the hour. If Country Multiplex Pricing is in use, this value includes the contributions from all CPCs on which the container ran.

If Country Multiplex Pricing is in use, each CPC on which the container was active is displayed on the hour row.

Tailored Fit Pricing for IBM Z and ISV reporting

By default, when ISV reporting is active—that is, the ISVLIB DD statement contains valid input and SCRT is processing a report for a specific ISV—SCRT ignores any containers that it finds in your environment.

You can enable container processing for ISV reporting by using the **ISV_SHOW_CONTAINERS** command (control statement). For more information, see [Chapter 23, “Reporting for independent software vendor programs,”](#) on page 413.

Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z

The following topics provides a detailed example of Tailored Fit Pricing for IBM Z using collocated Tailored Fit Pricing solutions in a standard sub-capacity environment.

Sample configuration for a collocated solution for Tailored Fit Pricing

The sample sub-capacity report for Tailored Fit Pricing for IBM Z is based on the following hardware and software configuration.

The sample configuration consists of a 2964 Model 7A1 CPC with a serial number of 02-01234. The CPC is configured with 4 LPARs, all running z/OS natively. The z/OS systems in SYS1, SYS2, SYS4, and SYS5 have z/OS SYSIDs of SYS1, SYS2, SYS4, and SYS5, respectively. Each LPAR is running a set of container tenant resource groups.

The z/OS systems are running a combination of the following z/OS products:

- Db2 11 for z/OS
- IBM MQ for z/OS V8
- CICS TS for z/OS V5*

*In this example, CICS TS for z/OS V5 is only active in container tenant resource groups.

Sample control statement input for a collocated solution for Tailored Fit Pricing

The sample sub-capacity report for Tailored Fit Pricing for IBM Z is based on the following control statement input to SCRT.

The following SPECIAL DD control statement was specified in the JCL for the SCRT job. You can request similar function when using SCRT for Windows and Linux.

```
UPDATE CONTAINER,
ID=D911111-N4ED9AD-DB3D215CAF-NTZZZZZZ-07F5-4EDF-AF34-E062EB-0ADCD7,
SET_NAME="ContainerName",
N089_PRODUCT_ID=(5698-BMP)
```

The **UPDATE CONTAINER** control statement is required for each solution ID that is found in the input SMF data or that is assigned using the **CONTAINER** control statement. The **UPDATE CONTAINER** control statement associates a container (represented by its unique solution ID) with a solution name.

The solution name associated with the container is assigned via the IBM LMS website and is available upon approval of the solution. You can use LMS to modify the solution name as you prefer; however, the solution name that you specify to SCRT must match the one defined in LMS.

Sample sub-capacity report for a collocated solution for Tailored Fit Pricing

A sub-capacity report for Tailored Fit Pricing for IBM Z has the same basic format and report sections as a sub-capacity report for a base z/OS system. However, a sub-capacity report with Tailored Fit Pricing solutions contains additional report sections that provide an overview of active containers, as well as container-specific report sections in both the billing-related and customer verification sections of the report.

The following sample report highlights the differences between a sub-capacity report for a base z/OS system and one that contains Tailored Fit Pricing solutions.

Billing-related section

The billing-related section of a sub-capacity report for Tailored Fit Pricing for IBM Z displays the following report sections:

- Customer information
- Tool Information
- Special Conditions
- Product Summary Information
- Missing LPAR Data Details
- Missing CPC Data Details
- Active Containers
- Container Billing Summary for each container
- Container Active Products for each container
- Container Product Metrics

There are no differences between the customer information, Tool Information, Special Conditions, Product Summary Information, Missing LPAR Data Details, and Missing CPC Data Details sections for a basic z/OS sub-capacity report and one for Tailored Fit Pricing, which are fully described in [“Billing-related section” on page 105](#).

Figure 231 on page 383 shows the customer information and Tool Information sections of the sample sub-capacity report for Tailored Fit Pricing.

```
==B5===== SCRT SUB-CAPACITY MVM REPORT - IBM Corp =====
```

```
Run Date/Time           06 Oct 2024 - 10:34
Name of Person Submitting Report:  Jane Customer
E-Mail Address of Report Submitter: jane@abc.com
Phone Number of Report Submitter:  8885550123
```

```
Customer Name           ABC Corp
Customer Number          US-S000000000
Machine Serial Number    02-01234
Machine Type and Model   2964-7A1
Machine Rated Capacity (MSUs) 10171
Purchase Order Number    (optional)
Customer Comments (255 chars max) (optional)
```

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

```
==C5=====
TOOL INFORMATION
```

```
Tool Release           30.1.0
Reporting Period        2 Sep, 2024 - 1 Oct, 2024 inclusive (30 days)
```

Figure 231. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Customer information and Tool Information sections

There is no Special Conditions section for this sample report.

Figure 232 on page 384 shows the Product Summary Information, Missing LPAR Data Details, and Missing CPC Data Details sections of the sample sub-capacity report for Tailored Fit Pricing.

==E5=====

PRODUCT SUMMARY INFORMATION

MLC Product Name	MLC Product ID	Tool MSUs
z/OS V2	5650-Z0S	137
DB2 11 for z/OS	5615-DB2	137
IBM MQ for z/OS V8	5655-W97	137

==H4=====

MISSING LPAR DATA DETAILS

Active LPARs with no SMF/SCRT89 records	Detected	Resolved (missing hours)	Est Peak MSU	Justification for low data collection (255 chars max)			
C08	25 Sep 2024 - 11:00	25 Sep 2024 - 16:00 (5 hours)	50	(required)			
C0A	25 Sep 2024 - 11:00	25 Sep 2024 - 16:00 (5 hours)	60	(required)			
C0B	25 Sep 2024 - 11:00	25 Sep 2024 - 16:00 (5 hours)	70	(required)			
C0D	25 Sep 2024 - 11:00	25 Sep 2024 - 16:00 (5 hours)	80	(required)			
Total Missing Hours		20					
OS Product Name	OS Product ID	Est Peak MSU	Date/Time	C08	C0A	C0B	C0D
z/OS V2	5650-ZOS	260	25 Sep 2024 - 11:00	50	60	70	80

==H6=====

MISSING CPC DATA DETAILS

CPC with no SMF/SCRT89 data from any LPAR	Detected	Resolved (missing hours)	Justification for low data collection (255 chars max)
2964-12345	02 Sep 2024 - 00:00	25 Sep 2024 - 11:00 (563 hours)	(required)
2964-12345	25 Sep 2024 - 16:00	02 Oct 2024 - 00:00 (152 hours)	(required)
Total Missing Hours		715	

Figure 232. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Product Summary Information, Missing LPAR Data Details, and Missing CPC Data Details sections

In the Product Summary Information section, the values listed in the **Tool MSUs** column represent the peak sub-capacity values that SCRT calculated for each product using the appropriate sub-capacity rules for that product. SCRT uses the Tailored Fit Pricing for IBM Z algorithm and SMF data to adjust the 4-hour rolling average sub-capacity MSU values for the sub-capacity eligible programs on a given CPC.

See “Product Max Contributors section” on page 389 for a deeper discussion about how the Tailored Fit Pricing for IBM Z algorithm affects the sub-capacity eligible programs that are listed in the Product Summary Information section. For general information about the Tool MSUs values, see “Product Summary Information section” on page 110.

The Missing LPAR Data Details section includes missing data for LPARs that are configured as part of a Tailored Fit Pricing solution using the **CONTAINER** control statement. As with non-container LPARs, 100 percent of all SMF data is required for Tailored Fit Pricing for IBM Z LPARs.

Active Containers section

The optional Active Containers section appears in the sub-capacity report because SCRT found one or more solution IDs in the configuration in one of the following ways:

- In SMF type 70 subtype 1 records associated with a tenant resource group
- Assigned using the **CONTAINER** control statement and associated with one or more dedicated LPARs

The Active Containers section provides a cross-reference of container identifiers (assigned by SCRT) to solution IDs.

In the sample configuration, there is one Tailored Fit Pricing solution active on the CPC. SCRT assigned the CPS1 container identifier to the container with solution ID Z894E15-F5F120B-D905DF9E5B-C5389F64-8E79-4B8E-81D6-ABC6A5-81B989.

The format of the SCRT container identifier always begins with the CPS prefix. SCRT then adds a numeric suffix for each container that it finds in the configuration (as identified by a unique solution ID). The SCRT container identifier is then used to identify the container in other sections of the sub-capacity report.

Figure 233 on page 385 shows the Active Containers section for the sample configuration.

```
==K5=====
ACTIVE CONTAINERS

SCRT Container Identifier      Solution ID
CPS1                         Z894E15-F5F120B-D905DF9E5B-C5389F64-8E79-4B8E-81D6-ABC6A5-81B989
```

Figure 233. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Active Containers section

Container Billing Summary section for container CPS1

SCRT generates a Container Billing Summary section for each Tailored Fit Pricing solution that it finds in the environment.

The header of each Container Billing Summary section contains the container ID assigned by SCRT within the header itself. For instance, the Container Billing Summary section for Tailored Fit Pricing solution CPS1 contains CPS1 within the header, as shown in Figure 234 on page 385.

```
==CPS1=====
Solution ID                Z894E15-F5F120B-D905DF9E5B-C5389F64-8E79-4B8E-81D6-ABC6A5-81B989
Solution Name              CICS-only Container 1
Peak Four Hour Rolling Average      14
Total MSU Consumption          255
```

Figure 234. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Billing Summary section

The Container Billing Summary section contains the following fields:

- Solution ID**
The 64-character IBM solution ID that represents the Tailored Fit Pricing for IBM Z qualified solution.
- Solution Name**
The solution name associated with the Tailored Fit Pricing solution in the LMS website.
- Peak Four Hour Rolling Average**
The peak 4-hour rolling average utilization for the Tailored Fit Pricing solution. This value is the sum of the 4-hour rolling average utilization for each tenant resource group and LPAR that has specified the container solution ID for the concurrent peak hour.
- For each standalone LPAR that is part of the container, the value is based on the weighted average SMF70LAC value for the hour.
 - For each tenant resource group that is part of the container, the value is based on the weighted average SMF70_TRG_LAC value for the hour.
- These values are added together to determine the sub-capacity MSU value for the container. The highest value during the reporting period is the displayed peak value.
- Total MSU Consumption**
The total MSUs used (consumed) by the Tailored Fit Pricing solution during the reporting period.
- To calculate this value, SCRT calculates the MSUs consumed for each tenant resource group and standalone LPAR that is part of the Tailored Fit Pricing solution.

- For tenant resource groups, SCRT takes the service units on general purpose processor value reported in the SMF70_TRG_SUCP field of SMF type 70 subtype 1 records and converts this value to MSUs on a per-hour basis for each applicable tenant resource group.
- For dedicated LPARs, SCRT takes the effective dispatch time of the LPAR reported in the SMF70EDT field of SMF type 70 subtype 1 records and converts this value to MSUs on a per-hour basis for each applicable LPAR.

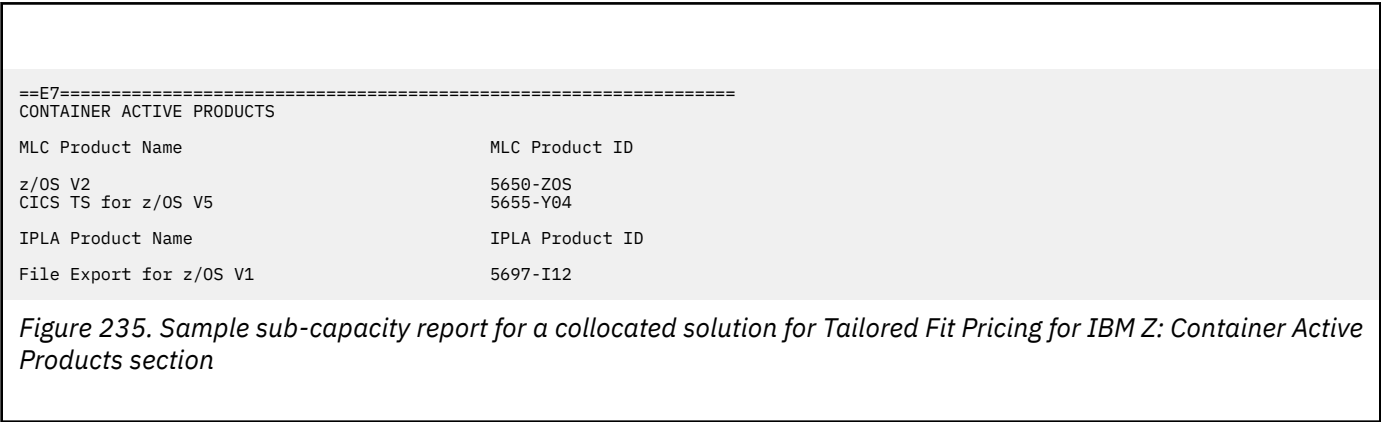
These calculated MSU values are then added together and reported as the total MSU consumption.

Container Active Products section for container CPS1

The Container Active Products section (E7) is similar to the Product Summary Information section of the sub-capacity report. This section lists all products that SCRT found to be running as part of a container tenant resource group or dedicated LPAR during the reporting period. Unlike the Product Summary Information section, product family roots are not reported in this section.

A key difference between the Container Active Products section and the Product Summary Information Section is that there are no tool MSU values reported for container products. This is a fundamental concept of Tailored Fit Pricing: The sub-capacity values of individual products are not reported for a container. Instead, Tailored Fit Pricing solutions are priced according to the container metric, such as the peak 4-hour rolling average utilization of the container or the total MSU consumption of the container.

[Figure 235 on page 386](#) shows the Container Active Products section for Tailored Fit Pricing solution CPS1.



MSU Consumption Data Collection section

The MSU Consumption Data Collection section (H7) tracks the completeness of data collection for Tailored Fit Pricing solutions that are billed based on an MSU consumption metric at sub-hour granularity. For each container, the section reports, on an LPAR-by-LPAR basis, the total possible number of minutes during which the container could have been running and the total number of minutes of MSU consumption data that was actually collected from SMF type 70 records. The difference between these two values is reported as the number of minutes of missing MSU consumption data.

Justification for missing data is required if the number of missing minutes exceeds the IBM-designated threshold.

[“MSU Consumption Data Collection section” on page 386](#) shows an example of the MSU Consumption Data Collection section for solution CPS1.


```

==H7=====
MSU CONSUMPTION DATA COLLECTION

LPAR   Total Possible Minutes   Minutes Collected   Missing Minutes   Justification
D0      3600                      3560                40   (optional)
D2      1440                      1420                20   (optional)
Total                                60   (required)

```

Figure 236. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: MSU Consumption Data Collection section

In “MSU Consumption Data Collection section” on page 386, the container could have been active for 3600 minutes on system D0 and 1440 minutes on system D2. Systems D0 and D2 collected MSU consumption data for 3560 minutes and 1420 minutes, respectively, during the reporting period, resulting in 40 and 20 missing minutes, respectively. A justification for the 60 missing minutes is required on the **Total** line. Optional justifications can be supplied per LPAR in order to provide additional detail.

When possible, SCRT does not report missing minutes that are near an hour in which an IPL or shutdown occurred. SCRT does not report LPARs that are designated to be part of a different container.

z/VM guest images that are not running with RMF Monitor I VMGUEST do not collect the required data to calculate MSU consumption. Guest images that are missing the required configuration are indicated with an (xi) footnote and the following footnote text at the end of the section:

(xi) VM guest image was running without RMF Monitor I VMGUEST data collection active.

When investigating missing time reported in this section, it can be helpful to refer to the Sub-Hour Data Collection Details section (V6), which you can enable by specifying the **GENERATE_DETAILED_DATA** control statement (or the CLI or GUI equivalent).

Container Interval Consumption section

The Container Interval Consumption section shows MSU consumption for the container with 5-minute granularity. This section is optional and is only produced when the **CONTAINER_INTERVAL_CONSUMPTION** control statement has been specified.

```

==I7=====
CONTAINER INTERVAL CONSUMPTION

Period                MSU Consumed   Footnotes
2024-09-02T00:00:00 UTC    21.43
2024-09-02T00:05:00 UTC    17.94
2024-09-02T00:10:00 UTC     7.21
2024-09-02T00:15:00 UTC     7.46

```

Figure 237. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Interval Consumption section

Container Product Metrics section

The Container Products Metrics section tracks products that have product-specific metrics that are not based on MSUs.

A value reported in the **Product Units** column is based on the product-specific algorithm, as applied to each tenant resource group and dedicated LPAR that is part of the Tailored Fit Pricing solution. This behavior is similar to how standard sub-capacity reporting aggregates product values for the metric outside of a container.

Since container CPS1 in the sample scenario does not use any products with a non-MSU metric, a separate example is shown in [Figure 238 on page 388](#).

```

==J7=====
CONTAINER PRODUCT METRICS

Product Name                Product ID   Product Units   Product Metric
IBM z/OS Connect EE V2     5655-CEE           40   Simultaneous Instances

```

Figure 238. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Product Metrics section

In Figure 238 on page 388, z/OS Connect EE V2 has a peak **Product Units** value of 40 simultaneous instances. Therefore, during the peak hour, the sum of all instances reported in tenant resource groups and dedicated LPARs related to the Tailored Fit Pricing solution is 40.

Customer verification section

The customer verification section presents information that supports and expands upon the information in the Product Summary Information section. The customer verification section of a sub-capacity report for Tailored Fit Pricing for IBM Z contains the following report sections:

- SMF / SCRT89 Input Data Statistics
- Detail LPAR Data Section
- Detail LPAR Usage Data
- Product Max Contributors
- Product Grid Snapshot
- Container Verification section for each container
- Container Product Grid Snapshot
- Container Max Contributors

Other sections that are not displayed for the sample configuration may be displayed under different conditions. For information about those sections, see [“Customer verification section ” on page 118](#).

There are no differences between the SMF / SCRT89 Input Data Statistics and Detail LPAR Data sections for a basic z/OS sub-capacity report and one for Tailored Fit Pricing, which are fully described in [“Customer verification section ” on page 118](#).

[Figure 239 on page 389](#) shows the SMF / SCRT89 Input Data Statistics and Detail LPAR Data sections of the sample sub-capacity report for Tailored Fit Pricing.

```
==M5=====
SMF/SCRT89 INPUT DATA STATISTICS
```

	SYSID	Input Data Start	Input Data End
SYS1	SYS1	25 Sep 2024 - 11:00	25 Sep 2024 - 16:00
SYS2	SYS2	25 Sep 2024 - 11:00	25 Sep 2024 - 16:00
SYS4	SYS4	25 Sep 2024 - 11:00	25 Sep 2024 - 16:00
SYS5	SYS5	25 Sep 2024 - 11:00	25 Sep 2024 - 16:00
CPC		25 Sep 2024 - 11:00	25 Sep 2024 - 16:00

```
==N5=====
Detail LPAR Data section
```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time
SYS1	214	1	25 Sep 2024 - 15:00	166	1	25 Sep 2024 - 14:00
SYS2	27	1	25 Sep 2024 - 15:00	25	1	25 Sep 2024 - 14:00
SYS4	28	1	25 Sep 2024 - 15:00	27	1	25 Sep 2024 - 14:00
SYS5	24	2	25 Sep 2024 - 14:00	23	1	25 Sep 2024 - 13:00
CPC	293	1	25 Sep 2024 - 15:00	242	1	25 Sep 2024 - 14:00

Figure 239. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: SMF / SCRT89 Input Data Statistics and Detail LPAR Data sections

Detail LPAR Usage Data section

There are no differences between the Detail LPAR Usage Data section for a basic z/OS sub-capacity report and a sub-capacity report generated using Tailored Fit Pricing for IBM Z. For details about this section, see [“Detail LPAR Usage Data section”](#) on page 120.

Product Max Contributors section

The MSU values that are listed in the Product Max Contributors section are the same for a basic z/OS sub-capacity report and for a report with Tailored Fit Pricing for IBM Z, with one exception: Whenever SCRT finds tenant resource groups for a Tailored Fit Pricing solution in the configuration and all Tailored Fit Pricing conditions have been met, SCRT generates a **Container MSU Reduction** column. For each standard sub-capacity product that is eligible to be reduced, this the **Container MSU Reduction** column reports the calculated reduction that was applied during the peak hour. The reduction only applies to products based on standard sub-capacity values.

For more information, see [“Product Max Contributors section”](#) on page 121.

[Figure 240 on page 389](#) shows the Product Max Contributors section for the sample configuration.

```
==P5=====
PRODUCT MAX CONTRIBUTORS
```

Product Name	Product ID	Highest	Date/Time	LPAR SYS1	LPAR SYS2	LPAR SYS4	LPAR SYS5	Container MSU Reduction
z/OS V2	5650-Z0S	137	25 Sep 2024 - 15:00	64	25	26	22	156
DB2 11 for z/OS	5625-DB2	137	25 Sep 2024 - 15:00	64	25	26	22	156
IBM MQ for z/OS V8	5655-M15	137	25 Sep 2024 - 15:00	64	25	26	22	156

Figure 240. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Product Max Contributors section

Product Grid Snapshot section

There are no differences between the Product Grid Snapshot section for a basic z/OS sub-capacity report and a sub-capacity report generated using Tailored Fit Pricing for IBM Z. For details about this section, see [“Product Grid Snapshot section” on page 123](#).

Figure 241 on page 390 shows the Product Grid Snapshot section for the sample configuration.

```
==Q5=====
PRODUCT GRID SNAPSHOT

Product Name                Product ID   SYS1    SYS2    SYS4    SYS5
z/OS V2                     5650-Z0S    0.60%   0.60%   0.60%   0.60%
DB2 11 for z/OS             5615-DB2    0.60%   0.60%   0.60%   0.60%
```

Figure 241. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Product Grid Snapshot section

Container Verification Section for container CPS1

SCRT generates a Container Verification section for each Tailored Fit Pricing solution that it finds in the environment. For each container, the Container Verification Section displays the Container Product Grid Snapshot section and Container Max Contributors section.

Container Product Grid Snapshot section

The Container Product Grid Snapshot section shows the products that ran in each container tenant resource group or dedicated LPAR for the container identified in the Container Verification section header. This section is similar to the Product Grid Snapshot section for a standard sub-capacity report and shows the percentage of time a product was active during the reporting period in each container tenant resource group or dedicated LPAR that is assigned to the container. (Note that products supported by the SCRT NO89 control statement are not listed in the Container Product Grid Snapshot section because this section is intended to report on the times when products actually ran.)

In the sample configuration, there are four tenant resource groups assigned to container CPS1. The Container Product Grid Snapshot section identifies these tenant resource groups first by the LPAR on which they ran, then by the tenant resource group name.

Figure 242 on page 390 shows the Container Product Grid Snapshot section for container CPS1. In this section, the TGCICS21 tenant resource group ran on four separate LPARs: SYS1, SYS2, SYS4, and SYS5. The percentage of the reporting period that each product was active in the TGCICS21 tenant resource group is reported for each LPAR on which TGCICS21 ran.

```
==Q7=====
Container Product Grid Snapshot

Product Name                Product ID   SYS1    SYS2    SYS4    SYS5
                        TGCICS21  TGCICS21  TGCICS21  TGCICS21
z/OS V2                     5650-Z0S    0.60%   0.60%   0.60%   0.60%
CICS TS for z/OS V5         5655-Y04    0.60%   0.60%   0.60%   0.60%
```

Figure 242. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Product Grid Snapshot section

Container Max Contributors section

The Container Max Contributors section shows the highest observed utilization values for the specific container identified in the Container Verification section header.

The highest hour depends on when the product ran in each of the tenant resource groups that specify the contain solution ID, along with any active, dedicated LPARs on the CPC that are configured with the **CONTAINER** control statement as part of the Tailored Fit Pricing solution.

Figure 243 on page 391 shows the Container Max Contributors section for container CPS1.

```

==T4=====
CONTAINER MAX CONTRIBUTORS

      TRG          Highest      Date/Time          Contribution to Highest
SYS2      TGCICS21                      0
SYS4      TGCICS21                      0
SYS1      TGCICS21                    14
SYS5      TGCICS21                      0
CPS1                      14      25 Sep 2024 - 15:00

```

Figure 243. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Max Contributors section

The Container Max Contributors section contains the following fields:

LPAR name (not labeled)

The name of the LPAR (displayed in the first column) that contributed to the Tailored Fit Pricing solution during the peak hour.

The named LPAR appears here either because a tenant resource group that ran on the LPAR specifies the associated solution ID, or it is a dedicated LPAR that was configured with the **CONTAINER** control statement as part of the Tailored Fit Pricing solution.

TRG

The name of the tenant resource group that ran on the LPAR named in the first column. The TRG is considered part of the Tailored Fit Pricing solution because it is defined with a solution ID.

If an asterisk (*) appears in this field, the LPAR is a dedicated container LPAR so there is no applicable TRG.

Highest

Displayed for a container row, this is the highest combined 4-hour rolling average utilization that was observed for the container, based on all of the dedicated LPARs and TRGs on which the container ran.

Date/Time

The date and time when the highest 4-hour rolling average utilization value was first reached.

Contribution to Highest

The amount of the contribution to the **Highest** value reported for the Tailored Fit Pricing solution, reported for each dedicated LPAR and TRG that on which the container ran.

In Figure 243 on page 391, container CPS1 had a 14 MSU peak on 25 September 2024 at 15:00. Although the TGCICS21 tenant resource group was active on four LPARs, the only recorded contribution came from the TGCICS21 running on LPAR SYS1.

Container MSU Consumption Contributors section

The Container MSU Consumption Contributors section shows the contribution to the total MSU consumption of the solution by each tenant resource group or LPAR.

Figure 244 on page 392 shows an example of the Container MSU Consumption Contributors section for a container named CPS1.

```

=====
==T6=====
CONTAINER MSU CONSUMPTION CONTRIBUTORS
CPS1      Z194E15-aaaaaaaa-aaaaaaaaaaaa-aaaaaaaa-aaaa-aaaa-aaaa-aaaaaa-aaaaaa
          TRG      Contribution to Total Consumed
LPAR1     TRG1      26712
LPAR1     TRG2      48081
LPAR1     TRG3      16027
LPAR2     TRG4      29383
LPAR2     TRG5      48081
LPAR3     *         16027
CPS1      184311

```

Figure 244. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container MSU Consumption Contributors section

The Container MSU Consumption Contributors section contains the following fields:

LPAR name (not labeled)

The name of the LPAR (displayed in the first column) that contributed to the solution's MSU consumption during the reporting period.

The named LPAR appears here either because a tenant resource group that ran on the LPAR specifies the associated solution ID, or it is a dedicated LPAR that was configured with the **CONTAINER** control statement as part of the Tailored Fit Pricing solution.

TRG

The name of the tenant resource group that ran on the LPAR named in the first column. The TRG is considered part of the Tailored Fit Pricing solution because it is defined with a solution ID.

If an asterisk (*) appears in this field, the LPAR is a dedicated container LPAR so there is no applicable TRG.

Contribution to Total Consumed

The amount of the contribution from the tenant resource group or LPAR to the total MSU consumption for the container.

In Figure 244 on page 392, container CPS1 consumed a total of 184311 MSU. LPAR1 had three tenant resource groups, TRG1, TRG2, and TRG3, contributing 26712 MSU, 48081 MSU, and 16027 MSU, respectively. LPAR2 had two tenant resource groups, TRG4 and TRG5, contributing 29383 MSU and 48081 MSU, respectively. LPAR3 is a dedicated container LPAR and contributed 16027 MSU.

Sample multiplex report for a standalone LPAR solution for Tailored Fit Pricing for IBM Z

This topic provides an example of Tailored Fit Pricing for IBM Z using standalone LPARs in a Tailored Fit Pricing solution with Country Multiplex Pricing (CMP) active.

Many of the multiplex report sections for Tailored Fit Pricing are the same as in a standard sub-capacity report with Tailored Fit Pricing active. This topic focuses on the specific differences that impact the multiplex report.

Sample configuration for a standalone solution for Tailored Fit Pricing

The sample multiplex report for Tailored Fit Pricing for IBM Z is based on the following hardware and software configuration.

The sample configuration consists of two CPCs:

- A 2965 Model J03 CPC with a serial number of 51-Z1Z11. The CPC is configured with 1 LPAR (LPAR1) running z/OS natively. The z/OS system in LPAR1 has a z/OS SYSID of SYS1.

- A 3906 Model 797 CPC with a serial number of 51-54321. The CPC is configured with 1 LPAR (LPARA) running z/OS natively. The z/OS system in LPARA has a z/OS SYSID of SYSA.

The z/OS systems are running a combination of the following z/OS products:

- Db2 10 for z/OS
- Db2 11 for z/OS
- WebSphere MQ for z/OS V7
- CICS TS for z/OS V4
- WebSphere Transformation Extended for z/OS V8

In this example, LPAR1 on CPC 51-Z1Z11 is configured to be part of a Tailored Fit Pricing solution.

Sample control statement input for a standalone solution for Tailored Fit Pricing

The sample multiplex report for Tailored Fit Pricing for IBM Z is based on the following control statement input to SCRT.

The following SPECIAL DD control statement was specified in the JCL for the SCRT job. You can request similar function when using SCRT for Windows and Linux.

```
CONTAINER CPC=2965-Z1Z11,IMAGE_ID=LPAR1,
ID=D9F4555-177C847-3DE17A3049-70AD6FC2-B263-4CBD-84EB-AA22CB-6915BB
```

The **CONTAINER** control statement indicates that an LPAR is dedicated to a Tailored Fit Pricing solution, and it associates one or more LPARs or z/VM systems with a container represented by its unique solution ID.

```
UPDATE CONTAINER,
ID=D9F4555-177C847-3DE17A3049-70AD6FC2-B263-4CBD-84EB-AA22CB-6915BB,
SET_NAME="Test and Dev"
```

The **UPDATE CONTAINER** control statement is required for each solution ID that is found in the input SMF data or is assigned by using the **CONTAINER** control statement. The **UPDATE CONTAINER** control statement associates a container (represented by its unique solution ID) with a solution name.

The solution name associated with the container is assigned via the IBM LMS website and is available upon approval of the solution. You can use LMS to modify the solution name as you prefer; however, the solution name that you specify to SCRT must match the one defined in LMS.

Sample multiplex report for a standalone solution for Tailored Fit Pricing

A multiplex report for Tailored Fit Pricing for IBM Z has the same basic format and report sections as a multiplex report for an environment without Tailored Fit Pricing. However, a multiplex report with Tailored Fit Pricing solutions contains additional report sections that provide an overview of active containers, as well as container-specific report sections in both the billing-related and customer verification sections of the report.

Only CPC M1C1 is shown in this sample report, as CPC M2C1 does not include any container tenant resource groups or LPARs.

Many of the report sections are the same as in a standard sub-capacity report with Tailored Fit Pricing; therefore, this topic focuses on the specific changes that impact the multiplex report.

The following sample report highlights the differences between a multiplex report for a base z/OS system and one that contains Tailored Fit Pricing solutions.

Multiplex summary section

The multiplex summary section contains information about the version of SCRT that was used to generate the multiplex report, customer information, a summary of the machine configuration that makes up the multiplex, and the multiplex MSU values for the reporting period.

Figure 245 on page 394 shows the multiplex summary section for the sample configuration.

```

==B5===== SCRT MULTIPLEX REPORT - IBM Corp =====
SCRT Tool Release                30.1.0
Name of Person Submitting Report: Jane Customer
E-Mail Address of Report Submitter: jane@abc.com
Phone Number of Report Submitter: 888-555-1212

Customer Name                    ABC Corp
Run Date/Time                    03 Dec 2024 - 14:13
Report Period                    2 Nov, 2024 - 1 Dec, 2024 inclusive (30 days)
Number of processors in Multiplex 2

```

Figure 245. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Multiplex summary section

For details about the fields shown in Figure 245 on page 394, see “Multiplex summary section” on page 347.

Machine summary section

Most of the fields in the machine summary section are unchanged when Tailored Fit Pricing is active in a multiplex environment.

The main difference is the appearance of two container summary information subsections. The first subsection displays the simultaneous multiplex peak value for the Tailored Fit Pricing solution and the contribution from each machine in the multiplex. The second subsection displays the total MSU consumption for the Tailored Fit Pricing solution and the contribution from each machine in the multiplex.

Figure 246 on page 394 shows the machine summary section for the sample configuration.

Machine identifier		M1C1	M2C1
Customer number		US-S000000000	US-S000000000
Machine Serial Number		51-Z1Z11	51-54321
Machine Type and Model		2965-J03	3906-797
Machine Rated Capacity (MSUs)		82	11115
Machine Model Changed		N	N
Exclude Data		N	N
Missing LPAR Data		Y	Y
Missing CPC Data		Y	Y

MLC Product Name	Number	MSU	Time	
z/OS V2	5650-Z05	84	27 Nov 2024 - 21:00	84
DB2 for z/OS	(All)	84	27 Nov 2024 - 21:00	84
DB2 11 for z/OS	5615-DB2	84	27 Nov 2024 - 21:00	84
DB2 10 for z/OS	5605-DB2	84	27 Nov 2024 - 21:00	84
CICS TS for z/OS V4	5655-S97	84	27 Nov 2024 - 21:00	84
WebSphere MQ for z/OS V7	5655-R36	84	27 Nov 2024 - 21:00	84

IPLA Product Name	Number	MSU	Time	
WebSphere Transformation Extender for z/OS V8	5655-R95	2	27 Nov 2024 - 20:00	2
IPLA z/OS-Based	(All)	84	27 Nov 2024 - 21:00	84

Container Identifier	Container Name	MSU	Time	
CPS1	DevTest Solution	56	27 Nov 2024 - 21:00	56

Container Identifier	Container Name	TOTAL MSU Consumption	
CPS1	DevTest Solution	5234	5234

Figure 246. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Machine summary section

In the machine summary section, the container summary information subsections immediately follow the product summary information and contain the data for the current reporting period for each container. Since the CMP algorithm analyzes data across a multiplex (not just a single CPC), the reported data has a different scope than in a standard sub-capacity report.

The first container summary information subsection for a multiplex report contains the following fields:

Container Identifier

The SCRT-assigned container identifier. This identifier is assigned by SCRT as it analyzes containers during report processing.

Container Name

The name assigned to this container via the **UPDATE CONTAINER** command (or the CLI or GUI equivalent).

MSU

Reports the concurrent sub-capacity peak utilization, in MSUs, across all CPCs in the multiplex. This differs from the value in a non-multiplex environment where separate peaks are calculated per CPC. SCRT calculates the peak concurrent sub-capacity MSU value for the multiplex from all of the sub-capacity eligible products listed on the report.

Time

Reports the date and time when the highest sub-capacity utilization peak in the multiplex was first reached.

Machine Identifier MSU fields

For each CPC in the multiplex, these fields report the contribution to the multiplex concurrent peak, in MSUs.

In the sample configuration, there is one Tailored Fit Pricing solution defined: CPS1. Tailored Fit Pricing solution CPS1 has a multiplex peak value of 56 MSUs. The peak occurred on 27 November 2017 at 21:00 local time. Machine M1C1 contributed all 56 MSUs to the peak.

In the sample environment, no non-container LPARs ran on machine M1C1; therefore, there is no contribution to non-container product peak values.

The second container summary information subsection for a multiplex report contains the following fields:

Container Identifier

The SCRT-assigned container identifier. This identifier is assigned by SCRT as it analyzes containers during report processing.

Container Name

The name assigned to this container via the **UPDATE CONTAINER** command (or the CLI or GUI equivalent).

TOTAL MSU Consumption

Reports the total MSU consumption of this Tailored Fit Pricing solution across all CPCs in the multiplex.

Machine Identifier MSU consumption fields

For each CPC in the multiplex, these fields report the contribution to the total MSU consumption of the Tailored Fit Pricing solution.

For more information about the product summary section, see [“Product Summary Information section” on page 110](#).

Machine summary section for CPC M1C1

SCRT generates a machine summary section for each CPC that it finds in the multiplex.

The header for each machine summary section contains the SCRT-generated machine identifier within the header itself. For instance, the machine summary section for CPC M1C1 contains the M1C1 value within the header, as shown in [Figure 247 on page 396](#).

Tailored Fit Pricing reporting

There are no differences in the machine information section, Product Summary Information section, Missing LPAR Data Details section, or Missing CPC Data Details section for a multiplex report with Tailored Fit Pricing solutions compared to a non-container multiplex report.

==M1C1=====

Machine identifier	M1C1
Customer number	US-S000000000
Machine serial number	02-12345
Machine Type and Model	3906-760
Machine Rated Capacity (MSUs)	82

==E5=====

PRODUCT SUMMARY INFORMATION

==H4=====

MISSING LPAR DATA DETAILS

Active LPARs with no SMF/SCRT89 records	Detected	Resolved (missing hours)	Est Peak MSU	Justification for low data collection (255 chars max)
LPAR1	27 Nov 2024 - 18:00	28 Nov 2024 - 06:00 (12 hours)	50	(required)
Total Missing Hours	12			
OS Product Name	OS Product ID	Multiplex Peak MSU	Est Peak MSU	Date/Time
z/OS V2	5650-Z0S	200	50	27 Nov 2024 - 18:00
				LPAR1 50

==H6=====

MISSING CPC DATA DETAILS

CPC with no SMF/SCRT89 data from any LPAR	Detected	Resolved (missing hours)	Justification for low data collection (255 chars max)
M1C1	02 Nov 2024 - 00:00	27 Nov 2024 - 15:00 (615 hours)	(required)
Total Missing Hours	615		

Figure 247. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Machine summary section for CPC M1C1

Since no LPARs on CPC M1C1 ran outside of a Tailored Fit Pricing solution during the reporting period, the Product Summary Information section for CPC M1C1 shows no product data.

Note: The Missing LPAR Data Details section continues to report on LPARs that are assigned to a Tailored Fit Pricing solution. There is no change in data collection requirements.

Active Containers section

The Active Containers section provides a cross reference of SCRT-assigned container identifiers to solution IDs.

There is no difference between an Active Containers section for Country Multiplex Pricing and one for a standard sub-capacity report.

When Country Multiplex Pricing is active, in cases where a single container runs on multiple CPCs (as identified by the solution ID), the SCRT-assigned container identifier is used to identify the container across CPCs in the multiplex.

shows the Active Containers section for CPC M1C1.

```

==K5=====
ACTIVE CONTAINERS

SCRT Container Identifier      Solution ID
CPS1                        D9F4555-177C847-3DE17A3049-70AD6FC2-B263-4CBD-84EB-AA22CB-6915BB

```

Figure 248. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Active Containers section

In the sample configuration, there is one Tailored Fit Pricing for IBM Z solution active on the CPC.

Container Billing Summary section for container CPS1

SCRT generates a Container Billing Summary section for each Tailored Fit Pricing solution that it finds in the environment for each CPC in the multiplex on which the container ran.

The header for each Container Billing Summary section contains the SCRT-assigned container identifier within the header itself. The Container Active Products section immediately follows the header, as shown in [Figure 249 on page 397](#).

```

==CPS1=====
Solution ID                D9F4555-177C847-3DE17A3049-70AD6FC2-B263-4CBD-84EB-AA22CB-6915BB
Solution Name              Test and Dev
Peak Four Hour Rolling Average      56
Total MSU Consumption        759

==E7=====
CONTAINER ACTIVE PRODUCTS

MLC Product Name          MLC Product ID
z/OS V2                   5650-Z0S
DB2 11 for z/OS           5615-DB2
CICS TS for z/OS V4       5655-S97

```

Figure 249. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Container Billing Summary and Container Active Products sections

There is no difference in format between the Container Billing Summary sections for Country Multiplex Pricing and the Container Billing Summary sections for standard sub-capacity pricing.

The data in the header of the Container Billing Summary section for a multiplex report differs from a standard sub-capacity report in that the reported peak 4-hour rolling average MSUs correspond to the contribution of the CPC to the multiplex peak value, instead of any individual peak that may have occurred on the CPC.

For more information about the Container Billing Summary section, see [“Container Billing Summary section for container CPS1” on page 385](#).

Customer verification section for CPC M1C1

The customer verification section of a multiplex report contains information that supports and expands on the information in the Product Summary Information section.

Other report sections that are not displayed for this sample configuration may be displayed under different conditions. For information about all of the report sections that can appear, see [“Customer verification section” on page 118](#).

The detail data sections for customer verification follow the machine information section (M7) shown in Figure 250 on page 398.

There are no differences between the Detail LPAR Data Section and Detail LPAR Usage Section of a multiplex report with Tailored Fit Pricing and a standard multiplex report.

```
==M7=====
Machine identifier      M1C1
Customer number        US-S000000000
Machine serial number  51-Z1Z11
Machine Type and Model  2965-J03
Machine Rated Capacity (MSUs) 82
```

Figure 250. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Customer verification section for CPC M1C1 - Machine information

Figure 251 on page 398 shows the Product Max Contributors and Product Grid Snapshot sections for CPC M1C1. No product peak information is displayed in the Product Max Contributors section for CPC M1C1 because no work ran in the standard sub-capacity (non-container) environment.

```
==P5=====
PRODUCT MAX CONTRIBUTORS

Product Name           Product ID  Highest  Date/Time

==Q5=====
PRODUCT GRID SNAPSHOT

Product Name           Product ID
z/OS V2                5650-Z0S
DB2 11 for z/OS        5615-DB2
CICS TS for z/OS V4    5655-S97
```

Figure 251. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Product Max Contributors and Product Grid Snapshot sections

Figure 252 on page 399 shows the Container Product Grid Snapshot and Container Max Contributors sections for solution CPS1. There are no differences between these sections for a multiplex report and a standard sub-capacity report.

```
==Q7=====
```

CONTAINER PRODUCT GRID SNAPSHOT

```
CPS1                      D9F4555-177C847-3DE17A3049-70AD6FC2-B263-4CBD-84EB-AA22CB-6915BB
```

Product Name	Product ID	LPAR1
		*
z/OS V2	5650-Z0S	1.90%
DB2 11 for z/OS	5615-DB2	1.90%
CICS TS for z/OS V4	5655-S97	1.90%

```
==T4=====
```

CONTAINER MAX CONTRIBUTORS

```
CPS1                      D9F4555-177C847-3DE17A3049-70AD6FC2-B263-4CBD-84EB-AA22CB-6915BB
```

LPAR1	TRG	Highest	Date/Time	Contribution to Highest
	*			84

```
CPS1                      84  27 Nov 2024 - 21:00
```

* indicates the entire LPAR was configured as part of the Container

Figure 252. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Container Product Grid Snapshot section

In the Container Max Contributors section for the sample environment, an asterisk (*) appears in the TRG (tenant resource group) column to indicate cases where the entire LPAR is part of the Tailored Fit Pricing solution.

Chapter 20. Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution

Tailored Fit Pricing for IBM Z - Software Consumption Solution (Software Consumption Solution) is a tailored usage-based pricing model where compute is measured on a per-MSU consumed basis. Charges are based on the total MSUs consumed annually.

Overview of Software Consumption Solution

- Committed per-MSU consumed licensing simplifies pricing and removes the need for capping.
- Annual MSU entitlements allow seasonal variations to be smoothed over a full 12-month period.
- Aggressive growth pricing is available for all MSU consumption above a committed baseline.
- The Software Consumption Solution is an alternative to a R4HA-based pricing model for all production workloads, whether new, existing, or growth. This greatly improves deployment flexibility, without the requirement for LPAR-level micromanagement of IBM product usage.
- The Software Consumption Solution offers price predictability and flexibility for clients who rely on IBM Z for their mission-critical workloads. It is offered in conjunction with the Application Development and Test Solution to deliver a comprehensive end-to-end pricing solution for all stages of the application lifecycle.

Requirements for Tailored Fit Pricing for IBM Z - Software Consumption Solution

The requirements for Tailored Fit Pricing for IBM Z - Software Consumption Solution are:

- Minimum hardware: All CPCs must be IBM z14[®] M01 - M05, or IBM z14 Model ZR1.
- Operating system: IBM z/OS V2.2, or later.
 - PTFs associated with the **IBM.Function.PricingInfrastructure** fix category (FIXCAT) must be applied.
 - PTFs for RMF APAR OA54916 are recommended.
- Sub-Capacity Reporting Tool (SCRT) V27.1, or later.
- SCRT is run for each sub-capacity reporting period and the resulting report is submitted to IBM on a monthly basis.
- One or more unique solution IDs is used to identify the Software Consumption solution. IBM will provide you with the solution IDs through the License Management Support (LMS) website upon approval of the solution. (See License Management Support (www.ibm.com/software/lms).)
- Production z/OS systems that run as z/VM guests must specify the VMGUEST RMF Monitor I option.

SCRT externals for Software Consumption reporting

SCRT provides the following externals for Software Consumption reporting:

- A report type specific to Software Consumption reporting: SCRT ENTERPRISE TAILORED FIT REPORT.
- The **ENTERPRISE_TAILORED_FIT_PRICING** control statement for requesting an Software Consumption report.
- All systems are required to be assigned to a container. SCRT issues a message and terminates if there are any unassigned systems in the SMF data. (See [“CONTAINER control statement” on page 41.](#))
- Software Consumption reports are always in UTC time.
- Existing SCRT externals apply to Software Consumption reporting, except for the following restrictions:

- The NO89 DD is optional and deprecated for use in Software Consumption reports. Designate NO89 products on a per-container basis via the **NO89_PRODUCT_ID** keyword of the **UPDATE CONTAINER** control statement, instead of on a per-LPAR basis as allowed via the NO89 DD. If the NO89 DD is present, SCRT will process it and apply it to the designated LPARs.
- The **COUNTRY_MULTIPLEX_PRICING** control statement is not allowed.
- The **REPORT_TIME** control statement is not allowed.
- The following SPECIAL DD control statements that apply to special IBM pricing offerings are not allowed:
 - INTEGRATED_WORKLOAD_PRICING
 - IGNORE_WLM_MWP_DATA
 - DIAGMSG=WLMZWPC
 - DIAGMSG=WLMMOB
 - ASSIGN ZCAP
 - ASSIGN ZWPC
- The following DD statements are not applicable; SCRT ignores them if they are specified:
 - INPUTCSV
 - IWPTRACE
 - GSTRACE
- The SCRT ISVLIB capability does not support the Enterprise Tailored Fit report type. Instead, use the SCRT report type required by the specific ISV. Either standard or multiplex reporting can be used for ISV reporting that includes container definitions.

Sample configuration for Software Consumption reporting

The sample report for Software Consumption is based on the following hardware and software configuration.

The sample configuration consists of a 3906-7E7 CPC with serial number 00001. The CPC is configured with 3 LPARs, all running z/OS natively. LPAR1 and LPAR2 are assigned to the Production1 container, which is based on a consumption metric. LPAR3 is assigned to the DevTest1 container, which is a capacity-based container.

Sample control statement input for Software Consumption reporting

The sample Software Consumption report is based on the following SPECIAL DD control statement input to SCRT:

```
CONTAINER CPC=3906-00001,IMAGE_ID=LPAR1,
ID=D911111-N285B66-88AF33BB8C-NSZZZZZZ-7203-417E-90DB-53931E-3788B8
CONTAINER CPC=3906-00001,IMAGE_ID=LPAR2,
ID=D911111-N285B66-88AF33BB8C-NSZZZZZZ-7203-417E-90DB-53931E-3788B8
CONTAINER CPC=3906-00001,IMAGE_ID=LPAR3,
ID=D911111-N6EE848-5A09B977D0-NTZZZZZZ-1DDE-4F4B-821D-3D8B30-301310
UPDATE CONTAINER,SET_NAME="Production1",
ID=D911111-N285B66-88AF33BB8C-NSZZZZZZ-7203-417E-90DB-53931E-3788B8
UPDATE CONTAINER,SET_NAME="DevTest1",
ID=D911111-N6EE848-5A09B977D0-NTZZZZZZ-1DDE-4F4B-821D-3D8B30-301310
```

Under Software Consumption reporting, all LPARs must be assigned to a container.

Sample report for Software Consumption reporting

A Software Consumption report has the same basic format and content as a Country Multiplex Pricing sub-capacity report with containers, except for the following points:

- Section B5 does not report sub-capacity 4-hour rolling average (4HRA) values for z/OS products.
- Section B5 does not report sub-capacity 4HRA peaks for containers that are based on a consumption metric.
- Sections E5, P5, and Q5 are not present in the report unless z/TPF or z/VSE or both are present in the reporting environment.
- The CPS header sections do not report the container's peak 4HRA for containers that are based on a consumption metric.

The following sample report highlights the differences between a sub-capacity Country Multiplex Pricing report and an Software Consumption Solution report.

Figure 253 on page 403 shows section B5 of the report. No product-level 4HRA information is presented in section B5.

```

==B5===== SCRT ENTERPRISE TAILORED FIT REPORT - IBM Corp =====

SCRT Tool Release                30.1.0
Name of Person Submitting Report: Joe Contact
E-Mail Address of Report Submitter: customer@example.com
Phone Number of Report Submitter: 888-555-5555

Customer Name                    Example Corp
Run Date/Time                    09 May 2024 - 11:12
Report Period                    2 Apr, 2024 - 1 May, 2024 inclusive (30 days)
Number of processors in Multiplex 1

Machine identifier                M1C1
Customer number                  US-S000000000
Machine Serial Number            02-00001
Machine Type and Model           3906-7E7
Machine Rated Capacity (MSUs)    15369
Machine Model Changed            N
Exclude Data                     N
Missing LPAR Data                N
Missing CPC Data                 N

Container Identifier  Container Name      MSU   Time
CPS2                 DevTest1          4139   06 Apr 2024 - 16:00 UTC      4139

Container Identifier  Container Name      TOTAL MSU Consumption
CPS1                 Production1          4185373      4185373
  
```

Figure 253. Sample report for Software Consumption Solution: Section B5

In section B5, no product-level 4HRA information is presented. The DevTest1 container is capacity-based and reports the 4HRA information for that container, but it does not report MSU consumption. The Production1 container is consumption-based and reports MSU consumption, but it does not report any 4HRA information.

Figure 254 on page 403 shows the CPS1 header section for the Production1 container.

```

==CPS1=====
Solution ID          D911111-N285B66-88AF33BB8C-NSZZZZZ-7203-417E-90DB-53931E-3788B8
Solution Name        Production1
Total MSU Consumption 4185373
  
```

Figure 254. Sample report for Software Consumption Solution: CPS1 header section

The solution associated with the CPS1 container is a consumption-based solution. No peak 4HRA is reported for the container.

Figure 255 on page 404 shows the CPS2 header section for the Production1 container.

```
==CPS2=====
Solution ID      D911111-N6EE848-5A09B977D0-NTZZZZZZ-1DDE-4F4B-821D-3D8B30-301310
Solution Name    DevTest1
Peak Four Hour Rolling Average    4139
```

Figure 255. Sample report for Software Consumption Solution: CPS2 header section

The solution associated with the CPS2 container is a capacity-based solution. No MSU consumption is reported for the container.

Chapter 21. Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution

SCRT supports Tailored Fit Pricing for IBM Z - Hardware Consumption Solution reporting based on interval rate data.

Hardware Consumption reporting identifies reporting period hours in which the MSU rate of a CPC exceeded the purchased capacity of the CPC within any 15-minute period. The MSU rate may exceed the purchased capacity as a result of additional active capacity from Capacity on Demand or other hardware offerings.

How SCRT calculates MSU rate data

SCRT MSU rate calculations are based on the reported effective dispatch time (EDT) on general purpose processors for all LPARs on the CPC. The EDT for all LPARs is summed for each 15-minute interval and converted to an hourly MSU rate based on the machine capacity factors at the time. All of the required information is included in the SMF 70 records.

SCRT requires SMF 70 records from at least one LPAR with SMF interval lengths of 1, 3, 5, or 15 minutes. When the interval lengths are less than 15 minutes, SCRT rolls up those intervals into 15-minute periods. (This occurs specifically for section I5 and related calculations only and is not generally true for other SCRT report values.)

Requirements for Hardware Consumption reporting

The requirements for Hardware Consumption reporting are:

- Sub-Capacity Reporting Tool (SCRT) V27.2, or later
- z/OS 2.1, or later
- At least one LPAR with an SMF interval length of 1, 3, 5, or 15 minutes that is active at all times

SCRT externals for Hardware Consumption reporting

SCRT provides the following externals in support of Hardware Consumption reporting:

- The **INTERVAL_RATE_DATA** control statement for requesting that SCRT include the interval rate data section I5 in an SCRT report.
- The **DETAIL_INTERVAL_RATE_DATA** control statement for requesting that SCRT include the detailed interval rate data section V9 in an SCRT report.
- A new CPC MSU Rate Above Permanent Capacity Rating section (I5) in SCRT reports. For details about this section, see [“Sample report for Hardware Consumption reporting”](#) on page 405.
- A new Detailed Interval Data section (V9) in SCRT reports. For details about this section, see [“Sample report for Hardware Consumption reporting”](#) on page 405.

Sample report for Hardware Consumption reporting

An SCRT report with interval rate data enabled for Hardware Consumption reporting (by specifying the **INTERVAL_RATE_DATA** control statement) displays an additional report section, section I5. Section I5 appears per CPC and reports every hour during which the MSU rate of any 15-minute interval within the hour exceeded the permanent (purchased) capacity. If multiple 15-minute intervals exceed the permanent capacity, section I5 reflects the highest excess over capacity.

[Figure 256 on page 406](#) shows an example of section I5.

```

==I5=====
CPC MSU RATE ABOVE PERMANENT CAPACITY RATING

Date/Time      MSU Rate      Perm Capacity      Excess Over Capacity
02 Mar 2024 -00:00      7766      5771      1995
02 Mar 2024 -01:00      6238      5771      467
02 Mar 2024 -21:00      6418      5771      647
02 Mar 2024 -22:00      6143      5771      372

```

Figure 256. Sample SCRT report: CPC MSU Rate Above Permanent Capacity Rating section (I5)

The column headings in section I5 are defined as follows:

Date/Time

The reporting period hour during which the interval occurred.

MSU Rate

The 15-minute MSU rate that resulted in the highest excess over capacity in the reporting period hour.

Perm Capacity

The permanent capacity of the CPC during the reporting period hour.

Excess Over Capacity

The difference between the MSU rate and the permanent capacity.

If the **DETAIL_INTERVAL_RATE_DATA** control statement is specified to include detailed interval data, the Detailed Interval Data section (V9) also appears in the report. Section V9 appears per CPC and provides finer granularity for interval rate data by showing a (typically) 15-minute view of the CPC MSU rate.

Figure 257 on page 406 shows an example of section V9.

```

==V9=====
DETAILED INTERVAL DATA

Date/Time      IntervalStart      IntervalEnd      IntLen      TypeMod      Rating      Hour R4HA      Hour MSU Consumed      MSU Rate      Perm Model Capacity      Perm Temp Model Capacity      Temp Var Model Capacity      Var Var Model Capacity
3906-6A477 02 Mar 2024 - 00:00 02 Mar 2024 - 00:00 02 Mar 2024 - 00:15 15 3906-763 7898 1122 796 7766 743 5771 763 7898 763 7898
3906-6A477 02 Mar 2024 - 00:00 02 Mar 2024 - 00:15 02 Mar 2024 - 00:30 15 3906-763 7898 1122 796 7642 743 5771 763 7898 763 7898
3906-6A477 02 Mar 2024 - 00:00 02 Mar 2024 - 00:30 02 Mar 2024 - 00:45 15 3906-763 7898 1122 796 6468 743 5771 763 7898 763 7898
3906-6A477 02 Mar 2024 - 00:00 02 Mar 2024 - 00:45 02 Mar 2024 - 01:00 15 3906-763 7898 1122 796 5603 743 5771 763 7898 763 7898
3906-6A477 02 Mar 2024 - 01:00 02 Mar 2024 - 01:00 02 Mar 2024 - 01:15 15 3906-763 7898 1063 1209 6238 743 5771 763 7898 763 7898
3906-6A477 02 Mar 2024 - 01:00 02 Mar 2024 - 01:15 02 Mar 2024 - 01:30 15 3906-763 7898 1063 1209 5903 743 5771 763 7898 763 7898
3906-6A477 02 Mar 2024 - 01:00 02 Mar 2024 - 01:30 02 Mar 2024 - 01:45 15 3906-763 7898 1063 1209 6032 743 5771 763 7898 763 7898
3906-6A477 02 Mar 2024 - 01:00 02 Mar 2024 - 01:45 02 Mar 2024 - 02:00 15 3906-763 7898 1063 1209 4962 743 5771 763 7898 763 7898

```

Figure 257. Sample SCRT report: Detailed Interval Data section (V9)

The column headers in section V9 are defined as follows:

Column 1 (untitled)

The CPC being reported

Data/Time

The reporting period hour during which the interval occurred.

IntervalStart

The start of the interval within the reporting period hour.

IntervalEnd

The end of the interval within the reporting period hour.

IntLen

The length of the interval, in minutes.

TypeMod

The active machine model and type.

Rating

The MSU rating of the active machine model.

Hour R4HA

The 4-hour rolling average CPU utilization of all LPARs that reported data for the report period hour. This value is for the entire hour, not the sub-hour interval.

Hour MSU Consumed

The total MSU consumed by all LPARs that reported data for the report period hour. This value is for the entire hour, not the sub-hour interval.

MSU Rate

The MSU rate of the CPC over the interval.

Perm Model

The permanent machine model during the interval.

Perm Capacity

The permanent machine capacity during the interval.

Temp Model

The temporary machine model during the interval.

Temp Capacity

The temporary machine capacity during the interval.

Var Model or Rep Model

If the machine type supports variable capacity, this field indicates the variable capacity machine model during the interval. If the machine type supports replacement capacity, this field indicates the replacement capacity machine model during the interval. If the machine type supports neither variable capacity nor replacement capacity, this field will not be present.

Var Capacity or Rep Capacity

If the machine type supports variable capacity, this field indicates the variable machine capacity during the interval. If the machine type supports replacement capacity, this field indicates the replacement machine capacity. If the machine type supports neither variable capacity nor replacement capacity, this field will not be present.

Chapter 22. Workload Classification Supplement

The Workload Classification Supplement is a supplementary report, separate from the SCRT report, that shows CPC utilization by workload class on a sub-hour basis (for instance, 15 minutes).

Overview of the Workload Classification Supplement

The Workload Classification Supplement is produced alongside the SCRT report during a single execution of SCRT. Only one Workload Classification Supplement is generated in a single execution of SCRT. This applies whether the regular SCRT report is a standard sub-capacity report (one report per CPC) or a Country Multiplex Pricing report (one report per multiplex) or an Enterprise Tailored Fit Pricing report (one report per enterprise).

A Workload Classification Supplement can be produced for any CPC type, but workload classification is only supported on IBM z17 or later. A Workload Classification Supplement for earlier processors will show 100% of processor time attributed to the default class for each processor type. The default class for general purpose processors is class 0, and the default class for zIIPs is class 1.

The Workload Classification Supplement is not submitted to IBM via the LMS portal. If you have an IBM offering that requires a Workload Classification Supplement to be submitted, the reporting and submission requirements will be specified in the offering terms. Contact your IBM representative for additional details.

Note: Use of the term *supplement* in this context refers to the supplementary report produced by SCRT, not to any contractual supplements associated with any IBM offerings.

How workload classification data is calculated

SCRT uses workload classification sample counts from the Workload Classification Data Section of SMF type 70 subtype 1 records and the effective dispatch time of general purpose processors and zIIPs to determine the percentage of processor time spent in each workload class. The dispatch time for each reporting system is apportioned according to the sampling data provided in the SMF records, then aggregated to form the CPC view.

If a system does not provide workload classification sample counts, all of the dispatch time for that system is attributed to the default class for each processor type. The default class for general purpose processors is class 0, and the default class for zIIPs is class 1.

SCRT externals for the Workload Classification Supplement

A Workload Classification Supplement is generated when the WCSUPP DD is present in your SCRT batch job. The WCSUPP DD identifies an output dataset to which the supplement is to be written and can be a sequential data set, a PDS, or a PDSE. Users running SCRT on Linux or Windows can use the **--wcsupp** command line argument or the **WCSUPP** checkbox and file name field on the **Trace Options** tab on the final GUI panel.

The Workload Classification Supplement uses the same information provided by the PARMS DD (or CLI or GUI equivalent) to identify the customer and submitter contact information.

Interaction with existing SCRT control statements

EXCLUDE control statements

- Product-specific excludes (that is, **EXCLUDE** control statements that specify the **PRODUCT_ID** keyword) do not have any impact on the Workload Classification Supplement.
- Non-product-specific excludes (that is, **EXCLUDE** control statements that do not specify the **PRODUCT_ID** keyword) impact the Workload Classification Supplement in the following ways:

If data is excluded from all systems on a CPC, the Workload Classification Supplement will not report any data for the excluded time periods. Otherwise, the following behaviors apply:

- Processor dispatch times from systems that are excluded continue to contribute to the total CPC dispatch times reported in the Workload Classification Supplement.
- Workload classification data from excluded systems is not considered when determining the percentage of time spent in each workload class. Instead, all dispatch time from an excluded system is attributed to the default class for each processor type. The default class for general purpose processors is class 0, and the default class for zIIPs is class 1.
- If you have specified an **EXCLUDE** control statement for software pricing purposes, consider whether or not that **EXCLUDE** is appropriate for any IBM offering that is based on the Workload Classification Supplement. If necessary, you can generate a Workload Classification Supplement by running SCRT a second time without the **EXCLUDE** statement.

The Workload Classification Supplement report

The Workload Classification Supplement consists of a header section followed by one classification section for each CPC. The header section provides much of the same information as section B5 of a regular SCRT report, such as SCRT version, customer information, and machine types and serial numbers.

There is one section for each CPC, identified by an M&C& header (see the description of the machine identifier for CMP reporting in [“Machine identifier” on page 348](#)), for example:

```
==M1C1=====
```

Each per-machine section consists of a header row and data rows that consist of the following column names and data:

Type

The CPC machine type.

Serial

The CPC serial number, including the plant of manufacturing (POM) code.

Interval Start

The start time corresponding to the data in the row.

Interval End

The end time corresponding to the data in the row.

Duration

The duration of the interval, in seconds.

Perm Model

The permanent capacity model of the CPC for this interval. If a change occurred during the interval, this is the model at the end of the interval.

Temp Model

The temporary capacity model of the CPC for this interval. If a change occurred during the interval, this is the model at the end of the interval.

Replacement Model

If supported by the processor type, the replacement capacity model of the CPC for this interval. If a change occurred during the interval, this is the model at the end of the interval. If not supported by the processor type, this value will be a dash (-) as a placeholder.

Active Model

The active capacity model of the CPC for this interval. If a change occurred during the interval, this is the model at the end of the interval.

Avg CPs Online

The average number of CPs online during the interval.

CP Dispatch Time

The total CPC dispatch time on general purpose processors during the interval.

CP Pct Util

General purpose processor utilization, as a percentage of total general purpose processor capacity.

CL0 CP

The percentage of processor time spent in workload class 0 for general purpose processors during the interval.

CL1 CP

The percentage of processor time spent in workload class 1 for general purpose processors during the interval.

CL2 CP

The percentage of processor time spent in workload class 2 for general purpose processors during the interval.

CL3 CP

The percentage of processor time spent in workload class 3 for general purpose processors during the interval.

CL4 CP

The percentage of processor time spent in workload class 4 for general purpose processors during the interval.

CL5 CP

The percentage of processor time spent in workload class 5 for general purpose processors during the interval.

CL6 CP

The percentage of processor time spent in workload class 6 for general purpose processors during the interval.

CL7 CP

The percentage of processor time spent in workload class 7 for general purpose processors during the interval.

Avg zIIPs Online

The average number of zIIPs online during the interval.

zIIP Dispatch Time

The total CPC dispatch time on zIIPs during the interval.

zIIP Pct Util

zIIP utilization, as a percentage of total zIIP capacity.

CL0 zIIP

The percentage of processor time spent in workload class 0 for zIIPs during the interval.

CL1 zIIP

The percentage of processor time spent in workload class 1 for zIIPs during the interval.

CL2 zIIP

The percentage of processor time spent in workload class 2 for zIIPs during the interval.

CL3 zIIP

The percentage of processor time spent in workload class 3 for zIIPs during the interval.

CL4 zIIP

The percentage of processor time spent in workload class 4 for zIIPs during the interval.

CL5 zIIP

The percentage of processor time spent in workload class 5 for zIIPs during the interval.

CL6 zIIP

The percentage of processor time spent in workload class 6 for zIIPs during the interval.

CL7 zIIP

The percentage of processor time spent in workload class 7 for zIIPs during the interval.

Important: If you are required to submit a Workload Classification Supplement to IBM for an IBM offering, do not edit the supplement in any way.

Sample Workload Classification Supplement report

This example shows the contents of a Workload Classification Supplement report.

Note: Because the Workload Classification Supplement report contains many more columns than a typical SCRT report, for space reasons, the column headers and rows in this example are shown here continued across multiple report segments.

===== WORKLOAD CLASSIFICATION SUPPLEMENT =====

SCRT Tool Release30.1.x
Name of Person Submitting ReportJoe Contact
E-Mail Address of Report Submitternobody@example.com
Phone Number of Report Submitter555-555-5555

Customer NameJoe Customer
Run Date/Time01 Apr 2025 - 15:37
Reporting Period2 Jan, 2025 - 1 Feb, 2025 inclusive (31 days)
Number of processors reported1

Machine IdentifierM1C1
Machine Serial Number51-77777
Machine Type8561

==M1C1=====

Type	Serial	Interval Start	Interval End	Duration	Perm Model	Temp Model	Replacement Model	Active Model	...
8561	51-77777	03 Jan 2025 - 00:00	03 Jan 2025 - 00:05	300	8561-725	8561-725	-	8561-725	...
8561	51-77777	03 Jan 2025 - 00:05	03 Jan 2025 - 00:10	300	8561-725	8561-725	-	8561-725	...
8561	51-77777	03 Jan 2025 - 00:10	03 Jan 2025 - 00:15	300	8561-725	8561-725	-	8561-725	...

...	Avg CPs Online	CP Dispatch Time	CP Pct Util	CL0 CP	CL1 CP	CL2 CP	CL3 CP	CL4 CP	CL5 CP	CL6 CP	CL7 CP	...
...	25	1636.61	21.82%	99.00%	0.01%	0.02%	0.03%	0.06%	0.13%	0.25%	0.51%	...
...	25	2927.03	39.03%	99.15%	0.01%	0.01%	0.03%	0.05%	0.11%	0.22%	0.43%	...
...	25	2702.14	36.03%	98.91%	0.01%	0.02%	0.03%	0.07%	0.14%	0.27%	0.55%	...

...	Avg zIIPs Online	zIIP Dispatch Time	zIIP Pct Util	CL0 zIIP	CL1 zIIP	CL2 zIIP	CL3 zIIP	CL4 zIIP	CL5 zIIP	CL6 zIIP	CL7 zIIP	
...	10	420.63	14.02%	0.00%	99.94%	0.01%	0.00%	0.00%	0.02%	0.00%	0.03%	
...	10	831.96	27.73%	0.00%	99.97%	0.00%	0.00%	0.00%	0.01%	0.00%	0.02%	
...	10	551.79	18.39%	0.00%	99.95%	0.01%	0.00%	0.00%	0.01%	0.00%	0.03%	

Chapter 23. Reporting for independent software vendor programs

SCRT supports the creation of reports for participating independent software vendor (ISV) programs.

This capability gives participating ISVs and their clients the ability to generate specific product utilization reports for each ISV product set. The ISV report is similar to an IBM sub-capacity report and features many of the same report sections.

Important:

- Other than for SCRT defect analysis purposes, the ISV report is **not** intended for use by IBM.
- The ISV reporting capability is **not** associated with any IBM licensing or pricing option.

Overview of reporting for ISV programs

SCRT support for ISV reporting provides a method for participating ISVs to understand how a client uses their software. Usage is analyzed and reported using the same approaches that SCRT uses for IBM products. This similarity provides the vendor and its clients with a familiar view of ISV product reporting; there are very few new concepts to learn. A goal of this support is to allow a client who is familiar with sub-capacity reporting for IBM products to easily adapt to reporting for ISV products.

The SCRT support for ISV program reporting does require ISV participation. There is no method to generate a report of ISV products without information provided by the ISV.

A key concept of this enhancement is the ISVLIB data set. This is an ISV-generated file that a participating ISV distributes to their clients as they deem appropriate. The ISV client then provides this ISVLIB data set as input to SCRT on z/OS via a new ISVLIB DD statement, or on Linux or Windows via the SCRT user interface (UI). The ISVLIB data set is the only input that is required to trigger report generation for the ISV product set.

When supplied with an ISVLIB data set, SCRT analyzes the customer data and reports on the specific utilization of ISV products represented in that data set. It will also report on the IBM operating system that is in use.

Note: The ISV report does not contain any IBM product information (beyond the operating system) nor any information about other ISV's programs. SCRT only accepts one ISVLIB input data set at a time. Clients who have programs from multiple participating ISVs must generate a separate report for each ISV, each requiring a separate execution of SCRT.

IBM urges (but does not require) ISVs to provide their clients with a sample job that demonstrates the appropriate JCL. The IBM SCRT sample job can also be modified for ISV reporting with guidance from the participating ISV.

SCRT externals for ISV program reporting

SCRT provides the following externals for ISV program reporting:

- New externals:
 - On z/OS systems, SCRT supports an optional ISVLIB DD statement which is used to refer to the ISV-provided library file.
 - On Windows and Linux systems, a separate launcher is installed for creating ISV reports. When you start the ISV version, the UI prompts you to specify the location of the ISVLIB data set. For more information, see [Chapter 4, “Using SCRT on Windows and Linux systems,”](#) on page 69.
- Changes to existing externals:
 - The NO89 DD statement supports ISV NO89 products when ISV reporting is active.

- The EXCLUDE DD statement supports ISV products when ISV reporting is active.
- The PARMS DD statement supports more generic customer identification parameters when ISV reporting is active.

ISVLIB DD statement

The default SCRT processing does not support ISV reporting. To report on program usage for a specific ISV, you must provide the ISVLIB data set on the ISVLIB DD statement.

The data in the data set pointed to by the ISVLIB DD statement is in a specific format, as required by SCRT. Participating ISVs supply their data in this format, which you should not modify.

If SCRT is unable to process the data in the ISVLIB DD statement, it issues an error message. The error description indicates the appropriate action. In many cases, it will be the responsibility of the ISV to provide a correction.

NO89 DD statement

When it finds a valid ISVLIB DD data set, SCRT does not accept the specification of IBM NO89 product usage. SCRT instead requires that you indicate the usage of any ISV NO89 programs.

Note that the syntax requirements for the NO89 product identifiers are relaxed for ISV report processing. Except for this modification, the syntax requirements and usage of the NO89 DD statement input is the same as for IBM product reporting. See [“NO89 DD statement” on page 61](#) for more information.

PARMS DD statement

When it finds a valid ISVLIB DD data set, SCRT does not perform syntax checking on the **CUSTOMER_NUMBER** parameter of the PARMS DD statement. Except for this modification, the syntax requirements and usage of the PARMS DD statement is not changed. See [“PARMS DD statement” on page 34](#) for more information.

SPECIAL DD statement

The default SCRT processing for ISV reporting does not use the Tailored Fit Pricing for IBM Z algorithm.

When you specify the **ISV_SHOW_CONTAINERS** control statement on the SPECIAL DD statement, SCRT uses the Tailored Fit Pricing for IBM Z algorithm to generate any Tailored Fit Pricing solutions. ISV products that utilize product registration or product usage services (and appear in SMF type 89 records) appear in a Tailored Fit Pricing solution under the same conditions as an IBM product in standard SCRT reporting.

Considerations for character set encoding

ISVs may optionally validate the integrity of an SCRT report submitted to them by using a utility provided by IBM. Successful report validation depends on correct character encoding at several steps of the process:

- SCRT reads parameters and control statements from the PARMS DD and the SPECIAL DD using the default character encoding of the JVM. You must ensure that the encoding of your PARMS and SPECIAL DD statements matches the default encoding of the JVM. For instance, if the default encoding of the JVM is IBM-1047, ensure that your parameters are encoded as IBM-1047.
- SCRT writes the report itself using the default encoding of the JVM. You might need to convert the report to the encoding expected by the ISV. Be sure to specify the correct input and output encodings for the conversion.

Restrictions for ISV program reporting

The following restrictions apply to ISV program reporting:

- MSU reductions from pricing options specific to IBM, such as Mobile Workload Pricing (MWP) or IBM Z Collocated Application Pricing (zCAP), are not applied to ISV products nor the IBM operating system.
- Multi-Version Measurement (MVM) is not supported for ISV products. Each version of an ISV product is reported independently, and no program family line is generated in the report.
- Tailored Fit Pricing for IBM Z is available for ISV program reporting; however, the default SCRT behavior for ISV reporting is to ignore any Tailored Fit Pricing solutions. You must use the **ISV_SHOW_CONTAINERS** command (control statement) to enable container reporting in ISV reports.

Sample SCRT JCL for ISV program reporting

Figure 258 on page 415 shows a portion of the JCL for SCRT that illustrates the specification of the ISVLIB DD statement and the use of an ISV-specific NO89 DD. In the example, the ISVLIB data supplied by the ISV has been placed in the EXAMPLE . SCRT . ISVLIB (VENDNXYX) data set. The ISV NO89 products with product IDs of 8997-XYZ and 6161-XYZ are specified in the NO89 DD. All other DDs must be specified as for a typical IBM sub-capacity report.

```
//VENDXYZ JOB MSGCLASS=H,NOTIFY=HLQ,REGION=0M,TIME=20
//*****
//* Sample JCL for invoking SCRT for vendor XYZ
//*****
//SCRT EXEC PROC=JVMPRC86,
// JAVACLS='com.ibm.scrt.SCRTe'
//STDENV DD ...
//STDOUT DD ...
//OUTPUT DD ...
//SMF DD ...
//ISVLIB DD DISP=SHR,DSN=EXAMPLE.SCRT.ISVLIB(VENDXYZ)
//*
//NO89 DD *
* XYZ LOREM IPSUM SUITE
8997-XYZ=LPAR2
* XYZ SAMPLE UTILITY
6161-XYZ=*ALL
```

Figure 258. Sample SCRT JCL for ISV program reporting

Sample SCRT product report for ISV programs

Figure 259 on page 416 and Figure 260 on page 416 are excerpts from an SCRT product report that highlight the differences from an IBM sub-capacity report.

Customer information section

In section B5, the header indicates the vendor name. In the example shown in Figure 259 on page 416, the vendor has identified themselves as XYZ COMPUTING in the supplied ISVLIB data set. A vendor-specific legal statement, also supplied by the vendor, also appears in section B5.

```
==B5===== SCRT PRODUCT REPORT FOR XYZ COMPUTING =====
```

```
Run Date/Time           07 Apr 2024 - 11:18
Name of Person Submitting Report:  John Customer
E-mail Address of Report Submitter: john@abc.com
Phone Number of Report Submitter:  800-555-1234
```

```
Customer Name           ABC Corp.
Customer Number          2000000000
Machine Serial Number    12-87654
Machine Type and Model   2817-707
Machine Rated Capacity (MSUs) 879
Purchase Order Number    (optional)
Customer Comments (255 chars max) (optional)
```

```
Sample vendor-specific legal statement...
Lorem ipsum dolor sit amet, consectetur
adipiscing elit, sed do eiusmod tempor incididunt ut
labore et dolore magna aliqua. Ut enim ad minim...
```

Figure 259. Sample SCRT product report: Customer Information section

Product Summary Information section

In the Product Summary Information section, SCRT reports on the IBM operating systems and the ISV products. In the example shown in [Figure 260 on page 416](#), both z/OS V1 and z/OS V2 are reported. The vendor products follow the IBM operating systems.

```
==E5=====
```

PRODUCT SUMMARY INFORMATION

Product Name	Product ID	Tool MSUs	Customer MSUs
z/OS	(All)	278	
z/OS V2	5650-ZOS	278	
z/OS V1	5694-A01	19	
XYZ Advanced Toolkit V17	1234-V17	97	
XYZ LOREM IPSUM SUITE	8997-XYZ	88	
XYZ SAMPLE UTILITY	6161-XYZ	278	

Figure 260. Sample ISV product report: Product Summary Information section

Part 4. Viewing, editing, and submitting a sub-capacity report

This part describes the tasks you perform on your sub-capacity report output, such as viewing and editing the report in a spreadsheet application, preparing to submit the report, and submitting the report to IBM. This part also describes how to correct some common errors that you might encounter.

The information in this part applies to sub-capacity reports generated on all platforms on which SCRT runs.

This part contains the following chapters:

- [Chapter 24, “Viewing and editing the sub-capacity report,” on page 419](#)
- [Chapter 25, “Preparing to submit the sub-capacity report to IBM,” on page 423](#)
- [Chapter 26, “Submitting sub-capacity reports to IBM,” on page 427](#)
- [Chapter 27, “Troubleshooting SCRT errors and problems,” on page 439](#)

Chapter 24. Viewing and editing the sub-capacity report

The person responsible for software billing (sometimes called the software asset manager) must analyze and, often, change and comment on the contents of a sub-capacity report prior to submitting it to IBM. This chapter describes how to perform these tasks using a spreadsheet application:

- View the contents of the sub-capacity report
- Change information in the sub-capacity report

Considerations for viewing and editing the sub-capacity report

For countries and regions where the License Management Support (LMS) web application is available, the recommended method for viewing, changing, and submitting your sub-capacity reports is described in [Chapter 26, “Submitting sub-capacity reports to IBM,” on page 427](#).

When LMS is available and when companies have multiple internal reviewers who assess and edit the sub-capacity report file using a spreadsheet application, IBM recommends uploading an *unaltered* copy of that sub-capacity report file into the LMS web application. Then, manually transfer (copy and paste) the reviewer comments that were previously entered into the spreadsheet file into the corresponding input fields in LMS. This approach has the benefit of allowing you to use the formatting capabilities of a spreadsheet application to view all sections of the sub-capacity report (including those sections not viewable in the LMS web interface) and to make changes or add comments to the sub-capacity report, but avoids the potential problem where a spreadsheet application might alter the sub-capacity report in unintended ways that can corrupt the data and cause LMS to reject the report file.

Both LMS and the spreadsheet method support double-byte character set (DBCS) data for certain fields, as described in [“Spreadsheet method for viewing and editing the sub-capacity report” on page 420](#). The LMS web interface does not show the Detail Data sections of the sub-capacity report, described in this chapter. You must use a spreadsheet application to view those sections.

For countries and regions where the LMS application is not available, use the spreadsheet method described in [“Spreadsheet method for viewing and editing the sub-capacity report” on page 420](#) to view and edit your sub-capacity reports before submitting them to IBM in an email.

Restriction for Solution Edition offerings

Although most of the sub-capacity report comments are accepted by the IBM business process as either single-byte character set (SBCS) data or, where supported, as DBCS data, the IBM business process does restrict the comments required for Solution Edition LPARs in the Special Conditions section and in the Detail Data Collection section to be the name of the Solution Edition offering in single-byte US English characters, such as:

Solution Edition for SAP Applications

For countries and regions where DBCS support is available for LMS and you are also using one or more Solution Edition offerings, you must use a non-LMS-based method to add the required comments for the **Exclude** control statements which are listed for Solution Edition LPARs in the Special Condition section of the sub-capacity report. Those comments must be added to the sub-capacity report file prior to uploading that file into LMS running in DBCS mode. There are several such non-LMS options for adding these comments:

- You could use the ISPF Edit panels under TSO/E on the z/OS system that created the sub-capacity report.
- You could use a simple text editor, such as Windows Notepad, running on a workstation.

In both cases, you must use extreme care to locate the appropriate place within sub-capacity report file to replace the (required) placeholder (see [“Using placeholder fields in the sub-capacity report” on page 420](#)) with the required **Exclude** comment, without changing the number of commas (,) or double quotation marks (") in the original file.

Although you can also use a spreadsheet application to add those same single-byte US English comments to the sub-capacity report file, be aware that spreadsheet applications also have the potential to alter the sub-capacity report in unintended ways that can corrupt the data and cause LMS to reject the report file. For details about using the **Exclude** statement, see [“EXCLUDE control statement” on page 44](#).

This same restriction applies to the comment in the **Detail Data Collection** section which is required for the same Solution Edition LPAR that is listed as having a value of 0.0% in the **Report Period % Data** column. The same techniques used for the **Exclude** comment will also work for this comment. However, the **LPAR_Comment** control statement can also be used to provide the required single-byte US English comment and has the added advantage of being reused without changes until the Solution Edition offering is no longer used. For details about using the **LPAR_Comment** control statement see [“PARMS parameters that apply to specific CPCs” on page 37](#).

Spreadsheet method for viewing and editing the sub-capacity report

The file produced by SCRT—the sub-capacity report—is a comma-separated values (CSV), flat text file. IBM recommends that you download the sub-capacity report (in text format with a .csv file extension) to a directory on your workstation.

You can review the contents of the sub-capacity report for accuracy and, if necessary, edit selected placeholder fields in the report using a standard text editor or a spreadsheet application, such as Microsoft Excel. Most spreadsheet applications support importing and exporting CSV-formatted files and can provide a convenient means to view and edit the sub-capacity report.

If you use a spreadsheet application to edit placeholder fields, you must save or export the report as a CSV file. Even when saved as a CSV file, the spreadsheet application might make unintentional changes to the format or structure of the report (such as adding additional empty columns, or adding or removing quotation marks). These changes might cause LMS to reject the report. If you are editing an ISV report, your ISV might reject the report. To avoid introducing such changes, instead of using a spreadsheet application, use any standard text editor to edit the placeholder fields.

Using placeholder fields in the sub-capacity report

The sub-capacity report uses special text values as placeholder fields to indicate where you can enter optional or required comments and where you can enter alternate product MSU values. Only these designated placeholder fields should be edited or changed after the report has been generated. Report changes outside of these fields will be ignored or will cause the report to be rejected by IBM business systems.

The sub-capacity report places one of the following three placeholder values in the specific fields where IBM accepts additional customer input (such as comments and MSU values) for processing:

(optional)

Indicates a field where additional, optional input will be accepted. The input could be an explanation for a condition or it could be an alternate MSU value in the **Customer MSUs** field. If it is an MSU value, the value must be a whole number with a maximum of four digits.

(conditional)

Follows the (optional) placeholder and indicates that if you provide input for the preceding (optional) placeholder, then you must also provide an explanation in the field indicated by the (conditional) placeholder. Otherwise, IBM will not accept the value in the preceding (optional) field. The (optional) and (conditional) pair of placeholders are used in the Product Summary Information section of the sub-capacity report.

(required)

Indicates a field where you must provide some form of input—typically, an explanation for some condition which is outside of the accepted sub-capacity rules. Failure to provide an acceptable explanation where requested may cause IBM to reject the entire sub-capacity report.

Edits to the placeholder fields have the following restrictions:

- Values are limited to 255 characters.
- Double quotation marks and commas are not permitted.

The LMS web application allows you to provide additional input for all fields that were originally marked by the (optional) and (conditional) placeholders and will prompt you to provide input for all fields that have the (required) placeholder for which no input was added prior to submission.

The LMS web application accepts double-byte character set (DBCS) data which has been entered in certain fields of a sub-capacity report file using a spreadsheet application for certain countries, regions, and code pages. [Table 56 on page 427](#) lists the countries, regions, and code pages for which LMS accepts DBCS data. LMS supports the normal ENGLISH_ENCODING = "ISO-8859-1" encoding for all other countries.

[Table 55 on page 421](#) lists the sub-capacity report fields in which LMS accepts DBCS values.

Table 55. Sub-capacity report fields in which LMS accepts DBCS data

Report section	Field/column name	Restrictions
First section (B5)	Customer Name	None.
First section (B5)	Customer Comments	None.
Special Conditions (D5)	Customer Comments	Except for Exclude statements for Solution Edition, which must match the title of the Solution Edition offering in SBCS US English text
Product Summary Information (E5)	Customer Comments	Only Customer Comments fields that contain the (conditional) placeholder
Missing LPAR Data Details	Justification for low data collection	None.
Missing CPC Data Details	Justification for low data collection	None.

Note: SCRT only supports Customer MSUs overrides for NO89 products. You must use the Exclude control statement to change the MSU values for products that are tracked by SMF type 89 records or SCRT89 records (that is, the products that have no placeholder in the Customer MSUs column). You must also provide a reason for using the Exclude control statement. IBM's business systems will ignore MSU override values in Customer MSUs fields that did not originally contain an (optional) placeholder.

If you must provide IBM with some information for a product that appears in the Product Summary Information section but for which SCRT does not provide any placeholder in the Customer Comments column for that product, enter that information at the (optional) placeholder for the Customer Comments row at the beginning of the sub-capacity report where general customer comments can be added.

Chapter 25. Preparing to submit the sub-capacity report to IBM

For a machine to qualify for sub-capacity pricing, a sub-capacity report must be officially submitted to IBM by the ninth day of the month following data collection. For example, if data is collected for 2 September through 1 October, you must submit the report to IBM by 9 October.

You must use the IBM License Management Support (LMS) application to submit your sub-capacity report to IBM, either by sending the report to LMS as an email attachment or by using the LMS web interface to upload and submit the report. See [Chapter 26, “Submitting sub-capacity reports to IBM,” on page 427](#) for instructions on using these submission processes.

Before submitting the sub-capacity report to IBM, review the items in [“Sub-capacity report checklist” on page 423](#) to help ensure a successful submission.

Sub-capacity report checklist

Verify the following items before you send a sub-capacity report to IBM:

1. Are you using the appropriate level of SCRT?

You are required to upgrade to a new SCRT version as soon as it is available. To determine the required version and to review the list of new products, SCRT fixes, and SCRT functional enhancements, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](#). Install the correct level of SCRT by applying the necessary PTFs or by downloading the appropriate deliverable.

2. Is the sub-capacity report for the correct reporting period?

Sub-capacity reports must be submitted to IBM by the ninth day of the month following data collection. If the sub-capacity report is for the wrong reporting period, check to see whether you are providing SCRT with SMF and SCRT89 records from the appropriate months. If you have provided the appropriate SMF data, you must use the **Report_Period** control statement to override the reporting date that SCRT automatically selects.

3. Does the message Machine Model change observed appear following the value in the **Machine Type/Model** field in section B5 of the sub-capacity report? If so, enter a comment to describe the reason for the machine model change. For details, see [“Machine software model changes” on page 441](#).

4. Do you need to use any of the SPECIAL SCRT control statements to:

- Exclude processing for a specific subset of the data collected during the SCRT reporting period?
- Suppress reporting of certain Getting Started Sub-capacity statistics?
- Request that SCRT use the Integrated Workload Pricing (IWP) option when generating SCRT reports?

If you use the **Integrated_Workload_Pricing** control statement to request IWP processing, you must also agree to (sign) the required IWP Addendum to the AWLC Attachment prior to submitting the sub-capacity report with Integrated Workload Pricing. Otherwise, IBM will reject sub-capacity reports with Integrated Workload Pricing.

- Request other reporting options?
5. Are you licensed for all the products for which a product defined capacity is specified on this report? If not, contact your IBM customer representative.
 6. Have you already submitted a report for this reporting period? If you submit multiple reports for the same CPC and reporting period only the first report is used.
 7. Has SCRT detected any conditions that would indicate that the z/OS systems on this CPC are not eligible for sub-capacity charges? For example:

- MVS and OS/390 must not be running on this CPC.
 - z/OS must not be running in 31-bit mode on this CPC.
8. Was the report generated using at least one month of SMF data (or SCRT89 data for z/TPF or z/VSE) as input?
 9. If you have multiple sub-capacity-eligible operating systems (z/OS, z/TPF, z/VSE) running on the same CPC, have you created a single sub-capacity report for that CPC that reports on all of the operating systems? See [“SMF DD statement” on page 65](#) for more information.
 10. Was data collected from the second day of the previous month through the first day of the current month?
 11. Have you reviewed all product utilization capacities generated by the SCRT and provided alternate values, where appropriate?

For example, did you account for any periods during which the collected SMF or SCRT89 data did not represent your normal or intended use of the products? (For example, for problem determination or disaster recovery testing.)

12. If you provided alternate **Customer MSUs** values for the product fields with the (optional) placeholder, were the MSU values whole numbers with a maximum of four digits? Did you also include an explanation for each alternate value in the corresponding **Customer Comments** field?

IBM will not accept alternate values that do not have an acceptable explanation in the corresponding **Customer Comments** field. IBM will also not accept alternate **Customer MSUs** values for products that do not have an (optional) placeholder in their **Customer MSUs** column. You must use an **Exclude** control statement to request SCRT to report an adjusted **Tool MSUs** value for products that do not have an (optional) placeholder in the corresponding **Customer MSUs** field. For details, see [“EXCLUDE control statement” on page 44](#).

13. If you needed to use one or more **Exclude** control statements to direct SCRT to exclude specific SMF and SCRT89 data from being processed, did you also provide an explanation for each **Exclude** control statement in the appropriate **Customer Comments** fields?

You can do this either by editing the .csv file generated by SCRT or, after sending or uploading the .csv file to LMS, by using the LMS web interface. IBM will *not* accept sub-capacity reports that use **Exclude** control statements but have no matching customer comments explaining the reason for using each **Exclude** control statement.

14. Do any data collection incidents appear in the Missing LPAR Data Details (H4) or Missing CPC Data Details (H6) section of the sub-capacity report?

If so, verify whether SMF or SCRT89 records were collected for those LPARs. If such records were collected, then rerun SCRT with the additional records, together with the SMF or SCRT89 files used to generate the original report. If no SMF or SCRT89 records were collected for that LPAR, then use the **Justification for low data collection** column to either explain that a non-sub-capacity eligible operating system was running (for instance, Linux) for the entire reporting period or explain why SMF or SCRT89 records were not collected for a sub-capacity eligible operating system.

See the description of the **LPAR_Comment** parameter control statement in [“PARMS parameters that apply to specific CPCs” on page 37](#) for the case where the same explanation is expected to apply to that LPAR (and CPC) for future reporting periods (for instance, the LPAR is dedicated to running Linux).

15. Was the **Generate_Detailed_Data** control statement used on the SPECIAL DD statement? If so, remove the **Generate_Detailed_Data** control statement and rerun SCRT.

After reviewing all the information in the sub-capacity report, the software asset manager must generate an output file by saving the edited .csv flat file. IBM accepts sub-capacity reports **only** in the .csv file format.

To qualify for sub-capacity pricing, the software asset manager must submit the sub-capacity report to IBM, as described in [Chapter 26, “Submitting sub-capacity reports to IBM,” on page 427](#).

How IBM handles the sub-capacity report

When a sub-capacity report is received, IBM sends a letter of acknowledgement to the report submitter, using the email address in the report or the email address of the IBM Registration ID that performed the LMS submission. In some cases, the letter will request additional information. If you do not receive the acknowledgment by the 15th of the month, or if you have concerns or experience problems during this process, contact the LMS support team by sending an email to **kmsweb@dk.ibm.com**.

When IBM processes your report, any MSU values that were changed as a result of one or more **Exclude** control statements accompanied by an acceptable explanation for the **Exclude** control statement will be accepted. Any **Customer MSUs** values that you specify for a product with an (optional) placeholder accompanied by an acceptable explanation will override the **Tool MSUs** values. If billing adjustments are needed, they are reflected in the billing period. For example, a September report submitted between 2 October and 9 October has an effective date of 1 November.

Chapter 26. Submitting sub-capacity reports to IBM

You must use the License Management Support (LMS) application to submit your sub-capacity reports to IBM, as described in [“LMS web interface”](#) on page 427.

However, before you can submit your sub-capacity report, you must get access to the LMS application, as described in [“Getting access to the License Management Support \(LMS\) application”](#) on page 427.

When you get access to the LMS application, you will automatically be subscribed to LMS informational email about SCRT. However, even if you do not have access to LMS, you can still elect to receive this informational email. For more information about subscribing or unsubscribing, see [“Subscribing to LMS informational email”](#) on page 437.

Getting access to the License Management Support (LMS) application

Before you can submit or process your sub-capacity reports, do the following to get access to the LMS application:

1. If you do not already have one, obtain an IBM ID and password at [IBM id registration \(https://www.ibm.com/account/profile?page=reg\)](https://www.ibm.com/account/profile?page=reg).
2. Complete the LMS authorization form online at [LMS Access Request \(www.ibm.com/software/lms/access/\)](http://www.ibm.com/software/lms/access/).
3. IBM processes the authorization form, grants the requested access, and sends an automated email confirmation to the email address shown on your IBM Registration ID, with a copy to your manager.
Tip: If you use an email spam filter, add the zSCM@dk.ibm.com email address as a trusted sender.
4. Once you receive this confirmation, you will be able to sign in at [License Management Support \(www.ibm.com/software/lms\)](http://www.ibm.com/software/lms).

LMS web interface

The LMS web interface enables you to upload and submit your sub-capacity report.

The LMS web application accepts double-byte character set (DBCS) data for some of the input fields supported by LMS. The fields in which LMS accepts DBCS values are identified in the descriptions of the LMS panels later in this section. This DBCS support applies *only* to the countries or regions listed in [Table 56](#) on page 427.

Table 56. Countries, regions, and code pages for which LMS accepts DBCS values

Country or region	Code	Double-byte code
China	641	CHINA_ENCODING = "GB18030"
Japan	760	JAPANESE_ENCODING = "Shift_JIS"
Korea	766	KOREA_ENCODING = "EUC-KR"
Taiwan	858	TAIWAN_ENCODING = "BIG5"

LMS technical support: For technical support with the LMS submission process (such as for LMS application outages or application support to resolve errors), use the following information to contact the Customer Support Center. The LMS Customer Support Center provides support in English 24 hours a day, 7 days a week.

Telephone

+44 2392 28 98 80 (worldwide)

Email
kmsweb@dk.ibm.com

Accessing the LMS application

To access the LMS application, go to [License Management Support \(www.ibm.com/software/lms\)](http://www.ibm.com/software/lms). After you log in to LMS using your IBM ID and password, click **Data submission**, as shown in [Figure 261](#) on [page 428](#).

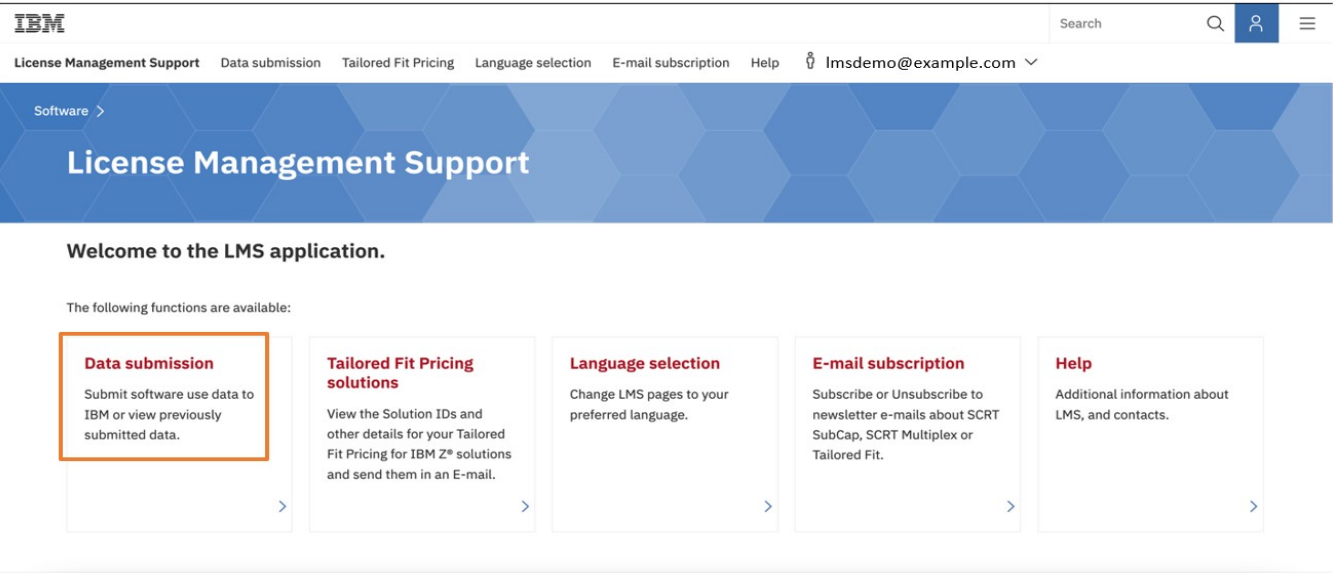


Figure 261. IBM License Management Support (LMS) web interface

The LMS Data Submission panel, shown in [Figure 262](#) on [page 428](#), lets you choose the LMS operation you wish to perform. You can view and, if necessary, modify uploaded reports, upload newly-created reports, or look at archived reports that you have previously submitted.

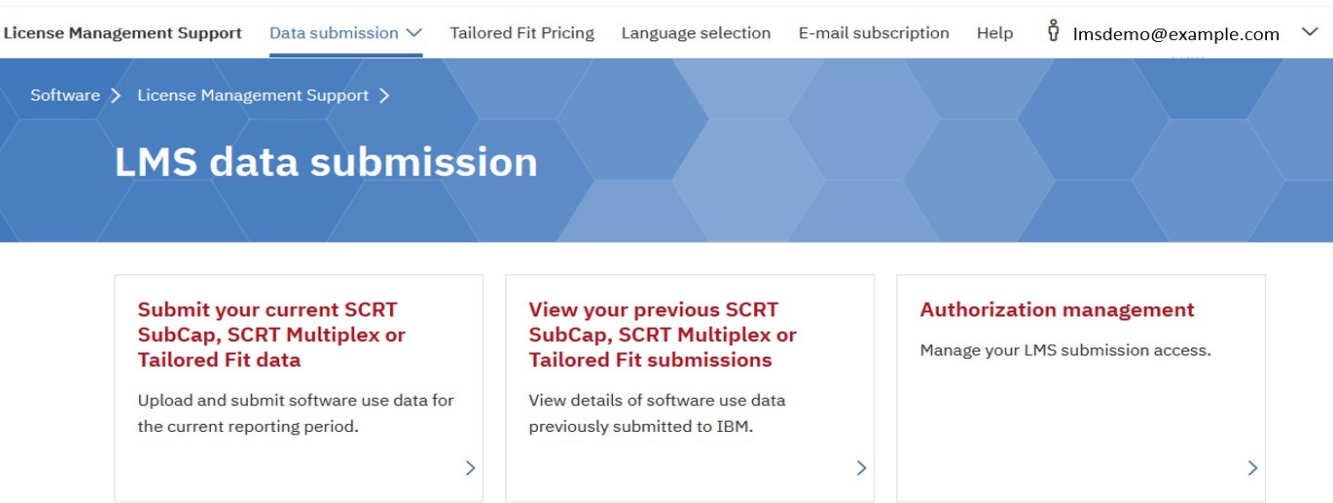


Figure 262. LMS web interface: LMS data submission panel

Viewing reports before submitting them

The system list panel shows you the reports that you have previously uploaded for your systems. If you have uploaded more than one sub-capacity report or if your report includes data for more than one CPC, you will see more than one system listed on this panel. The **Status** field shows reports that have already been submitted, reports that are available for submission, and reports that have errors and cannot be submitted. For information to help resolve errors, click **Help** in the navigation bar.

Systems with a status of **Work in progress** indicate reports that you previously saved as a draft. These reports have not yet been submitted to IBM.

Figure 263 on page 429 shows an example of the LMS system list panel.

You can select any system not yet submitted by 'clicking' on its number. You will have the opportunity to update your sub-capacity information before submitting it. After each report submission has been confirmed, the status field for that system will change to 'Submitted'.

System	Customer	Customer name	Status	Time stamp
8561-710 02-DEMO1	S012345678	LMS DEMO	Work in progress	3Apr2022 07:17 UTC
8561-710 02-DEMO2	S012345678	LMS DEMO	Ready for submission	
8561-710 02-DEMO3	S012345678	LMS DEMO	Submitted	3Apr2022 07:19 UTC

At the bottom of the table are two buttons: 'Work on new file' and 'Save as draft'.

Figure 263. LMS web interface: System list panel

To upload a new sub-capacity report, click **Work on new file**. The upload panel appears, as shown in Figure 264 on page 429. You can click **Browse** to find sub-capacity reports in spreadsheet format (comma-separated value format with a file extension of .csv) that are on your workstation hard drive or a network-attached drive accessible to your browser.

To upload your SCRT SubCap, SCRT Multiplex or Tailored Fit reports, use 'Browse' button.

After successful file upload, you will have the opportunity to review the data, and to provide additional information before final submission.

Please note:

- LMS submission window is open between 2nd and 9th day of the month inclusive.
- Comments for compliance conditions (i.e LPAR with 0% data collection) have to be provided during submission.

SCRT SubCap, SCRT Multiplex or Tailored Fit report file*

Buttons: **Browse** and **Continue**

Figure 264. LMS web interface: Uploading a new sub-capacity report

Submitting reports

After the sub-capacity report files have been uploaded, select and submit each report individually. If you select **Save as draft**, the uploaded report files are stored for later submission and are not considered submitted by IBM.

You can use the OUTPUT DD statement on the SCRT JCL to put multiple reports into a single output file, so there is only one file to upload. See [Chapter 3, “Using SCRT on z/OS systems,” on page 21](#) for information on directing SCRT output to a single sequential data set.

If you did not modify the SCRT JCL to put multiple reports in a single file, you can still upload the separate reports as a single file by putting them in a single directory on your workstation and issuing the following command (from a Windows command prompt) to concatenate them:

```
copy /b *csv AllinOne.csv
```

where *AllinOne* is any name you give to the concatenated file. You can now upload the *AllinOne.csv* file.

After you upload your reports, click **Continue**. The **Status** field shows Ready for submission on the system list panel, as shown in [Figure 265 on page 430](#).

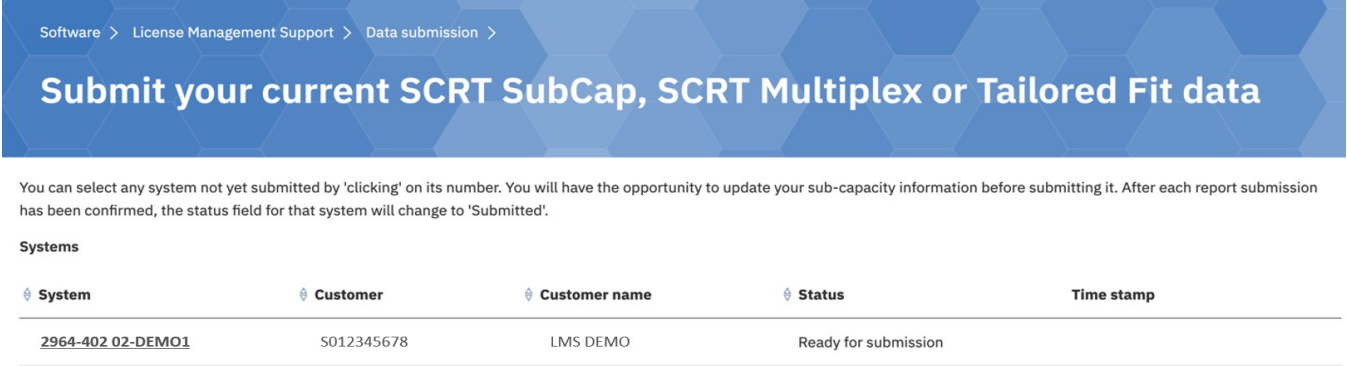


Figure 265. LMS web interface: System list panel - Uploaded report ready for submission

After the sub-capacity report files have been uploaded, verify and submit each of the reports individually. If you click **Save as draft**, the uploaded report files are stored for submission at a later time; they are not yet considered submitted by IBM.

To verify and submit a report for a system that you have not yet submitted, click its system ID in the system list (see [Figure 265 on page 430](#)). The report is displayed, as shown in [Figure 266 on page 431](#).

Submit your current SCRT SubCap, SCRT Multiplex or Tailored Fit data

Please edit / update as needed for the machine 2964-402 02-DEMO1 and select 'continue'.

The 'Customer/Submitter' information requested below is required to process your sub-capacity transactions with IBM. This data will be used by us (or by others involved in fulfillment) to fulfill your requests. It may also be used to contact you for customer satisfaction surveys, market research or in connection with certain other related transactions such as notification of updates to the tool or process.

By submitting SCRT SubCap, SCRT Multiplex and/or Tailored Fit data to IBM, you agree that IBM will collect and use certain Personal information (PI) about you. This information includes email address and/or IBM ID used for the submission, your name, email address and phone number as included in SCRT reports. Collected PI will be used to identify your data submission, should a problem occur and you reach out to IBM for support. More information on Data privacy at IBM, please visit <https://www.ibm.com/privacy/details/us/en/>

PI collected will not be shared or discussed with any third-parties, nor with anyone inside of IBM other than those who have a business need for the information, subject to appropriate confidentiality obligations.

Should you disagree with this practice, do NOT submit any SCRT SubCap, SCRT Multiplex or Tailored Fit report.

2964-402 02-DEMO1

Customer/Submitter information

Submitter name

LMS DEMO

Submitter e-mail

lmsdemo@example.com

Submitter phone

800-555-0123

Customer name

LMS DEMO

Customer number

US-S012345678

Machine details

Figure 266. LMS web interface: Viewing and editing customer and machine information in the sub-capacity report

This section of the report shows the customer and configuration data you entered in the JCL used to run SCRT, along with some fields that were generated by the tool. You can change the data in the **Submitter name** and **Submitter phone** fields. The **Submitter name** field supports DBCS data for the countries, regions, and code pages listed in Table 56 on page 427. The **Submitter phone** field only accepts SBCS numeric digits. (For details about all the fields shown in the report, see “Sample sub-capacity report for single-version z/OS native systems” on page 105.)

Figure 267 on page 431 shows the Machine details and Tool information section of the sub-capacity report.

Machine details

Machine type and model

2964-402

Machine serial

02-DEMO1

Machine rated capacity

60

Machine Peak Utilization

14

Machine MSU Consumed

5718

Tool information

Tool release

28.2.0

Reporting period

2 Sep, 2022 - 1 Oct, 2022 inclusiv

Figure 267. LMS web interface: Viewing and editing machine details and tool information in the sub-capacity report

The next section of the report, shown in Figure 268 on page 432, contains the Product Summary Information section. This section lists all the products from this sub-capacity report.

Product summary information				
MLC Product name	Product ID	Tool MSU	Customer MSU	Customer comment
z/OS V2	5650-ZOS	14		
CICS TS for z/OS V5	5655-Y04	12		
IBM MQ for z/OS V9	5655-MQ9	14		
IBM Enterprise Cobol for z/OS V4	5655-S71	14	<input type="text"/>	<input type="text"/>
IPLA Product name	Product ID	Tool MSU	Customer MSU	Customer comment
z/OS Unreduced	(All)	14		

Figure 268. LMS web interface: Viewing and editing product summary information in the sub-capacity report

Under the special circumstances for the **Customer MSUs** field, as described in “Customer MSUs” on page 111, you can enter a whole-number value with a maximum of 4 digits in the **Customer MSU** field for one or more products. If you do so and the value you enter is smaller than the value in the **Tool MSU** field, you must explain the reason for each value you entered in the corresponding **Customer Comment** field. The **Customer Comment** field supports DBCS data for the countries, regions, and code pages listed in Table 56 on page 427.

Figure 269 on page 432 shows the Special conditions section of the sub-capacity report.

Special conditions					
SMF/SCRT89 Exclusion Control Statements	Image ID	Product ID	Start	Resume	*Customer Comments (255 chars max)
Exclude Control Statement	*ALL	*ALL	02 Sep 2022 - 00:00	11 Sep 2022 - 08:00	<input type="text"/>
SMF/SCRT89 Exclusion Results	Image ID	Product ID	Start	Resume	Hours
SMF/SCRT89 records excluded	*ALL	*ALL	02 Sep 2022 - 00:00	11 Sep 2022 - 08:00	0

Figure 269. LMS web interface: Viewing and editing special conditions information in the sub-capacity report

Whenever you use one or more SCRT **Exclude** control statements, the resulting sub-capacity report will contain a Special Conditions section which lists each **Exclude** statement and the results for that statement. You must provide a customer comment explaining the reason for using each **Exclude** statement. LMS will not submit this sub-capacity report until you provide an explanation for every input field listed in this section of the report.

Whenever such **Exclude** related comments are for using Solution Edition LPARs, those comments must match the title of the Solution Edition that is being supported. For that case, the comments must be in single-byte US English. For all other uses of the **Exclude** statement, the comment field on this panel supports double-byte characters for the countries, regions, and code pages listed in Table 56 on page 427. If you need to supply a Solution Edition explanation but you are using LMS in DBCS mode, see “Considerations for viewing and editing the sub-capacity report” on page 419.

Figure 270 on page 433 shows the Detailed Data Collection section of the sub-capacity report.

Missing LPAR data			
LPARs with unmatched SMF70 and SMF89 records	Detected	Resolved (missing hours)	*Justification for low data collection (255 chars max)
SYS1	03 Sep 2022 - 20:00	03 Sep 2022 - 21:00 (1 hours)	1 hour missing data (no SMF70 data collected)
Total Missing Hours		3	
Active LPARs with no SMF/SCRT89 records	Detected	Resolved (missing hours)	*Justification for low data collection (255 chars max)
SYS3	03 Sep 2022 - 20:00	03 Sep 2022 - 22:00 (2 hours)	Example comment
Total Missing Hours		3	
Estimated utilization per Operating system family		Product ID	Est Peak MSU
z/OS V2		5650-ZOS	14
Missing CPC data			
CPC with no SMF/SCRT89 data from any LPAR	Detected	Resolved (missing hours)	*Justification for low data collection (255 chars max)
2817-99B76	02 Sep 2022 - 00:00	03 Sep 2022 - 20:00 (44 hours)	Example comment
Total Missing Hours		44	

Figure 270. LMS web interface: Viewing and editing missing LPAR data and missing CPC data in the sub-capacity report

Whenever the Missing LPAR data section or Missing CPC data section indicates that a data collection incident was detected, you must provide a comment to explain the missing SMF or SCRT89 data for that LPAR or CPC. LMS will not submit this sub-capacity report until you provide an explanation for each LPAR and CPC for which a data gap was detected. These comment fields support double-byte characters for the countries, regions, and code pages listed in [Table 56 on page 427](#). For the case where the LPAR is a Solution Edition LPAR, see the **LPAR_Comment** control statement in [“PARMS parameters that apply to specific CPCs” on page 37](#).

In the comment field under the **General** section of the panel, you can add overall comments. This comment field supports DBCS data for the countries, regions, and code pages listed in [Table 56 on page 427](#). When you add overall comments, enter comments that relate to the overall sub-capacity report, such as:

- Information that this report refers to an upgraded CPC (include any change in type-serial number)
- Information that this report refers to a system that consolidates two previously owned CPCs (include the type-serial numbers from the previous CPCs, as well as the type-serial number for the current CPC)
- Information about any special conditions that require an explanation, as described in [“Special Conditions section” on page 107](#)

All comment fields are limited to 255 characters.

After you have verified that the sub-capacity report contains all the necessary information, click **Continue** to continue the submission process, or click **Save as draft** to save your changes to the report without submitting it to IBM. You can click **Cancel** at any time to cancel your changes and close the report.

Confirming and submitting sub-capacity reports

Before the report is submitted, it is displayed again, showing any changes you made. You can select **Submit**, **Back**, **Modify distribution list**, or **Save as draft**.

Figure 271 on page 434, Figure 272 on page 435, Figure 273 on page 436, and Figure 274 on page 436 show the LMS confirm and submit panel.

License management data submission - confirm page

Please confirm that the following data is to be submitted for the machine 2964-402 02-DEMO1

Control number 1665498381612 2964-402 02-DEMO1

Customer/Submitter information

Submitter name	Submitter e-mail	Submitter phone
LMS DEMO	lmsdemo@example.com	800-555-0123
Customer name	Customer number	
LMS DEMO	US-S012345678	

Machine details

Machine type and model	Machine serial number	Machine rated capacity
2964-402	02-DEMO1	60
Machine Peak Utilization	Machine MSU Consumed	
14	5718	

Tool information

Tool release	Reporting period
28.2.0	2 Sep, 2022 - 1 Oct, 2022 inclusive (30 days)

Figure 271. LMS web interface: Confirm and submit panel (Part 1 of 4)

Special conditions					
SMF/SCRT89 Exclusion Control Statements	Image ID	Product ID	Start	Resume	Customer Comments (255 chars max)
Exclude Control Statement	*ALL	*ALL	02 Sep 2022 - 00:00	11 Sep 2022 - 08:00	IBM submission no customer comment provided
SMF/SCRT89 Exclusion Results	Image ID	Product ID	Start	Resume	Hours
SMF/SCRT89 records excluded	*ALL	*ALL	02 Sep 2022 - 00:00	11 Sep 2022 - 08:00	0
Product summary information					
MLC Product name	Product ID		Tool MSU	Customer MSU	Customer comment
z/OS V2	5650-ZOS		14		
CICS TS for z/OS V5	5655-Y04		12		
IBM MQ for z/OS V9	5655-MQ9		14		
IBM Enterprise Cobol for z/OS V4	5655-S71		14		
IPLA Product name	Product ID		Tool MSU	Customer MSU	Customer comment
z/OS Unreduced	(All)		14		

Figure 272. LMS web interface: Confirm and submit panel (Part 2 of 4)

Submitting reports

Missing LPAR data

LPARs with unmatched SMF70 and SMF89 records	Detected	Resolved (missing hours)	Justification for low data collection (255 chars max)
SYS1	03 Sep 2022 - 20:00	03 Sep 2022 - 21:00 (1 hours)	1 hour missing data (no SMF70 data collected)

Total Missing Hours 3

Active LPARs with no SMF/SCRT89 records	Detected	Resolved (missing hours)	Justification for low data collection (255 chars max)
SYS3	03 Sep 2022 - 20:00	03 Sep 2022 - 22:00 (2 hours)	Example comment

Total Missing Hours 3

Estimated utilization per Operating system family	Product ID	Est Peak MSU
z/OS V2	5650-ZOS	14

Missing CPC data

CPC with no SMF/SCRT89 data from any LPAR	Detected	Resolved (missing hours)	Justification for low data collection (255 chars max)
2817-99B76	02 Sep 2022 - 00:00	03 Sep 2022 - 20:00 (44 hours)	Example comment

Total Missing Hours 44

Figure 273. LMS web interface: Confirm and submit panel (Part 3 of 4)

Other information

For recurring charge (MLC) products, the data supplied in this report will be used to adjust the billable MSUs in inventory for all MLC Products listed under the MLC Product Name column on this report. In accordance with our Agreement, IBM will treat a change in product licensed capacity as an order. If the MSUs have changed since the last report, software billing based on inventory MSUs will increase or decrease accordingly.

For One Time Charge (IPLA) products, the data supplied in this report will be used to bill those IPLA products listed under the IPLA Product Name column in this report which exceed your entitled capacity. In accordance with our agreement, IBM will treat the use of a product in excess of its entitled capacity as an order and you will be billed for the amount in excess of your entitlement.

This report is prepared by the System z customer identified above ("Customer") or its authorized designee, and such Customer is solely responsible for the contents and accuracy of this report. Specifically, IBM makes no representations or warranties regarding the contents or accuracy of this report. Any questions concerning this report should be directed to the customer.

☒ Check to receive Confirmation receipt via email

☐ Check to attach submitted report with the Confirmation receipt

Privacy Notice: Pursuant to General Data Protection Regulation (GDPR) requirements, IBM will attach submitted report to Confirmation receipt only upon client request/authorization

☐ Check to copy additional recipients on Confirmation receipt. To create or modify recipient list, please select 'Distribution List' below.

NOTE: The distribution list will be in effect for the current session only. It will not be saved after you log off.

Distribution List

Back

Submit

Figure 274. LMS web interface: Confirm and submit panel (Part 4 of 4)

If you see a mistake, click **Back**. The report is displayed in edit mode again, with all your previous additions and corrections still in place.

If you checked the box to include recipients from your distribution list, when the email receipt is sent, copies will be sent to the addresses on the distribution list. Click **Modify distribution list** to add recipients to the distribution list. The distribution list is valid for this session only.

When the report is correct and complete, click **Submit** to submit the report to IBM. You return to the system list panel, where the status of the system for this report changes to Submitted.

If you click **Save as draft**, you return to the system list panel, and the status of this system displays Work in progress and the report is not submitted to IBM.

Viewing previously submitted reports

You can click **View your previous SCRT SubCap, SCRT Multiplex or Tailored Fit submissions** from the LMS data submission panel to review the sub-capacity reports that you previously submitted using the LMS web interface. This ability can help you check trends or compare fields in the current month's report to the same fields in previous reports. You can either view all the reports for a particular reporting period, or all the reports for a particular machine. Select the appropriate tab (**By Reporting Period** or **By Machine**, as shown in [Figure 275 on page 437](#)), and then select the period or machine from the list presented on that tab.

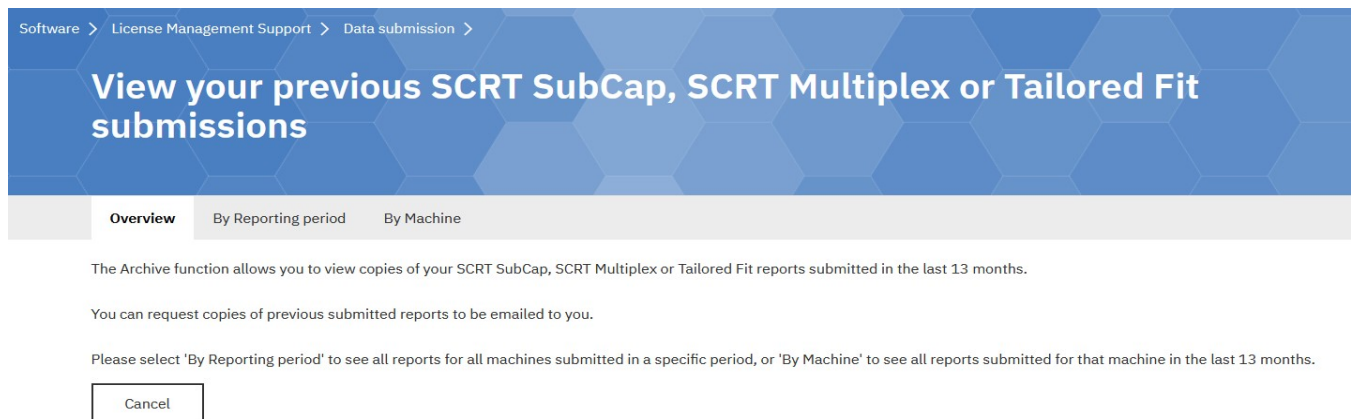


Figure 275. LMS web interface: Viewing previously submitted sub-capacity reports

Subscribing to LMS informational email

The email subscription function within the IBM LMS web application allows you to subscribe to or unsubscribe from informational email about SCRT from IBM, such as news about the target date for, availability of, and content in the next SCRT release.

About this task

If you are authorized to submit sub-capacity reports through the LMS web application, you are automatically subscribed to SCRT informational email. If you wish, you may unsubscribe.

If you are not authorized to submit sub-capacity reports through LMS but would like to receive SCRT informational email, you may use this procedure to subscribe.

LMS email subscriptions are available to anyone with an IBM id (formerly known as an IBM Registration ID or IBM ID). If you do not yet have an IBM id, you can get one by filling out a profile at [IBM id registration](https://www.ibm.com/account/profile?page=reg) (<https://www.ibm.com/account/profile?page=reg>).

Procedure

Perform these steps to subscribe to or unsubscribe from SCRT informational email.

1. Once you have your IBM id, go to [License Management Support \(www.ibm.com/software/lms\)](https://www.ibm.com/software/lms).
2. Click the **Sign in** button.

3. On the Sign in page, enter your **IBM id** and **Password**, then click **Sign in**.

If you are an IBM employee, click **Sign in with your intranet ID**, instead of **Sign in**.

When you are successfully signed in, you will see the IBM License Management Support page, where you can select your language, set your email subscription preferences, and get help with the LMS application. If you are authorized to submit sub-capacity reports, you will also see the data submission option.

4. Click **Email Subscription** to check your subscription settings and to subscribe to or unsubscribe from LMS informational email.
5. On the E-mail subscription management page, select the **SCRT newsletter subscription** check box to begin receiving informational email about the Sub-Capacity Reporting Tool, or clear the check box to stop receiving email.

If you are authorized to submit sub-capacity reports, the **SCRT newsletter subscription** check box should already be selected, indicating that you were automatically subscribed.

6. Click **Continue**.

You will return to the IBM License management support page. If you selected the **SCRT newsletter subscription** check box, you will receive a confirmation email and you will also see the following message:

Thank you for subscribing to the SCRT newsletter

7. If you are authorized, you can click **Data submission** to continue using the LMS application, or, if you are finished, click **Leave the LMS application**.

Chapter 27. Troubleshooting SCRT errors and problems

This chapter discusses the following SCRT troubleshooting topics:

- [“Obtaining technical support from IBM” on page 439](#)
- [“SCRT limitations” on page 439](#)
- [“Limitations that have been removed” on page 440](#)
- [“Machine replacements and machine software model changes” on page 441](#)
- [“z/OS issues message IWM063I” on page 442](#)
- [“Common SCRT problem scenarios” on page 442](#)
- [“Environments in which you can and cannot generate a full report” on page 447](#)
- [“SCRT return codes” on page 448](#)
- [“SCRT fixes” on page 439](#)

SCRT fixes

For information about APARs and other software fixes in the current version of SCRT, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](#).

Obtaining technical support from IBM

About this task

If you have a technical question or problem about SCRT, use the following procedure to obtain technical support from IBM. Also use this procedure whenever the SCRT documentation directs you to open a service request.

Procedure

1. Go to [IBM Support \(https://www.ibm.com/mysupport\)](https://www.ibm.com/mysupport).
2. Sign in with your IBMid.
If you don't have an IBMid, you can register for one.
3. Under **Support Cases**, click **Open a case**.
4. Specify a product/component, as follows:
 - **z/OS customers:** Enter one of the following keywords to locate the Sub-Capacity Reporting Tool:
 - SCRT
 - 5752SCRT2
 - sub-capacity reporting tool
 - **z/VSE customers:** Enter VSE.
5. Provided all other necessary information to complete the request.

SCRT limitations

In some situations, SCRT either does not have enough information to generate a precise report or is not designed to reflect a particular option. These situations and how the SCRT handles them are described here.

LPARs with z/OS V1R6 and earlier systems running as z/VM guests

In configurations with **any** z/OS guests at V1.6 or earlier levels, the SCRT is not able to determine the LPAR's 4-hour rolling average utilization for any z/OS guest system, including guest systems at z/OS V1.7. Instead, the SCRT substitutes the LPAR's maximum potential capacity for the LPAR's 4-hour rolling average utilization. The maximum capacity of an LPAR is equal to the MSU capacity of the machine divided by the number of CPs corresponding to that machine model, times the number of CPs configured for that LPAR.

SMF type 89 record-related limitations

If you run a product that does not produce SMF type 89 records at only certain times during the reporting period, or if you move such a product between LPARs during a reporting period, SCRT will not be able to reflect your utilization of that product with complete accuracy. You may have to enter a value for Customer MSUs in the sub-capacity report, as described under [“Product Summary Information section” on page 110](#). Use the Customer Comments field to enter a short explanation of the reason for the value you entered for Customer MSUs.

Products with features that can run in different LPARs

Licensed priced features of sub-capacity eligible products are not reported by SCRT. IBM calculates charges for licensed priced features based on the MSU value reported for the product base.

Multi-Version Measurement

When you install a new version of a licensed program, IBM allows you to run both the new and old versions of that product simultaneously for an unlimited duration. This is known as Multi-Version Measurement (MVM) for z/OS and z/VSE. The SCRT report output reflects the impact of MVM by showing a combined concurrent peak MSU value for multiple program versions in an eligible program family.

For more information about MVM for multiple program versions, see [Chapter 17, “Sub-capacity reporting and Multi-Version Measurement \(MVM\),” on page 367](#).

Single Version charging

In some cases, when you are installing a new version of a licensed program IBM allows you to run both the new and old versions of that product for a migration period and be charged via the new version. This is known as Single Version Charging (SVC). The SCRT output does not reflect the impact of SVC; it shows MSU values for both versions of the product. IBM must adjust your bill to reflect SVC. Contact your IBM representative for more information about products that are eligible for SVC.

Note that Multi-Version Measurement (MVM) for z/OS and z/VSE replaced SVC starting in SCRT V24.2.0 and V24.11.0. For more information, see [Chapter 17, “Sub-capacity reporting and Multi-Version Measurement \(MVM\),” on page 367](#).

Limitations that have been removed

This section lists limitations of previous versions of SCRT that have been removed.

Support for up to 128 native and guest systems

The total number of native plus guest systems that SCRT can process has been increased from 64 to 128, per CPC. For instance, a single execution of SCRT can process a total of 128 native plus guest systems from one CPC, and process a total of 128 native plus guest systems from a different CPC.

This support requires SCRT to use a larger **LRECL** value for the .csv file. See the description of the OUTPUT DD statement in [“OUTPUT DD statement” on page 64](#) for details about the change to the **LRECL** value.

LPARs with z/OS V1R7 and higher systems running as z/VM guests

As long as all of the z/OS guests running under z/VM in an IBM Z environment are at z/OS V1.7 or higher, SCRT reports the 4-hour rolling average utilization for each z/OS guest system.

Dedicated CPs and Wait Completion=Yes

In early versions of SCRT, configurations where some physical CPs were dedicated to certain workloads, rather than being shared, or where physical CPs had `Wait Completion=Yes` set, SCRT was not able to determine the LPAR's 4-hour rolling average utilization. Instead, SCRT substituted the LPAR's maximum potential capacity for the 4-hour rolling average utilization.

This limitation was lifted in version 6.1; SCRT now calculates the 4-hour rolling average utilization for LPARs with dedicated CPs or `Wait Completion=Yes`.

Machine replacements and machine software model changes

SCRT makes a distinction between machine (CPC) replacements and machine (CPC) model upgrades and will automatically create the appropriate sub-capacity reports whenever those conditions occur. SCRT automatically handles these conditions, so there is no need and no means to tell SCRT what to do for these situations. Note that the machine replacement condition does require you to add a specific set of customer comments to both the sub-capacity report that was generated for the machine being replaced and the sub-capacity report that was generated for the replacement machine.

Review each of the conditions to determine when you must add the specified customer comments to the sub-capacity reports for the replaced machine and for the replacement machine.

Machine software model changes

If the capacity of a CPC is upgraded or downgraded during a reporting period (that is, the type-serial number of the CPC remains the same, but its model number changes), SCRT generates a single sub-capacity report for that CPC, which lists the highest machine model observed during the reporting period. The machine model is followed by the message `Machine Model change observed`. The report reflects the highest 4-hour rolling average utilization for each product for all models of that CPC during that reporting period. Your billing will reflect the highest 4-hour rolling average utilization as shown in that report, regardless of whether that capacity was available to you for the full month or not.

Capacity changes can be temporary, as when an OOCOD or CBU action is taken. Capacity changes can also be permanent, as when the CPC software model change is permanently upgraded—that is, there is no intention to re-configure the CPC back to the prior CPC software model number.

If the message `Machine Model change observed` appears following the value in the **Machine Type/Model** field in section B5 of your sub-capacity report, you should add a comment in the optional **Customer Comments** field in that section to describe the reason for the model change. This is voluntary; however, doing so will help IBM to process your sub-capacity report as accurately and efficiently as possible.

- If the machine model change is temporary due to CBU or OOCOD, the suggested wording for your comment is (choose one of the bracketed reasons):

This machine was changed from model *xxx* to *yyy* for [CBU] [OOCOD] [CBU and OOCOD]. The permanent model remains *zzz*.

Example: This machine was changed from model 708 to 710 for CBU and OOCOD. The permanent model remains 708.

Note that either the *from* or *to* model values could be the same as the permanent model value. However, there may be cases when multiple upgrades occur over different reporting periods such that neither the *from* nor the *to* model value is the same as the permanent model value.

- If the machine model change is permanent, the suggested wording for your comment is:

This machine was permanently upgraded from model *xxx* to *yyy*.

Machine replacements

By definition, a machine (CPC) replacement occurs whenever the combination of CPC type (for example, 2097) and CPC serial number (for example, 02-12345) of the replaced (before) and replacement (after) CPCs are different. [Table 57 on page 442](#) shows examples of CPCs which have been replaced by other CPCs, since the combination of CPC type and CPC serial number of the before and after CPCs are different.

Table 57. Examples of conditions where CPCs have been replaced

Before (replaced) CPC Type and serial number	After (replacement) CPC Type and serial number	Replacement differences
2097 02-12345	2097 02-22222	CPC serial number changed
2097 02-12345	2817 02-12345	CPC type changed
2097 02-12345	2817 02-22222	CPC type and serial number changed

A CPC replacement can occur even when the same physical box (frame) is used by both the before and after CPCs. Note that the machine (CPC) software model (for instance, 710 for a 2097-710 CPC with serial number 02-11111) plays no role in determining whether a CPC has been replaced.

If a CPC is replaced by another CPC during a reporting period, SCRT generates one report for the before CPC and one report for the after CPC, regardless of whether the serial number of the CPC changes or not. Your billing will reflect the highest 4-hour rolling average utilization for the combined reports, regardless of whether that capacity was available to you for the full month or not.

Note that replacing a CPC usually means that one or both of the CPCs will indicate a data collection incident in report sections H4 or H6 for that reporting period. You must replace the (required) placeholder in the **Justification for low data collection** column in sections H4 or H6 of the sub-capacity reports for both CPCs with a comment, such as: CPC type-serial *tttt-xxsssss* was replaced by CPC type-serial *tttt-xxsssss* on *yyyy/mm/dd*.

z/OS issues message IWM063I

A z/OS system issues the following message:

```
IWM063I WLM POLICY WAS REFRESHED DUE TO A PROCESSOR SPEED CHANGE
```

whenever WLM detects a processor speed change for either of the following reasons:

- The processor is being dynamically upgraded by the customer to a processor type with a faster or slower processor speed. The processor type will change accordingly.
- The processor is experiencing a cooling problem and has reduced its speed in order to reduce its cooling requirements. The processor type and software model remain unchanged.

If you receive message IWM063I and believe that a cooling incident lasted longer than one hour, open a service request, as described in [“Obtaining technical support from IBM” on page 439](#), describing the condition and its duration and requesting guidance on how to generate the sub-capacity report for the affected CPC.

Common SCRT problem scenarios

The following topics describe some of the problem scenarios and error situations that frequently come up in the service requests that we receive, along with suggestions for how to resolve them.

Product MSUs exceed defined capacity value set for an LPAR

If you specify `Enforce relative weights = yes` on the HMC when you set up your LPARs, the defined capacity you set is ignored and the MSU values shown on the sub-capacity report can exceed the defined capacity you set up. For more information, see [“Capping individual LPARs” on page 497](#).

Data from the same CPC reported on more than one sub-capacity report

To generate a valid report, the SCRT needs to have all the SMF records for all the LPARs defined for a single CPC (one type-serial number). If you have collected the records for different LPARs in different data sets, all the data sets must be listed on the JCL that runs the SCRT (see [Chapter 3, “Using SCRT on z/OS systems,”](#) on page 21.)

If you do not do this and you accidentally produce multiple sub-capacity reports for one CPC, your input will be invalid and will not be accepted for billing purposes.

Changing values in the Tool MSUs field of the sub-capacity report or deleting SMF records

If you have an unusual situation, such as a disaster recovery, that causes your usage of a product to be very different from your normal usage, you should use the appropriate method (described in [“Providing alternate MSU values”](#) on page 443) to adjust the MSU value for each affected product.

If you change the **Tool MSUs** field instead of using the appropriate method, your report is invalidated and cannot be accepted for billing purposes.

If you delete SMF or SCRT89 records that you feel reflect exceptional usage from the input to SCRT, your report is invalidated and cannot be accepted for billing purposes.

Providing alternate MSU values

If you have an unusual situation, such as a disaster recovery, that causes your usage of a product to differ significantly from your normal usage, you should adjust your product usage in one of the following two ways, depending on the type of product:

- Whenever the product is instrumented on an SMF type 89 record (that is, it is not a NO89 product), use one or more SCRT **Exclude** control statements to exclude the appropriate product hours from the sub-capacity report and replace the (required) placeholder in the **Customer Comments** field with an explanation for the matching **Exclude** control statement in the Special Conditions section of the report. For more information, see [“Special Conditions section”](#) on page 107.
- Whenever the product is a NO89 product, you should replace the product's (optional) placeholder in the **Customer MSUs** field of the sub-capacity report with a corrected MSU value and replace the (conditional) placeholder in the matching **Customer Comments** field with a description of the situation that caused you to provide a value in the **Customer MSUs** field. For more information, see the descriptions of these fields in [“Product Summary Information section”](#) on page 110.

If you describe the situation that caused the unusual usage in an email or letter to IBM instead of editing the **Customer MSUs** and **Customer Comments** fields, your bill will not be adjusted to correct the unusual usage situation. Thus, it is very important to review your sub-capacity report each month.

Testing Capacity BackUp, disaster recovery, or OOCOD

Reporting intervals that have increased 4-hour rolling average utilization because of non-production testing of Capacity BackUp (CBU), non-production testing of disaster recovery procedures, or non-production testing of On/Off Capacity on Demand (OOCOD) (24-hour test period) are not subject to sub-capacity software charging.

Non-production testing is when you test a copy of production systems, but those copied systems are not processing useful workloads; in other words, those copied systems are not doing work. For instance, doing work includes, but is not limited to, production, development, program maintenance or testing, and data synchronization. For details about how IBM defines the term *doing work* under the backup use policies for the *IBM Customer Agreement* (ICA) and the *IBM International Program License Agreement* (IPLA), see [IBM Software Licenced Under the ICA - Backup Use Defined \(www.ibm.com/support/customer/csol/terms/?id=ICA-SW-Backup-DR&lc=en\)](#).

You may use **Exclude** control statements to exclude from SCRT processing those intervals during which such non-production testing was conducted. Note that the **Resume** value on the **Exclude** control

statement should be *three hours after* the testing was completed due to the continuing contribution of the testing activity to the 4-hour rolling average utilization over the next three hours. For a description of the **Exclude** control statement, see [“EXCLUDE control statement” on page 44](#).

IBM will accept the exclusion provided that a valid explanation has been given in the Customer Comments field for each **Exclude** control statement. You can enter such comments by editing the .csv report file or by using the LMS application.

CBU and Country Multiplex Pricing

The rules for excluding CBU testing differ slightly for customers who use Country Multiplex Pricing (CMP):

- CMP customers *may* use **Exclude** control statements to remove periods of non-production backup testing and non-production CBU testing from their multiplex reports.
- CMP customers *may not* use **Exclude** control statements to remove production-move CBU testing or any similar special-bid production-move testing from their multiplex reports.

Example scenarios

Example 1: Assume that CBU, disaster recovery testing, or OOCOD testing was done on a 2097 with a serial number of 45678 in two LPARs that are also used for purposes other than CBU, disaster recovery, or OOCOD testing:

- LPAR DEV1A, from 24 October 2024 at 7:30 AM to 25 October 2024 at 7:30 AM
- LPAR DEV2C, on 24 October 2024, between 9:00 AM and 11:45 PM

Then, appropriate **Exclude** control statements would be:

```
* Exclude hours +3 for CBU testing for DEV1A
Exclude CPC=2097-45678,Image_Id=DEV1A,
  Start=2024/10/24/07,Resume=2024/10/25/11
*
* Exclude hours +3 for CBU testing for DEV2C
Exclude CPC=2097-45678,Image_Id=DEV2C,
  Start=2024/10/24/09,Resume=2024/10/25/03
```

The **Resume** value for LPAR DEV1A is 11:00 AM, based on the fact that the testing ended before the start of the 8:00 AM hour, so the 4-hour rolling average utilization for hours 8:00 AM, 9:00 AM, and 10:00 AM needs to be excluded due to the testing that took place in that LPAR. The **Resume** value for LPAR DEV2C is 3:00 AM, based on the fact that the testing ended before 00:00 on 25 October, so the 4-hour rolling average utilization for the hour starting at 00:00 and for hours 1:00 AM and 2:00 AM needs to be excluded due to the testing that took place in that LPAR.

Example 2: Assume that the same CBU, disaster recovery, or OOCOD testing that was described in Example 1 was done on LPARs that were created for such testing and only used for such testing:

- LPAR DRTEST1 was used from 24 October 2024 at 7:30 AM to 25 October 2024 at 7:30 AM.
- LPAR DRTEST2 was used on 24 October 2024, between 9:00 AM and 11:45 PM.

While the **Exclude** control statements used in Example 1 could also be used for this example (provided that the **Image_Id** values of DEV1A and DEV2C are replaced by DRTEST1 and DRTEST2, respectively), there is an easier way to handle the case where the LPARs are dedicated to being used only for CBU, disaster recovery, and OOCOD testing. Because those LPARs are only used for this special purpose, you can simply set the values for the **Start** and **Resume** parameters to cover the entire reporting period, as shown in the following example **Exclude** control statements:

```
* Exclude for all CBU testing on DRTEST1
Exclude CPC=2097-45678,Image_Id=DRTEST1,
  Start=2024/10/02/00,Resume=2024/11/02/00
*
* Exclude for all CBU testing on DRTEST2
Exclude CPC=2097-45678,Image_Id=DRTEST2,
  Start=2024/10/02/00,Resume=2024/11/02/00
```

Unexpected product reported on a sub-capacity report

An unexpected product might be reported on a sub-capacity report for any of the following reasons:

The product was used.

SCRT correctly reported a product that was used. The reporting of this product on this CPC might be unexpected if the product was recently licensed and deployed, and this is the first sub-capacity report on which the product appears. Verify your usage of this product.

Another possibility is that the product is not licensed on this CPC but was run there as a result of a backup recovery situation. In that case, see [“Testing Capacity BackUp, disaster recovery, or OOCOD” on page 443](#) for the appropriate actions to take.

The product was not used but was reported because of incorrect information on an SCRT control statement.

The product is a NO89 product which has been reported on a sub-capacity report as being used, even though the product was not used. This is due to an error on the NO89 control statement that was customized for that product. That control statement indicated that the product was used during the reporting period. If that product was never used on any LPAR on that CPC during the reporting period, correct the NO89 control statement for that product to use a value of *NONE and rerun SCRT to create a corrected sub-capacity report. Submit the corrected sub-capacity report to IBM. See [“NO89 DD statement” on page 61](#) for a description of the NO89 control statement.

The product used a function that was licensed under a different product.

A number of IBM products support the same IBM debug function and must be properly configured so that the shared debug function identifies itself under the appropriate product ID. These products include: C/C++, Cobol, PL/I, VAJAVA, VAPLI, and the Debug Tool. The installation instructions for each product describe how to configure the IFAPRDxx member of SYS1.PARMLIB so that the debug function correctly identifies itself as being a part of that product. Failure to correctly configure the IFAPRDxx member either prevents the debug function from running or causes the debug tool to run but to identify itself as having run under one of the other products that share the debug function. To correct this problem, configure the IFAPRDxx member according to the installation instructions for the particular licensed product that uses the shared debug function. Note that SCRT will continue to report the incorrect product ID until the first full month that the IFAPRDxx member is correctly configured.

The product is an old version of a product that should no longer be used.

One or more workloads are still using a function provided by an older version of an IBM product that should no longer be in use. This usually happens because a workload's JOBLIB, STEPLIB, or LNKST product library is pointing to an older version of the product, usually an older version of CICS, Db2, or IMS. As a result, SCRT reports that an old version of CICS, Db2, or IMS is being used because those older product versions are being used by some jobs.

The only way to address this problem is to identify the workloads (that is, the z/OS jobs) that are using the older versions of these products and change the JCL and/or LNKST for those jobs to point to the libraries for the appropriate versions of the products. Note that SCRT will continue to report the old product until the first full month in which the old version of the product is no longer used.

To identify the jobs that are using the old version of a product, use the information found in SMF type 30 subtype 4 (SMF30-4) records. The **SMF30UPI** field contains the product ID. Look in this field for the product ID of the old version of the product that should no longer be used. Once you find an SMF30-4 record whose **SMF30UPI** field contains the product ID of an old product, look in the **SMF30JBN**, **SMF30STM**, and **SMF30PGM** fields on that record to identify the job name, step name, and program name, respectively, that invoked the old product. The **SMF30STD** and **SMF30SIT** fields on that record contain the date and time that the step started. After you identify all of the jobs that are invoking old versions of products, you can correct the problem by updating the JCL (JOBLIB, STEPLIB) and/or LNKST used by those jobs to point to the appropriate versions of the products.

If you have access to the MXG tool, you should be able to use it to carry out the technique described earlier to identify all of the jobs that are using the old versions of products simply by using the original SMF files that were input to SCRT. If you do not have access to the MXG tool, you can use the following manual method to locate all of the jobs:

1. In order to reduce the volume of data to be searched and separate it into more manageable files, use the Product Grid Snapshot section of the sub-capacity report to determine the LPARs in which the old product ran and then use the IFASMFDP utility to create separate files for the SMF records for each of those LPARs. Use the utility to place only the SMF30-4 records in each of those files in order to minimize the file size for the steps to follow.
2. Use the IBM IDCAMS utility to create a separate file with the hexadecimal printout (SYSPRINT output) for each of the SMF files created in step 1.
3. Log on to TSO/E and use the ISPF **Browse** function to open a file that contains the hexadecimal printout and use the Find command to locate each of the SMF30-4 records that contains the eight-character product ID (*xxxx-yy*) of the old product.
4. Repeat step 3 for each of the files created in step 2 so that you can identify all of the jobs that need to be changed to use the appropriate product versions.

SCRT generates sub-capacity reports instead of a multiplex report

If the output from SCRT is a set of sub-capacity reports instead of the single multiplex report that you expected, it is likely that the **Country_Multiplex_Pricing** control statement was not recognized by SCRT. Any SPECIAL DD control statements must immediately follow the SPECIAL DD line. JCL will not pass control statements to SCRT if they follow JCL comment lines (lines that start with */**). If the JCL comment lines are still there after uncommenting the **Country_Multiplex_Pricing** control statement, the statement will not take effect.

SCRT JAR file not found

If the correct location to the SCRT JAR file is not specified in the APP_HOME variable in the SCRT JCL, the JZOS batch launcher issues a class-not-found exception, as in the following example:

```
java.lang.NoClassDefFoundError: com.ibm.scrt.SCRTe
Caused by: java.lang.ClassNotFoundException: com.ibm.scrt.SCRTe
  at java.net.URLClassLoader.findClass(URLClassLoader.java:600)
  at java.lang.ClassLoader.loadClassHelper(ClassLoader.java:777)
  at java.lang.ClassLoader.loadClass(ClassLoader.java:750)
  at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:326)
  at java.lang.ClassLoader.loadClass(ClassLoader.java:731)
```

If you encounter this error, perform the following steps before contacting IBM Support:

1. Ensure that the z/OS UNIX System Services file system is accessible from the system where you ran SCRT.

To ensure that the SCRT job runs on a system where the file system is accessible, you might need to set the system affinity (**SYSAFF**) in the JCL to target the job to the necessary system.

Example: To ensure that the SCRT job runs on system SY1, specify:

```
/*JOBPARM SYSAFF=SY1
```

2. Ensure that the **APP_HOME** environment variable specifies the fully-qualified directory path where SCRT is installed.

The installation path is the location of the `scrt.jar` file. To display the fully-qualified path, navigate to the location of the `scrt.jar` file and issue the **pwd** command. Ensure that the displayed path matches the **APP_HOME** value.

Out-of-memory conditions

If the value specified for the JVM maximum heap size is not large enough for SCRT to process data for the reporting period, and out-of-memory exception may occur, as in the following example:

```
IBM Sub-Capacity Reporting Tool v.r.m
Build id: yyyyymmdd-tttt
Out of memory: Java heap space
```

If you are using the 31-bit SCRT JCL, switch to using the 64-bit SCRT JCL instead.

If you are using the 64-bit SCRT JCL and you encounter this condition, increase the value of the **-Xmx** parameter in the STDENV statement.

Ensure that the JOB statement in the JCL specifies the REGION=0M parameter.

Another common cause of an out-of-memory condition is that an installation policy is restricting the amount of available memory. To correct this, the z/OS UNIX **MEMLIMIT** parameter might need to be adjusted. Consult your system programmer for assistance.

Using a Java version earlier than the minimum required

If you are using a Java version that is earlier than the minimum version required by SCRT, you may see one of the following messages in the job output:

- Java 8 or higher is required to run SCRT
- java.lang.UnsupportedClassVersionError: JVMCFRE003 bad major version;
class=com/ibm/scrt/SCRTe, offset=6

You can issue the following command from the z/OS UNIX command line to display the Java version that is in use:

```
java -version
```

See [“Requirements for SCRT” on page 21](#) for more information.

Environments in which you can and cannot generate a full report

SCRT does *not* generate a sub-capacity report under the following circumstances:

- Required SCRT control statements are missing or if SCRT control statements have errors
- SMF type 70 subtype 1 records are present for an LPAR but SMF type 89 subtype 1 and subtype 2 records are missing for that LPAR, or SMF type 89 subtype 1 and subtype 2 records are present for an LPAR but SMF 70 subtype 1 records are missing for that LPAR
- For z/TPF-only systems, SCRT89 records are not present in the input data set
- SCRT encounters an input data set containing data only from a non-IBM Z CPC
- A CPC is running in a bimodal environment
- A CPC is running in basic mode (no LPARs are defined on it).

Note that if you accidentally submit the same data twice, you will still get a correct report. The SCRT discards duplicate records.

See [Table 58 on page 447](#) for other possible data collection environments and whether or not SCRT will generate a sub-capacity report.

Table 58. SCRT reporting behavior in various hardware and operating system environments

Operating system	IBM Z hardware, normal environment	Non-IBM Z hardware	Multiple mainframe CPCs	IBM Z hardware, bimodal environment
z/OS	analyze data, produce full report	no report generated	analyze data, produce report for each IBM Z CPC in input data set	analyze data, produce abbreviated report (cannot be used for billing)

Table 58. SCRT reporting behavior in various hardware and operating system environments (continued)

Operating system	IBM Z hardware, normal environment	Non-IBM Z hardware	Multiple mainframe CPCs	IBM Z hardware, bimodal environment
z/TPF	analyze data, produce full report	no report generated	analyze data, produce report for each IBM Z CPC in input data set	analyze data, produce abbreviated report (cannot be used for billing)
z/VSE	analyze data, produce full report	not applicable	analyze data, produce report for each IBM Z CPC in input data set	not applicable
Mixed z/OS and OS/390	analyze data, produce abbreviated report (cannot be used for billing)	no report generated	analyze data, produce abbreviated report (cannot be used for billing) for each IBM Z CPC in input data set	analyze data, produce abbreviated report (cannot be used for billing)
OS/390	analyze data, produce error report without the Product Summary Information section (cannot be used for billing)	no report generated	analyze data, produce error report without the Product Summary Information section (cannot be used for billing)	not applicable
MVS	analyze data, produce error report without the Product Summary Information section (cannot be used for billing)	no report generated	analyze data, produce error report without the Product Summary Information section (cannot be used for billing)	not applicable

SCRT return codes

The following return codes are associated with SCRT processing:

Return code

Description

0

A report was successfully produced for each CPC in the input.

4

A report was successfully produced for each CPC in the input, but at least one warning message was issued for one or more of those reports. For instance, corrupted SMF records were detected and deleted for at least one CPC. (See error message “[SCRTTOOL068](#)” on [page 456](#) for more information.)

In other cases, follow the instructions in the accompanying SCRTTOOLxxx message.

8

At least one report was successfully produced and at least one report could not be produced.

16

No reports could be produced.

If there is any valid data in your input, SCRT creates a report showing results based on the data it had. To ensure you have enough valid data for a full report, check the information in sections H4 and H6 of the sub-capacity report.

Appendix A. SCRT messages

This appendix contains an explanation of all messages issued by the Sub-Capacity Reporting Tool (SCRT) and the SCRT89 Record Utility. This includes:

- “SCRT messages for z/OS systems” on page 449
- “SCRT messages for internal SCRT restrictions” on page 494
- “SCRT89 Record Utility messages” on page 495

SCRT messages for z/OS systems

This section lists the messages that may be issued from the SCRT deliverable that runs on a z/OS system.

SCRT VERSION *vv.r.m* *****

Explanation:

Indicates the version (*vv*), release (*r*) and modification level (*m*) of the SCRT deliverable. This is the first message written to SYSPRINT.

User response:

None. This message is informational only.

SCRTTOOL001

ERROR: CUSTOMER NAME PARAMETER OMITTED

Explanation:

You did not specify a customer name in the PARMS DD statement of the SCRT JCL. This is a mandatory parameter. The SCRT exits with return code 16. A sub-capacity report is not generated.

User response:

Specify a customer name in the PARMS DD statement and then rerun SCRT.

SCRTTOOL002

ERROR: UNABLE TO OPEN SMF DD STATEMENT

Explanation:

The SCRT could not open the SMF input data set. The SMF DD statement may not exist, or it may exist but the data set specified on it may not be cataloged.

User response:

Make sure the SMF DD statement exists and is correctly cataloged. Then rerun SCRT.

SCRTTOOL003

PROCESSING COMPLETE

Explanation:

The SCRT has completed its processing. This message is issued even if this SCRT run was not successful.

User response:

None. This message is informational only.

SCRTTOOL004

ERROR: UNABLE TO OPEN OUTPUT DD STATEMENT

Explanation:

The SCRT was unable to open the output data set. The OUTPUT DD statement may not exist, or it may exist but the data set specified on it may not be valid.

User response:

Make sure the OUTPUT DD statement exists and is correct. Then rerun SCRT.

SCRTTOOL005

PROCESSED *type-serial* INTO MEMBER #*serial*

Explanation

The SCRT has processed the CPC identified by *type-serial* into the output data set (PDS member) identified by #*serial*.

In the message text:

type

The four-digit CPC machine type. For instance, 9175, 3932, 3931, 8562, 8561, 3907, 3906, 2965, 2964, 2828, 2827, 2817, 2818, 2098, 2097, 2096, 2094, 2086, 2084, 2066, or 2064.

serial

The seven-digit serial number of this CPC. The first two digits represent the plant (location) where the CPC was manufactured.

User response:

None. This message is informational only.

SCRTTOOL006	SYSID <i>sysname</i> IS RUNNING AN INELIGIBLE LEVEL OF THE Z/OS FAMILY: <i>type</i>-<i>serial</i>
--------------------	--

Explanation

The system indicated by SYSID *sysname* is running MVS or OS/390. The CPC indicated by *type*-*serial* is not eligible for sub-capacity pricing for z/OS and its middleware.

In the message text:

sysname

The system name from the **SID** parameter in the SMFPRMxx member for this LPAR.

type

The four-digit CPC machine type. For instance, 9175, 3932, 3931, 8562, 8561, 3907, 3906, 2965, 2964, 2828, 2827, 2817, 2818, 2098, 2097, 2096, 2094, 2086, 2084, 2066, or 2064.

serial

The seven-digit serial number of this CPC. The first two digits represent the plant (location) where the CPC was manufactured.

User response:

To be eligible for sub-capacity pricing for z/OS on this CPC, the CPC must not have an MVS or OS/390 license or run MVS or OS/390 for the entire reporting period.

SCRTTOOL007	MACHINE NOT FOUND IN TABLE. MODEL TYPE <i>type</i> MODEL NUMBER <i>model</i>
--------------------	---

Explanation

This CPC was not recognized by SCRT. One or more of the SMF or SCRT89 records processed by SCRT were created on a CPC with the machine type and model listed in the message. However, the level of SCRT that you are using does not support records created on that CPC machine type and model.

In the message text:

type

The four-digit CPC machine type. For instance, 9175, 3932, 3931, 8562, 8561, 3907, 3906, 2965, 2964, 2828, 2827, 2817, 2818, 2098, 2097, 2096, 2094, 2086, 2084, 2066, or 2064.

model

The three-digit software model number which determines the software capacity of the machine.

User response:

If a newer level of SCRT is available from the SCRT home page, download and use that level to process this data.

SCRTTOOL009	ERROR: UNABLE TO OPEN PARMS DD STATEMENT.
--------------------	--

Explanation:

The SCRT could not read the PARMS DD statement.

User response:

Correct the PARMS DD statement and rerun the SCRT.

SCRTTOOL012	SYNTAX ERROR IN VMGUESTS PARAMETER FILE.
--------------------	---

Explanation:

There is a syntax error in the parameters you coded in the VMGUESTS parameter file.

User response:

Correct the error and rerun the SCRT.

SCRTTOOL015***nn type* RECORDS PROCESSED**

Explanation

This message is issued at the end of all SCRT processing to indicate the number of input records that were processed.

In the message text:

nn

The number of *type* records that were processed by SCRT

type

One of the following record types:

SMF 70

z/OS SMF type 70, subtype 1 records

SMF 89

z/OS SMF type 89 records

SCRT89

z/TPF or z/VSE SCRT89 records

For z/OS systems, SCRT issues this message twice, one for SMF type 70, subtype 1 records and one for SMF type 89 records.

User response:

None. This message is informational only.

SCRTTOOL017**NO DATA FOUND TO PROCESS**

Explanation:

SCRT did not find any SMF records or SCRT89 records to process. No report is generated, and the SCRT exits with return code 16.

User response

Make sure your input data set has SMF type 70 and type 89 records for z/OS systems, or SCRT89 records for z/TPF and z/VSE systems.

If the optional PARMS **Report_Period** control statement was used, make sure the input data contains SMF or SCRT89 records for the specified reporting period.

If one or more SPECIAL **Exclude** control statements were used, make sure those statements do not exclude all of the data either within the default reporting period or the reporting period specified by the **Report_Period** control statement.

If z/OS and z/TPF run on the same machine, you must provide the SMF records and optionally also provide SCRT89 records. SCRT uses the PR/SM data from the SMF records to calculate the MSU values for z/TPF and z/TPFDF for those LPARs that do not provide SCRT89 records.

SCRTTOOL025**VMIMAGE NAME OVER 8 CHARACTERS IN VMGUESTS PARAMETER FILE**

Explanation:

The value you specified for a VMIMAGE in the VMGUESTS parameter list is longer than eight characters.

User response:

Specify a value that is eight characters or less in the VMGUESTS parameter list in the SCRT JCL and rerun SCRT.

SCRTTOOL026**VMIMAGE NAME DUPLICATED IN VMGUESTS PARAMETER FILE.**

Explanation:

You specified two identical names for separate **VMIMAGE** names in the VMQUESTS parameter list.

User response:

Specify a different **VMIMAGE** name in the VMQUESTS parameter list in the SCRT JCL and rerun SCRT.

SCRTTOOL030**CONTINUATION NOT RECEIVED IN VMQUESTS PARAMETER FILE.****Explanation:**

A VMQUESTS statement contained a comma as the last character on the line, indicating the statement would continue to the next line, but the continued line was not found.

User response:

Correct the VMQUESTS statement in the SCRT JCL and then rerun SCRT.

SCRTTOOL031**SYSID *sysname* IN VMQUESTS NOT FOUND IN SMF DATA.****Explanation**

On the VMQUESTS DD statement you specified an SMF ID that does not exist.

In the message text:

sysname

The system name from the **SID** parameter in the SMFPRM~~xxx~~ member for this LPAR.

User response:

Specify a correct SMF ID in the VMQUESTS DD statement in the SCRT JCL and then rerun SCRT.

SCRTTOOL032**SYSID *sysname* IN VMQUESTS NOT RUNNING VM.****Explanation**

On the VMQUESTS DD statement you specified an SMF ID that is not associated with a z/VM LPAR.

In the message text:

sysname

The system name from the **SID** parameter in the SMFPRM~~xxx~~ member for this LPAR.

User response:

Specify a correct SMF ID in the VMQUESTS DD statement in the SCRT JCL and then rerun SCRT.

SCRTTOOL033**SYSID *sysname* IS VMQUEST. REQUIRED IN VMQUESTS PARAMETER FILE.****Explanation**

You are running at least one z/VM LPAR which you did not specify in the VMQUESTS parameter file.

In the message text:

sysname

The system name from the **SID** parameter in the SMFPRM~~xxx~~ member for this LPAR.

User response:

Make sure all z/VM LPARS are specified in the VMQUESTS DD statement in the SCRT JCL and then rerun SCRT.

SCRTTOOL036**VMQUESTS DD MUST BE SPECIFIED IF VM IS PRESENT IN INPUT.****Explanation**

Your SMF input included information for a CPC on which z/OS runs as a guest under z/VM. However, you did not specify the VMQUESTS DD statement.

User response

Specify the VMQUESTS DD statement in the SCRT JCL and then rerun SCRT.

Explanation

An invalid value has been specified for an SCRT parameter.

User response

Review the SCRT output for accompanying messages that indicate the source of the error. Correct the invalid parameter and rerun SCRT.

Explanation

Running SCRT under the version of Java indicated by *version* is deprecated. SCRT will continue processing as normal, but future versions of SCRT may require a newer Java runtime environment.

User response:

Update your Java runtime environment to a more recent level. The minimum and recommended Java levels are documented in [“Requirements for SCRT” on page 21](#) and [“Requirements for SCRT for Windows and Linux” on page 69](#).

Explanation

This message is issued in conjunction with SCRTTOOL046. It indicates a syntax error in the line of JCL given in SCRTTOOL046.

User response

Check the syntax in the JCL line given in message SCRTTOOL046 (see the sample JCL in [“Sample JCL to run the 64-bit SCRT on a z/OS system” on page 25](#)). Correct the JCL and rerun SCRT.

Explanation

LPAR names and guest system IDs specified on the NO89 DD statement in the SCRT JCL must either have data in the input stream being processed by this execution of SCRT, or the LPAR must be defined on the CPC and be inactive for the month. Either the LPAR name listed in the error message did not meet either of those requirements or the input stream did not contain any data for the guest system ID. If the product ID in the error message is for a z/OS product, no SMF records were provided for that LPAR or z/OS guest system. If the product ID in the error message is for a z/VSE product, no SCRT89 records were provided for that LPAR or z/VSE guest system. Your configuration may have changed, or there may be a typographical error in the NO89 DD statement.

In the message text:

lparname

The name of an LPAR or system ID of a guest system that you specified in a NO89 DD statement in the SCRT JCL.

productid

The seven-character product ID, in the form *nnnn-nnn*. See the sample JCL in Figure 1 on page 16 for the complete list of products that do not produce SMF type 89 records or SCRT89 records.

User response:

Correct the NO89 DD statement in the SCRT JCL either providing the correct LPAR name or guest system ID, or by removing the LPAR name or guest system ID from the NO89 control statement indicated in the error message.

If that LPAR name or guest system ID is the only *lparname* specified on that NO89 control statement, replace it with a value of *NONE. Then, rerun SCRT.

SCRTTOOL048**TYPE-SERIAL: *type-serial* IN NO89 DD DATA NOT FOUND IN SMF DATA**

Explanation

You have a product that does not produce SMF 89 records running in two or more LPARs with the same LPAR name on different CPCs. The LPARs are identified on the NO89 DD statement in the SCRT JCL by the type-serial numbers of the CPCs. You have entered a type-serial number that does not exist in your configuration. Your configuration may have changed or there may be a typographical error in the NO89 DD statement.

In the message text:

type

The four-digit CPC machine type. For instance, 9175, 3932, 3931, 8562, 8561, 3907, 3906, 2965, 2964, 2828, 2827, 2817, 2818, 2098, 2097, 2096, 2094, 2086, 2084, 2066, or 2064.

serial

The six-digit serial number of this CPC. The first two digits represent the plant (location) where the CPC was manufactured.

User response:

Correct the NO89 DD statement in the SCRT and rerun SCRT.

SCRTTOOL050***productid* IS NOT A NO89 PRODUCT. CHECK NO89 DD**

Explanation

A product with the product ID *productid* was listed in the NO89 DD statement in the SCRT JCL, but there is no NO89 product with this product ID. There might be a typographical error in the product ID, which you should correct. Or, you might be attempting to add a product ID that is not already defined by IBM as a NO89 product, which is not allowed. Check to make sure the product ID you are trying to indicate in the NO89 DD statement is a valid NO89 sub-capacity-eligible product.

In the message text:

productid

The seven-character product ID, in the form *nnnn-nnn*. See the sample JCL in [“Sample JCL to run the 64-bit SCRT on a z/OS system” on page 25](#) for a list of valid products that do not produce SMF type 89 records.

User response:

Correct the NO89 DD statement in the JCL and rerun SCRT.

SCRTTOOL051**ERROR IN PARMS: SYNTAX ERROR IN CPC= KEYWORD *control-statement***

Explanation:

There is an error in one of the control statements. The control statement in error (*control-statement*), or a portion of it, follows the message.

User response:

Review the SCRT input and correct the incorrect control statement.

SCRTTOOL052**ERROR IN PARMS: SYNTAX ERROR IN LPAR= KEYWORD *control-statement***

Explanation:

There is an error in one of the control statements. The control statement in error (*control-statement*), or a portion of it, follows the message.

User response:

Review the SCRT input and correct the incorrect control statement.

SCRTTOOL054***type-serial* PROCESSING *month, year text***

Explanation

This message is informational and tells you what records were processed in this invocation of the SCRT. The records being processed are from the CPC identified by *type-serial* and were created for the period identified by *month, year*.

In the message text:

type

The four-digit CPC machine type. For instance, 9175, 3932, 3931, 8562, 8561, 3907, 3906, 2965, 2964, 2828, 2827, 2817, 2818, 2098, 2097, 2096, 2094, 2086, 2084, 2066, or 2064.

serial

The seven-digit serial number of this CPC. The first two digits represent the plant (location) where the CPC was manufactured.

month

The month for which this report is being generated.

year

The year for which this report is being generated.

text

Optional text, (CUSTOMER OVERRIDE), indicating that the month and year were specified for this report using the **Report_Period** parameter of the PARMS DD statement.

User response:

This message is informational; no action is necessary.

SCRTTOOL055***type-serial* NO DATA TO PROCESS**

Explanation

This message usually indicates a mismatch between the reporting interval and the SMF records that are available to be processed.

In the message text:

type

The four-digit CPC machine type. For instance, 9175, 3932, 3931, 8562, 8561, 3907, 3906, 2965, 2964, 2828, 2827, 2817, 2818, 2098, 2097, 2096, 2094, 2086, 2084, 2066, or 2064.

serial

The seven-digit serial number of this CPC. The first two digits represent the plant (location) where the CPC was manufactured.

User response

Check that the reporting interval given in message SCRTTOOL54 is the correct interval. If it is not, check that your SMF records are from the correct interval. Correct your input and rerun the SCRT.

If you do not expect a report for the indicated CPC (for example, the CPC was retired), disregard the message. SCRT may have produced a report or reports for other CPCs. Check that reports have been generated for all expected CPCs.

SCRTTOOL056**ERROR READING SMF OR SCRT89 RECORDS FOR *filename***

Explanation

A malformed SMF or SCRT89 file was supplied to SCRT.

In the message text:

filename

The name of the file that SCRT failed to read.

User response

Ensure that the SMF or SCRT89 records were collected and transferred correctly, and rerun SCRT.

SCRTTOOL065	ALL VM GUESTS RUNNING z/OS 1.7 OR GREATER VMGUESTS DD IGNORED
--------------------	--

Explanation

Since all the z/OS guest systems running under VM on this CPC are at z/OS V1R7 or later releases, the VMGUESTS DD statement is not needed and the SCRT is ignoring it.

User response

No action is needed; you can remove the VMGUESTS DD statement from the JCL to avoid seeing this message.

SCRTTOOL068	nnnnnn INVALID SMF xx RECORDS DROPPED
--------------------	--

Explanation

The SCRT detected that an SMF record or records were shorter than the record length indicated by their RDW. The SCRT dropped the records from the report to prevent any problems in further processing.

In the message text:

nnnnnn

The number of SMF records that were shorter than the expected record length and dropped from further processing.

xx

The type of SMF record that the SCRT found to be invalid and dropped from the report.

User response:

Examine your SMF records and correct any that are not the expected record length

SCRTTOOL069	nnnnnn DUPLICATE SMF xx RECORDS DROPPED
--------------------	--

Explanation

SCRT detected a duplicate SMF record.

In the message text:

nnnnnn

The number of SMF records that were found to be duplicates and dropped from further processing.

xx

The type of SMF record that the SCRT found to be duplicated and dropped from the report.

User response:

Review your data collection procedures and input SMF specification to supply only unique SMF records to SCRT.

SCRTTOOL070	MULTIPLE MONITORING PRODUCTS WERE DETECTED FOR IMAGE <i>lparname</i> ON CPC type-serial
--------------------	--

Explanation

SCRT detected that SMF records from more than one monitoring facility (such as RMF and a third-party product) were supplied for one or more hours. For each reporting period hour, SCRT accepts records from the first monitoring facility processed. Subsequent records for the same hour from a different monitoring facility are dropped.

In the message text:

lparname

Then ame of an LPAR or system ID for which records from multiple monitoring products were detected.

type-serial

The machine type and serial number of the CPC running the image indicated by *lparname*.

User response:

Supply records from only a single monitoring facility as input to SCRT.

SCRTTOOL072	ERROR IN PARMS: REQUIRED KEYWORD (CONTACT_NAME) IS MISSING
--------------------	---

Explanation:

The **Contact_Name** keyword is missing from the PARMS DD statement. That keyword is required.

User response:

Examine your parameters on the PARMS DD statement and include the required **Contact_Name** keyword. For details, see “PARMS DD statement” on page 34.

SCRTTOOL073	ERROR IN PARMS: REQUIRED KEYWORD (CONTACT_EMAIL) IS MISSING
--------------------	--

Explanation:

The **Contact_Email** keyword is missing from the PARMS DD statement. That keyword is required.

User response:

Examine your parameters on the PARMS DD statement and include the required **Contact_Email** keyword. For details, see “PARMS DD statement” on page 34.

SCRTTOOL074	ERROR IN PARMS: REQUIRED KEYWORD (CONTACT_PHONE) IS MISSING
--------------------	--

Explanation:

The **Contact_Phone** keyword is missing from the PARMS DD statement. That keyword is required.

User response:

Examine your parameters on the PARMS DD statement and include the required **Contact_Phone** keyword. For details, see “PARMS DD statement” on page 34.

SCRTTOOL075	ERROR IN PARMS: REQUIRED KEYWORD (CUSTOMER_NAME) IS MISSING
--------------------	--

Explanation:

The **Customer_Name** keyword is missing from the PARMS DD statement. That keyword is required.

User response:

Examine your parameters on the PARMS DD statement and include the required **Customer_Name** keyword. For details, see “PARMS DD statement” on page 34.

SCRTTOOL076	ERROR IN PARMS: REQUIRED KEYWORD (CUSTOMER_NUMBER) IS MISSING
--------------------	--

Explanation:

The **Customer_Number** keyword is missing from the PARMS DD statement. That keyword is required.

User response:

Examine your parameters on the PARMS DD statement and include the required **Customer_Number** keyword. For details, see “PARMS DD statement” on page 34.

SCRTTOOL077	ERROR IN PARMS: REPORT_PERIOD DATE IS INVALID
--------------------	--

Explanation:

The date specified for the **Report_Period** keyword on the PARMS DD statement is not correct. That keyword is ignored.

User response:

Examine your parameters on the PARMS DD statement and provide a supported date for the specified **Report_Period**. The date must specify a year and month in the format *yyyymm*, or specify the value *Last_Month* to indicate that the report covers the previous month. For details, see “PARMS DD statement” on page 34.

SCRTTOOL080**SPECIAL CONTROL STATEMENT ERROR****Explanation:**

One or more of the SPECIAL DD control statements is in error. The specific error will be listed on one of the error messages that follows this message. In some cases, message SCRTTOOL119 will list the control statement in error.

User response:

Read the error message(s) that follow this one to determine the specific error and follow the User Response for the error message(s).

SCRTTOOL091**CPC= MUST BE SPECIFIED****Explanation:**

The **CPC=** keyword parameter, which is required for the **Exclude** control statement, is missing. Message SCRTTOOL119 lists the control statement in error.

User response:

Add the **CPC=** keyword and a valid CPC value to the **Exclude** control statement in error.

SCRTTOOL092**IMAGE_ID= MUST BE SPECIFIED****Explanation:**

The **Image_Id=** keyword parameter, which is required for the **Exclude** control statement, is missing. Message SCRTTOOL119 lists the control statement in error.

User response:

Provide the required **Image_Id=** keyword parameter together with a valid value. Note that a value of *ALL can be used to specify all of the images on this CPC. The use of the *ALL value for this parameter will likely result in a data collection incident being reported in section H4 or H6 of the sub-capacity report for the specified CPC.

SCRTTOOL093**BOTH START= AND RESUME= MUST BE SPECIFIED****Explanation:**

The **Start=** and **Resume=** keyword parameters are required for the **Exclude** control statement. Either or both keywords are missing from an **Exclude** control statement. Message SCRTTOOL119 lists the control statement in error.

User response:

Provide the required **Start=** and **Resume=** keywords together with valid values for each.

SCRTTOOL094**CPC= ALREADY SPECIFIED****Explanation:**

An **Exclude** control statement contains more than one instance of the **CPC=** keyword parameter. Each **Exclude** control statement supports only one **CPC=** parameter. Message SCRTTOOL119 lists the control statement in error.

User response:

Remove the additional instances of the **CPC=** keyword parameter. If multiple CPCs must be specified, use a separate **Exclude** control statement for each CPC. Note that the same CPC can appear on separate **Exclude** control statements.

SCRTTOOL095**CPC= VALUE LENGTH MUST BE 10, TTTT-SSSS****Explanation:**

The value for the **CPC=** keyword parameter on an **Exclude** control statement was not 10 characters long. The **CPC** value must be 10 characters long, in the form *tttt-ssss*. Message SCRTTOOL119 lists the control statement in error.

User response:

Provide a valid CPC value that is 10 characters long.

SCRTTOOL099**CPC= NOT FOUND IN THE DATA****Explanation:**

The value for the **CPC=** keyword parameter on an **Exclude** control statement was not found on any of the SMF or SCRT89 records processed by this execution of SCRT. Message SCRTTOOL119 lists the control statement in error.

User response:

Change the **CPC** value on the **Exclude** control statement so that it matches a CPC whose SMF or SCRT89 records are being processed by this execution of SCRT. To determine the CPCs that are processed by this execution of SCRT, run SCRT with the same SMF and SCRT files but without any **Exclude** control statements.

SCRTTOOL100 **IMAGE_ID= ALREADY SPECIFIED**

Explanation:

An **Exclude** control statement contains more than one instance of the **Image_Id=** keyword parameter. Each **Exclude** control statement supports only one **Image_Id=** parameter. Message SCRTTOOL119 lists the control statement in error.

User response:

Remove the additional instances of the **Image_Id=** keyword parameter. If multiple image IDs must be specified on this **Exclude** control statement, use the list form — that is, (*image_id1, image_id2, . . . , image_idn*) — to provide multiple **Image_Id=** values.

SCRTTOOL101 **IMAGE_ID= VALUE MUST BE 1 TO 8 CHARS**

Explanation:

The value for the **Image_Id=** keyword parameter is greater than eight characters long. The **Image_Id=** value provides the name of the operating system image to which the **Exclude** control statement applies and that name must be from one to eight characters long. Message SCRTTOOL119 lists the control statement in error.

User response:

Change the **Image_Id=** value so that it is from one to eight characters long and matches an image ID on the SMF or SCRT89 records that are being processed by this execution of SCRT. To determine the image IDs that are being processed by this execution of SCRT, run SCRT with the same SMF and SCRT files but without any **Exclude** control statements. All of the image IDs for each CPC will be listed in the first column of the Detail Data Collection section of the sub-capacity report for that CPC.

SCRTTOOL102 **IMAGE_ID= IN LIST SUBVALUE MUST BE 1 TO 8 CHARS**

Explanation:

One of the values in the list format of the **Image_Id=** keyword parameter is greater than eight characters long. The list of **Image_Id=** values provides the names of the operating system images to which the **Exclude** control statement applies; those names must be from one to eight characters long. Message SCRTTOOL119 lists the control statement in error.

User response:

Change the **Image_Id=** value so that it is from one to eight characters long and matches an image ID on the SMF or SCRT89 records that are being processed by this execution of SCRT. To determine the image IDs that are being processed by this execution of SCRT, run SCRT with the same SMF and SCRT files but without any **Exclude** control statements. All of the image IDs for each CPC will be listed in the first column of the Detail Data Collection section of the sub-capacity report for that CPC.

SCRTTOOL103 **IMAGE_ID= NOT FOUND ON THIS CPC**

Explanation:

The value for the **Image_Id=** keyword parameter on an **Exclude** control statement was not found on any of the SMF or SCRT89 records for this CPC being processed by this execution of SCRT. The **Image_Id=** value provides the name of the operating system image to which the **Exclude** control statement applies. Message SCRTTOOL119 lists the control statement in error.

User response:

Change the **Image_Id=** value on the **Exclude** control statement so that it matches an image ID on the SMF or SCRT89 records being processed for this CPC by this execution of SCRT. To determine the image IDs that are being processed by this execution of SCRT, run SCRT with the same SMF and SCRT files but without any **Exclude** control statements. All of the image IDs for each CPC will be listed in the first column of the Detail Data Collection section of the sub-capacity report for that CPC.

SCRTTOOL104**IMAGE_ID= IN LIST NOT FOUND ON THIS CPC****Explanation:**

One or more of the values in the list format of the **Image_Id=** keyword parameter was not found on any of the SMF or SCRT89 records for this CPC being processed by this execution of SCRT. The list of **Image_Id=** values provides the names of the operating system images to which the **Exclude** control statement applies. Message SCRTTOOL119 lists the control statement in error.

User response:

Change one or more of the **Image_Id=** values on the **Exclude** control statement so that they match an image ID on the SMF or SCRT89 records being processed for this CPC by this execution of SCRT. To determine the image IDs that are being processed by this execution of SCRT, run SCRT with the same SMF and SCRT files but without any **Exclude** control statements. All of the image IDs for each CPC will be listed in the first column of the Detail Data Collection section of the sub-capacity report for that CPC.

SCRTTOOL105**START= ALREADY SPECIFIED****Explanation:**

An **Exclude** control statement contains more than one instance of the **Start=** keyword parameter. Each **Exclude** control statement supports only one **Start=** parameter. Message SCRTTOOL119 lists the control statement in error.

User response:

Remove the additional instances of the **Start=** keyword parameter. If multiple **Start** or **Resume** date or time ranges must be specified, use a separate **Exclude** control statement for each range.

SCRTTOOL106**START= ERROR IN DATE FORMAT****Explanation:**

The value for the **Start=** keyword parameter on an **Exclude** control statement contains a formatting error. The **Start** value must be in the format *yyyy/mm/dd/hh*. Message SCRTTOOL119 lists the control statement in error. This message will be followed by a message that lists the specific format error.

User response:

Review the additional error message and follow the instructions provided in the user response for that message.

SCRTTOOL107**RESUME DATETIME MUST BE GREATER THAN START DATETIME****Explanation:**

The value for the **Resume=** keyword parameter on an **Exclude** control statement is earlier than the corresponding value for the **Start=** keyword. The **Resume** value must be greater (later) than the corresponding **Start** value. Message SCRTTOOL119 lists the control statement in error.

User response:

Change the value of the **Resume=** keyword so that it is at least one hour greater (later) than the value for the **Start=** keyword. Note that in order to exclude a single hour the **Resume** hour must be one hour greater than the **Start** hour.

SCRTTOOL109**RESUME= ALREADY SPECIFIED****Explanation:**

An **Exclude** control statement contains more than one instance of the **Resume=** keyword parameter. Each **Exclude** control statement supports only one **Resume=** parameter. Message SCRTTOOL119 lists the control statement in error.

User response:

Remove the additional instances of the **Resume=** keyword parameter. If multiple **Start** or **Resume** date or time ranges must be specified, use a separate **Exclude** control statement for each range.

SCRTTOOL110**RESUME= ERROR IN DATE FORMAT****Explanation:**

The value for the **Resume=** keyword parameter on an **Exclude** control statement contains a formatting error. The **Resume** value must be in the format *yyyy/mm/dd/hh*. Message SCRTTOOL119 lists the control statement in error. This message will be followed by a message that lists the specific format error.

User response:

Review the additional error message and follow the instructions provided in the user response for that message.

SCRTTOOL119**COMPLETE EXCLUDE STATEMENT FOLLOWS****Explanation:**

All of the parameters and values on the **Exclude** control statement in error are listed immediately following this message. Each of the statements (lines) that comprise the complete **Exclude** control statement will be listed separately. Message SCRTTOOL119 will be issued for each such statement.

User response:

Review the additional error message and follow the instructions provided in the user response for that message.

SCRTTOOL120**PRODUCT_ID= ALREADY SPECIFIED****Explanation:**

An **Exclude** control statement contains more than one instance of the **Product_Id=** keyword parameter. Each **Exclude** control statement supports only one **Product_Id=** parameter. Message SCRTTOOL119 lists the control statement in error.

User response:

Remove the additional instances of the **Product_Id=** keyword parameter. If multiple product IDs must be specified on this **Exclude** control statement, use the list form — that is, (*xxxx-xxx*, . . . , *xxxx-xxx*) — to provide multiple **Product_Id** values.

SCRTTOOL121**PRODUCT_ID= LENGTH MUST BE 8****Explanation:**

The value for the **Product_Id=** keyword parameter on an **Exclude** control statement is not eight characters long. The **Product_Id** value must be eight characters long, in the form *xxxx-xxx*. Message SCRTTOOL119 lists the control statement in error.

User response:

Provide a valid **Product_Id** value that is eight characters long.

SCRTTOOL123**PRODUCT_ID= NOT SUPPORTED BY SCRT****Explanation:**

The value for the **Product_Id=** keyword parameter on an **Exclude** control statement is not supported by the level of SCRT that was run. Message SCRTTOOL119 lists the control statement in error.

User response:

Verify that the **Product_Id** value does not contain a typographical error. To determine the product IDs that are reported by this execution of SCRT, run SCRT with the same SMF and SCRT files but without any **Exclude** control statements. The Product Grid Snapshot section of the sub-capacity report lists the valid product IDs. Note, however, that the product IDs for the operating system products that are also listed in this section are not supported as valid values for the **Product_Id** parameter.

SCRTTOOL124**PRODUCT_ID= CANNOT BE A NO89 PRODUCT****Explanation:**

The value for the **Product_Id=** keyword parameter on an **Exclude** control statement is for a NO89 product. The **Exclude** control statement does not support NO89 products. Message SCRTTOOL119 lists the control statement in error.

User response:

Only specify **Product_Id** values that SCRT supports.

SCRTTOOL125**PRODUCT_ID= CANNOT BE AN OPERATING SYSTEM PRODUCT****Explanation:**

The value for the **Product_Id=** keyword parameter on an **Exclude** control statement is for an operating system product. The **Exclude** control statement does not support operating system products. Message SCRTTOOL119 lists the control statement in error.

User response:

The products z/OS (5694-A01), z/TPF (5748-T15, 5748-F15), and z/VSE (5686-CF8, 5686-CF9) are all considered operating system products and are not accepted as valid values for the **Product_Id=** parameter on the **Exclude** control statement because there is no valid reason to exclude those operating system products without also excluding all of the middleware products that were running at the same time on the same system. If you intended to exclude all of the products, including operating system products, then simply remove the **Product_Id=** parameter from the **Exclude** control statement. For examples of such usage for z/OS systems, see “EXCLUDE control statement” on page 44.

SCRTTOOL127 **PRODUCT_ID= LIST ENTRY LENGTH MUST BE 8****Explanation:**

One of the values in the list format of the **Product_Id=** keyword parameter on an **Exclude** control statement is greater than eight characters long. Each product ID in the list of **Product_Id** values must be eight characters long, in the form *xxxx-xxx*. Message SCRTTOOL119 lists the control statement in error.

User response:

Provide a list of valid **Product_Id** values that are each eight characters long.

SCRTTOOL129 **PRODUCT_ID= LIST ENTRY NOT SUPPORTED BY SCRT****Explanation:**

One of the values in the list format of the **Product_Id=** keyword parameter on an **Exclude** control statement is not supported by the level of SCRT that was run. Message SCRTTOOL119 lists the control statement in error.

User response:

Verify that there are no typographical errors in any of the **Product_Id** values in the list. To determine the product IDs that are reported by this execution of SCRT, run SCRT with the same SMF and SCRT files but without any **Exclude** control statements. The Product Grid Snapshot section of the sub-capacity report lists the valid product IDs. Note, however, that the product IDs for the operating system products that are also listed in this section are not supported as valid values for the **Product_Id** parameter.

SCRTTOOL130 **PRODUCT_ID= LIST ENTRY CANNOT BE A NO89 PRODUCT****Explanation:**

One of the values in the list format of the **Product_Id=** keyword parameter on an **Exclude** control statement is for a NO89 product. The **Exclude** control statement does not support NO89 products. Message SCRTTOOL119 lists the control statement in error.

User response:

Only specify **Product_Id** values that SCRT supports.

SCRTTOOL131 **PRODUCT_ID= LIST ENTRY CANNOT BE AN OPERATING SYSTEM PRODUCT****Explanation:**

One of the values in the list format of the **Product_Id=** keyword parameter on an **Exclude** control statement is for an operating system product. The **Exclude** control statement does not support operating system products. Message SCRTTOOL119 lists the control statement in error.

User response:

Only specify **Product_Id** values that SCRT supports.

SCRTTOOL134 **EXPECTED KEYWORD IS MISSING****Explanation:**

SCRT did not recognize a control statement in the SPECIAL DD.

User response:

Update the SPECIAL DD to contain only valid SCRT SPECIAL control statements. For descriptions of the SCRT SPECIAL control statements and their supported keywords, see “Customizing the SCRT JCL” on page 25.

SCRTTOOL135 **SPECIAL CONTROL STATEMENT IN ERROR FOLLOWS, LINE = nnn**

Explanation

All of the parameters and values on the SPECIAL DD control statement in error are listed immediately following this message. Each of the statements (lines) that comprise the complete SPECIAL DD control statement will be listed separately. Message SCRTTOOL135 will be issued for each such statement.

nnn

The line number within the SPECIAL DD control statement that contains the error. Comment lines are included in the count.

User response:

Review the accompanying error message and follow the instructions provided in the user response for that message.

SCRTTOOL136

SUPPRESSED SECTION NAME IS INVALID

Explanation:

The parameter value supplied on the **Suppress_Report_Section** control statement is not recognized by SCRT.

User response:

Currently, SCRT only supports the value GSSP. Verify that the value GSSP was spelled correctly and rerun SCRT.

SCRTTOOL137

START.UTC NE RESUME.UTC

Explanation:

There is a mismatch in the formats of the date/time values specified for the **Start** and **Resume** parameters of the **Exclude** control statement. Both the **Start** and **Resume** parameter values must use either the local format or the UTC format to express the date/time values.

User response:

Change the values for the **Start** and **Resume** parameters of the **Exclude** control statement so that both use the local time format (*yyyy/mm/dd/hh*) or both use the UTC time format (*yyyy/mm/dd/hh/UTC*). Note that there are specific conditions under which you *must* use the UTC format.

SCRTTOOL138

ONE OF THE IMAGE_IDS HAS MULTIPLE UTC OFFSETS DURING THE INTERVAL

Explanation:

The local date/time format (*yyyy/mm/dd/hh*) was used for the **Start** and **Resume** parameters on the **Exclude** control statement. One or more of the **Image_Ids** for that control statement has multiple UTC offsets during the **Start** and **Resume** interval that is being excluded. The difference in UTC offsets is due to time zone changes and typically occurs in the spring and fall, when some locations change between Standard Time and Daylight Savings Time.

User response:

Convert the date/time values for both the **Start** and **Resume** parameters from local time to UTC time and use the UTC time format (*yyyy/mm/dd/hh/UTC*) on the **Exclude** control statement. The converted date/time values must be the actual UTC values for those dates and times. Note that the date can change when converting a local time to a UTC time. You might need to contact your operations manager or system programmer to determine which UTC offsets were used on which dates for each **Image_Id** being excluded.

SCRTTOOL139

MULTIPLE IMAGE_IDS HAVE DIFFERENT UTC OFFSETS

Explanation:

The local date/time format (*yyyy/mm/dd/hh*) was used for the **Start** and **Resume** parameters on an **Exclude** control statement that specifies more than one **Image_Id**. Two or more of the **Image_Ids** specified on that control statement have UTC offsets that differ from each other. This is one of the few conditions that requires the date/time values for the **Start** and **Resume** parameters to be in the UTC format (*yyyy/mm/dd/hh/UTC*).

User response

You have two options:

- **Option 1:** Split the single **Exclude** control statement into separate **Exclude** control statements for each **Image_Id**. When you do this, the date/time values for the **Start** and **Resume** parameters on each of the

Exclude statements can remain in the local format (*yyyy/mm/dd/hh*), unless there is a UTC offset change for the date/time range being excluded for that **Image_Id**. In that case, the UTC format (*yyyy/mm/dd/hh/UTC*) is required.

- **Option 2:** Convert the date/time values for the **Start** and **Resume** parameters on the single **Exclude** control statement into their corresponding UTC values and use the UTC format (*yyyy/mm/dd/hh/UTC*) on the control statement. The converted date/time values must be the actual UTC values for those dates and times. Note that the date can change when converting a local time to a UTC time. You might need to contact your operations manager or systems programmer to determine which UTC offsets were used on which dates for each **Image_Id** being excluded.

SCRTTOOL140**ERROR IN PARMS: SYNTAX ERROR IN COMMENT= KEYWORD****Explanation:**

There is an error on one of the PARMS DD control statements. The value provided for the **Comment=** parameter on the **LPAR_Comment** statement has a syntax error. Message SCRTTOOL046 lists the control statement in error.

User response:

Verify that the text assigned to the **Comment=** parameter, which is listed in the SCRTTOOL046 message, ends with double quotation marks (") before position 72 on the control statement.

SCRTTOOL141**ERROR IN PARMS: REQUIRED KEYWORD (CPC=) IS MISSING****Explanation:**

There is an error on one of the PARMS DD control statements. The required **CPC=** keyword is missing from an **LPAR_Comment** control statement. Message SCRTTOOL046 lists the last control statement being processed when SCRT detected the error.

User response:

Review the **LPAR_Comment** control statement and verify that the required **CPC=** parameter is present and precedes the **Comment=** parameter. Note that the **CPC=** parameter can appear on its own statement as part of a continuation of an **LPAR_Comment** control statement. Whenever multiple sets of **LPAR_Comment** control statements are being submitted, use the statement listed by message SCRTTOOL046 to determine which set of **LPAR_Comment** control statements was being processed when the error occurred. Whenever the **LPAR_Comment** parameters are split across multiple control statements, the last statement in that set is listed by message SCRTTOOL046.

SCRTTOOL142**ERROR IN PARMS: REQUIRED KEYWORD (LPAR=) IS MISSING****Explanation:**

There is an error on one of the PARMS DD control statements. The required **LPAR=** keyword is missing from a **LPAR_Comment** control statement. Message SCRTTOOL046 lists the last control statement being processed when SCRT detected the error.

User response:

Review the **LPAR_Comment** control statement and verify that the required **LPAR=** parameter is present and precedes the **Comment=** parameter. Whenever multiple sets of **LPAR_Comment** control statements are being submitted, use the statement listed by message SCRTTOOL046 to determine which set of **LPAR_Comment** control statements was being processed when the error occurred. Whenever the **LPAR_Comment** parameters are split across multiple control statements, the last statement in that set is listed by message SCRTTOOL046.

SCRTTOOL143**ERROR IN PARMS: REQUIRED KEYWORD (COMMENT=) IS MISSING****Explanation:**

There is an error on one of the PARMS DD control statements. The required **Comments=** keyword is missing from a **LPAR_Comment** control statement. Message SCRTTOOL046 lists the last control statement being processed when SCRT detected the error.

User response:

Review the **LPAR_Comment** control statement and verify that the required **Comments=** parameter is present and is the last parameter on the **LPAR_Comment** control statement. Note that the **Comments=** parameter can appear on its own control statement provided that the last parameter on the preceding statement is followed by a comma. Whenever multiple sets of **LPAR_Comment** control statements are being submitted, use the statement listed by SCRTTOOL046 to determine which set of **LPAR_Comment** control statements was being processed

when the error occurred. Whenever the **LPAR_Comment** parameters are split across multiple control statements, the last statement in that set is listed by SCRTTOOL046.

SCRTTOOL148 **ERROR IN PARMS: CPC= ALREADY SPECIFIED****Explanation:**

There is an error on one of the PARMS DD control statements. There are one or more duplicate **CPC=** parameters on this set of **LPAR_Comment** statements. Message SCRTTOOL046 lists the control statement in error.

User response:

Review the **LPAR_Comment** control statement listed in the SCRTTOOL046 message and remove the duplicate **CPC=** parameters and their associated values. Each **LPAR_Comment** control statement supports a single **CPC=** parameter.

SCRTTOOL149 **ERROR IN PARMS: LPAR= ALREADY SPECIFIED****Explanation:**

There is an error on one of the PARMS DD control statements. There are one or more duplicate **LPAR=** parameters on this set of **LPAR_Comment** statements. Message SCRTTOOL046 lists the control statement in error.

User response:

Review the **LPAR_Comment** control statement listed in the SCRTTOOL046 message and remove the duplicate **LPAR=** parameters and their associated values. Each **LPAR_Comment** control statement supports a single **LPAR=** parameter.

SCRTTOOL150 **ERROR IN PARMS: COMMENT= ALREADY SPECIFIED****Explanation:**

There is an error on one of the PARMS DD control statements. There are one or more duplicate **Comment=** parameters on this set of **LPAR_Comment** statements. Message SCRTTOOL046 lists the control statement in error.

User response:

Review the **LPAR_Comment** control statement listed in the SCRTTOOL046 message and remove the duplicate **Comment=** parameters. Each **LPAR_Comment** control statement supports a single **Comment=** parameter. Note, however, that the text for the **Comment=** parameter can be split across multiple control statements in order to support up to 255 characters of comment text. Whenever the comments span multiple control statements, the **Comment=** keyword is required on the first statement which starts the comment. The **Comment=** keyword should not be used for the text that is split onto the following control statements.

SCRTTOOL153 **ERROR IN PARMS: COMMENT MISSING BEGINNING QUOTE****Explanation:**

There is an error on one of the PARMS DD control statements. The text value provided for the **Comment=** parameter on the **LPAR_Comment** statement is not preceded by the double quotation mark character ("), but must be. Message SCRTTOOL046 lists the control statement in error.

User response:

Correct the syntax error in the **Comment=** value by inserting a double quotation mark character (") before the actual comment text. Also be sure to use a double quotation mark character (") after the end of the comment text on this control statement.

SCRTTOOL154 **ERROR IN PARMS: COMMENT EXCEEDS 255 CHARACTERS****Explanation:**

There is an error on one of the PARMS DD control statements. The text value provided for the **Comment=** parameter on the **LPAR_Comment** statement exceeds the limit of 255 characters. Message SCRTTOOL046 lists the control statement in error.

User response:

Correct the syntax error in the **Comment=** value by limiting the text value to a maximum of 255 characters. The SCRTTOOL045 message lists the specific statement that contains the text that exceeds the 255-character maximum.

SCRTTOOL155 **ERROR IN PARMS: T/S type-serial LPAR *lparname* NOT FOUND IN DATA**

Explanation:

There is an error on one of the PARMs DD control statements. The PR/SM Partition Data Section on SMF or SCRT89 records did not report any busy time for the reporting period for the LPAR named *lparname* on the CPC specified by *type-serial*.

User response:

Verify that there are no typographical errors in the values for the **CPC=** and **LPAR=** parameters on the **LPAR_Comment** control statement listed in message SCRTTOOL155. Because that LPAR is not listed in the Detail Data Collection section (H5) of the sub-capacity report, SCRT does not see any activity for this LPAR during the reporting period for the specified CPC and will not accept an **LPAR_Comment** control statement for this LPAR..

SCRTTOOL158**ERROR IN PARMS: T/S *type-serial* LPAR *lparname* HAS SMF DATA****Explanation:**

There is an error on one of the PARMs DD control statements. SMF or SCRT89 records were found for the LPAR named *lparname* on the CPC specified by *type-serial*.

User response:

The **LPAR_Comment** control statement does not support comments for LPARs which have provided SMF or SCRT89 records that have been processed by this execution of SCRT. Verify that there are no typographical errors in the values for the **CPC=** and **LPAR=** parameters on the **LPAR_Comment** control statement listed in message SCRTTOOL158.

SCRTTOOL161**POTENTIAL ISSUE WITH POWER SAVE MODE DETECTED****Explanation:**

SCRT has detected a condition that might affect the sub-capacity values of products that were running while the CPC may have been running in power save mode.

User response

This is a warning message which is intended to highlight the need to take additional actions prior to running a CPC in static power save mode (for those CPCs that support this mode). This message does not necessarily mean that the CPC whose SMF data was being processed by SCRT was running in power save mode during the reporting period. This message is issued for either or both of the following conditions:

- One or more of the operating systems that were running on a CPC that supports static power save mode during the reporting period were either not at the appropriate z/OS release level or did not have the required service (APAR) installed to properly support power save mode.
- One or more z/OS systems did not have its active SMFPRMxx member of PARMLIB optimally configured for collecting SMF data when the CPC was running in power save mode.

The first of these conditions makes it impossible for SCRT to even determine whether the CPC is running in power save mode. You must perform the actions in the following paragraph to determine how to properly prepare for running in power save mode and to determine how SCRT handled the measurement data on the SMF and SCRT89 records that were generated on a CPC that was running in power save mode.

Power save mode and the SCRT support for power save mode are described in [“Static power save mode” on page 506](#). The Detail LPAR Data Section on one or more of the sub-capacity reports generated by this execution of SCRT displays a group of columns in that report section to categorize the SCRT view of power save mode. This view is based entirely on the data provided by the SMF or SCRT89 records from the LPAR or guest system listed in the first column on the same line of the report. [“Static power save mode” on page 506](#) contains tables of explanations of these columns specific to each operating system. Locate the table of explanations that matches the operating system and APAR level that has hour counts listed in the power save mode columns of this sub-capacity report and use those explanations to determine the impact, if any, on sub-capacity pricing and to determine any actions to take to improve the reporting of power save mode. Implement the suggested actions for the applicable conditions.

SCRTTOOL171**ERROR IN PARMS: CUSTOMER_NUMBER VALUE IS INVALID****Explanation:**

The value used for the **Customer_Number** parameter does not match one of the formats supported by SCRT. This message is followed by an SCRTTOOL046 message which lists the value specified for the customer number. SCRT terminates with return code 16 without generating any sub-capacity reports.

User response:

Check for typing errors and verify that the value specified for the customer number conforms to the format and value described for the **Customer_Number** parameter in [“PARMS parameters that apply to all CPCs” on page 34](#). Rerun SCRT with the correct customer number.

SCRTTOOL173

REPORT_TIME ERROR - VALUE IS INVALID

Explanation:

The option keyword supplied on the **Report_Time** control statement used by this execution of SCRT is not supported. SCRT terminates without generating any sub-capacity reports. Message SCRTTOOL135 lists the control statement in error.

User response:

Review the **Report_Time** control statement listed in message SCRTTOOL135 and rerun SCRT, replacing the invalid option value with one of the supported options: **Special_Local_Time** or **UTC_Time**.

SCRTTOOL174

REPORT_TIME ERROR - MULTIPLE CONTROL STATEMENTS

Explanation:

More than one **Report_Time** control statement was specified, using either the same or different option values. SCRT only supports one **Report_Time** control statement at a time. SCRT terminates without generating any sub-capacity reports.

User response:

Review the **Report_Time** control statements and rerun SCRT using a single **Report_Time** control statement to request either the **Special_Local_Time** or **UTC_Time** option.

SCRTTOOL175

REPORT_TIME ERROR - ALL HARDWARE TIME OFFSETS NOT EQUAL

Explanation:

The hardware timer offsets to UTC time reported in the SMF70HOF and SMF89HOF fields are not the same on all of the SMF records processed by this execution of SCRT, including the SMF records before and after the reporting period. The hardware timer offsets are set by the Sysplex Timer Offset, LPAR Timer Offset, and the z/VM virtual clock (VTOD). This condition violates one of the restrictions for using the **Special_Local_Time** option. SCRT terminates without generating any sub-capacity reports.

User response

Try to correct the problem by rerunning SCRT using the **Exclude** control statement to exclude the SMF records before and after the reporting period for all of the CPCs being processed, in case those are the only records whose SMF70HOF and SMF89HOF fields differ from those on the SMF records during the reporting period. Note that the times you specify for the **Start** and **Resume** parameters on the **Exclude** control statement must be in local time (without a UTC suffix).

If the problem persists, the data in the SMF records violates one of the restrictions for using the **Special_Local_Time** option on the **Report_Time** control statement. SCRT does not support the use of the **Special_Local_Time** option for this condition and cannot generate a sub-capacity report. However, SCRT does support the **UTC_Time** option for this condition and can generate a sub-capacity report with dates and times displayed in UTC time (with the UTC suffix added to the time values).

SCRTTOOL176

REPORT_TIME ERROR - HARDWARE TIME OFFSET TOO LARGE

Explanation:

The hardware timer offsets to UTC time reported in the SMF70HOF and SMF89HOF fields had values greater than 16 hours or less than -16 hours on one or more of the SMF records processed by this execution of SCRT, including the SMF records before and after the reporting period. The hardware timer offsets are set by the Sysplex Timer Offset, LPAR Timer Offset, and the z/VM virtual clock (VTOD). This condition violates one of the restrictions for using the **Special_Local_Time** option. SCRT terminates without generating any sub-capacity reports.

User response

Try to correct the problem by rerunning SCRT using the **Exclude** control statement to exclude the SMF records before and after the reporting period for all of the CPCs being processed, in case those are the only records whose SMF70HOF and SMF89HOF values are too high or too low. Note that the times you specify for the **Start** and **Resume** parameters on the **Exclude** control statement must be in local time (without a UTC suffix).

If the problem persists, the data in the SMF records violates one of the restrictions for using the **Special_Local_Time** option on the **Report_Time** control statement. SCRT does not support the use of the **Special_Local_Time** option for this condition and cannot generate a sub-capacity report. However, SCRT does support the **UTC_Time** option for this condition and can generate a sub-capacity report with dates and times displayed in UTC time (with the UTC suffix added to the time values).

SCRTTOOL177

REPORT_TIME ERROR - ALL Z/OS TIME OFFSETS NOT EQUAL

Explanation:

The z/OS offset to UTC time reported in the SMF70LGO and SMF89DTO fields was not the same on all of the SMF records processed by this execution of SCRT, including the SMF records before and after the reporting period. The z/OS offset to UTC time is set by the **TIMEZONE** value in the **CLOCKxx** member of **SYS1.PARMLIB** or can be set at IPL time. This condition violates one of the restrictions for using the **Special_Local_Time** option. SCRT terminates without generating any sub-capacity reports.

User response

Try to correct the problem by rerunning SCRT using the **Exclude** control statement to exclude the SMF records before and after the reporting period for all of the CPCs being processed, in case those are the only records whose SMF70LGO and SMF89DTO fields have this problem. Note that the times you specify for the **Start** and **Resume** parameters on the **Exclude** control statement must be in local time (without a UTC suffix).

If the problem persists, the data in the SMF records violates one of the restrictions for using the **Special_Local_Time** option on the **Report_Time** control statement. SCRT does not support the use of the **Special_Local_Time** option for this condition and cannot generate a sub-capacity report. However, SCRT does support the **UTC_Time** option for this condition and can generate a sub-capacity report with dates and times displayed in UTC time (with the UTC suffix added to the time values).

SCRTTOOL178

REPORT_TIME ERROR - z/OS TIME OFFSET TOO LARGE

Explanation:

The value of the z/OS offset to UTC time reported in the SMF70LGO and SMF89DTO was greater than 16 hours or less than -16 hours on one or more of the SMF records processed by this execution of SCRT, including the SMF records before and after the reporting period. The z/OS offset to UTC time is set by the **TIMEZONE** value in the **CLOCKxx** member of **SYS1.PARMLIB** or can be set at IPL time. This condition violates one of the restrictions for using the **Special_Local_Time** option. SCRT terminates without generating any sub-capacity reports.

User response

Try to correct the problem by rerunning SCRT using the **Exclude** control statement to exclude the SMF records before and after the reporting period for all of the CPCs being processed, in case those are the only records whose SMF70LGO and SMF89DTO fields have this problem. Note that the times you specify for the **Start** and **Resume** parameters on the **Exclude** control statement must be in local time (without a UTC suffix).

If the problem persists, the data in the SMF records violates one of the restrictions for using the **Special_Local_Time** option on the **Report_Time** control statement. SCRT does not support the use of the **Special_Local_Time** option for this condition and cannot generate a sub-capacity report. However, SCRT does support the **UTC_Time** option for this condition and can generate a sub-capacity report with dates and times displayed in UTC time (with the UTC suffix added to the time values).

SCRTTOOL179

REPORT_TIME ERROR - CANNOT USE SPEC_LOCAL_TIME AND EXCLUDE WITH UTC

Explanation:

SCRT processed a **Report_Time** control statement with the **Special_Local_Time** option and an **Exclude** control statement that uses the UTC format (times specified with a UTC suffix) for the **Start** and **Resume**

parameters. This condition violates one of the restrictions for using the `Special_Local_Time` option. SCRT terminates without generating any sub-capacity reports.

User response:

Correct the **Exclude** control statement by using the local time format (times specified without a UTC suffix) for the **Start** and **Resume** parameters, then rerun SCRT. Be sure to determine the correct local dates and times to be excluded.

SCRTT00L180

REPORT_TIME ERROR - CANNOT USE UTC_TIME AND EXCLUDE IN LOCAL TIME

Explanation:

SCRT processed a **Report_Time** control statement with the `UTC_Time` parameter and an **Exclude** control statement that uses the local time format (times specified without a UTC suffix) for the **Start** and **Resume** parameters. This condition violates one of the restrictions for using the `UTC_Time` option. SCRT terminates without generating any sub-capacity reports.

User response:

Correct the **Exclude** control statement by using the UTC time format (times specified with a UTC suffix) for the **Start** and **Resume** parameters, then rerun SCRT. Be sure to determine the correct UTC dates and times to be excluded.

SCRTT00L200

ASSIGN statement is missing CPC= parameterASSIGN parameters follow:

Explanation:

The **ASSIGN** control statement has been used in the **PARMS DD** statement, but the required **CPC** parameter is missing.

User response:

Examine the parameters on the **ASSIGN** control statement in your **PARMS DD** statement and include the required **CPC** parameter. For details, see [“Customizing the SCRT JCL”](#) on page 25.

SCRTT00L201

ASSIGN statement is missing CUSTOMER_NUMBER= parameterASSIGN parameters follow:

Explanation:

The **ASSIGN** control statement has been used in the **PARMS DD** statement, but the required **CUSTOMER_NUMBER** parameter is missing.

User response:

Examine the parameters on the **ASSIGN** control statement in your **PARMS DD** statement and include the required **CUSTOMER_NUMBER** parameter. For details, see [“Customizing the SCRT JCL”](#) on page 25.

SCRTT00L202

ASSIGN statement CPC= parameter is invalidASSIGN parameters follow:

Explanation:

The **ASSIGN** control statement has been used in the **PARMS DD** statement, but the value specified for the **CPC** parameter is not valid.

User response:

Examine the parameters on the **ASSIGN** control statement in your **PARMS DD** statement and correct the error in the **CPC** parameter. For details, see [“Customizing the SCRT JCL”](#) on page 25.

SCRTT00L203

ASSIGN statement CUSTOMER_NUMBER= parameter is invalidASSIGN parameters follow:

Explanation:

The **ASSIGN** control statement has been used in the **PARMS DD** statement, but the value specified for the **CUSTOMER_NUMBER** parameter is not valid.

User response:

Examine the parameters on the **ASSIGN** control statement in your **PARMS DD** statement and correct the error in the **CUSTOMER_NUMBER** parameter. For details, see [“Customizing the SCRT JCL”](#) on page 25.

SCRTTOOL204	ASSIGN statement CPC= parameter already specified parameters follow:
Explanation: The ASSIGN control statement has been used in the PARMS DD statement, but the CPC parameter is specified multiple times.	
User response: Examine the parameters on the ASSIGN control statement in your PARMS DD statement and remove duplicate instances of the CPC parameter. For details, see “Customizing the SCRT JCL” on page 25.	
SCRTTOOL205	ASSIGN statement CUSTOMER_NUMBER= parameter already specified parameters follow:
Explanation: The ASSIGN control statement has been used in the PARMS DD statement, but the CUSTOMER_NUMBER parameter is specified multiple times.	
User response: Examine the parameters on the ASSIGN control statement in your PARMS DD statement and remove duplicate instances of the CUSTOMER_NUMBER parameter. For details, see “Customizing the SCRT JCL” on page 25.	
SCRTTOOL210	NO89 statement syntax errorCPC:LPARlist format is incorrect:NO89 statement follows:
Explanation: The specified NO89 DD statement contains an invalid LPAR name in the CPC:LPAR list. Either the LPAR name exceeds 8 characters, or it could not be found in the input SMF data. The LPAR name in question is displayed in the message, as well as the NO89 statement on which it was specified.	
User response: Correct the NO89 DD statement in the SCRT JCL.	
SCRTTOOL211	NO89 statement syntax errorLPARlist format is incorrect:NO89 statement follows:
Explanation: The specified NO89 DD statement contains an invalid LPAR name in the LPAR list. Either the LPAR name exceeds 8 characters, or it could not be found in the input SMF data. The LPAR name in question is displayed in the message, as well as the NO89 statement on which it was specified.	
User response: Correct the NO89 DD statement in the SCRT JCL.	
SCRTTOOL212	Product nnnnnn specification is missing from NO89 DD data
Explanation The specified product ID (PID) was not found in the NO89 DD statement. In the message text: nnnnnn The product ID that is missing from the NO89 data.	
User response: Add the missing PID to the NO89 DD statement in the SCRT JCL.	
SCRTTOOL220	Country Multiplex Pricing and Report Time are mutually exclusive
Explanation: The Country_Multiplex_Pricing and Report_Time SPECIAL DD control statements cannot be used together. When Country Multiplex Pricing is requested, SCRT determines whether the resulting report will be in UTC or local time in the same manner as with traditional sub-capacity reports. An override of this function is not allowed.	
User response:	

If Country Multiplex Pricing is desired, remove the **Report_Time** SPECIAL DD control statement.

SCRTT00L221 **All Multiplex CPCs must be located in the same country.****Explanation:**

When the **Country_Multiplex_Pricing** control statement is specified on the SPECIAL DD, SCRT performs additional checking to ensure that the **Customer_Number** values assigned to the input CPCs all contain the same country code.

User response:

Examine the **Customer_Number** values specified in the PARMS DD control statement for both the **Customer_Number** parameter and the **ASSIGN Customer_Number** control statement. Correct the parameters and rerun SCRT, specifying only CPCs and customer numbers that are contained within the same country multiplex.

SCRTT00L230 **PARMS syntax errorPARMS statement follows:****Explanation:**

The PARMS DD statement specified in the SCRTT00L230 message is not recognized by SCRT.

User response:

Correct the PARMS DD statement.

SCRTT00L240 **xxxxxxx DD is missing****Explanation:**

The DD statement identified by xxxxxxxx was not found in the JCL.

User response:

Specify the missing DD statement and any required inputs associated with it.

SCRTT00L241 **Failure when checking for DD xxxxxxxx****Explanation**

SCRT encountered an error while checking for the DD statement identified in the message text.

In the message text:

xxxxxxx

The name of the DD statement that was being checked.

User response

Check the SCRT job output for additional error messages or exceptions, which may be helpful in determining the reason for the failure.

Correct the error with the specified DD statement.

SCRTT00L242 **xxxxxx ReadFile FileException****Explanation**

SCRT encountered an error while attempting to read from the data set associated with the DD statement identified in the message text.

In the message text:

xxxxxx

The name of the DD statement from which SCRT was attempting to read.

User response

Check the SCRT job output for additional error messages or exceptions, which may be helpful in determining the reason for the failure.

Correct the error with the specified DD statement.

SCRTT00L243 **xxxxxx ReadFile RcException**

Explanation

SCRT encountered an error while attempting to read from the data set associated with the DD statement identified in the message text.

In the message text:

xxxxxx

The name of the DD statement from which SCRT was attempting to read.

User response

Check the SCRT job output for additional error messages or exceptions, which may be helpful in determining the reason for the failure.

Correct the error with the specified DD statement.

SCRTTOOL244

xxxxxx Closing FileException

Explanation

SCRT encountered an error while attempting to close the data set associated with the DD statement identified in the message text.

In the message text:

xxxxxx

The name of the DD statement that SCRT was attempting to close.

User response

Check the SCRT job output for additional error messages or exceptions, which may be helpful in determining the reason for the failure.

Correct the error with the specified DD statement.

SCRTTOOL245

ddname I/O ERROR ENCOUNTERED. text

Explanation

SCRT encountered an I/O error on a z/OS file.

In the message text:

ddname

The name of the DD statement for the file that encountered an I/O error.

text

A description of the I/O error condition.

User response:

Determine the cause of the I/O error for the file referenced by *ddname* and take corrective action.

SCRTTOOL250

Error found while processing INPUTCSV

Explanation:

SCRT encountered an error while processing the INPUTCSV DD data sets. One or more of the statements in the INPUTCSV data is incorrect.

User response:

See the SCRTTOOL254 message for information about the specific error that was encountered. Identify and correct the error that caused the failure.

SCRTTOOL251

Error found while verifying INPUTCSV data.

Explanation:

SCRT encountered an error while processing the INPUTCSV DD data sets. The syntax is correct, but the INPUTCSV files contain a combination of parameters that is not allowed. The SCRTTOOL254 message reports the specific condition that was encountered.

User response:

See the SCRTTOOL254 message for information about the specific error that was encountered. Identify and correct the error that caused the failure.

SCRTTOOL252 **Error encountered opening INPUTCSV**
Explanation:

SCRT encountered a critical error while opening a data set specified on the INPUTCSV DD statement.

User response

Check the SCRT job output for additional error messages or exceptions, which may be helpful in determining the reason for the failure.

Identify and correct the error that caused the failure.

SCRTTOOL253 **INPUTCSV DD member not found.**
Explanation:

SCRT could not find the PDS member specified on the INPUTCSV DD statement.

User response:

Correct the error that caused the failure.

SCRTTOOL254 ***text***
Explanation:

In the message, *text* reports the specific error encountered while parsing a data set specified on the INPUTCSV DD statement.

User response:

Examine the reported condition and correct it.

SCRTTOOL255 **CPC: *cpu* IMAGE: *image* PRODUCT: *product* TIME: *time* DIAG1: *diag1***
Explanation

SCRT is processing one or more data sets specified on the INPUTCSV DD statement. When SCRT detects that the Mobile Workload Pricing (MWP) or IBM Z Collocated Application Pricing (zCAP) usage value exceeds the actual usage time for the respective defining program during a given hour, SCRT issues the SCRTTOOL255 message to highlight that hour.

SCRT also issues the SCRTTOOL256 or SCRTTOOL257 message to indicate whether the instances of the SCRTTOOL255 message represent MWP or zCAP processing.

In the message text:

cpu

The type-serial of the CPC being processed

image

The LPAR name or virtual machine name where the mobile or zCAP product ran

product

The product or product family that the MWP or zCAP algorithm is processing

time

The time in the INPUTCSV data set being processed

diag1

A diagnostic indicator for IBM use only.

User response:

Correct the data during the interval specified in the message. The MWP or zCAP values cannot exceed the products' actual usage.

SCRTTOOL256 **INPUTCSV MOBILE CPU DATA EXCEEDS PRODUCT USAGE**
Explanation:

SCRT is processing one or more data sets specified on the INPUTCSV DD statement and has detected that the Mobile Workload Pricing (MWP) data for a given product family exceeds the actual usage of the products in the product family. See the instances of the SCRTTOOL255 message for details.

User response:

Correct the data for the intervals highlighted by the associated SCRTTOOL255 messages. The MWP values cannot exceed the products' actual usage.

SCRTTOOL257	INPUTCSV ZCAP CPU DATA EXCEEDS PRODUCT USAGE
--------------------	---

Explanation:

SCRT is processing one or more data sets specified on the INPUTCSV DD statement and has detected that the IBM Z Collocated Application Pricing (zCAP) data for a given product exceeds the product's actual usage. See the instances of the SCRTTOOL255 message for details.

User response:

Correct the data for the intervals highlighted by the associated SCRTTOOL255 messages. The zCAP values cannot exceed the product's actual usage.

SCRTTOOL258	INPUTCSV TOTAL ZCAP WORKLOAD DATA EXCEEDS USAGE OF THE PARTITION.
--------------------	--

Explanation

SCRT is processing one or more data sets specified on the INPUTCSV DD statement and has detected that the values for a zCAP workload exceed the usage of the partition.

See the instances of the SCRTTOOL255 message for details.

User response:

Examine the instances of the SCRTTOOL255 messages to determine the intervals during which the zCAP workload data exceeds the LPAR usage. Correct the data in the INPUTCSV DD data sets, as needed.

SCRTTOOL259	INPUTCSV ZWPC CPU DATA EXCEEDS PRODUCT USAGE
--------------------	---

Explanation:

SCRT is processing one or more data sets specified on the INPUTCSV DD statement and has detected that the IBM Z Workload Pricing for Cloud (zWPC) data for a given product family exceeds the actual usage of the products in the product family. See the instances of the SCRTTOOL255 message for details.

User response:

Correct the data for the intervals identified by the associated SCRTTOOL255 messages. The zWPC values cannot exceed the products' actual usage.

SCRTTOOL260	Exclude CPC= parameter invalid value
--------------------	---

Explanation:

The specified **CPC** parameter does not meet the expected syntax for the **Exclude** control statement.

User response:

See the SCRTTOOL135 message for more information and correct the **CPC** parameter accordingly.

SCRTTOOL261	EXCLUDE IMAGE_ID= PARAMETER INVALID VALUE
--------------------	--

Explanation:

The specified **IMAGE_ID** parameter does not meet the expected syntax for the **Exclude** control statement.

User response:

See the SCRTTOOL135 message for more information and correct the **IMAGE_ID** parameter accordingly.

SCRTTOOL262	EXCLUDE PRODUCT_ID= PARAMETER INVALID VALUE
--------------------	--

Explanation:

The specified **PRODUCT_ID** parameter does not meet the expected syntax for the **Exclude** control statement.

User response:

See the SCRTTOOL135 message for more information and correct the **PRODUCT_ID** parameter accordingly.

SCRTTOOL263 **UNRECOGNIZED COMMAND IN SPECIAL DD INPUT****Explanation:**

An unrecognized command was found in the SPECIAL DD control statement.

User response:

Correct the error in the SPECIAL DD statement.

SCRTTOOL264 **EXCLUDE COMMENT= ALREADY SPECIFIED****Explanation:**

An **EXCLUDE** control statement contains more than one instance of the **COMMENT** keyword parameter. Each **EXCLUDE** control statement supports only one **COMMENT** instance.

User response:

Remove the extra instances of the **COMMENT** parameter.

SCRTTOOL265 **EXCLUDE COMMENT= EXCEEDS MAX LENGTH OF *length*****Explanation:**

The comment specified via the **COMMENT** keyword parameter of an **EXCLUDE** control statement exceeds the maximum allowable length (*length*).

User response:

Update the comment so that it does not exceed the allowable length.

SCRTTOOL275 **UNRECOGNIZED ASSIGN PARAMETER SPECIFIED FOR *offering*****Explanation**

SCRT SPECIAL DD processing does not recognize the **ASSIGN** command parameter specified for the listed offering. The valid values that can be assigned to an offering are:

CATEGORYA

CATEGORYB

In the message text:

offering

The offering that is assigned via the **ASSIGN** command (for instance, ZWPC).

User response:

Correct the **ASSIGN** statement to reflect a valid parameter.

SCRTTOOL276 **DUPLICATE ASSIGN KEYWORD SPECIFIED FOR *offering*****Explanation**

SCRT SPECIAL DD processing has detected that an offering has been assigned multiple times using the **ASSIGN** statement.

In the message text:

offering

The offering that has been assigned multiple times via the **ASSIGN** command (for instance, ZWPC).

User response:

Correct the applicable **ASSIGN** commands to only assign a given offering once.

SCRTTOOL277 ***offering* HAS NOT BEEN ASSIGNED USING THE SPECIAL DD INPUT****Explanation**

The DIAGMSG=WLMZWPC function was requested, but ZWPC has not been assigned to a WLM classification (or category) using the **ASSIGN** statement.

In the message text:

offering

The offering that has not been assigned via the **ASSIGN** statement (for instance, ZWPC).

User response:

If WLM-based zWPC is not desired, remove the DIAGMSG=WLMZWPC statement. If it is desired, you must specify the **ASSIGN** statement to request the function.

SCRTTOOL278

category HAS BEEN ASSIGNED TO MORE THAN ONE OFFERING USING THE SPECIAL DD INPUT

Explanation

The specified category can only be assigned to one offering.

In the message text:

category

The category name that has been assigned to multiple offerings.

User response:

Correct the category assignments so that the specified category is only assigned to one offering.

SCRTTOOL280

INPUTCSV CPC AND/OR LPAR VALUE CANNOT BE LOCATED: *lpar*

Explanation

A specified INPUTCSV data set contains an LPAR name that cannot be located in the SMF input data.

In the message text:

lpar

The CPC:LPAR or LPAR name that SCRT cannot locate in the SMF input data.

User response:

Correct the CPC:LPAR or LPAR name. Note that the specified name must exist in the SMF input data, and the value specified in the INPUTCSV file, between double quotation marks, must not contain blanks.

SCRTTOOL281

INPUTCSV ZCAP WORKLOAD CPU DATA EXCEEDS PARTITION

Explanation:

The IBM Z Collocated Application Pricing (zCAP) workload MSUs derived from the CPU time reported in the INPUTCSV file exceed the LAC value for the LPAR.

User response:

Correct the workload CPU time so that it does not exceed the LAC value for the hosting LPAR.

SCRTTOOL282

INPUTCSV ZCAP WORKLOAD DATA SET DOES NOT CONTAIN REQUIRED CPU TIME INPUT

Explanation:

For IBM Z Collocated Application Pricing (zCAP) workloads that do not use WLM, you must provide valid CPU time data in the INPUTCSV file.

User response:

Update the INPUTCSV file with the CPU times for the zCAP workload.

SCRTTOOL295

SCRT PROCESSING ISV REPORT FOR: *isvname*

Explanation

SCRT is processing an independent software vendor (ISV) report for the specified vendor.

In the message text:

isvname

The name of the ISV.

User response:

None. This message is informational only.

SCRTTOOL296 **ERROR PROCESSING ISVLIB DD STATEMENT.****Explanation:**

SCRT cannot build its internal tables from the information in the ISVLIB file. Other accompanying messages, such as SCRTTOOL297, provide details about the cause of the error.

User response:

Contact the vendor who supplied the ISVLIB file to obtain a corrected copy.

SCRTTOOL297 ***text*****Explanation**

SCRT cannot build its internal tables from the information in the ISVLIB file.

In the message text, *text* provides a description of the error.

User response:

Contact the vendor who supplied the ISVLIB file to obtain a corrected copy.

SCRTTOOL300 **ERROR PROCESSING COMMAND LINE ARGUMENT: *argument*****Explanation**

A command-line argument was incorrectly specified. If the argument requires a parameter, the parameter value might be missing.

In the message text:

argument

The command-line argument in error.

User response:

Correctly specify the indicated argument and rerun SCRT.

SCRTTOOL301 **MISSING COMMAND LINE ARGUMENT: *argument*****Explanation**

A required command-line argument has not been specified.

In the message text:

argument

The required command-line argument.

User response:

Specify the indicated argument and rerun SCRT.

SCRTTOOL302 ***type* FILE COULD NOT BE FOUND OR IS INCORRECTLY SPECIFIED
PATH SPECIFIED = *path*****Explanation**

The command line specified a path to a file that is incorrect. The path might not exist, might be unreadable (for input files), might be unwritable (for output files), or is otherwise incorrect. The message indicates the path that SCRT attempted to use.

In the message text:

type

The type of file, either N089, OUTPUT, or SMF.

path

The path that SCRT attempted to use.

User response:

Check and correct the specified path. Then, rerun SCRT.

SCRTTOOL303 **REPORT GENERATED: *report-path***

Explanation

A sub-capacity report was generated.

In the message text:

report-path

The absolute path to the report.

User response:

None.

SCRTTOOL319

**SOLUTION ID IS NOT VALID OR ITS FORMAT WAS NOT RECOGNIZEDID
= *solutionID***

Explanation

A Tailored Fit Pricing for IBM Z solution ID was encountered that is not valid. The solution ID either does not match the format of a valid solution ID, or otherwise appears to not be validly formed.

In the message text:

solutionID

The solution ID that SCRT identified as not meeting syntax requirements.

User response:

Validate that the solution ID correctly matches the one provided by IBM in License Management Support (LMS) for the container solution. If the solution ID appears to be incorrectly defined in the WLM tenant resource group definition, use the **UPDATE SOLUTION_ID** command (control statement) in SCRT to correct the solution ID in existing SMF data. Also update the WLM tenant resource group definition with the correct value. If you cannot determine why the solution ID is incorrect, contact IBM Support.

SCRTTOOL320

CONTAINER COMMAND CPC= MUST BE SPECIFIED

Explanation:

The **CPC** keyword parameter, which is required for the **CONTAINER** control statement, is missing. Message SCRTTOOL135 lists the control statement in error.

User response:

Add the **CPC** keyword and a valid CPC value to the indicated **CONTAINER** control statement.

SCRTTOOL321

CONTAINER COMMAND CPC= PARAMETER INVALID VALUE

Explanation:

The **CPC** keyword parameter, which is required for the **CONTAINER** control statement, does not contain a valid value. Message SCRTTOOL135 lists the control statement in error.

User response:

Correct the **CPC** keyword value, ensuring it meets the syntax requirements of the **CONTAINER** control statement.

SCRTTOOL322

CONTAINER COMMAND CPC= ALREADY SPECIFIED

Explanation:

A **CONTAINER** control statement contains more than one instance of a **CPC** keyword parameter. Each **CONTAINER** control statement supports only one **CPC** instance. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the extra instances of the **CPC** parameter. If you need to specify more than one CPC, use a separate **CONTAINER** control statement for each CPC.

SCRTTOOL323

CONTAINER COMMAND CPC= VALUE LENGTH MUST BE 10, TTTT-SSSS

Explanation:

The value for the **CPC** keyword parameter on a **CONTAINER** control statement was not 10 characters long. The **CPC** value must be 10 characters long, in the form *tttt-ssss*. Message SCRTTOOL135 lists the control statement in error.

User response:

Provide a **CPC** value that is 10 characters long.

SCRTTOOL324 **CONTAINER COMMAND CPC=tttt-sssss NOT FOUND IN THE DATA****Explanation:**

The value for the **CPC** keyword parameter on a **CONTAINER** control statement was not found on any of the SMF records processed by this execution of SCRT. Message SCRTTOOL135 lists the control statement in error.

User response:

Change the **CPC** value on the **CONTAINER** control statement so that it matches a CPC whose SMF records are being processed by this execution of SCRT. To determine the CPCs that are processed by this execution of SCRT, run SCRT without any **CONTAINER** control statements and review the generated report.

SCRTTOOL325 **CONTAINER COMMAND IMAGE_ID= MUST BE SPECIFIED****Explanation:**

The **IMAGE_ID** keyword parameter, which is required for the **CONTAINER** control statement, is missing.

User response:

Add the **IMAGE_ID** keyword and a valid image ID to the **CONTAINER** control statement.

SCRTTOOL326 **CONTAINER COMMAND IMAGE_ID= PARAMETER INVALID VALUE****Explanation:**

The **IMAGE_ID** keyword parameter, which is required for the **CONTAINER** control statement, does not contain a valid value. Message SCRTTOOL135 lists the control statement in error.

User response:

Correct the **IMAGE_ID=** keyword value, ensuring it meets the syntax requirements of the **CONTAINER** control statement.

SCRTTOOL327 **CONTAINER COMMAND IMAGE_ID= ALREADY SPECIFIED****Explanation:**

A **CONTAINER** control statement contains more than one instance of the **IMAGE_ID** keyword parameter. Each **CONTAINER** control statement supports only one **IMAGE_ID** instance. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the extra instances of the **IMAGE_ID** parameter. If you need to specify more than one CPC, use a separate **CONTAINER** control statement for each CPC.

SCRTTOOL328 **CONTAINER COMMAND IMAGE_ID= VALUE MUST 1 TO 8 CHARS****Explanation:**

The value for the **Image_Id** keyword parameter exceeds 8 characters in length. The **Image_Id** value provides the name of the operating system image to which the **CONTAINER** control statement applies, and that name must be 1 - 8 characters long. Message SCRTTOOL135 lists the control statement in error.

User response:

Change the **Image_Id** value so that it is 1 - 8 characters long and matches an image ID in the SMF records that are being processed by this execution of SCRT. To determine the image IDs that are being processed by this execution of SCRT, run SCRT with the same SMF files but without any **CONTAINER** control statements. All the image IDs for each CPC will be listed in the first column of the Detail Data Collection section of the sub-capacity report for that CPC.

SCRTTOOL329 **CONTAINER COMMAND IMAGE_ID=imageID NOT FOUND ON THIS CPC=tttt-sssss****Explanation**

The value for the **IMAGE_ID** keyword parameter on a **CONTAINER** control statement was not found in any of the SMF records processed by this execution of SCRT. Message SCRTTOOL135 lists the control statement in error.

In the message text:

imageID

The image ID that was not found in the SMF data.

tttt-sssss

The type and serial number of the CPC that you specified on the **CONTAINER** command (control statement) for which the specified image ID was not found in the SMF data.

User response:

Change the **IMAGE_ID** value on the **CONTAINER** control statement so that it matches a CPC whose SMF records are being processed by this execution of SCRT. To determine the image IDs that are processed by this execution of SCRT, run SCRT without any **CONTAINER** control statements and review the generated report.

SCRTTOOL330**CONTAINER COMMAND INVALID FORMAT****Explanation:**

There is an error in the **CONTAINER** control statement.

User response:

Review the **CONTAINER** control statement. There must be a **CPC** parameter specifying the CPC, an **IMAGE_ID** parameter specifying one or more images that are being selected, and an **ID** parameter specifying a valid solution ID.

SCRTTOOL334**CONTAINER COMMAND ID= MUST BE SPECIFIED****Explanation:**

A **CONTAINER** control statement was found for which the required **ID** keyword parameter was not specified. The **ID** keyword parameter is used to specify the 64-character solution ID. Message SCRTTOOL135 lists the control statement in error.

User response:

Modify the **CONTAINER** command to specify the **ID** keyword parameter. You can obtain the solution ID from the IBM LMS web portal.

SCRTTOOL335**CONTAINER COMMAND ID= PARAMETER INVALID VALUE****Explanation:**

A **CONTAINER** control statement was found for which the **ID** keyword parameter is not valid. The **ID** keyword parameter is used to specify the 64-character solution ID. Message SCRTTOOL135 lists the control statement in error.

User response:

Modify the **CONTAINER** control statement to specify a correct **ID** keyword parameter. You can obtain the solution ID from the IBM LMS web portal.

SCRTTOOL336**CONTAINER COMMAND ID= ALREADY SPECIFIED****Explanation:**

A **CONTAINER** control statement contains more than one instance of the **ID** keyword parameter. Each **CONTAINER** control statement supports only one **ID** instance. The **ID** parameter is used to specify the 64-character Solution ID. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the extra instances of the **ID** control statement. If you need to specify more than one solution ID, use a separate **CONTAINER** control statement for each CPC.

SCRTTOOL337**CONTAINER COMMAND ID= VALUE MUST BE 64 CHARS****Explanation:**

The value for the **ID** keyword parameter on a **CONTAINER** control statement is not 64 characters long. The solution ID value must be 64 characters long. Message SCRTTOOL135 lists the control statement in error.

User response:

Provide an **ID** value that is 64 characters long.

SCRTTOOL338**CONTAINER COMMAND THE CPC AND IMAGE_ID CAN ONLY BE ASSIGNED ONCE *cpc* , *image***

Explanation

SCRT has encountered an instance of a **CONTAINER** and **IMAGE_ID** combination being assigned to multiple containers. An image (LPAR or VM) can only be assigned to one container during a reporting period.

In the message text:

cpc

The CPC machine type and serial number, in the form *tttt-sssss*, that was specified on the **CONTAINER** command.

image

The image ID that was specified on the **CONTAINER** command.

User response:

Correct any **CONTAINER** commands, ensuring that no CPC and image ID pair is assigned to more than one container.

SCRTTOOL343

CONTAINER COMMAND IMAGE_ID=*image* MUST BE Z/OS

Explanation

A **CONTAINER** control statement contains an image that is running an operating system other than z/OS. Tailored Fit Pricing for IBM Z only supports the z/OS operating system.

In the message text:

image

The image on which a non-z/OS operating system was detected.

User response:

Remove the specified image from the **CONTAINER** control statement.

SCRTTOOL344

UPDATE SOLUTION_ID DUPLICATE COMMAND ENCOUNTERED

Explanation:

SCRT has encountered a solution ID that is specified on more than one **UPDATE SOLUTION_ID** control statements as the **OLD** keyword parameter. A solution ID cannot be reassigned to multiple values in the same execution of SCRT. Message SCRTTOOL135 lists the control statement in error.

User response:

Review any instances of the **UPDATE SOLUTION_ID** command. Ensure each **OLD** parameter value is only assigned once for this execution of SCRT.

SCRTTOOL345

UPDATE SOLUTION_ID OLD= MUST BE SPECIFIED

Explanation:

The **OLD** keyword parameter, which is required for the **UPDATE SOLUTION_ID** control statement, is missing. Message SCRTTOOL135 lists the control statement in error.

User response:

Specify the **OLD** parameter with a valid solution ID value on the indicated **UPDATE SOLUTION_ID** control statement.

SCRTTOOL346

UPDATE SOLUTION_ID NEW= MUST BE SPECIFIED

Explanation:

The **NEW** keyword parameter, which is required for the **UPDATE SOLUTION_ID** control statement, is missing. Message SCRTTOOL135 lists the control statement in error.

User response:

Specify the **NEW** parameter with a valid solution ID value on the indicated **UPDATE SOLUTION_ID** control statement.

SCRTTOOL347

UPDATE SOLUTION_ID OLD= ALREADY SPECIFIED

Explanation:

An **UPDATE SOLUTION_ID** control statement contains more than one instance of the **OLD** keyword parameter. Each **UPDATE SOLUTION_ID** control statement supports only instance of the **OLD** parameter. The **OLD**

parameter specifies the 64-character solution ID that is to be replaced. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the extra instances of the **OLD** parameter. If you need to specify more than one solution ID, use a separate **UPDATE SOLUTION_ID** control statement for each CPC.

SCRTTOOL348

UPDATE SOLUTION_ID NEW= ALREADY SPECIFIED

Explanation:

An **UPDATE SOLUTION_ID** control statement contains more than one instance of the **NEW** keyword parameter. Each **UPDATE SOLUTION_ID** control statement supports only one instance of the **NEW** parameter. The **NEW** parameter specifies the 64-character solution ID that is to replace the **OLD** solution ID. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the extra instances of the **NEW** parameter. If you need to specify more than one solution ID, use a separate **UPDATE SOLUTION_ID** control statement for each CPC.

SCRTTOOL349

UPDATE SOLUTION_ID NEW= INVALID VALUE

Explanation:

An **UPDATE SOLUTION_ID** control statement was found for which the **NEW** keyword parameter is not valid. The **NEW** parameter specifies the 64-character solution ID. Message SCRTTOOL135 lists the control statement in error.

User response:

Modify the **UPDATE SOLUTION_ID** control statement to specify a correct **NEW** parameter. You can obtain the solution ID from the IBM LMS web portal.

SCRTTOOL350

UPDATE SOLUTION_ID OLD= INVALID VALUE

Explanation:

An **UPDATE SOLUTION_ID** command was found for which the **OLD** keyword parameter is not valid. The **OLD** parameter specifies the 64-character solution ID. Message SCRTTOOL135 lists the control statement in error.

User response:

Modify the **UPDATE SOLUTION_ID** control statement to specify a correct **OLD** parameter. You can obtain the solution ID from the IBM LMS web portal.

SCRTTOOL351

UPDATE CONTAINER ID= MUST BE SPECIFIED

Explanation:

The **ID** keyword parameter, which is required for the **UPDATE CONTAINER** control statement, is missing. Message SCRTTOOL135 lists the control statement in error.

User response:

Specify the **ID** keyword and a valid solution ID value on the indicated **UPDATE CONTAINER** control statement.

SCRTTOOL352

UPDATE CONTAINER SET_NAME= MUST BE SPECIFIED

Explanation:

The **SET_NAME** keyword parameter, which is required for the **UPDATE CONTAINER** control statement, is missing. Message SCRTTOOL135 lists the control statement in error.

User response:

Specify the **SET_NAME** keyword and a valid solution name value on the indicated **UPDATE CONTAINER** control statement.

SCRTTOOL353

UPDATE CONTAINER ID= ALREADY SPECIFIED

Explanation:

An **UPDATE CONTAINER** control statement contains more than one instance of the **ID** keyword parameter. Each **UPDATE CONTAINER** control statement supports only one instance of the **ID** parameter. The **ID** parameter specifies the 64-character solution ID that will be associated with the **SET_NAME** keyword parameter. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the extra instances of the **ID** parameter. If you need to specify more than one solution ID, use a separate **UPDATE CONTAINER** control statement for each CPC.

SCRTTOOL354**UPDATE CONTAINER SET_NAME= ALREADY SPECIFIED****Explanation:**

An **UPDATE CONTAINER** control statement contains more than one instance of the **SET_NAME** keyword parameter. Each **UPDATE CONTAINER** control statement supports only one instance of the **SET_NAME** parameter. The **SET_NAME** parameter specifies a solution name of up to 40 characters that will be associated with the **ID** keyword parameter. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the extra instances of the **SET_NAME** parameter from the **UPDATE CONTAINER** control statement.

SCRTTOOL355**UPDATE CONTAINER ID= INVALID VALUE****Explanation:**

An **UPDATE CONTAINER** command (control statement) was found for which the **ID** keyword parameter is not valid. The **ID** parameter specifies the 64-character solution ID. Message SCRTTOOL135 lists the control statement in error.

User response:

Modify the **UPDATE CONTAINER** command to specify a correct **ID** parameter. You can obtain the solution ID from the IBM LMS web portal.

SCRTTOOL356**UPDATE CONTAINER SET_NAME= INVALID VALUE****Explanation:**

An **UPDATE CONTAINER** command (control statement) was found for which the **SET_NAME** keyword parameter is not valid. The **SET_NAME** parameter specifies a solution name of up to 40 characters which must match the solution name specified in LMS. Message SCRTTOOL135 lists the control statement in error.

User response:

Modify the **UPDATE CONTAINER** command to specify a correct **SET_NAME** parameter. You can obtain the solution name from the IBM LMS web portal.

SCRTTOOL357**UPDATE CONTAINER SET_NAME= EXCEEDS 40 CHARACTERS****Explanation:**

An **UPDATE CONTAINER** command (control statement) was found for which the **SET_NAME** keyword parameter exceeds the maximum allowed length. The **SET_NAME** parameter specifies a solution name of up to 40 characters which must match the solution name specified in LMS. Message SCRTTOOL135 lists the control statement in error.

User response:

Modify the **UPDATE CONTAINER** command to specify a correct **SET_NAME** parameter. You can obtain the solution name from the IBM LMS web portal.

SCRTTOOL358**UPDATE CONTAINER DUPLICATE COMMAND ENCOUNTERED****Explanation:**

An **UPDATE CONTAINER** command (control statement) was found for which the **ID** keyword parameter contains a solution ID which already has a solution name associated with it. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the duplicate **UPDATE CONTAINER** command. Note that the **SET_NAME** parameter specifies a solution name of up to 40 characters which must match the one specified in LMS.

SCRTTOOL359**UPDATE CONTAINER SOLUTION ID NOT FOUND IN SMF DATA ID=
solutionID NAME= solutionName**

Explanation

An **UPDATE CONTAINER** command (control statement) was specified for a container that was not found in the SMF input data. No tenant resource group was found having the solution ID, and no image was assigned to the solution ID using a **CONTAINER** control statement.

In the message text:

solutionID

The 64-character solution ID for the container that was not found in the SMF data.

solutionName

The solution name identified in the **UPDATE CONTAINER** command that is associated with *solutionID*.

User response:

Check the solution ID identified in the SCRTTOOL359 message. Verify that this solution ID exists in the environment or is assigned to a dedicated image, as appropriate.

SCRTTOOL360

CONTAINER SOLUTION NAME NOT FOUND FOR CONTAINERS ID=*id*

Explanation

One or more Tailored Fit Pricing for IBM Z solutions were encountered for which no solution name is assigned.

In the message text:

id

One or more solution IDs without an assigned solution name.

User response:

For each solution ID, use the **UPDATE CONTAINER** command (control statement) to specify a valid solution name. You can obtain the solution name from the IBM LMS web portal.

SCRTTOOL361

IGNORE CONTAINER ID= MUST BE SPECIFIED

Explanation:

The **ID** keyword parameter, which is required for the **IGNORE CONTAINER** control statement, is missing. Message SCRTTOOL135 lists the control statement in error.

User response:

Specify the **ID** keyword and a valid solution ID value on the indicated **IGNORE CONTAINER** control statement.

SCRTTOOL362

IGNORE CONTAINER START= MUST BE SPECIFIED

Explanation:

The **START** keyword parameter, which is required for the **IGNORE CONTAINER** control statement, is missing. Message SCRTTOOL135 lists the control statement in error.

User response:

Specify the **START** keyword and a valid start time on the indicated **IGNORE CONTAINER** control statement.

SCRTTOOL363

IGNORE CONTAINER RESUME= MUST BE SPECIFIED

Explanation:

The **RESUME** keyword parameter, which is required for the **IGNORE CONTAINER** control statement, is missing. Message SCRTTOOL135 lists the control statement in error.

User response:

Specify the **RESUME** keyword and a valid start time on the indicated **IGNORE CONTAINER** control statement.

SCRTTOOL364

IGNORE CONTAINER ID= INVALID VALUE

Explanation:

An **UPDATE SOLUTION_ID** command (control statement) was found for which the **OLD** keyword parameter is not valid. The **OLD** parameter specifies the 64-character solution ID. Message SCRTTOOL135 lists the control statement in error.

User response:

Modify the **UPDATE SOLUTION_ID** command to specify a correct **OLD** parameter. You can obtain the solution ID from the IBM LMS web portal.

SCRTTOOL365**IGNORE CONTAINER RESUME=, ERROR IN DATE FORMAT (date)**

Explanation

The value for the **START** keyword parameter on an **IGNORE CONTAINER** control statement contains a formatting error. The **START** value must be in the format *yyyy/mm/dd/hh*. Message SCRTTOOL135 lists the control statement in error.

In the message text:

date

The **START** value that is incorrectly formatted.

User response:

Review the format of the **START** date, and modify the **IGNORE CONTAINER** control statement to specify the correct value for the **START** parameter.

SCRTTOOL366**IGNORE CONTAINER RESUME=, ERROR IN DATE FORMAT (date)**

Explanation

The value for the **RESUME** keyword parameter on an **IGNORE CONTAINER** control statement contains a formatting error. The **RESUME** value must be in the format *yyyy/mm/dd/hh*. Message SCRTTOOL135 lists the control statement in error.

In the message text:

date

The **RESUME** value that is incorrectly formatted.

User response:

Review the format of the **RESUME** date, and modify the **IGNORE CONTAINER** control statement to specify the correct value for the **RESUME** parameter.

SCRTTOOL367**IGNORE CONTAINER BOTH START AND RESUME MUST BE EITHER UTC OR LOCAL**

Explanation:

There is a mismatch in the formats of the date/time values specified for the **START** and **RESUME** parameters of the **IGNORE CONTAINER** control statement. You must specify both the **START** and **RESUME** parameter values in either the local format or the UTC format.

User response:

Change the values for the **START** and **RESUME** parameters of the **IGNORE CONTAINER** control statement so that both use the local time format (*yyyy/mm/dd/hh*) or both use the UTC time format (*yyyy/mm/dd/hh/UTC*). Note that there are specific conditions under which you must use the UTC format.

SCRTTOOL368**IGNORE CONTAINER DUPLICATE COMMAND ENCOUNTERED**

Explanation:

An **IGNORE CONTAINER** command (control statement) was found for which the one of the keyword parameters is duplicated. Message SCRTTOOL135 lists the control statement in error.

User response:

Remove the duplicate parameter from the indicated **IGNORE CONTAINER** command.

SCRTTOOL369**IGNORE CONTAINER SOLUTION ID NOT FOUND IN SMF DATA ID=id**

Explanation

An **UPDATE CONTAINER** command (control statement) was specified for a container that was not found in the SMF input data. No tenant resource group was found having the solution ID, and no image that provided SMF data was assigned to the solution ID using the **CONTAINER** control statement.

In the message text:

id

The 64-character solution ID for the container that was not found in the SMF data.

User response:

Check the solution ID (*id*) identified in the message. Validate that this solution ID exists in the environment or is assigned to a dedicated image, as appropriate.

SCRTTOOL370

SOLUTION ID IS NOT ELIGIBLE TO BE IGNORED*ID=id*

Explanation:

The solution ID (indicated by *id* in the message text) specified on an **IGNORE CONTAINER** control statement (or the CLI or GUI equivalent) is not eligible to be ignored. Certain solution IDs cannot be ignored according to IBM business rules.

User response:

Remove or update the **IGNORE CONTAINER** control statement to generate your report.

SCRTTOOL371

SOLUTION ID IS NOT ELIGIBLE FOR CONTAINER COMMAND*ID=id*

Explanation:

The solution ID (indicated by *id* in the message text) specified on an **CONTAINER** control statement (or the CLI or GUI equivalent) is not eligible to be assigned to a whole LPAR. The solution ID is valid only for containers derived from Tenant Resource Groups or is otherwise not eligible under IBM business rules.

User response:

Remove or update the **CONTAINER** control statement to generate your report.

SCRTTOOL372

**SOLUTION ID OLD= IS NOT ELIGIBLE TO BE UPDATED TO
NEW=OLD=*old*NEW=*new***

Explanation:

An **UPDATE SOLUTION_ID** control statement (or the CLI or GUI equivalent) attempted to update a solution ID (indicated by *old* in the message text) to another solution ID (indicated by *new* in the message text). Certain solution ID updates are not permitted under IBM business rules.

User response:

Remove or update the **UPDATE SOLUTION_ID** control statement to generate your report.

SCRTTOOL373

UPDATE CONTAINER NO89_PRODUCT_ID= ALREADY SPECIFIED

Explanation:

There is an error in an **UPDATE CONTAINER** control statement. There are one or more duplicate **NO89_PRODUCT_ID** parameters.

User response:

Review the **UPDATE CONTAINER** control statement and remove the duplicate keyword. A single **NO89_PRODUCT_ID** parameter can specify multiple product ID values.

SCRTTOOL374

UPDATE CONTAINER NO89_PRODUCT_ID= SYNTAX ERROR

Explanation:

There is an error in an **UPDATE CONTAINER** control statement. The product ID value for the **NO89_PRODUCT_ID** parameter is incorrectly specified.

User response:

Review the **UPDATE CONTAINER** control statement and correct the product ID value.

SCRTTOOL375

UPDATE CONTAINER NO89_PRODUCT_ID= LENGTH MUST BE 8 (*pid*)

Explanation:

There is an error in an **UPDATE CONTAINER** control statement. The product ID value (*pid*) specified by the **NO89_PRODUCT_ID** parameter must be 8 characters in length.

User response:

Review the **UPDATE CONTAINER** control statement and correct the product ID value.

SCRTTOOL376	UPDATE CONTAINER NO89_PRODUCT_ID= NOT A NO89 PRODUCT (<i>pid</i>)
--------------------	--

Explanation:

There is an error in an **UPDATE CONTAINER** control statement. The product ID value (*pid*) specified by the **NO89_PRODUCT_ID** parameter is not a NO89 product.

User response:

Review the **UPDATE CONTAINER** control statement and correct the product ID value.

SCRTTOOL377	UPDATE CONTAINER NO89_PRODUCT_ID= LIST ENTRY LENGTH MUST BE 8 (<i>pid</i>)
--------------------	---

Explanation:

There is an error in an **UPDATE CONTAINER** control statement. Each product ID value (*pid*) specified in the list of values on the **NO89_PRODUCT_ID** parameter must be 8 characters in length.

User response:

Review the **UPDATE CONTAINER** control statement and correct the product ID value.

SCRTTOOL378	UPDATE CONTAINER NO89_PRODUCT_ID= LIST ENTRY NOT A NO89 PRODUCT (<i>pid</i>)
--------------------	---

Explanation:

There is an error in an **UPDATE CONTAINER** control statement. Each product ID value (*pid*) specified in the list of values on the **NO89_PRODUCT_ID** parameter must be a valid NO89 product.

User response:

Review the **UPDATE CONTAINER** control statement and correct the product ID value.

SCRTTOOL380	INCLUDE CPC= MUST BE SPECIFIED
--------------------	---------------------------------------

Explanation:

The **CPC** keyword parameter, which is required for the **INCLUDE** control statement, is missing.

User response:

Add the **CPC** keyword and a valid CPC value to the **INCLUDE** control statement.

SCRTTOOL381	INCLUDE CPC= ALREADY SPECIFIED
--------------------	---------------------------------------

Explanation:

An **INCLUDE** control statement contains more than one instance of the **CPC** keyword parameter.

User response:

Remove the additional instances of the **CPC** keyword parameter. If multiple CPCs must be specified, use a separate **INCLUDE** control statement for each CPC.

SCRTTOOL382	INCLUDE CPC= INVALID VALUE
--------------------	-----------------------------------

Explanation:

The value for the **CPC** keyword parameter on an **INCLUDE** control statement is not valid. The **CPC** value must be 10 characters long, in the form *tttt-sssss*.

User response:

Correct the **CPC** specification on the **INCLUDE** control statement.

SCRTTOOL383	INCLUDE IMAGE_ID= MUST BE SPECIFIED
--------------------	--

Explanation:

The **IMAGE_ID** keyword parameter, which is required for the **INCLUDE** control statement, is missing.

User response:

Add the **IMAGE_ID** keyword and a valid value or list of values to the **INCLUDE** control statement.

SCRTTOOL384	INCLUDE IMAGE_ID= ALREADY SPECIFIED
--------------------	--

Explanation:

An **INCLUDE** control statement contains more than one instance of the **IMAGE_ID** keyword parameter.

User response:

Remove the additional instances of the **IMAGE_ID** keyword parameter. If necessary, you can specify multiple images by using the list form of the **IMAGE_ID** parameter.

SCRTTOOL385 **INCLUDE IMAGE_ID= INVALID VALUE****Explanation:**

The value for the **IMAGE_ID** keyword parameter on an **INCLUDE** control statement is not valid.

User response:

Correct the **IMAGE_ID** specification on the **INCLUDE** control statement to specify a valid image name or list of image names.

SCRTTOOL386 **INCLUDE COMMAND INVALID FORMAT****Explanation:**

The **INCLUDE** control statement is not valid.

User response:

Review and correct the **INCLUDE** control statement.

SCRTTOOL387 **INCLUDED IMAGES NOT FOUND ON CPC=*cpc* IN DATAIMAGES=(*image-list*)****Explanation:**

The images (*image-list*) specified on an **INCLUDE** control statement for CPC *cpc* are not found in the input data.

User response:

Review the **INCLUDE** control statements to ensure that all systems that are intended to be included are specified correctly. Review the input data to ensure that all data is provided.

SCRTTOOL388 **INCLUDED CPC=*cpc* NOT FOUND IN DATA****Explanation:**

The indicated CPC (*cpc*) specified on an **INCLUDE** control statement is not found in the input data.

User response:

Review the **INCLUDE** control statements to ensure the CPC is specified correctly. Review the input data to ensure that all data is provided.

SCRTTOOL389 **INCLUDE COMMENT= ALREADY SPECIFIED****Explanation:**

An **INCLUDE** control statement contains more than one instance of the **COMMENT** keyword parameter. Each **INCLUDE** control statement supports only one **COMMENT** instance.

User response:

Remove the extra instances of the **COMMENT** parameter.

SCRTTOOL390 **INCLUDE COMMENT= EXCEEDS MAX LENGTH OF *length*****Explanation:**

The comment specified via the **COMMENT** keyword parameter of an **INCLUDE** control statement exceeds the maximum allowable length (*length*).

User response:

Update the comment so that it does not exceed the allowable length.

SCRTTOOL398 **EXCLUDE MSU_CONSUMPTION CPUTIME= INVALID VALUE****Explanation:**

The **CPUTIME** value specified on an **EXCLUDE MSU_CONSUMPTION** control statement is not valid.

User response:

Specify the **CPUTIME** value as a whole number of seconds.

SCRTTOOL399 **EXCLUDE MSU_CONSUMPTION CPC= MUST BE SPECIFIED****Explanation:**

The **CPC** keyword parameter, which is required for the **EXCLUDE MSU_CONSUMPTION** control statement, is missing.

User response:

Add the **CPC** keyword and a valid CPC value to the **EXCLUDE MSU_CONSUMPTION** control statement.

SCRTT00L400 **EXCLUDE MSU_CONSUMPTION CPC= PARAMETER INVALID VALUE**

Explanation:

The value for the **CPC** keyword parameter on an **EXCLUDE MSU_CONSUMPTION** control statement is not valid. The CPC value must be 10 characters long, in the form *tttt-ssss*.

User response:

Correct the **CPC** specification on the **EXCLUDE MSU_CONSUMPTION** control statement.

SCRTT00L401 **EXCLUDE MSU_CONSUMPTION CPC= ALREADY SPECIFIED**

Explanation:

An **EXCLUDE MSU_CONSUMPTION** control statement contains more than one instance of the **CPC** keyword parameter.

User response:

Remove the additional instances of the **CPC** parameter. If multiple CPCs must be specified, use a separate **EXCLUDE MSU_CONSUMPTION** control statement for each CPC.

SCRTT00L402 **EXCLUDE MSU_CONSUMPTION CPC= VALUE LENGTH MUST BE 10, TTTT-SSSS**

Explanation:

The value for the **CPC** keyword parameter on an **EXCLUDE MSU_CONSUMPTION** control statement is not valid. The **CPC** value must be 10 characters long, in the form *tttt-ssss*.

User response:

Correct the **CPC** specification on the **EXCLUDE MSU_CONSUMPTION** control statement.

SCRTT00L403 **MSU CAPACITY INFORMATION MISSING FROM SMF 70s**

Explanation:

An **EXCLUDE MSU_CONSUMPTION** control statement was specified, but the SMF records do not contain the data necessary to calculate the MSU consumption based on the CPU time specified.

User response:

z/OS 2.1 and later includes the necessary data in the SMF type 70 records. If you need to exclude MSU consumption from an earlier z/OS release, work with your IBM representative.

SCRTT00L404 **EXCLUDE MSU_CONSUMPTION ERROR CALCULATING MSU CONSUMPTION TO EXCLUDE**

Explanation:

SCRT encountered an error while calculating the MSUs to exclude based on an **EXCLUDE MSU_CONSUMPTION** control statement.

User response:

Try specifying a smaller amount of CPU time to exclude on the **EXCLUDE MSU_CONSUMPTION** control statement. If this does not resolve the error, contact IBM Support and open a Service Request.

SCRTT00L405 **SOLUTION ID IS NOT APPLICABLE FOR EXCLUDE MSU_CONSUMPTION**

Explanation:

The solution ID specified on an **EXCLUDE MSU_CONSUMPTION** control statement is not associated with a consumption-based offering and is not eligible for excluding MSU consumption.

User response:

Review the **EXCLUDE MSU_CONSUMPTION** control statement and correct or update the specified solution ID. For non-consumption-based offerings, use another form of the **EXCLUDE** control statement, if appropriate.

SCRTT00L406 **EXCLUDE MSU_CONSUMPTION ID= MUST BE SPECIFIED**

Explanation:

The **ID** keyword parameter, which is required for the **EXCLUDE MSU_CONSUMPTION** control statement, is missing.

User response:

Add the **ID** keyword and a valid consumption-based solution ID value to the **EXCLUDE MSU_CONSUMPTION** control statement.

SCRTTOOL407**EXCLUDE MSU_CONSUMPTION DUPLICATE ID= SPECIFIED****Explanation:**

An **EXCLUDE MSU_CONSUMPTION** control statement contains more than one instance of the **ID** keyword parameter.

User response:

Remove the additional instances of the **ID** parameter. If multiple solution IDs must be specified, use a separate **EXCLUDE MSU_CONSUMPTION** control statement for each ID.

SCRTTOOL408**EXCLUDE MSU_CONSUMPTION ID= INVALID VALUE****Explanation:**

The value for the **ID** keyword parameter on an **EXCLUDE MSU_CONSUMPTION** control statement is not valid. The ID value must be a valid Tailored Fit Pricing solution ID.

User response:

Correct the **ID** parameter specification on the **EXCLUDE MSU_CONSUMPTION** control statement.

SCRTTOOL409**EXCLUDE MSU_CONSUMPTION HOUR= MUST BE SPECIFIED****Explanation:**

The **HOUR** keyword parameter, which is required for the **EXCLUDE MSU_CONSUMPTION** control statement, is missing.

User response:

Add the **HOUR** keyword and a valid hour value to the **EXCLUDE MSU_CONSUMPTION** control statement.

SCRTTOOL410**EXCLUDE MSU_CONSUMPTION DUPLICATE HOUR= SPECIFIED****Explanation:**

An **EXCLUDE MSU_CONSUMPTION** control statement contains more than one instance of the **HOUR** keyword parameter.

User response:

Remove the additional instances of the **HOUR** parameter. If an incident occurs over multiple hours, use a separate **EXCLUDE MSU_CONSUMPTION** control statement for each hour.

SCRTTOOL411**EXCLUDE MSU_CONSUMPTION HOUR=, ERROR IN DATE FORMAT (*date*)****Explanation:**

An **EXCLUDE MSU_CONSUMPTION** control statement specifies an invalid value (*date*) for the **HOUR** parameter.

User response:

Specify the value for the **HOUR** parameter in the format *yyyy/mm/dd/hh*.

SCRTTOOL412**EXCLUDE MSU_CONSUMPTION CPUTIME= MUST BE SPECIFIED****Explanation:**

The **CPUTIME** keyword parameter, which is required for the **EXCLUDE MSU_CONSUMPTION** control statement, is missing.

User response:

Add the **CPUTIME** keyword and the CPU seconds to the **EXCLUDE MSU_CONSUMPTION** control statement.

SCRTTOOL413**EXCLUDE MSU_CONSUMPTION INVALID FORMAT****Explanation:**

The **EXCLUDE MSU_CONSUMPTION** control statement is not valid.

User response:

Review and correct the **EXCLUDE MSU_CONSUMPTION** control statement.

SCRTTOOL414 **EXCLUDE MSU_CONSUMPTION INVALID HOUR SPECIFIED****Explanation:**

No data was found for the hour specified by the **HOUR** parameter on an **EXCLUDE MSU_CONSUMPTION** control statement. SCRT cannot compute the MSU consumption without valid data from the CPC for the specified hour.

User response:

Review the **EXCLUDE MSU_CONSUMPTION** control statement and verify that the correct hour was specified. Verify that you have supplied all the relevant SMF records and that the records contain data for the specified hour.

SCRTTOOL415 **EXCLUDE MSU_CONSUMPTION DUPLICATE CPUTIME= SPECIFIED****Explanation:**

An **EXCLUDE MSU_CONSUMPTION** control statement contains more than one instance of the **CPUTIME** keyword parameter.

User response:

Remove the additional instances of the **CPUTIME** parameter. If you need to exclude multiple incidents, sum the CPU time or use a separate **EXCLUDE MSU_CONSUMPTION** control statement for each hour.

SCRTTOOL416 **EXCLUDE MSU_CONSUMPTION SOLUTION NOT ACTIVE ON
CPC=*cpcID=id*****Explanation:**

An **EXCLUDE MSU_CONSUMPTION** control statement was specified, but there were no systems or tenant resource groups on the specified CPC (*cpc*) that were associated with the specified solution ID (*id*).

User response:

Correct the **EXCLUDE MSU_CONSUMPTION** control statement to specify the proper CPC or solution ID or both.

SCRTTOOL425 **NOT ALL IMAGES WERE ASSIGNED TO A CONTAINER****Explanation:**

For an Enterprise Tailored Fit Pricing report, all z/OS images must be assigned to a container by using the **CONTAINER** control statement. SCRT found one or more images that were not assigned to a container. The SCRTTOOL426 message is issued for each image that was unassigned.

User response:

Review the input to SCRT and the SCRTTOOL426 messages. Use the **CONTAINER** control statement to assign all images to a container.

SCRTTOOL426 **CPC=*cpc* IMAGE_ID=*image* IS NOT ASSOCIATED WITH A CONTAINER****Explanation:**

The image specified by *image* on the CPC specified by *cpc* is not assigned to a container. SCRT issues this message only when an Enterprise Tailored Fit Pricing report is requested.

User response:

Review the input to SCRT and the SCRTTOOL426 messages. Use the **CONTAINER** control statement to assign all images to a container.

SCRTTOOL427 **IGNORE CONTAINER IS NOT SUPPORTED WHEN GENERATING AN
ENTERPRISE TAILORED FIT REPORT****Explanation:**

The **IGNORE CONTAINER** control statement is not valid when an Enterprise Tailored Fit Pricing report is requested.

User response:

Remove the **IGNORE CONTAINER** control statement from your SCRT input.

SCRTTOOL428 **ENTERPRISE TAILORED FIT PRICING IS INCOMPATIBLE WITH:
*function*****Explanation:**

The indicated SCRT control statement or function (*function*) is not valid when an Enterprise Tailored Fit Pricing report is requested.

User response:

Remove the incompatible control statement from your SCRT input.

SCRTTOOL429	COMMAND LINE ARGUMENT <i>arg</i> IS INCOMPATIBLE WITH ENTERPRISE TAILORED FIT PRICING
--------------------	--

Explanation:

The indicated SCRT command line option (*arg*) is not valid when an Enterprise Tailored Fit Pricing report is requested.

User response:

Remove the incompatible command line argument from your SCRT invocation.

SCRTTOOL601	ENVIRONMENT INFORMATION: <i>text</i>
--------------------	---

Explanation

This message displays diagnostic information about the environment for use by IBM in diagnosing SCRT failures.

In the message text:

text

The environment diagnostic information.

User response:

Provide this information to IBM Support when reporting SCRT problems.

SCRTTOOL602	UNSUPPORTED FUNCTION: <i>xxxxxxx</i> DD STATEMENT
--------------------	--

Explanation

SCRT encountered a DD statement that is not supported by the version of SCRT that you are running.

In the message text:

xxxxxxx

The name of the unsupported DD statement.

User response:

Review your requirement for the function provided by the specified DD statement. You must either remove the DD statement from the JCL, or use a version of SCRT that supports that function.

SCRTTOOL603	SCRT PROCESSING HAS FAILED DUE TO A SEVERE ERROR.
--------------------	--

Explanation:

SCRT has encountered a severe error and cannot continue; processing has been halted.

User response:

Look for accompanying messages and attempt problem determination and resolution using the information provided in those messages.

SCRTTOOL605	DIAGNOSTICS: <i>text</i>
--------------------	---------------------------------

Explanation

A ZFile exception has occurred.

In the message text, *text* provides a description of the error.

User response:

Correct the condition that is causing the file error and rerun SCRT.

SCRTTOOL610	SMF RECORDS DO NOT CONTAIN REQUIRED MOBILE DATA FOR CPC: <i>cpc_name</i> LPAR: <i>lpar_name</i>
--------------------	--

Explanation

This message displays diagnostic information related to the SCRT support for WLM mobile workload classification. When the DIAGMSG=WLMMOB control statement is specified in the SPECIAL DD, the SCRTTOOL610 message is issued for each CPC:LPAR that does not have the required APAR OA48466 applied and accepted.

In the message text:

cpc_name

The name of the CPC on which *lpar_name* resides.

lpar_name

The name of the LPAR that requires that APAR OA48466 be applied and accepted.

User response:

Apply and accept APAR OA48466 on the specified LPAR.

SCRTTOOL611	SMF RECORDS DO NOT CONTAIN REQUIRED MOBILE DATA FOR CPC: <i>cpc_name</i> LPAR: <i>lpar_name</i>
--------------------	--

Explanation

This message displays diagnostic information related to the SCRT support for WLM mobile workload classification. When the DIAGMSG=WLMMOB control statement is specified in the SPECIAL DD, the SCRTTOOL611 message is issued for each CPC:LPAR that reports a zero value for the mobile 4-hour rolling average CPU consumption field during the reporting period.

In the message text:

cpc_name

The name of the CPC on which *lpar_name* resides.

lpar_name

The name of the LPAR that reported a zero value for the mobile 4-hour rolling average utilization.

User response:

See the WLM documentation for mobile workload classification in [z/OS MVS Planning: Workload Management](#).

SCRTTOOL612	SMF RECORDS DO NOT CONTAIN REQUIRED DATA FOR CPC: <i>cpc</i> LPAR: <i>image</i>
--------------------	--

Explanation

The SMF records that were collected and provided as input for WLM for Cloud are missing a category field.

In the message text:

cpc

The name of the CPC whose SMF records are missing the category field.

image

The image ID on *cpc* whose SMF records are missing the category field.

User response:

Correct your data collection by ensuring that your workload is correctly classified in z/OS workload management (WLM) to use the proper category.

SCRTTOOL613	SMF RECORDS DO NOT CONTAIN A NON-ZERO VALUE FOR CPC: <i>cpc</i> LPAR: <i>image</i>
--------------------	---

Explanation

The SMF records that were collected and provided as input for WLM for Cloud always contain a zero value in the category fields.

In the message text:

cpc

The name of the CPC whose SMF records always contain a zero value in the category fields.

image

The image ID on *cpc* whose SMF records always contain a zero value in the category field.

User response:

Correct your data collection by ensuring that your workload is correctly classified in z/OS workload management (WLM) to use the proper category.

SCRTT00L620

**ASSIGN ZCAP COMMAND DOES NOT CONTAIN REQUIRED ZCAPNAME
PARAMETER: *text***

Explanation

The IBM Z Collocated Application Pricing (zCAP) workload assigned to this WLM category must be named.

In the message text, *text* contains the **Assign ZCAP** control statement, as presented to SCRT in the SPECIAL DD.

User response:

Correct the **Assign ZCAP** control statement in the SPECIAL DD and rerun SCRT.

SCRTT00L621

**INPUTCSV WORKLOAD NAME DOES NOT MATCH NAME SPECIFIED ON
ASSIGN ZCAP COMMAND**

Explanation:

The zCAP workload name in the INPUTCSV file does not match the name specified on the **Assign ZCAP** control statement in the SPECIAL DD.

User response:

Correct the input so that the workload name in the INPUTCSV file matches the name specified on the **Assign ZCAP** control statement in the SPECIAL DD.

SCRTT00L650

ERROR OPENING CSV FILE NAME: *text*

Explanation

SCRT encountered an error upon attempting to open the specified INPUTCSV file. (This message applies to SCRT running in a non-z/OS environment.)

In the message text, *text* indicates the name of the file that failed to open.

User response:

Correct the file name and try again.

SCRT messages for internal SCRT restrictions

SCRT messages in the format SCRTT00L5xx report internal SCRT restrictions. Whenever SCRT generates a message with a message ID from the following list, contact the SCRT support team for assistance by sending an email to scrt@us.ibm.com with the error message that you received.

- **SCRTT00L500**
- **SCRTT00L501**
- **SCRTT00L502**
- **SCRTT00L503**
- **SCRTT00L504**
- **SCRTT00L505**
- **SCRTT00L506**
- **SCRTT00L507**
- **SCRTT00L508**
- **SCRTT00L509**
- **SCRTT00L510**

SCRT89 Record Utility messages

This section contains explanations of all messages issued by the SCRT89 Record Utility.

SCRTRU001	UNABLE TO OPEN SMFUSFS DD STATEMENT.
------------------	---

Explanation:

The SCRT89 record utility could not open the SMF UNIX file system.

User response:

Make sure the SMFUSFS DD statement exists and is correctly coded. Then rerun the job.

SCRTRU002	UNABLE TO OPEN SMFVB DD STATEMENT.
------------------	---

Explanation:

The SCRT89 record utility could not open the SMF output data set. The SMFVB DD statement may not exist, or it may exist but the data set specified on it may not be cataloged.

User response:

Make sure the SMFVB DD statement exists and is correctly cataloged. Then rerun the job.

SCRTRU003	FINAL RECORD IS INCOMPLETE. RECORD OMITTED.
------------------	--

Explanation:

Not all of the data for the final record was present in the input data set.

User response:

Check your input and run the job again.

SCRTRU004	NEGATIVE RECORD DESCRIPTOR WORD ENCOUNTERED.
------------------	---

Explanation:

The input data is not SCRT data, or the data has been corrupted.

User response:

Check your input data.

SCRTRU005	INPUT FILE NOT CONTAINED IN A UNIX FILE SYSTEM.
------------------	--

Explanation:

The SMFUSFS DD statement must point to a z/OS UNIX file system.

User response:

Correct the JCL and rerun the job.

Appendix B. Advanced topics in sub-capacity pricing and reporting

The following are some advanced topics in sub-capacity pricing and reporting and might not apply to every installation:

- [“Capping individual LPARs” on page 497](#)
- [“Group capacity” on page 498](#)
- [“Static power save mode” on page 506](#)
- [“Sending TSAD data” on page 514](#)

Capping individual LPARs

There are several techniques for controlling an LPAR's 4-hour rolling average utilization in a sub-capacity software pricing environment. We refer to these capping techniques as *individual LPAR capping* to distinguish them from LPAR group capping (described in [“Group capacity” on page 498](#)). These individual LPAR capping techniques include:

- PR/SM hard capping
- LPAR physical capacity limit
- Defined capacity

If you have set a cap for the capacity of an individual LPAR but the 4-hour rolling average utilization of the LPAR does not exceed that capping value, the software charges will be based on the highest observed 4-hour rolling average utilization for that LPAR.

If the 4-hour rolling average utilization of the LPAR exceeds the capping value that has been set for that LPAR, SCRT reports the capped value for that LPAR and IBM charges for the software at the capped level (rather than at the higher 4-hour rolling average utilization).

You can use the Hardware Management Console (HMC) to set the capping values for each of the individual capping techniques for a given LPAR, even though some of those capping values might cause other capping values to be ignored. The HMC does not issue any warning messages when that condition occurs.

PR/SM hard capping

PR/SM hard capping has the following characteristics:

- PR/SM hard capping is applied to all z/OS, z/TPF, and z/VSE products, including the operating system products themselves, running on hard-capped LPARs. If multiple versions of z/OS or z/VSE are running on hard-capped LPARs, hard capping applies to each version of z/OS individually (which are summed and reported in the z/OS (A11) entry) and each version of z/VSE individually (which are summed and reported in the z/VSE (A11) entry) running on those LPARs.
- PR/SM hard capping is supported on all sub-capacity eligible hardware.
- The LPAR capping value is optional and is based on the weight assigned to a specific LPAR relative to the weights assigned to all of the other LPARs.
- The LPAR capping value is set via the HMC.
- PR/SM enforces the capping for an LPAR independently of the level of the operating system running on that LPAR.
- PR/SM hard capping causes z/OS (WLM) and SCRT to ignore any defined capacity or LPAR group capacity values that might also be set for the hard-capped LPAR.

Group capacity

- SCRT enforces the lesser of the PR/SM hard cap value and the LPAR physical capacity limit capping value whenever PR/SM is not precise enough to maintain the capped value.

LPAR physical capacity limit

LPAR physical capacity limit has the following characteristics:

- LPAR physical capacity limit is applied to all z/OS, z/TPF, and z/VSE products, including the operating system products themselves, running on LPARs with a physical capacity limit. If multiple versions of z/OS or z/VSE are running on LPARs with a physical capacity limit, the physical capacity limit applies to each version of z/OS individually (which are summed and reported in the z/OS (A11) entry) and each version of z/VSE individually (which are summed and reported in the z/VSE (A11) entry) running on those LPARs.
- LPAR physical capacity limit is supported by zEC12 hardware starting with the GA2 level microcode, zBC12 hardware, and later processor families.
- The LPAR capping value is set in terms of hundredths of a CP and is independent of the other LPARs on the same CPC.
- The LPAR capping value is set via the HMC.
- PR/SM enforces this capping for an LPAR independently of the level of the operating system running on that LPAR.
- SCRT enforces the lesser of the PR/SM hard cap value and the LPAR physical capacity limit capping value whenever PR/SM is not precise enough to maintain the capped value.
- SCRT applies this capping value to the LPAR first, if necessary, prior to applying any LPAR group capping value that would occur if the LPAR is not hard capped, is part of an LPAR group, and the LPARs in that group need to be capped.

Defined capacity

Defined capacity has the following characteristics:

- Defined capacity is only supported for z/OS products, including the z/OS operating system product itself, running on LPARs under control of defined capacity. If multiple versions of z/OS are running on LPARs under defined capacity, the defined capacity applies to each version of z/OS individually (which are summed and reported in the z/OS (A11) entry) running on those LPARs.
- Defined capacity is supported on all sub-capacity eligible hardware.
- The LPAR capping value is set in terms of the MSUs used for software billing.
- The LPAR capping value is set via the HMC.
- If an LPAR is hard capped, z/OS (WLM) and SCRT ignore any defined capacity or LPAR group capacity values that might be set for that LPAR.
- PR/SM and the WLM component of z/OS enforce the capping for an LPAR starting with z/OS V1.1.
- SCRT enforces the lesser of the defined capacity value and the LPAR physical capacity limit capping value whenever the LPAR is not hard capped.
- SCRT applies this capping value to the LPAR first, if necessary, prior to applying any LPAR group capping value that would occur if the LPAR is not hard capped, is part of an LPAR group, and the LPARs in that group need to be capped.

Guideline: It is generally preferable to use defined capacity to cap individual LPARs, instead of the other individual LPAR capping techniques, because the defined capacity value is set in terms of billing MSUs and enables LPARs to also be capped as members of an LPAR group.

Group capacity

Group capacity allows soft capping across multiple LPARs defined in the same LPAR group.

Group capacity is only supported for z/OS products, including the z/OS operating system product itself, running on LPARs in an LPAR group. If multiple versions of z/OS are running on LPARs in an LPAR group, the group capacity applies to each version of z/OS individually (which are summed and reported in the z/OS (A11) entry) running on those LPARs.

Group capacity is available to LPARs in an IBM z14, IBM z13, zEnterprise, System z10, or System z9 environment that:

- Run native z/OS V1R8 (and higher) systems. z/OS guest systems running under z/VM do not support group capacity.
- Use z/OS workload management (WLM) to manage the 4-hour rolling average utilization of all of the LPARs within the same group
- Do not enforce relative weight (no PR/SM hard cap set)
- Use the Hardware Management Console (HMC) to define the capacity group

Note:

1. The planning tool does not cap LPARs that are defined to an LPAR group. However, SCRT provides full capping support for the LPAR defined to a group, as explained in [“How SCRT uses group capacity”](#) on page 499.
2. Although PR/SM and WLM will enforce the group capacity for the environment described earlier, SCRT has an additional requirement that needs to be met before it can perform the SCRT-specific group capping and group reporting described in [“How SCRT uses group capacity”](#) on page 499. SCRT requires that the SMF type 70 records be generated by the RMF feature shipped with z/OS V1R8 (or higher) or that comparable SMF records be generated by a competitive ISV product. The SMF type 70 records for z/OS V1R8 (and higher) have additional fields which provide data for group capping.

How SCRT uses group capacity

For the first four hours after an IPL and for the hours immediately after adding LPARs to a group or changing the group's capacity, the capacities of those LPARs may exceed either the lesser of the LPAR defined capacity value and LPAR physical capacity limit enforcement value (if either or both are set), or exceed the capacity of the LPAR group to which the LPAR has been defined (if part of a group), until WLM and PR/SM finally cap the LPARs to your provided capping values. Whenever SCRT finds SMF data that indicates that the LPARs within the same LPAR group have exceeded the capacity cap set for that group, SCRT assigns the capacity of the group to the products running in the LPARs in the group, instead of assigning the actual LPAR capacities. SCRT applies the lesser of any defined capacity capping or LPAR physical capacity limit capping (when either is applicable) to individual LPARs before assessing group capping.

The Product Max Contributors section of the sub-capacity report indicates when the peak hour for the product has been capped by SCRT at the MSU cap set for the LPAR group. The Sub-Capacity Capping Enforced by SCRT section of the sub-capacity report indicates the number of hours that SCRT was required to set the MSU value to the cap value set for the LPAR group because PR/SM and WLM could not react fast enough to enforce the capped value. The hours reported in the Sub-Capacity Capping Enforced by SCRT section do not include the number of hours that were successfully capped by PR/SM and WLM. The Group Capacity LPARs section lists the LPARs defined to each group during the reporting period.

Group capacity and Tailored Fit Pricing for IBM Z

In cases where a capacity group contains LPARs where the entire LPAR participates in a Tailored Fit Pricing solution (container LPARs), SCRT considers such LPARs separately for purposes of evaluating group capacity. That is, for all LPARs within a capacity group, group capping is applied independently for traditional, non-container LPARs and for container LPARs on a per-solution ID basis.

Examples of capping enforcement by SCRT

Note: This topic describes the types of capping that SCRT enforces or in which SCRT may otherwise participate. Other forms of capping that do not require SCRT involvement (such as absolute group capping) are not included here.

Because group capping, with or without the capping of individual LPARs within a group, can be complicated, the following set of examples illustrate how SCRT assesses the MSU values for the products in various sample configurations. These examples describe how SCRT handles capping for individual LPARs due to hard capping, LPAR physical capping, and defined capacity capping, and how SCRT handles group capping relative to the individual LPAR capping. All of the examples include LPAR physical capping, defined capacity capping, and group capping, and one example ([“Capping enforcement on a z/OS V1 and z/OS V2 configuration” on page 505](#)) also includes hard capping.

Note: All of the following examples were deliberately configured in order to describe the SCRT capping logic for all capping conditions. Although the combination of capping conditions shown in each example can technically be set using the Hardware Management Console (HMC), it is unlikely that any customer would use all of the capping combinations shown in any of the examples.

Capping enforcement on a traditional z/OS configuration

Table 59 on page 500 shows a sample traditional z/OS configuration with individual LPAR capping and group capping.

- LPAR2 has two individual LPAR capping values:
 - A defined capacity value set to 28 MSUs, and
 - An LPAR physical capacity limit value set to the equivalent of 30 MSUs
- LPAR3 has two individual LPAR capping values:
 - A defined capacity value set to 50 MSUs, and
 - An LPAR physical capacity limit value set to the equivalent of 44 MSUs
- LPAR2 and LPAR3 are part of an LPAR group that has been capped at 70 MSUs.

Each LPAR in the configuration also shows the MSU values for hours 51 and 80 of the sample reporting period.

CPC: 2096-S04 (serial = 02-12345)			
LPAR1	LPAR2	LPAR3	LPAR4
	Defined cap = 28 MSUs LPAR physical cap = 30 MSUs Group = GRPA Group cap = 70 MSUs	Defined cap = 50 MSUs LPAR physical cap = 44 MSUs Group = GRPA Group cap = 70 MSUs	
z/OS V1R7 (SYS1) CICS TS V3 IMS V8 File Manager V6	z/OS V1R8 (SYS2) DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8 File Manager V6 WebSphere Application Server V6 WebSphere MQ Workflow V3	z/OS V1R8 (SYS3) DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8 Lotus Domino V7 WebSphere Application Server V6 WebSphere MQ Workflow V3	z/VSE

Table 59. Sample z/OS traditional configuration with individual LPAR capping and group capping (continued)			
CPC: 2096-S04 (serial = 02-12345)			
hour 51 = 27 MSUs hour 80 = 25 MSUs	hour 51 = 32 MSUs hour 80 = 30 MSUs	hour 51 = 39 MSUs hour 80 = 45 MSUs	

Table 60 on page 501 shows the MSUs for hour 51 for the sample configuration shown in Table 59 on page 500.

Table 60. MSUs for hour 51 in the sample z/OS traditional configuration				
Product	MSUs contributed by each LPAR for hour 51 (actual values for each LPAR are as shown in the configuration in Table 59 on page 500)			Reported MSUs for hour 51
	LPAR1	LPAR2	LPAR3	
z/OS	27	28 ^{1,3}	39 ^{2,3}	94
DB2 UDB V8		28 ^{1,3}	39 ^{2,3}	67
CICS TS V3	27	28 ^{1,3}	39 ^{2,3}	94
MQSeries V6		28 ^{1,3}	39 ^{2,3}	67
IMS V8	27	28 ^{1,3}	39 ^{2,3}	94
Lotus Domino V7			39 ²	39
File Manager V6	27	28 ¹		55
IBM WebSphere Application Server for z/OS V6		28 ^{1,3}	39 ^{2,3}	67
IBM WebSphere MQ Workflow for z/OS V3		28 ^{1,3}	39 ^{2,3}	67
Footnotes: <ol style="list-style-type: none"> 1. The cap value for this LPAR is 28 MSUs because that is the lesser of the defined capacity value of 28 MSUs and the LPAR physical capacity limit value of 30 MSUs. SCRT assigned a sub-capacity value of 28 MSUs for this hour because the actual sub-capacity value was 32 MSUs. 2. The cap value for this LPAR is 44 MSUs because that is the lesser of the defined capacity value of 50 MSUs and the LPAR physical capacity limit value of 44 MSUs. SCRT assigned a sub-capacity value of 39 MSUs for this hour because the actual sub-capacity value was 39 MSUs, which is less than the capped value. 3. The total MSUs for the LPARs in this group was 67 MSUs, which is less than the group cap of 70 MSUs, so no group capping took place. 				

Summary analysis for hour 51 shown in Table 60 on page 501:

- LPAR1 and LPAR4 can not participate in group capping, even if LPAR1 and LPAR4 were defined to be part of the LPAR group on the LPAR definition panel, because z/OS V1R7 (running in LPAR1), does not support group capping and z/VSE (running in LPAR4) does not support LPAR capping. Group capping support was introduced in z/OS V1R8.
- SCRT capped all products running in LPAR2 for hour 51 because LPAR2 exceeded its individual LPAR cap of 28 MSUs.
- SCRT did not apply group capping to the LPARs in this group (LPAR2 and LPAR3) because the MSUs for the LPARs in this group totaled 67 MSUs, which is less than the group cap of 70 MSUs.

If hour 51 was the earliest peak hour for all of the products in this sample configuration for the reporting period, Figure 276 on page 502 shows what the corresponding Product Max Contributors section of the sub-capacity report would look like.

Capping examples

```
==P5=====
```

PRODUCT MAX CONTRIBUTORS									
Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR2	LPAR LPAR3	LPAR LPAR4	GROUP LGRPA	
z/OS V1	5694-A01	94	04 Oct 2024 - 03:00	27	28	39	0	0	
DB2 UDB for z/OS V8	5625-DB2	67	04 Oct 2024 - 03:00	0	28	39	0	0	
CICS TS for z/OS V3	5655-M15	94	04 Oct 2024 - 03:00	27	28	39	0	0	
MQSeries for z/OS V6	5655-L82	67	04 Oct 2024 - 03:00	0	28	39	0	0	
IMS V8	5655-C56	94	04 Oct 2024 - 03:00	27	28	39	0	0	
Lotus Domino for z/OS V7	5655-N14	39	04 Oct 2024 - 03:00	0	0	39	0	0	
File Manager for z/OS V6	5655-P17	55	04 Oct 2024 - 03:00	27	28	0	0	0	
IBM WebSphere Application Server for z/OS V6	5655-N01	67	04 Oct 2024 - 03:00	0	28	39	0	0	
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	67	04 Oct 2024 - 03:00	0	28	39	0	0	
VSE Central Functions	5686-CF8	11	04 Oct 2024 - 03:00	0	0	0	11	0	

Figure 276. Sample Product Max Contributors section for hour 51 in the sample z/OS traditional configuration

Table 61 on page 502 shows the MSUs for hour 80 for the sample configuration shown in Table 59 on page 500.

Table 61. MSUs for hour 80 in the sample z/OS traditional configuration				
Product	MSUs contributed by each LPAR for hour 80 (actual values for each LPAR are as shown in the configuration in Table 59 on page 500)			Reported MSUs for hour 80
	LPAR1	LPAR2	LPAR3	
z/OS	25	28 ^{1,3}	44 ^{2,3}	95
DB2 UDB V8		28 ^{1,3}	44 ^{2,3}	70
CICS TS V3	25	28 ^{1,3}	44 ^{2,3}	95
MQSeries V6		28 ^{1,3}	44 ^{2,3}	70
IMS V8	25	28 ^{1,3}	44 ^{2,3}	95
Lotus Domino V7			44 ³	44
File Manager V6	25	28 ¹		53
IBM WebSphere Application Server for z/OS V6		28 ^{1,3}	44 ^{2,3}	70
IBM WebSphere MQ Workflow for z/OS V3		28 ^{1,3}	44 ^{2,3}	70
Footnotes: <ol style="list-style-type: none"> 1. The cap value for this LPAR is 28 MSUs because that is the lesser of the defined capacity value of 28 MSUs and the LPAR physical capacity limit value of 30 MSUs. SCRT assigned a sub-capacity value of 28 MSUs for this hour because the actual sub-capacity value was 32 MSUs. 2. The cap value for this LPAR is 44 MSUs because that is the lesser of the defined capacity value of 50 MSUs and the LPAR physical capacity limit value of 44 MSUs. SCRT assigned a sub-capacity value of 44 MSUs for this hour because the actual sub-capacity value was 45 MSUs. 3. The total MSUs for the LPARs in this group was 72 MSUs (28 + 44), which is more than the group cap of 70 MSUs. For hour 80, SCRT assigned a value of 70 MSUs to products that ran in both LPAR2 and LPAR3, instead of the sum of the actual MSUs for those LPARs. 				

Summary analysis for hour 80 shown in Table 61 on page 502:

- LPAR1 can not participate in group capping, even if LPAR1 was defined to be part of the LPAR group on the LPAR definition panel, because z/OS 1.7, which is running in LPAR1 does not support group capping. Group capping support was introduced by z/OS 1.8.
- SCRT capped all products running in LPAR2 for hour 80, because LPAR2 exceeded its individual LPAR cap of 28 MSUs.
- SCRT capped all products running in LPAR3 for hour 80 because LPAR 3 exceeded its individual LPAR cap of 44 MSUs.
- SCRT set the group cap for all products that were running in both of the LPARs in the group (LPAR2 and LPAR3) during hour 80, because the capacities of both LPARs in the group (72 MSUs) exceed the group cap.
- SCRT was not required to cap any product that ran in only one of the LPARs in the group because neither LPAR alone exceeded the capacity of the group.

If hour 80 was the earliest peak hour for all of the products in this sample configuration for the reporting period, [Figure 277 on page 503](#) shows what the corresponding Product Max Contributors section of the sub-capacity report would look like.

```

==P5=====
PRODUCT MAX CONTRIBUTORS

```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR2	LPAR LPAR3	LPAR LPAR4	GROUP LGRPA
z/OS V1	5694-A01	95	05 Oct 2024 - 08:00	25	0	0	0	70
DB2 UDB for z/OS V8	5625-DB2	70	05 Oct 2024 - 08:00	0	0	0	0	70
CICS TS for z/OS V3	5655-M15	95	05 Oct 2024 - 08:00	25	0	0	0	70
MQSeries for z/OS V6	5655-L82	70	05 Oct 2024 - 08:00	0	0	0	0	70
IMS V8	5655-C56	95	05 Oct 2024 - 08:00	25	0	0	0	70
Lotus Domino for z/OS V7	5655-N14	44	05 Oct 2024 - 08:00	0	0	44	0	0
File Manager for z/OS V6	5655-P17	53	05 Oct 2024 - 08:00	25	28	0	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	70	05 Oct 2024 - 08:00	0	0	0	0	70
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	70	05 Oct 2024 - 08:00	0	0	0	0	70
VSE Central Functions	5686-CF8	13	05 Oct 2024 - 08:00	0	0	0	13	0

Figure 277. Sample Product Max Contributors section for hour 80 in the sample z/OS traditional configuration

Capping enforcement on a z/OS traditional and z/OS zNALC configuration

z/OS zNALC and z/OS traditional are treated the same way as any other product for individual LPAR capping and group capping. This can be demonstrated by changing the configuration to contain a zNALC LPAR. LPAR3 was replaced by ZNAL3, which is a zNALC LPAR. Note that some of the products were dropped from the zNALC LPAR because zNALC LPARs must be dedicated to zNALC workloads and the applications which support those zNALC workloads. See [Chapter 7, “Sub-capacity reporting for System z New Application License Charges \(zNALC\),” on page 145](#) for a description of how SCRT reports on configurations with zNALC LPARs.

Table 62 on page 503 shows a sample z/OS traditional and z/OS zNALC configuration with individual LPAR capping and group capping.

CPC: 2096-S04 (serial = 02-12345)			
LPAR1	LPAR2	ZNAL3	LPAR4
	Defined cap = 28 MSUs LPAR physical cap = 30 MSUs Group = GRPA Group cap = 70 MSUs	Defined cap = 50 MSUs LPAR physical cap = 44 MSUs Group = GRPA Group cap = 70 MSUs	
z/OS V1R7 (SYS1)	z/OS V1R8 (SYS2)	z/OS V1R8 (SYS3)	z/VSE
CICS TS V3	DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8	DB2 UDB V8	
IMS V8		Lotus Domino V7	
File Manager V6	File Manager V6 WebSphere Application Server V6 WebSphere MQ Workflow V3		
(hour 80 = 25 MSUs)	(hour 80 = 30 MSUs)	(hour 80 = 45 MSUs)	

[Table 63 on page 504](#) shows the MSUs for hour 80 for the sample configuration shown in [Table 62 on page 503](#).

Capping examples

Table 63. MSUs for hour 80 in the sample z/OS traditional and z/OS zNALC configuration

Product	MSUs contributed by each LPAR for hour 80 (actual values for each LPAR are as shown in the configuration in Table 62 on page 503)			Reported MSUs for hour 80
	LPAR1	LPAR2	ZNAL3	
z/OS	25	28 ^{1,3}	44 ^{2,3}	95
z/OS (Traditional)	25	28 ¹		53
z/OS (zNALC)			44 ²	44
DB2 UDB V8		28 ^{1,3}	44 ^{2,3}	70
CICS TS V3	25	28 ¹		53
MQSeries V6		28 ¹		28
IMS V8	25	28 ¹		53
Lotus Domino V7			44 ²	44
File Manager V6	25	28 ¹		53
IBM WebSphere Application Server for z/OS V6		28 ¹		28
IBM WebSphere MQ Workflow for z/OS V3		28 ¹		28

Footnotes:

1. The cap value for this LPAR is 28 MSUs because that is the lesser of the defined capacity value of 28 MSUs and the LPAR physical capacity limit value of 30 MSUs. SCRT assigned a sub-capacity value of 28 MSUs for this hour because the actual sub-capacity value was 32 MSUs.
2. The cap value for this LPAR is 44 MSUs because that is the lesser of the defined capacity value of 50 MSUs and the LPAR physical capacity limit value of 44 MSUs. SCRT assigned a sub-capacity value of 44 MSUs for this hour because the actual sub-capacity value was 45 MSUs.
3. The total MSUs for the LPARs in this group was 72 MSUs (28 + 44), which is more than the group cap of 70 MSUs. For hour 80, SCRT assigned a value of 70 MSUs to products that ran in both LPAR2 and LPAR3, instead of the sum of the actual MSUs for those LPARs.

If hour 80 was the earliest peak hour for all of the products in this sample mixed z/OS traditional and zNALC configuration for the reporting period, Figure 278 on page 504 shows what the corresponding Product Max Contributors section of the sub-capacity report would look like.

==P5=====

PRODUCT MAX CONTRIBUTORS

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR2	LPAR ZNAL3	LPAR LPAR4	GROUP LGRPA
z/OS V1	5694-A01	95	05 Oct 2024 - 08:00	25	0	0	0	70
z/OS V1 (Traditional_	5694-A01	53	05 Oct 2024 - 08:00	25	28	0	0	0
z/OS V1 (zNALC)	5694-A01	44	05 Oct 2024 - 08:00	0	0	44	0	0
DB2 UDB for z/OS V8	5625-DB2	70	05 Oct 2024 - 08:00	0	0	0	0	70
CICS TS for z/OS V3	5655-M15	53	05 Oct 2024 - 08:00	25	28	0	0	0
MQSeries for z/OS V6	5655-L82	28	05 Oct 2024 - 08:00	0	28	0	0	0
IMS V8	5655-C56	53	05 Oct 2024 - 08:00	25	28	0	0	0
Lotus Domino for z/OS V7	5655-N14	44	05 Oct 2024 - 08:00	0	0	44	0	0
File Manager for z/OS V6	5655-P17	53	05 Oct 2024 - 08:00	25	28	0	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	28	05 Oct 2024 - 08:00	0	28	0	0	0
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	28	05 Oct 2024 - 08:00	0	28	0	0	0
VSE Central Functions	5686-CF8	12	05 Oct 2024 - 08:00	0	0	0	12	0

Figure 278. Sample Product Max Contributors section for hour 80 in the sample z/OS traditional and z/OS zNALC configuration

For more information, see the topics on:

- Defining group capacity in the chapter on “Workload Management Participants of WLM Planning” in *z/OS MVS Planning: Workload Management*.
- “Workload Charging by Soft-Capping to a Defined Capacity” in *PR/SM Planning Guide*, available from the IBM Systems Support website for your hardware model.

Capping enforcement on a z/OS V1 and z/OS V2 configuration

This example is for a CPC where z/OS V1 and z/OS V2 systems ran during the reporting period and is based on the same z/OS Traditional sample configuration shown in [Table 59 on page 500](#), with the following changes:

- The z/OS V1 system in LPAR3 was replaced with a z/OS V2 system.
- Several individual LPAR capping values were set for LPAR1, including LPAR hard capping.
- A group capping value was set for LPAR1.
- The actual sub-capacity value for LPAR1 was changed from 27 MSUs to 30 MSUs.

[Table 64 on page 505](#) shows the sample z/OS traditional configuration with z/OS V1 and z/OS V2.

CPC: 2096-S04 (serial = 02-12345)			
LPAR1	LPAR2	LPAR3	LPAR4
Defined cap = 20 MSUs LPAR hard cap = 25 MSUs LPAR physical cap = 40 MSUs Group = GRPA Group cap = 70 MSUs	Defined cap = 28 MSUs LPAR hard cap = none LPAR physical cap = 30 MSUs Group = GRPA Group cap = 70 MSUs	Defined cap = 50 MSUs LPAR hard cap = none LPAR physical cap = 44 MSUs Group = GRPA Group cap = 70 MSUs	
z/OS V1 (SYS1) CICS TS V3 IMS V8 File Manager V6 hour 80 = 30 MSUs	z/OS V1 (SYS2) DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8 File Manager V6 WebSphere Application Server V6 WebSphere MQ Workflow V3 hour 80 = 30 MSUs	z/OS V2 (SYS3) DB2 UDB V8 CICS TS V3 MQSeries V6 IMS V8 Lotus Domino V7 WebSphere Application Server V6 WebSphere MQ Workflow V3 hour 80 = 45 MSUs	z/VSE

[Table 65 on page 505](#) shows the MSUs for the sample configuration shown in [Table 64 on page 505](#).

Table 65. MSUs for the sample z/OS z/OS V1 and z/OS V2 configuration				
Product	MSUs contributed by each LPAR for hour 80 (actual values for each LPAR are as shown in the configuration in Table 64 on page 505)			Reported MSUs for hour 80
	LPAR1	LPAR2	LPAR3	
z/OS	25 ⁴	28 ^{1,3}	44 ^{2,3}	95
z/OS V2			44 ³	44
z/OS V1	25 ⁴	28 ¹		53
DB2 UDB V8		28 ^{1,3}	44 ^{2,3}	70
CICS TS V3	25 ⁴	28 ^{1,3}	44 ^{2,3}	95
MQSeries V6		28 ^{1,3}	44 ^{2,3}	70
IMS V8	25 ⁴	28 ^{1,3}	44 ^{2,3}	95
Lotus Domino V7			44 ³	44
File Manager V6	25 ⁴	28 ¹		53
IBM WebSphere Application Server for z/OS V6		28 ^{1,3}	44 ^{2,3}	70
IBM WebSphere MQ Workflow for z/OS V3		28 ^{1,3}	44 ^{2,3}	70

Table 65. MSUs for the sample z/OS z/OS V1 and z/OS V2 configuration (continued)

Product	MSUs contributed by each LPAR for hour 80 (actual values for each LPAR are as shown in the configuration in Table 64 on page 505)			Reported MSUs for hour 80
	LPAR1	LPAR2	LPAR3	
Footnotes:				
<div>1. The cap value for this LPAR is 28 MSUs because that is the lesser of the defined capacity value of 28 MSUs and the LPAR physical capacity limit value of 30 MSUs. SCRT assigned a sub-capacity value of 28 MSUs for this hour because the actual sub-capacity value was 32 MSUs.</div> <div>2. The cap value for this LPAR is 44 MSUs because that is the lesser of the defined capacity value of 50 MSUs and the LPAR physical capacity limit value of 44 MSUs. SCRT assigned a sub-capacity value of 44 MSUs for this hour because the actual sub-capacity value was 45 MSUs.</div> <div>3. The total MSUs for the LPARs in this group was 72 MSUs (28 + 44), which is more than the group cap of 70 MSUs. For hour 80, SCRT assigned a value of 70 MSUs to products that ran in both LPAR2 and LPAR3, instead of the sum of the actual MSUs for those LPARs.</div> <div>4. The cap value for this LPAR is 25 MSUs because that is the lesser of the hard cap value of 25 MSUs and the LPAR physical capacity limit value of 40 MSUs. The even lesser defined capacity value of 20 MSUs is ignored because defined capacity is not supported on an LPAR that is hard capped. However, an LPAR physical capacity limit can coexist on an LPAR that is hard capped. Although this LPAR is a member of the same LPAR group as LPAR2 and LPAR3, group capping is not supported for this LPAR because it is hard capped.</div>				

Whenever z/OS V1 and z/OS V2 are running on the same CPC during the reporting period, SCRT adds a z/OS (A11) entry in order to report on both versions of z/OS as if they were the same product, so Table 65 on page 505 contains the z/OS entry to match the corresponding z/OS (A11) entry reported in the Product Max Contributors section. SCRT also reports on z/OS V1 and z/OS V2 as separate products.

Note that LPAR2 and LPAR3 for the z/OS (A11) entry are capped as part of the GRPA LPAR group, but none of the LPARs for either the z/OS V1 or z/OS V2 entries are capped as part of an LPAR group. Also note that there are no capping differences between middleware products in this example and the example in Table 61 on page 502 because the z/OS version does not affect the capping for middleware products. See “Sub-capacity reporting for multiple versions of z/OS” on page 127 for information about how z/OS V1 and z/OS V2 are reported in other sub-capacity report sections.

If hour 80 was the earliest peak hour for all of the products for the reporting period, Figure 279 on page 506 shows what the corresponding Product Max Contributors section of the sub-capacity report would look like.

```
==P5=====
PRODUCT MAX CONTRIBUTORS
```

Product Name	Product ID	Highest	Date/Time	LPAR LPAR1	LPAR LPAR2	LPAR LPAR3	GROUP LGRPA
z/OS	(A11)	98	05 Oct 2024 - 08:00	25	0	0	70
z/OS V2	5650-Z0S	44	05 Oct 2024 - 08:00	0	0	44	0
z/OS V1	5694-A01	53	05 Oct 2024 - 08:00	25	28	0	0
DB2 UDB for z/OS V8	5625-DB2	70	05 Oct 2024 - 08:00	0	0	0	70
CICS TS for z/OS V3	5655-M15	95	05 Oct 2024 - 08:00	25	0	0	70
MQSeries for z/OS V6	5655-L82	70	05 Oct 2024 - 08:00	0	0	0	70
IMS V8	5655-C56	95	05 Oct 2024 - 08:00	25	0	0	70
Lotus Domino for z/OS V7	5655-N14	44	05 Oct 2024 - 08:00	0	0	44	0
File Manager for z/OS V6	5655-P17	53	05 Oct 2024 - 08:00	25	28	0	0
IBM WebSphere Application Server for z/OS V6	5655-N01	70	05 Oct 2024 - 08:00	0	0	0	70
IBM WebSphere MQ Workflow for z/OS V3	5655-BPM	70	05 Oct 2024 - 08:00	0	0	0	70

Figure 279. Sample Product Max Contributors section for hour 80 in the sample z/OS V1 and z/OS V2 configuration

Static power save mode

The IBM zEnterprise 196 (z196) and later families of servers support a capability known as *static power save mode*. This mode is designed to reduce power consumption on supporting CPCs whenever full performance is not required. The mode can be switched on and off without disruption to currently running workloads, except for the change in performance.

Note: The IBM z13s, zEnterprise BC12 (zBC12), and zEnterprise 114 (z114) families of servers do *not* support static power save mode.

You can use static power save mode in the following situations:

- Periods of lower utilization, such as weekends or other non-prime shifts
- Systems that are used for backup capacity, such as for emergency backup. Such systems run with reduced energy consumption, but can quickly be brought back to full performance.
- CPCs that are under sub-capacity pricing whose machine capacities are much higher than needed for the current workloads and whose rated capacities are not needed until some time in the future

This documentation refers to static power save mode simply as *power save mode*.

Power save mode applies to all engine types—CPs and specialty engines—and is simultaneously effective for all engines and engine types within the same CPC. The reduction in power consumption is achieved by elongating the cycle time of all of the engines. The change in cycle time has a direct impact on that CPC's rated capacity (that is, an elongated cycle time means that it takes more CPU time to perform the same amount of work). This documentation uses the term *high performance mode* to refer to the mode when the CPC is not running in power save mode.

Because sub-capacity pricing is based on the CP utilization of the operating system image (typically the entire LPAR) and that utilization is dependent on the cycle time of the engines being used, the sub-capacity pricing MSU values, which are calculated by the operating system and by SCRT, in some cases, need to be properly adjusted for the difference in engine cycle time that occurs under power save mode. The adjustment is done by applying the appropriate factor when converting busy time to MSUs. For example, when the appropriate conversion factors are applied to a job which is run once in power save mode and again without power save mode, the calculated MSU value for each job should be the same, even though the job that ran in power save mode used more CPU time.

Each operating system has different dependencies on the collection, use, and reporting of engine cycle time; thus, each operating system supports power save mode in its own way, as described in the following topics.

If power save mode is not used during the reporting period, you need not perform any special actions and SCRT need not perform any special processing in order to correctly report on all sub-capacity products using the existing SCRT algorithms.

SCRT support for power save mode

SCRT provides special support for power save mode by:

1. Characterizing the level of power save mode support provided by each of the operating systems that are supplying SMF or SCRT89 records to SCRT
2. Adjusting the MSU values that are calculated for Getting Started Sub-Capacity Pricing (GSSP) products, when present

The Detail LPAR Data Section of the sub-capacity report has been enhanced to support power save mode by displaying additional columns to categorize the level of power save mode support provided by the SMF and SCRT89 records that were processed by SCRT. The additional columns are only displayed when SCRT determines that a CPC ran in power save mode for one or more hours during the reporting period or when SCRT could not determine whether a CPC ran in power save mode.

SCRT provides these enhancements to the Detail LPAR Data Section to help you recognize when you have not applied the required APARs or have not correctly configured your operating systems to fully support power save mode for sub-capacity pricing. These enhancements are not intended to count the number of hours that the CPC was in power saving mode because SCRT cannot reliably determine all of the hours during which the CPC was in power save mode during the reporting period, so such an hour count has not been included.

[Figure 280 on page 508](#) shows a sample Detail LPAR Data Section for a CPC that collected SMF and SCRT89 records while that CPC was running in power save mode. The configuration upon which the report is based consists of several native z/OS, z/TPF, and z/VSE systems and several guest systems. While such a configuration might be unusual for a single CPC, it is technically valid and serves to explain the differences in support for power save mode for each of the operating system conditions.

```

=====
DETAIL LPAR DATA SECTION

```

	Highest	Hour Count	Date/Time	2nd Highest	Hour Count	Date/Time	PSM Hours Full Support	PSM Hours Partial Support	Hours Unknown PSM Support	Hours No PSM Support
ZOS7A	134	6	03 Oct 2024 - 14:00	132	11	03 Oct 2024 - 13:00			697	
ZOS10A	151	8	04 Oct 2024 - 11:00	150	12	04 Oct 2024 - 10:00	20	2		
ZOS10B(S10B)	131	5	02 Oct 2024 - 06:00	128	8	02 Oct 2024 - 05:00	22			
ZOS11(zNALC)	148	9	17 Oct 2024 - 12:00	147	15	17 Oct 2024 - 01:00	22			
TPFP1	93	16	11 Oct 2024 - 01:00	89	22	11 Oct 2024 - 02:00	22			
TPFP2(TEST)	33	11	05 Oct 2024 - 09:00	31	14	05 Oct 2024 - 08:00	22			
TPFP3	56	2	09 Oct 2024 - 08:00	54	21	09 Oct 2024 - 07:00				454
TPFT4*	72	8	03 Oct 2024 - 13:00	70	9	03 Oct 2024 - 12:00	22			
ZVSE1	27	9	07 Oct 2024 - 11:00	26	12	07 Oct 2024 - 10:00	22			
ZVSE2(VSEPROD2)	14	5	02 Oct 2024 - 15:00	11	9	02 Oct 2024 - 14:00	22			
CPC	831	5	03 Oct 2024 - 14:00	827	11	03 Oct 2024 - 13:00				

Figure 280. Sample sub-capacity report: Detail LPAR Data Section for static power save mode

The existing first seven columns shown in the sample Detail LPAR Data Section in Figure 280 on page 508 are fully described in “Detail LPAR Data section” on page 119. Four new, additional columns appear whenever SCRT determines that a CPC ran in power save mode for one or more hours during the reporting period or whenever SCRT cannot determine whether a CPC ran in power saving mode during the reporting period. The four new columns are:

PSM Hours Full Support

The number of hours, including partial hours, during which SCRT detected that the LPAR or guest system was running in power save mode and that all required APARs and configuration changes have been applied. For more operating system-specific information, see the topic that applies to that operating system.

PSM Hours Partial Support

The number of hours, including partial hours, during which SCRT detected that the LPAR or guest system was running in power save mode and that some of the required APARs or configuration changes have not been applied. For more operating system-specific information, see the topic that applies to that operating system.

Hours Unknown PSM Use

The number of hours, including partial hours, during which SCRT was unable to determine whether the LPAR or guest system was running in power save mode because the required APARs or configuration changes have not been applied. For more operating system-specific information, see the topic that applies to that operating system.

Hours No PSM Support

The number of hours, including partial hours, during which SCRT was unable to support power save mode on the LPAR. For more operating system-specific information, see the topic that applies to that operating system.

Whenever any hours appear under the **PSM Hours Partial Support**, **Hours Unknown PSM Use**, or **Hours No PSM Support** columns, SCRT also issues the SCRTTOOL161 message to SYSPRINT to indicate that you might be able to take actions to obtain full SCRT support in power save mode.

Note: There are conditions, such as changing the CPC model either by making a permanent upgrade or by using On Off Capacity on Demand (OOCOD) or Capacity BackUp (CBU), when SCRT might count the hour when that change occurred as an hour during which power save mode was used, even though power save mode was not used. These conditions are explained under the specific support topics for z/OS, z/TPF, and z/VSE systems, which follow.

z/OS support for power save mode

The introduction of power save mode requires z/OS and SCRT to track if and when a CPC is running in power save mode and to take appropriate actions to calculate the correct sub-capacity values based on CPC mode. If power save mode is not used during the reporting period, you need not perform any special actions and SCRT need not perform any special processing in order to correctly report on all sub-capacity products based on its existing algorithms.

Except for GSSP products, the sub-capacity values for all other z/OS sub-capacity products will always be correctly reported based on the operating mode (engine cycle time) when that z/OS system image was running, regardless of whether the required power saving mode APARs have been applied and the configuration for power save mode has been properly set in the SMFPRMxx member of SYS1.PARMLIB. Although the z/OS sub-capacity values will be correct for all products (except Getting Started Sub-Capacity Pricing products), SCRT will be unable to determine and report the number of hours during which the z/OS system was running in power save mode unless the power save mode APARs have been applied. The sub-capacity values will be correct because:

- The MSU conversion factor that applies to power save mode will be used to convert the busy times to MSUs for those times that were collected while the CPC was running in power save mode
- The MSU conversion factor that applies to the rated capacity of the CPC will be used to convert the busy times to MSUs for those times that were collected while the CPC was not running in power save mode

The MSU conversion uses the appropriate conversion factor because the WLM component of z/OS continually monitors the mode in which the CPC is running, performs the 4-hour rolling average MSU calculations, and passes those MSU values by way of SMF type 70 records on to SCRT for sub-capacity reporting.

Unlike other sub-capacity products, the accuracy of the sub-capacity values reported for GSSP products that are running while the CPC is in power save mode are highly dependent on whether the required power save mode APARs have been applied and whether you have chosen appropriate settings in the SMFPRMxx member. As long as all of the required actions have been performed, SCRT will accurately report the sub-capacity values for GSSP products. Otherwise, whenever the product's peak hour occurs while the product was running with the CPC in power save mode, the sub-capacity value reported for one or more GSSP products might be higher than it would have been had the required actions been taken.

The following APARs are required for running z/OS on a CPC in power save mode and must be applied before SCRT can provide full support for power save mode:

- SMF APAR OA30563 (for z/OS V1.7 through V1.12)
- WLM APAR OA30968 (for z/OS V1.10 through V1.12)
- RMF APAR OA29530 (for z/OS V1.10 through V1.12)

Note that z/VM APARs VM64798 and VM64879 for z/VM 5.4 and 6.1 are required to run z/VM on zEC12, zBC12, z196, or z114 CPCs. (The z13s, zBC12 and z114 processor families do not support power save mode.)

SMF APAR OA30563 supports the new **MAXEVENTINTRECS** parameter in the SMFPRMxx member. This parameter is used to indicate the maximum number of event-driven SMF type 89 interval records that are allowed to be created during the regular SMF89 interval. Whenever the **MAXEVENTINTRECS** value is set high enough so that no SMF89 measurement interval contains data collected for both power save mode and high performance mode, then SCRT accurately calculates the sub-capacity values for the GSSP products reported on those records because SCRT applies the appropriate MSU conversion factors to the busy times that were collected during the respective modes. See *z/OS MVS Initialization and Tuning Reference* for more information.

Whenever the **MAXEVENTINTRECS** value is set too low so that an SMF89 measurement interval contains data collected for both power saving mode and high performance mode, SCRT applies the MSU conversion factor for the rated capacity of the CPC to all of the busy time reported for that SMF interval, including the busy time collected in power save mode. The result is that the MSU values calculated for that interval might be higher than they would have been if the **MAXEVENTINTRECS** value had not been set too low.

The default value for the **MAXEVENTINTRECS** parameter is 0 (zero), which means that no event-driven records will be written at the instant when power save mode is entered or exited.

Table 66 on page 510 provides explanations for the sub-capacity reporting implications for the power save mode hours that are reported in the Detail LPAR Data Section of the sub-capacity report, along with a list of corrective actions. See “SCRT support for power save mode” on page 507 for descriptions of each of the types of hours listed in Table 66 on page 510.

Static power save mode

Table 66. Explanations for power save mode values listed in Detail LPAR Data Section for z/OS systems

Hours reported in this column of the Detail LPAR Data Section	z/OS release	Sub-capacity reporting impact and suggested actions
PSM Hours Full Support	V1.10 and higher	All of the MSU values for the sub-capacity products, including GSSP products, have been appropriately adjusted for power save mode. No customer action is necessary.
PSM Hours Partial Support	V1.10 and higher	The MSU values for all sub-capacity products, except GSSP products, have been appropriately adjusted for power save mode. The sub-capacity values that were calculated for any GSSP products that were running during these hours might be higher than they otherwise would have been because SCRT applied the conversion factor for high performance mode to all of the busy time (power save mode and high performance mode) reported for that SMF interval. Note that the sub-capacity values for GSSP products were correctly calculated with the appropriate conversion factor for all other power save mode intervals and for all other high performance mode intervals. This condition can be corrected for SMF type 89 records that are collected during power save mode in the future by setting the SMF MAXEVENTINTRECS value high enough to allow SMF to report the busy times for both modes in different SMF records.
Hours Unknown PSM Use	V1.10 and higher	The MSU values for all sub-capacity products, except GSSP products, have been appropriately adjusted for power save mode, when necessary. The sub-capacity values that were calculated for any GSSP products that were running during these hours might be higher than they otherwise would have been because SCRT applied the conversion factor for high performance mode to all of the busy time (power save mode and high performance mode) reported for that SMF interval because the hardware was unable to report whether the CPC was running in power save mode or high performance mode. You can correct this condition by applying all of the z/OS related APARs that were listed earlier in this section. However, if this condition continues to occur for the SCRT reporting period following the month in which the APARs were applied and the z/OS system was IPLed, then check for any new microcode drivers that might be needed to address this condition.
	V1.7 through V1.9	The MSU values for all sub-capacity products, except GSSP products, have been appropriately adjusted for power save mode, when necessary. SCRT cannot determine whether power save mode was actually used for the hours reported under this column. Therefore, every hour during the reporting period when the associated LPAR or guest system ran will be counted and reported under this column, regardless of whether or not power save mode was actually used and regardless of whether or not APAR OA30563 was applied. The impact to the sub-capacity values for GSSP products depends on one of the following specific conditions: <ul style="list-style-type: none"> If the CPC was never in power save mode during the reporting period, then all of the sub-capacity values assigned to GSSP products are correct because none of those hours required any adjustment for power save mode. If all of the following conditions are true, then the sub-capacity values for GSSP products running during any SMF interval during which the CPC entered or exited power save mode might be reported as being higher than they otherwise would have been: <ul style="list-style-type: none"> The CPC was in power save mode during the reporting period. SMF APAR OA30563 was applied. The SMF MAXEVENTINTRECS value was not high enough to allow SMF to report the busy times for power save mode and high performance mode on separate SMF records. If both of the following conditions are true, then none of the sub-capacity values for GSSP products running during power save mode have been adjusted for power save mode because the SMF records do not contain the necessary data with which to make such adjustments: <ul style="list-style-type: none"> The CPC was in power save mode during the reporting period. SMF APAR OA30563 was not applied. <p>You can correct this condition by applying SMF APAR OA30563 and setting the SMF MAXEVENTINTRECS value high enough to enable SMF to report the busy times for power save mode and high performance mode on separate SMF records.</p> <p>The best way to address this condition is to migrate to z/OS V1.10 or higher and to apply the required APARs to support power save mode, as described earlier.</p>

Figure 280 on page 508 shows a sample Detail LPAR Data Section for a CPC that ran in static power save mode for 22 hours during the reporting period. Each of the z/OS systems reported their own view of power save mode based on their release level and the effect of the SMF **MAXEVENTINTRECS** parameter. The reporting behavior of each of the following z/OS systems, as shown in the figure, is based on the SCRT processing rules described earlier in this topic.

ZOS7A represents a z/OS V1.7 system. z/OS V1.7 systems are unable to determine whether they are running in power save mode. Thus, SCRT reports all of the hours collected for any z/OS V1.7 system during the reporting period (697 hours) under the **Hours Unknown PSM Use** column. SCRT reports

the hours that way regardless of whether SMF APAR OA30563 was applied to that system. Note that, although SCRT does not know whether power save mode was being used, it would have always correctly calculated the sub-capacity values for the GSSP products, even in power save mode, provided that the SMF **MAXEVENTINTRECS** parameter was set high enough. However, that is not the case if APAR OA30563 was not applied. The same result would have occurred for z/OS V1.8 and V1.9 systems. The sub-capacity values reported for non-GSSP products are always correctly calculated based on the mode that was being used because those calculations are performed by the WLM component of z/OS.

ZOS10A represents a z/OS V1.10 system where the SMF **MAXEVENTINTRECS** parameter was not set high enough. The result is that the busy times reported on two SMF records were the sum of the busy time for power save mode and for high performance mode. Because these times were not reported separately, SCRT applied the conversion factor for high performance mode to that time and counted that hour under the **PSM Hours Partial Support** column, as the sub-capacity values for GSSP products might be higher than if the busy times for power save mode and high performance mode had been reported on separate SMF records.

ZOS10B represents a z/OS V1.10 guest system whose SMF **MAXEVENTINTRECS** parameter was set high enough to always report the busy times for power save mode and high performance mode on separate SMF records. This separation of busy times enabled SCRT to accurately calculate the sub-capacity values for GSSP products for ZOS10B for every hour during the reporting period. The hours during which SCRT provided that level of support are reported under the **PSM Hours Full Support** column.

ZOS11A represents a z/OS V1.11 system running a zNALC workload with the SMF **MAXEVENTINTRECS** parameter set high enough to always report the busy times for power save mode and high performance mode on separate SMF records, so SCRT was able to handle power save mode in the same manner as for ZOS10B.

z/TPF support for power save mode

The introduction of power save mode requires z/TPF and SCRT to track if and when a CPC is running in power save mode and to take appropriate actions to calculate the correct sub-capacity values based on CPC mode.

The support for power save mode for z/TPF requires z/TPF APAR PJ38004 (z/TPF PUT 07) and z/TPF APAR PJ38530. This support will continue to create SCRT89 records every 60 minutes and those records will contain sub-capacity values which have been calculated by applying the appropriate conversion factor—for any busy time collected in power save mode and for any busy time collected in high performance mode—during that 60-minute interval.

SCRT supports power save mode for z/TPF LPARs that do not provide SCRT89 records only when that CPC also supports a native z/OS system (with the required power save mode APARs applied) that collects the required SMF type 70 and SMF type 89 records, and which runs at the same time as those z/TPF LPARs. SCRT does not support power save mode for z/TPF LPARs that do not provide SCRT89 records and instead rely entirely on the SCRT89 records collected by other native z/TPF systems to supply the required power save mode information (and no native z/OS system is running or meets the conditions listed earlier). For more information about the SCRT support for this type of z/TPF LPAR, see [Chapter 13, “Sub-capacity reporting for z/TPF systems,”](#) on page 281.

Note that z/VM APARs VM64798 and VM64879 for z/VM 5.4 and 6.1 are required to run z/VM on zEC12, zBC12, z196, and z114 CPCs. (The zBC12 and z114 processor families do not support power save mode.)

[Table 67 on page 512](#) provides explanations for the sub-capacity reporting implications for the power save mode hours that are reported in the Detail LPAR Data Section of the sub-capacity report, along with a list of corrective actions. See [“SCRT support for power save mode” on page 507](#) for descriptions of each of the types of hours listed in [Table 67 on page 512](#).

Table 67. Explanations for power save mode values listed in Detail LPAR Data Section for z/TPF systems

Hours reported in this column of the Detail LPAR Data Section	z/TPF level	Sub-capacity reporting impact and suggested actions
PSM Hours Full Support	With APAR PJ38004	<p>All of the MSU values for all sub-capacity products have been appropriately adjusted for power save mode.</p> <p>No customer action is necessary.</p> <p>Note that changing the number of CPs on a CPC, either due to a permanent model upgrade or due to a temporary change via On Off Capacity on Demand (OOCOD), might cause that hour to be counted and reported under this column.</p>
	Without APAR PJ38004	<p>The MSU values for all sub-capacity products have been calculated as if the CPC was running in high performance mode because the hardware was unable to report whether the CPC was running in power save mode or high performance mode.</p> <p>If this condition persists, check for any new microcode drivers that might be needed to address this condition.</p>
Hours No PSM Support	Any	<p>None of the MSU values for the sub-capacity products have been adjusted for power save mode because the hours being reported are for a z/TPF LPAR that has not provided SCRT89 records and whose reporting is dependent on either the SMF data collected by a z/OS system without the required power save APARs running on the same CPC, or on the SCRT89 data collected by another z/TPF system running on the same CPC. SCRT does not support power save mode for such non-reporting z/TPF LPARs.</p> <p>This condition will be corrected in a future level of SCRT.</p>

Figure 280 on page 508 shows a sample Detail LPAR Data Section for a CPC that ran in static power save mode for 22 hours during the reporting period. Each of the z/TPF systems reported their own view of power save mode based on whether the required z/TPF APAR had been applied and whether the z/TPF system supplied SCRT89 records. The reporting behavior of each of the following z/TPF systems, as shown in the figure, is based on the SCRT processing rules described earlier in this topic.

TPFP1 and TPFP2 represent a native z/TPF system and a z/TPF guest system, respectively, with the required power save mode APAR applied. SCRT accurately reported the sub-capacity values for both of these systems for the entire reporting period, including any periods of time that were run in power save mode. SCRT reported all the hours or portions of an hour during which the systems ran in power save mode under the **PSM Hours Full Support** column.

TPFP3 represents an LPAR where a z/TPF system ran for 454 hours during the reporting period without the required power save mode APAR applied, and for the remaining 231 hours of the reporting period with the APAR applied. The hours without the APAR applied are reported under the **Hours No PSM Support** column because neither z/TPF nor SCRT could determine whether power save mode was in use. The 454

hours without the APAR applied happen to include the 22 hours during which power save mode was in use but was not detected. z/TPF did not use the conversion factor for power save mode when calculating the sub-capacity values for those 22 hours because the level of z/TPF that was running at the time did not contain the logic to support power save mode.

TPFT4 represents an LPAR that has not supplied any SCRT89 records, possibly because the LPAR contains a large number of z/TPF development guest systems or test guest systems. SCRT used the busy times that were collected for that LPAR by one of the z/OS systems with the required power save mode APARs and which was running at the same time on that CPC to calculate and report the sub-capacity values for that LPAR. Since one or more of the z/OS systems running on this CPC met all of the required conditions, SCRT was able to support power save mode for the z/TPF guest systems running in LPAR TPFT4 and post a count of 22 hours under the **PSM Hours Full Support** column.

z/VSE support for power save mode

The introduction of power save mode requires z/VSE and SCRT to track if and when a CPC is running in power save mode and to take appropriate actions to calculate the correct sub-capacity values based on CPC mode.

The following APARs are required for z/VSE to collect SCRT89 records on zEC12, ZBC12, z196, and z114 CPCs. (The zBC12 and z114 processor families do not support power save mode.) Without the required APAR, the CMT will be unable to collect SCRT89 records and SCRT will be unable to generate any sub-capacity reports for the z/VSE systems that ran on those CPCs.

- z/VSE V4.1: APAR DY47110
- z/VSE V4.2: APAR DY47111
- z/VSE V4.3: No APARs required; support for power save mode is already included

Note that z/VM APARs VM64798 and VM64879 for z/VM V5.4 and V6.1 are required to run z/VM on zEC12, zBC12, z196, and z114 CPCs. (The zBC12 and z114 processor families do not support power save mode.)

The z/VSE support for power save mode will continue to create SCRT89 records every 30 minutes and those records will contain sub-capacity values which have been calculated by applying the appropriate conversion factor—for any busy time collected in power save mode and for any busy time collected in high performance mode—during that 30-minute interval. Note that the support for power save mode in z/VSE V4.1 and V4.2 cannot detect power save mode under the unusual case where the CPC enters and exits power save mode totally within the 30 minutes reported on an SCRT89 record. When that happens, z/VSE treats all of the busy time collected during that 30-minute interval as if the CPC was in the mode indicated by the CPC at the end of the measurement interval, and such an hour is not counted in [Table 68 on page 514](#). The same limitation does not apply to z/VSE 4.3, where z/VSE is able to apply the appropriate conversion factors to the appropriate modes, even when the CPC enters and exits power save mode within the same 30-minute interval reported on an SCRT89 record.

[Table 68 on page 514](#) provides explanations for the sub-capacity reporting implications for the power save mode hours that are reported in the Detail LPAR Data Section of the sub-capacity report, along with a list of corrective actions. See [“SCRT support for power save mode” on page 507](#) for descriptions of each of the types of hours listed in [Table 68 on page 514](#).

Table 68. Explanations for power save mode values listed in Detail LPAR Data Section for z/VSE systems

Hours reported in this column of the Detail LPAR Data Section	z/VSE level	Sub-capacity reporting impact and suggested actions
PSM Hours Full Support	z/VSE V4.1 with APAR DY47110	All of the MSU values for all sub-capacity products have been appropriately adjusted for power save mode. No customer action is necessary.
	z/VSE V4.2 with APAR DY47111	Note that changing the number of CPs on a zEC12 or z196 CPC, either due to a permanent model upgrade or due to a temporary change via On Off Capacity on Demand (OOCOD), might cause that hour to be counted and reported under this column.
	z/VSE V4.3	
Hours Unknown PSM Use	z/VSE V4.1 with APAR DY47110	The MSU values for all sub-capacity products have been calculated as if the CPC was running in high performance mode because the hardware was unable to report whether the CPC was running in power save mode or high performance mode.
	z/VSE V4.2 with APAR DY47111	If this condition persists, check for any new microcode drivers that might be needed to address this condition.
	z/VSE V4.3	

Figure 280 on page 508 shows a sample Detail LPAR Data Section for a CPC that ran in static power save mode for 22 hours during the reporting period. Each of the z/VSE systems reported their own view of power save mode based on the support provided by the required z/VSE APAR. Note that the CMT will not collect SCRT89 records for the CPC unless the required APAR is applied. The reporting behavior of each of the following z/VSE systems, as shown in the figure, is based on the SCRT processing rules described earlier in this topic.

ZVSE1 and ZVSE2 represent a native z/VSE system and a z/VSE guest system, respectively. Both systems had the required APAR for power save mode applied. (Otherwise, the CMT would have issued an error message and would have been unable to collect SCRT89 records, which are required for sub-capacity reporting.) SCRT reported all of the hours or portions of an hour during which the systems ran in power save mode under the **PSM Hours Full Support** column.

Sending TSAD data

Sending weekly Transmit System Availability Data (TSAD) is part of the terms and conditions for sub-capacity pricing. The data is required for audit purposes. You can send TSAD by one of the following methods:

- Using IBM Z Remote Support Facility (RSF)
- Mailing a diskette, DVD, or USB flash memory drive to IBM
- Emailing the data to IBM

Typically, an IBM service representative enables RSF for you when your CPC is installed; most customers will find that their CPC is already set up for RSF. If RSF is not enabled, contact your IBM service representative.

If you are not going to be using RSF, see the procedure described in the Transmit Service Data task for the Support Element (SE) in *Hardware Management Console Operations Guide*. When you perform this task, select the **System Availability Data** button. This will create several files of data for this CPC on your chosen target device.

Submit the data to IBM in one of the following ways:

- If you intend to mail a diskette, DVD, or USB flash memory drive to IBM, copy all the files created by the SE to your chosen media. Label the media with the date, CPC type, model, and serial number, and mail it to the following address:

IBM Z RAS Engineering
B64A/707-2B86 (MS P317)
2455 South Road
Poughkeepsie, NY 12601

- If you intend to email the data to IBM, create a .zip file containing all the files created by the SE for the CPC. Name the .zip file SADsssss.zip, where sssss is the five-digit serial number of the CPC. Send an email, including the date, CPC type, and model with the .zip file as an attachment to: tsadmail@us.ibm.com

Appendix C. SCRT country code list

The SCRT program and the LMS authorization request form each require a country code to be used in conjunction with your customer number.

You can find the 2-character or 3-digit country code for your country or region in the second column of [Table 69 on page 517](#).

Table 69. SCRT country code list

Country or region	Code
Albania	AL
Argentina	613
Armenia	AM
Australia	616
Austria	AT
Azerbaijan	AZ
Bahamas	619
Bahrain	BH
Barbados	621
Belarus	BY
Belgium	BE
Bermuda	627
Bolivia	629
Bosnia-Herzegovina	BA
Brazil	631
Bulgaria	BG
Canada	CA
Cayman Islands	647
Chile	655
Colombia	661
Costa Rica	663
Croatia	HR
Cyprus	CY
Czech Republic	CZ
Denmark	DK
Dominican Republic	681
Ecuador	683
Egypt	EG

Country codes

Table 69. SCRT country code list (continued)

Country or region	Code
El Salvador	829
Finland	FI
France	706
Georgia	GE
Germany	DE
Greece	GR
Guatemala	731
Guyana	640
Honduras	735
Hong Kong	738
Hungary	HU
Iceland	742
India	744
Indonesia	749
Ireland	IE
Israel	IL
Italy	758
Jamaica	759
Japan	760
Jordan	JO
Kazakhstan	KZ
Korea	766
Kuwait	KW
Kyrgistan	KG
Lesotho	711
Lithuania	638
Luxembourg	BE
Macau	736
Macedonia	MK
Malaysia	778
Malta	MT
Mexico	781
Moldavia	MD
Morocco	MA

Table 69. SCRT country code list (continued)

Country or region	Code
Mozambique	782
Namibia	682
Neth Antilles	791
Netherlands	NL
New Zealand	796
Nicaragua	799
Norway	NO
Oman	OM
Pakistan	PK
Panama	811
Paraguay	813
Peru	815
Philippines	818
Poland	PL
Portugal	PT
Qatar	QA
Romania	RO
Russian Federation	RU
Saudi Arabia	SA
Serbia	RS
Singapore	834
Slovakia	SK
Slovenia	SI
South Africa	ZA
Spain	838
Suriname	843
Sweden	SE
Switzerland	CH
Taiwan	858
Tajikistan	360
Thailand	856
Trinidad/Tobago	859
Turkiye	TR
Turkmenistan	TM

Country codes

Table 69. SCRT country code list (continued)

Country or region	Code
Ukraine	UA
United Arab Emirates	AE
United Kingdom	GB
United States	US
Uruguay	869
Uzbekistan	UZ
Venezuela	871
Yemen	YE

Appendix D. Changes and fixes in prior SCRT versions

This appendix summarizes the changes and fixes that have been made in prior versions of SCRT.

Summary of changes in prior SCRT versions

This section summarizes the changes made to this document in prior versions of SCRT.

For a summary of the changes made in the current version of SCRT, see [“Summary of changes” on page xix](#).

Summary of changes for SC23-6845-21, SCRT 29.1.0

This document contains information previously presented in SC23-6845-20 and includes terminology, maintenance, and editorial changes.

New and changed information

Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Notice: SCRT 29.1.0 is the last level for which PTFs will be provided for z/OS 2.3. End of service for z/OS 2.3 occurred on 30 September 2022. Customers on z/OS 2.3 may continue to install the downloadable version of SCRT for z/OS or use SCRT for Windows and Linux.

Documentation changes: The following list highlights the technical changes made to this documentation in this release of SCRT. For the complete list of functional and technical changes in this release of SCRT, including changes in any interim modification levels since the last release, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](#).

- As previously announced, SCRT now requires a Java 8 environment.
- The SCRT installer for Windows environments now supports Windows 11.
- Support has been added for new z/OS sub-capacity products, as described in "Product support changes."
- The sample SCRT JCL has been updated in [“Sample JCL to run the 64-bit SCRT on a z/OS system” on page 25](#).
- A **Comment** parameter is added to the **Exclude**, **EXCLUDE MSU_CONSUMPTION**, and **Include** control statements in [“SPECIAL DD statement” on page 40](#), and to the panel in [“Exclude panel” on page 85](#).
- The **--report-time** command line option is added in [“Specifying command line options” on page 91](#).
- Additional guidance is provided to ensure an accurate calculation of the 4-hour rolling average utilization for the first hour of the report period when providing INPUTCSV data for Mobile Workload Pricing and IBM Z Collocated Application Pricing. See [“Mobile workload reporting using an INPUTCSV file” on page 200](#) and [“INPUTCSV DD statement for zCAP” on page 234](#).
- Information about the LMS email submission process has been removed. All users must use the LMS web interface to submit reports to IBM, as described in [Chapter 26, “Submitting sub-capacity reports to IBM,” on page 427](#).
- The figures depicting the LMS web interface have been updated to match the current tool in [“LMS web interface” on page 427](#).

SCRT message changes: The following new messages have been added in [Appendix A, “SCRT messages,” on page 449](#):

SCRTTOOL264–SCRTTOOL265

SCRTTOOL389–SCRTTOOL390

Product support changes: SCRT 29.1.0 adds sub-capacity support for the products listed in Table 70 on page 522. The NO89 products are marked with an asterisk (*), and updates to the NO89 DD of the SCRT JCL for the z/OS deliverable are required to self-report usage or non-usage of these products.

Table 70. Additional and revised product support in SCRT 29.1.0

Operating system	Product ID	Product name
z/OS	5655-EP6	IBM Enterprise PL/I VUE for z/OS V6
	5655-G75 ¹	File Manager for z/OS and OS/390 V2* ¹
	5655-PL6	IBM Enterprise PL/I for z/OS V6
	5655-NOE	IBM Open Enterprise SDK for Node.js 16
	5655-YA1	IBM CICS TS for z/OS 6
	5696-VTP	IBM Z Virtual Test Platform V2
	5698-ANA	IBM Z Anomaly Analytics with Watson 5
	5698-DB2	Db2 13 for z/OS
	5698-DBV	Db2 13 for z/OS Value Unit Edition
	5698-LDA	IBM Z Operational Log and Data Analytics 5
	5698-T09	IBM Z Workload Scheduler 10*
	5698-ZWG	IBM Z Distribution for Zowe 2
	5698-020	IBM Z JCL Expert*
	5698-5A1	IBM Observability by Instana® Application Performance Monitoring on z/OS
	5722-DV1	IBM CICS TS VUE 6
	5755-A01	IBM Application Delivery Foundation for z/OS 4
	5755-A02	IBM Fault Analyzer for z/OS 15
	5755-A03	IBM File Manager for z/OS 15
	5755-A04	IBM Application Performance Analyzer for z/OS 15
	5755-A05	IBM Developer for z/OS Enterprise Edition 16
	5755-A06	IBM Debug for z/OS 16

Table notes:

1. 5655-G75 is not a new product; it is an older product whose name has been corrected from "File Analyzer for z/OS and OS/390 V2" to "File Manager for z/OS and OS/390 V2". The name has been updated in the internal SCRT product catalog and in the NO89 DD in the sample JCL.

SCRT 29.1.1 adds limited sub-capacity support for IBM z/VSE products listed in [Figure 166 on page 321](#) that run on 21st Century Software VSEⁿ 6.3. Note that SCRT does not report on the 21st Century Software VSEⁿ operating system nor on any 21st Century Software VSEⁿ products. Neither does SCRT reflect any 21st Century Software VSEⁿ operating system nor any 21st Century Software VSEⁿ products in the optional (A11) entries used for Multi-Version Measurement (MVM) reporting. For more information, see [Reporting of IBM z/VSE products on 21st Century Software VSEⁿ](#).

Software fixes in SCRT: For information about APARs and other software fixes in this release of SCRT, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](#).

Summary of changes for SC23-6845-19, SCRT Version 28.2.0

This document contains information previously presented in SC23-6845-18 and includes terminology, maintenance, and editorial changes.

New and changed information

Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Documentation changes: The following list highlights the technical changes made to this documentation in this release of SCRT. For the complete list of functional and technical changes in this release of SCRT, including changes in any interim modification levels since the last release, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](https://ibm.biz/SCRTCurrentLevels).

- India has been added to the list in [Appendix C, “SCRT country code list,”](#) on page 517. (Updated 14 Oct 2021)
- Notice is given that SCRT Version 28 is the last version of SCRT that will run in a Java 7 environment. See [“Requirements for SCRT”](#) on page 21 and [“Requirements for SCRT for Windows and Linux”](#) on page 69.
- The sample SCRT JCL has been updated in [“Sample JCL to run the 64-bit SCRT on a z/OS system”](#) on page 25.
- SCRT supports a new non-MSU metric type, Number of CPCs, as described in [“Non-MSU metric types”](#) on page 371 and [“SCRT externals for reporting of non-MSU metrics”](#) on page 371.
- SCRT allows you to specify a Tailored Fit Pricing for IBM Z solution ID at the system level, as described in [“Requirements for Tailored Fit Pricing for IBM Z”](#) on page 375 and [“Interactions with traditional sub-capacity offerings”](#) on page 377.
- The name "Enterprise Consumption Solution" has been updated to "Software Consumption Solution" in [Chapter 20, “Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution,”](#) on page 401.
- The name "Interval rate data" has been updated to "Tailored Fit Pricing for IBM Z - Hardware Consumption Solution" in [Chapter 21, “Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution,”](#) on page 405.
- A new topic has been added, [“Considerations for character set encoding”](#) on page 414.
- Support has been added for new z/OS sub-capacity products, as described in "Product support changes."

Product support changes: SCRT V28.2.0 adds sub-capacity support for the products listed in [Table 71](#) on page 523. The NO89 products are marked with an asterisk (*), and updates to the NO89 DD of the SCRT JCL for the z/OS deliverable are required to self-report usage or non-usage of these products.

Table 71. Additional and revised product support in SCRT V28.2.0

Operating system	Product name	Product ID
z/OS	IBM Debug for z/OS V14	5755-Q50
	IBM Developer for z/OS Enterprise Edition V15	5755-AC5
	IBM Migration Utility for z/OS V5	5698-MG5*
	IBM Tivoli Workload Scheduler distributed Agent for z/OS V8	5698-AA3*
	IBM Workload Scheduler distributed Agent for z/OS V9	5698-AAR*
	IBM Z Distribution for Zowe V1	5698-ZWE

Software fixes in SCRT: For information about APARs and other software fixes in this release of SCRT, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](https://ibm.biz/SCRTCurrentLevels).

Summary of changes for SC23-6845-18, SCRT Version 28.1.0

This document contains information previously presented in SC23-6845-17 and includes terminology, maintenance, and editorial changes.

New and changed information

Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Documentation changes: The following list highlights the technical changes made to this documentation in this release of SCRT. For the complete list of functional and technical changes in this release of SCRT, including changes in any interim modification levels since the last release, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](http://ibm.biz/SCRTCurrentLevels).

- Notice is given that SCRT Version 28 is the last version of SCRT that will run in a Java 7 environment. See [“Requirements for SCRT” on page 21](#) and [“Requirements for SCRT for Windows and Linux” on page 69](#).
- Two new System Recovery Boost report sections have been added in [“System Recovery Boost Usage section” on page 117](#) and [“System Recovery Boost Usage Details section” on page 126](#).
- Information about the format of a solution ID for Tailored Fit Pricing for IBM Z has been added in [“Format of a Tailored Fit Pricing solution ID” on page 376](#).
- Three new sample report sections have been added in [“Product Features section” on page 352](#), [“System Recovery Boost Activity section” on page 353](#), and [“Detail LPAR Usage Data section for CPC M1C1” on page 359](#). These sections also appear for CPC M2C1.
- Information about the Machine Peak Utilization and Machine MSU Consumed fields has been added in [“Customer information” on page 105](#) and [“Machine summary section” on page 348](#).
- Minor updates have been made in [Chapter 23, “Reporting for independent software vendor programs,” on page 413](#).
- Updates and clarifications have been made in [“Spreadsheet method for viewing and editing the sub-capacity report” on page 420](#).
- The sample SCRT JCL has been updated in [“Sample JCL to run the 64-bit SCRT on a z/OS system” on page 25](#).
- Support has been added for new z/OS sub-capacity products, as described in [“Product support changes.”](#)

SCRT message changes: The following new messages have been added in [Appendix A, “SCRT messages,” on page 449](#):

SCRTTOOL038

Product support changes: SCRT V28.1.0 adds sub-capacity support for the products listed in [Table 72 on page 524](#).

Table 72. Additional and revised product support in SCRT V28.1.0

Operating system	Product name	Product ID
z/OS	IBM Connect:Direct for z/OS V6	5655-X12
	IBM Open Enterprise Python for z/OS	5655-PYT
	IBM SDK for Node.js - z/OS 14	5655-NOD
	IBM Wazi Virtual Test Platform	5697-VTP
	IBM Z Common Data Provider V2	5698-ABN
	IBM Z Open Automation Utilities	5698-PA1
	IBM Z Operations Analytics V4	5698-AAV

SCRT V28.1.0 also adds support for the following z/OS priced features:

- IBM z/OS Authorized Code Scanner
- IBM z/OS Workload Interaction Correlator

Software fixes in SCRT: For information about APARs and other software fixes in this release of SCRT, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](#).

Summary of changes for SC23-6845-17, SCRT Version 27.2.0

This document contains information previously presented in SC23-6845-16 and includes terminology, maintenance, and editorial changes.

New and changed information

Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Documentation changes: The following list highlights the technical changes made to this documentation in this release of SCRT. For the complete list of functional and technical changes in this release of SCRT, including changes in any interim modification levels since the last release, see [SCRT Current Levels](https://ibm.biz/SCRTCurrentLevels) (ibm.biz/SCRTCurrentLevels).

- Information about the format of SCRT customer numbers has been updated.
- Support has been added for reference product lines that indicate simultaneous peak 4-hour rolling average utilization across related product families, such as between MLC and VUE product variants. These new lines appear in the IPLA product groups in sections E5 and B5 and in supporting sections. For an example, see [“Product Summary Information section” on page 110](#).
- In Tailored Fit Pricing reports, the Container Active Products section (E7) no longer reports product family roots. (Product family roots are reported in the Product Summary Information section (E5).)
- The Product Features section (F5) has been added, as described in [“Product Features section” on page 115](#).
- Report sections H4 and H6 now have a line that reports the total number of missing hours, as described in [“Missing LPAR Data Details section” on page 112](#) and [“Missing CPC Data Details section” on page 114](#).
- Report section H7 no longer requires individual justification for missing minutes for each LPAR. Instead, a single justification is required for each H7 section that indicates missing time. Optionally, you can also provide justification per LPAR.
- Guidance is provided about SCRT NO89 products not being listed in section Q7 in [“Customer verification section” on page 388](#).
- Information about optional interval rate data reporting has been added in Chapter 21, [“Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution,” on page 405](#). This report section is disabled by default and should only be enabled at the direction of IBM.
- Support has been added for IBM z15 CPCs, machine type 8561.
- Updates have been made in [“Obtaining technical support from IBM” on page 439](#).
- The sample SCRT JCL has been updated in [“Sample JCL to run the 64-bit SCRT on a z/OS system” on page 25](#).
- Appendix C, [“SCRT country code list,” on page 517](#) has been added.
- Support has been added for new z/OS sub-capacity products, as described in [“Product support changes.”](#)

SCRT message changes:

- The following new messages have been added in [Appendix A, “SCRT messages,” on page 449](#):

SCRTTOOL037
SCRTTOOL056
SCRTTOOL069
SCRTTOOL070

- The following message are no longer issued and have been removed:

SCRTTOOL037
SCRTTOOL038

Product support changes: SCRT V27.2.0 adds sub-capacity support for the products listed in [Table 73](#) on page 526.

Table 73. Additional and revised product support in SCRT V27.2.0

Operating system	Product name	Product ID
z/OS	IBM Operational Decision Manager for z/OS V8 (Existing product, renamed from WebSphere Business Rules for z/OS V8.)	5655-Y31
	IBM SDK for Node.js on z/OS, V8	5655-DKN
	IBM SDK for Node.js on z/OS, V12	5655-NJS
	IBM Z Service Automation Suite V1	5698-SA1
	IBM Z Table Accelerator	5698-TA1

Software fixes in SCRT: For information about APARs and other software fixes in this release of SCRT, see [SCRT Current Levels](#) ([ibm.biz/SCRTCurrentLevels](#)).

Summary of changes for SC23-6845-16, SCRT Version 27.1.0

This document contains information previously presented in SC23-6845-15 and includes terminology, maintenance, and editorial changes.

New and changed information

Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Documentation changes: The following list highlights the technical changes made to this documentation in this release of SCRT. For the complete list of functional and technical changes in this release of SCRT, including changes in any interim modification levels since the last release, see [SCRT Current Levels](#) ([ibm.biz/SCRTCurrentLevels](#)).

- The term "Container Pricing for IBM Z" has been changed to "Tailored Fit Pricing for IBM Z."
- "Advanced image naming" on page 104 has been added.
- The CONTAINER control statement now supports the *ALL value for the **CPC** and **IMAGE_ID** parameters, as described in "CONTAINER control statement" on page 41.
- Information about the **INCLUDE** control statement and command line option has been added in "INCLUDE control statement" on page 54 and "Specifying INCLUDE statements" on page 94.
- Information about the **EXCLUDE MSU_CONSUMPTION** control statement has been added in "Unusual situations that affect utilization" on page 17 and "EXCLUDE MSU_CONSUMPTION control statement" on page 51.
- The **UPDATE CONTAINER** control statement now allows you to specify a list of NO89 products for a container, as described in "UPDATE CONTAINER control statement" on page 59 and "Updating container properties" on page 95.
- A new report section, Capacity Changes (G5), has been added as described in "Capacity Changes section" on page 116.
- A new report section, SCRT Softcap Enforcement (G6), has been added as described in "SCRT Softcap Enforcement section" on page 116.
- A new pricing offering, Tailored Fit Pricing for IBM Z - Software Consumption Solution, has been added as described in Chapter 20, "Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution," on page 401.
- Support has been added for new z/OS sub-capacity products, as described in "Product support changes."

SCRT message changes:

- The following new messages have been added in [Appendix A, "SCRT messages,"](#) on page 449:

SCRTTOOL051 - SCRTTOOL052
 SCRTTOOL373 - SCRTTOOL378
 SCRTTOOL380 - SCRTTOOL388
 SCRTTOOL398 - SCRTTOOL399
 SCRTTOOL400 - SCRTTOOL415*
 SCRTTOOL425 - SCRTTOOL429*

*In past versions of SCRT that ran on z/VSE systems, messages in the SCRTTOOL4xx range were issued only on z/VSE systems. Since recent versions of SCRT no longer run on z/VSE systems, this message number range is now being re-purposed.

- The following existing messages have been updated:

SCRTTOOL325
 SCRTTOOL330

- The following message are no longer issued and have been removed:

SCRTTOOL053
 SCRTTOOL331 - SCRTTOOL333
 SCRTTOOL342
 SCRTTOOL400 - SCRTTOOL412**

**These messages were formerly issued by past versions of SCRT that ran on z/VSE systems. Since recent versions of SCRT no longer run on z/VSE systems, these old messages have been removed from this documentation.

Product support changes: SCRT V27.1.0 adds sub-capacity support for the products listed in [Table 74 on page 527](#).

Table 74. Additional and revised product support in SCRT V27.1.0

Operating system	Product name	Product ID
z/OS	IBM Z Batch Resiliency	5698-BR1

Software fixes in SCRT: For information about APARs and other software fixes in this release of SCRT, see [SCRT Current Levels \(ibm.biz/SCRTCurrentLevels\)](#).

Appendix E. Accessibility

Publications for this product are offered in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties when using PDF files, send an email to scrt@us.ibm.com or write to:

IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
USA

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. SCRT relies on the major accessibility features available in z/OS and z/VSE systems which enable users to:

- Use assistive technologies such as screen readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

Notices

This information was developed for products and services offered in the USA or elsewhere.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, HCD, JES2, JES3, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Minimum supported hardware

The minimum supported hardware for z/OS releases identified in z/OS announcements can subsequently change when service for particular servers or devices is withdrawn. Likewise, the levels of other software products supported on a particular release of z/OS are subject to the service support lifecycle of those products. Therefore, z/OS and its product publications (for example, panels, samples, messages, and product documentation) can include references to hardware and software that is no longer supported.

- For information about software support lifecycle, see: [IBM Lifecycle Support for z/OS \(www.ibm.com/software/support/systemsz/lifecycle\)](http://www.ibm.com/software/support/systemsz/lifecycle)
- For information about currently-supported IBM hardware, contact your IBM representative.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at [Copyright and Trademark information \(www.ibm.com/legal/copytrade.shtml\)](http://www.ibm.com/legal/copytrade.shtml).

Adobe is a trademark or registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Red Hat is a registered trademark of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Index

Special Characters

- container command line option [95](#)
- customer-number command line option [93](#)
- enterprise-tailored-fit-pricing command line option [96](#)
- exclude command line option [94](#)
- ignore-container command line option [96](#)
- include command line option [94](#)
- inputcsv command line option [94](#)
- isv command line option [96](#)
- isv-show-containers command line option [96](#)
- lpar-comment command line option [95](#)
- no89 command line option [93](#)
- update-container command line option [95](#)
- update-solution-id command line option [96](#)
- zwpc command line option [94](#)
- .csv format
 - reports in [4](#)
 - working with [420](#)
- % data collected field of the sub-capacity report
 - how determined [16](#)

Numerics

- 21st Century Software VSEⁿ
 - reporting IBM z/VSE products on [307](#)
- 2nd Highest field
 - Detail LPAR Data section [120](#)
- 4-hour rolling average utilization [16](#)
- 5655-R95 (WebSphere Transformation Extender V8),
 - considerations for
 - NO89 considerations [61](#)

A

- accessibility
 - contact IBM [529](#)
 - features [529](#)
- Active Containers section of multiplex report [397](#)
- Active Containers section of sub-capacity report [385](#)
- Advanced Entry Workload License Charges (AEWLC) [5](#)
- advanced image naming [104](#)
- Advanced Workload License Charges (AWLC) [5](#)
- alternate MSU values, placeholder fields for [420](#)
- analyzing the sub-capacity report [420](#)
- ASCII character set [39](#)
- ASSIGN ZWPC statement [40](#), [41](#)
- assistive technologies [529](#)
- audits of sub-capacity reports [17](#)

B

- billing effective date [425](#)
- billing related section of sub-capacity report [105](#)
- billing-related section of sub-capacity report [142](#)
- bimodal environment [447](#)

C

- calculating MSUs
 - for mixed z/VM guest systems [133](#)
 - for z/OS V1.6 or earlier guest systems [139](#)
- Capacity BackUp (CBU) testing [443](#)
- Capacity changes section of sub-capacity report [116](#)
- Capacity Changes section of sub-capacity report [116](#)
- Capacity Measurement Tool (CMT)
 - z/VSE APAR levels [306](#), [307](#)
- capacity-based pricing metrics [5](#)
- capping enforcement by SCRT, examples
 - traditional z/OS configuration [500](#)
 - z/OS traditional and z/OS zNALC configuration [503](#)
 - z/OS V1 and z/OS V2 configuration [505](#)
- capping individual LPARs
 - defined capacity [498](#)
 - LPAR physical capacity limit [498](#)
 - PR/SM hard capping [497](#)
- changes, machine software model [441](#)
- checklist for sub-capacity report [423](#)
- CLI [90](#)
- clock, CPC
 - set to future or past time [15](#)
- CMT
 - preparing z/VSE [313](#)
 - skeleton [310](#)
 - verifying [314](#)
- CMT files
 - how to configure [310](#)
- CMTLAB [311](#)
- CMTREPO
 - defining [316](#)
- CMTSTART [312](#)
- collecting SMF records [13](#)
- comma-separated-value file
 - format
 - reports in [4](#)
 - working with [420](#)
- command line interface (CLI) [90](#)
- command line options
 - container [95](#)
 - customer-number [93](#)
 - enterprise-tailored-fit-pricing [96](#)
 - exclude [94](#)
 - ignore-container [96](#)
 - include [94](#)
 - inputcsv [94](#)
 - isv [96](#)
 - isv-show-containers [96](#)
 - lpar-comment [95](#)
 - no89 [93](#)
 - update-container [95](#)
 - update-solution-id [96](#)
 - zwpc [94](#)

- COMMENT parameter, EXCLUDE control statement [49](#)
- common problem scenarios
 - testing Capacity BackUp, disaster recovery, or OOCod [443](#)
- common problems
 - data for same CPC on more than one sub-capacity report [443](#)
 - product MSUs exceed defined capacity for an LPAR [442](#)
 - providing alternate MSU values [443](#)
 - SMF records deleted [443](#)
 - Tool MSUs values changed [443](#)
- considerations
 - for Getting Started Sub-Capacity Pricing [165](#)
- Container Active Products section of multiplex report [397](#)
- Container Active Products section of sub-capacity report [386](#)
- Container Billing Summary section of multiplex report [397](#)
- Container Billing Summary section of sub-capacity report [385](#)
- CONTAINER control statement [41](#)
- Container Interval Consumption section of sub-capacity report [387](#)
- Container Max Contributors section of sub-capacity report [391](#), [399](#)
- Container MSU Consumption Contributors section of sub-capacity report [392](#)
- Container Product Grid Snapshot section of sub-capacity report [390](#), [399](#)
- Container Product Metrics section of sub-capacity report [388](#)
- CONTAINER statement [40](#)
- CONTAINER_INTERVAL_CONSUMPTION statement [40](#)
- control statements
 - SPECIAL DD
 - Suppress_Report_Section [170](#)
- control statements, customizing [320](#)
- Coordinated Universal Time (UTC) [14](#)
- country code [35](#)
- Country Multiplex Pricing (CMP)
 - multiplex [343](#)
 - multiplex report [343](#)
 - overview [343](#)
 - qualifications [343](#)
 - reporting for z/VSE systems [336](#)
 - sample
 - configuration [345](#)
 - control statement input [346](#)
 - customer verification section [339](#), [357](#)
 - detail data sections [358](#), [360](#)
 - detailed interval data section [362](#)
 - machine summary section [337](#), [348](#)
 - multiplex report [347](#)
 - multiplex summary section [337](#), [347](#)
- SCRT externals
 - CMP control statement [344](#)
 - Generate_Detailed_Data control statement [345](#)
- COUNTRY_MULTIPLEX_PRICING statement [40](#)
- CPC changes during reporting period [441](#)
- CPC MSU Rate Above Permanent Capacity Rating section of sub-capacity report [405](#)
- CPC parameter, EXCLUDE control statement [45](#)
- CPSTRACE DD statement [66](#)
- csv format
 - reports in [4](#)
 - working with [420](#)

- current levels, SCRT [439](#)
- Customer Comments field
 - customer information section [107](#)
 - Product Summary Information section [111](#)
- customer comments, placeholder fields for [420](#)
- customer information section of SCRT product report [416](#)
- customer information section of sub-capacity report [105](#), [107](#), [147](#), [155](#), [159](#), [162](#), [175](#), [190](#), [209](#), [219](#), [240](#), [262](#), [271](#), [291](#), [295](#), [324](#), [328](#), [331](#), [334](#), [383](#)
- customer MSUs field
 - Product Summary Information section [191](#)
- Customer MSUs field
 - Product Summary Information section [111](#), [176](#)
- Customer Name field
 - customer information section [106](#)
- customer number [35](#)
- Customer Number field
 - customer information section [106](#)
- customer verification section of sub-capacity report [118](#), [142](#)
- customer verification section of the sub-capacity report for native systems [118](#)
- customizing
 - NO89 control statements [320](#)
 - PARMS control statements [320](#)
 - SPECIAL control statements [321](#)
- customizing SCRT control statements [320](#)

D

- data collection incidents [17](#)
- data collection reporting [17](#)
- data gaps [17](#)
- date billing changes become effective [425](#)
- date to submit sub-capacity reports [423](#)
- Date/Time field
 - Detail LPAR Data section [120](#)
 - Missing LPAR Data Details section [114](#)
 - Product Max Contributors section [122](#)
- date/time stamps [14](#)
- DD statements
 - CPSTRACE [66](#)
 - GSTRACE [66](#)
 - INPUTCSV [60](#)
 - ISVLIB [61](#)
 - IWPTRACE [66](#)
 - NO89 [61](#)
 - OUTPUT [64](#)
 - PARMS [34](#)
 - SMF [65](#)
 - SPECIAL
 - ASSIGN ZWPC [40](#), [41](#)
 - CONTAINER [40](#), [41](#)
 - CONTAINER_INTERVAL_CONSUMPTION [40](#), [43](#)
 - COUNTRY_MULTIPLEX_PRICING [40](#), [43](#)
 - DETAIL_INTERVAL_RATE_DATA [40](#), [43](#)
 - DIAGMSG [40](#), [43](#)
 - ENTERPRISE_TAILORED_FIT_PRICING [40](#), [44](#)
 - EXCLUDE [40](#), [44](#)
 - EXCLUDE MSU_CONSUMPTION [51](#)
 - GENERATE_DETAILED_DATA [40](#), [52](#)
 - IGNORE CONTAINER [40](#), [53](#)
 - IGNORE_WLM_MWP_Data [54](#)
 - IGNORE_WLM_MWP_DATA [40](#)

DD statements (*continued*)

SPECIAL (*continued*)

- INCLUDE [40, 54](#)
- INTEGRATED_WORKLOAD_PRICING [40, 55](#)
- INTERVAL_RATE_DATA [40, 56](#)
- IPLA_ALL [40, 56](#)
- ISV_SHOW_CONTAINERS [40, 56](#)
- MLC_ALL [40, 56](#)
- REPORT_TIME [40, 56](#)
- SUPPRESS_REPORT_SECTION [40, 58](#)
- UPDATE_CONTAINER [40, 59](#)
- UPDATE SOLUTION_ID [41, 59](#)
- z/TPF_Transformation_Engines_Offering [60](#)
- z/TPF_TRANSFORMATION_ENGINES_OFFERING [41](#)
- STDENV [34](#)
- VMGUESTS [67](#)

dedicated CPUs

- limitation in prior versions [441](#)

defined capacity

- exceeded [442](#)
- in subcapacity report section [124](#)
- used instead of 4-hour rolling average [124](#)
- See also* capping individual LPARs

defined capacity, examples of, *See* capping enforcement by SCRT, examples

DELETE EXEC statement [33](#)

deleting SMF records [18](#)

Detail LPAR Data section of sub-capacity report [119, 120, 140, 143, 148, 156, 160, 163, 178, 196, 212, 222, 243, 264, 274, 293, 296, 325, 329, 332, 335, 389, 508](#)

Detail LPAR Usage Data section of sub-capacity report [120, 121](#)

DETAIL_INTERVAL_RATE_DATA statement [40, 43](#)

Detailed Interval Data section of multiplex report [363–365](#)

Detailed Interval Data section of sub-capacity report [406](#)

Detected field

- Missing CPC Data Details section [115](#)

- Missing LPAR Data Details section [113](#)

DIAGMSG statement [40, 43](#)

disaster recovery testing [443](#)

disaster recovery, handling [17](#)

DLBL statement [310](#)

double-byte character set (DBCS)

- countries, regions, and code pages acceptable by LMS [427](#)

- LMS support for [427, 431–433](#)

- values in report fields acceptable by LMS [421](#)

duplicate SMF records [13](#)

E

EBCDIC character set [39](#)

editing the sub-capacity report

- as a spreadsheet [420](#)

ENTERPRISE_TAILORED_FIT_PRICING statement [40, 44](#)

Entry Workload License Charges (EWLC) [5](#)

environments for generating reports [447](#)

error conditions

- displayed on the sub-capacity report [107](#)

Est Peak MSU field

- Missing LPAR Data Details section [114](#)

example

- 4-hour rolling average utilization [16](#)

- Country Multiplex Pricing

example (*continued*)

Country Multiplex Pricing (*continued*)

- customer verification section [339, 357](#)

- detail data sections [358, 360](#)

- detailed interval data section [362](#)

- machine summary section [337, 348](#)

- multiplex summary section [337, 347](#)

DLBL and EXTENT statements [312](#)

FTP transfer from z/VSE [318](#)

FTP transfer from z/VSE to z/OS [318](#)

FTP transfer to z/OS [319](#)

Getting Started Sub-Capacity Pricing [173, 174](#)

IBM Z Collocated Application Pricing [237, 239](#)

IBM Z Workload Pricing for Cloud, INPUTCSV-based [260](#)

IBM Z Workload Pricing for Cloud, WLM classification-based [269](#)

Integrated Workload Pricing [188, 189](#)

license capacity calculation [15](#)

Mobile Workload Pricing, INPUTCSV-based [207](#)

Mobile Workload Pricing, WLM classification-based [217](#)

multiple LPARS with mixed z/OS guests [136](#)

multiple LPARS with mixed z/OS guests and a z/OS native system [137](#)

multiple LPARs with z/OS V1.7 guests and native systems [137](#)

multiplex report for z/OS and z/VSE systems [336](#)

native z/OS systems [104](#)

native z/OS systems, only zNALC workloads [153](#)

native z/OS systems, traditional and zNALC workloads [145](#)

NO89 DD statement [64](#)

NO89 values for z/VSE [321](#)

single LPAR with mixed z/OS guests [135](#)

single LPAR with z/OS guests at V1.6 or lower [134](#)

single LPAR with z/OS guests at V1.7 or higher [134](#)

single-version z/OS native systems [105](#)

sub-capacity report for z/OS and z/VSE [333](#)

sub-capacity report for z/OS V1.6 guest systems [139](#)

sub-capacity report for z/OS V1.7 and later guest systems [141](#)

sub-capacity report for z/VSE only [323](#)

sub-capacity report for z/VSE V4 and V5 [327](#)

sub-capacity report for z/VSE V6, V5, and V4 [330](#)

z/OS and z/TPF systems [294](#)

z/OS and z/VSE systems [333](#)

z/OS V1 and z/OS V2 native systems [128](#)

z/OS V1 and z/OS V2 systems [127](#)

z/OS V1 and z/OS V2 traditional and zNALC workloads [149](#)

z/OS V1.6 guests, traditional and zNALC workloads, different LPARs [154](#)

z/OS V1.6 guests, traditional and zNALC workloads, same LPARs [157](#)

z/OS V1.7 guests, traditional and zNALC workloads [161](#)

z/TPF systems [290](#)

z/VSE systems only [323](#)

z/VSE V4 and V5 systems [327](#)

z/VSE V6, V5, and V4 systems [330](#)

EXCLUDE control statement

- COMMENT parameter [49](#)

- CPC parameter [45](#)

- IMAGE_ID parameter [46](#)

- PRODUCT_ID parameter [46](#)

- RESUME parameter [48](#)

EXCLUDE control statement (*continued*)

START parameter [47](#)

using UTC time [47](#)

EXCLUDE MSU_CONSUMPTION statement [51](#)

EXCLUDE statement [40](#), [44](#)

Exclude statements

displaying on the sub-capacity report [108](#)

EXEC statement [33](#)

execution-based IPLA licensing [6](#)

EXTENT statement [310](#)

F

features, product [440](#)

fixes, SCRT [439](#)

Flat WLC (FWLC) [5](#)

footnotes field

Product Summary Information section [191](#), [210](#), [220](#),
[241](#), [263](#), [272](#)

Footnotes field

Product Summary Information section [177](#)

full capacity pricing

metrics

Parallel Sysplex License Charges (PSLC) [5](#)

zSeries Entry License Charges (zELC) [5](#)

FWLC [5](#)

G

GENERATE_DETAILED_DATA statement [40](#), [52](#)

Getting Started Sub-Capacity Pricing [6](#)

Getting Started Sub-Capacity Pricing (GSSP)

considerations

for WebSphere Transformation Extender V8 (5655-
R95) [166](#)

externals

footnote [169](#)

report sections [169](#)

Suppress_Report_Section control statement [170](#)

trace records [171](#)

restrictions

for partially supported products [165](#)

for power save mode [166](#)

for z/OS 1.6 guests [165](#)

sample

configuration [173](#)

control statement input [174](#)

sample sub-capacity report [174](#)

understanding [165](#)

Getting Started Sub-Capacity Pricing Assigned Hours section
of sub-capacity report [214](#), [224](#), [245](#), [266](#), [276](#)

Getting Started Sub-Capacity Pricing Assigned section of
sub-capacity report [179](#), [180](#), [214](#), [224](#), [246](#), [266](#), [276](#)

Getting Started Sub-capacity Pricing Equal to Traditional
Sub-Capacity section of sub-capacity report [180](#)

Getting Started Sub-Capacity Pricing Equal to Traditional
Sub-Capacity section of sub-capacity report [179](#)

Getting Started Sub-Capacity Pricing Not Assigned section of
sub-capacity report [180](#)

GETVIS

requirements [310](#)

graphical user interface (GUI) [73](#)

group capacity LPARs

group capacity LPARs (*continued*)

assigning capacity [499](#)

availability [498](#)

requirements [499](#)

Group Capacity LPARs section of sub-capacity report [124](#),
[125](#), [149](#), [181](#), [197](#)

group capping, examples of, See capping enforcement by
SCRT, examples

group names and values fields

Product Max Contributors section [122](#)

GSSP

considerations

for WebSphere Transformation Extender V8 (5655-
R95) [166](#)

externals

footnote [169](#)

report sections [169](#)

Suppress_Report_Section control statement [170](#)

trace records [171](#)

restrictions

for partially supported products [165](#)

for power save mode [166](#)

for z/OS 1.6 guests [165](#)

sample

configuration [173](#)

control statement input [174](#)

sample sub-capacity report [174](#)

understanding [165](#)

GSTRACE DD statement [66](#)

guest systems at z/OS V1.6, sample sub-capacity report [139](#)

guest systems at z/OS V1.7 and later, sample sub-capacity
report [141](#)

guest z/OS V1.6 systems, traditional and zNALC workloads,
different LPARs, sample sub-capacity report [154](#)

guest z/OS V1.6 systems, traditional and zNALC workloads,
same LPARs, sample sub-capacity report [157](#)

guest z/OS V1.7 systems, traditional and zNALC workloads,
sample sub-capacity report [161](#)

GUI [73](#)

H

handling the sub-capacity report (IBM) [425](#)

Hardware Consumption Solution reporting [405](#)

Highest field

Detail LPAR Data section [120](#)

Product Max Contributors section [122](#)

Highest field of the sub-capacity report [120](#)

highest utilization

for native systems [119](#)

Hour Count field

Detail LPAR Data section [120](#)

hypervisor offset [14](#)

I

IBM country code [35](#)

IBM customer number [35](#)

IBM ID [427](#)

IBM support, obtaining [439](#)

IBM Z Collocated Application Pricing (zCAP)

data collection requirements [233](#)

explanation [244](#)

IBM Z Collocated Application Pricing (zCAP) *(continued)*

- monthly reporting process [233](#)
- overview [229](#)
- qualifications [232](#)
- sample
 - configuration [237](#)
 - control statement input [238](#)
 - sub-capacity report [239](#)
- SCRT externals
 - Generate_Detailed_Data control statement [234](#)
 - IBM Z Collocated Application Pricing footnote [234](#)
 - INPUTCSV DD statement [234](#)
 - INPUTCSV input data format [235](#)
 - Mobile/zCAP Pricing Detailed Data section [235](#)
- zCAP defining programs [229](#)
- IBM z Integrated Information Processor (zIIP) [16](#)
- IBM Z International Product License Agreement (IPLA) [6](#)
- IBM Z New Application License Charges (zNALC) [7](#)
- IBM Z Software Pricing Reference Guide [8](#)
- IBM Z Workload Pricing for Cloud (zWPC)
 - data collection requirements [254](#)
 - explanation [265](#), [275](#)
 - monthly reporting process [254](#)
 - overview [251](#)
 - qualifications [253](#)
 - sample
 - configuration [260](#), [269](#)
 - control statement input for INPUTCSV-based reporting [260](#)
 - control statement input for WLM-based reporting [270](#)
 - sub-capacity report [261](#), [271](#)
 - sample for INPUTCSV-based reporting [260](#)
 - sample for WLM classification-based reporting [269](#)
- SCRT externals
 - Assign ZWPC control statement [258](#)
 - DIAGMSG=WLMZWPC control statement [259](#)
 - Generate_Detailed_Data control statement [256](#), [259](#)
 - IBM Z Workload Pricing for Cloud footnote [255](#), [259](#)
 - INPUTCSV DD statement [255](#)
 - INPUTCSV input data format [256](#)
 - Mobile/zCAP Pricing Detailed Data section [256](#), [259](#)
 - zWPC_Input_From_WLM message [258](#)
- SCRT externals for CSV input [255](#)
- SCRT externals for WLM input [258](#)
- zWPC defining programs [251](#)
- IBM zEnterprise Application Assist Processor (zAAP) [16](#)
- ID and password for submitting reports [427](#)
- identities, z/OS native systems [103](#)
- IGNORE CONTAINER statement [40](#), [53](#)
- IGNORE_WLM_MWP_Data statement [54](#)
- IGNORE_WLM_MWP_DATA statement [40](#)
- image naming, advanced [104](#)
- IMAGE_ID parameter, EXCLUDE control statement [46](#)
- INCLUDE statement [40](#), [54](#)
- INCLUDE statements
 - displaying on the sub-capacity report [109](#)
- independent software vendor (ISV) programs
 - character set encoding [414](#)
 - externals for reporting [413](#)
 - overview [413](#)
 - restrictions for reporting [414](#)
 - sample JCL [415](#)

independent software vendor (ISV) programs *(continued)*

- sample report
 - customer information section [416](#)
 - Product Summary Information section [416](#)
- individual LPAR capping, examples of, *See* capping
- enforcement by SCRT, examples
- input for SCRT
 - NO89 control statements [310](#)
 - PARMS control statements [310](#)
 - SCRT89 records [309](#)
 - SPECIAL control statements [310](#)
- INPUTCSV DD statement [60](#)
- installing the SCRT89 record utility [287](#)
- Integrated Workload Pricing (IWP)
 - explanation [193](#)
 - externals
 - control statement [186](#)
 - footnote [187](#)
 - Special Conditions section [186](#)
 - trace records [187](#)
 - IWP adjusted programs [183](#), [184](#)
 - IWP defining programs [183](#), [184](#)
 - overview [183](#)
 - qualifications for [185](#)
 - sample
 - configuration [188](#)
 - control statement input [189](#)
 - sub-capacity report [189](#)
- INTEGRATED_WORKLOAD_PRICING statement [40](#), [55](#)
- International Product License Agreement (IPLA) licensing [6](#)
- interval rate data, reporting of [405](#)
- INTERVAL_RATE_DATA statement [40](#), [56](#)
- interval, maximum utilization
 - for native systems [119](#)
- IPLA CICS-Based virtual product [112](#)
- IPLA COBOL-Based virtual product [112](#)
- IPLA DB2-Based virtual product [112](#)
- IPLA IMS-Based virtual product [112](#)
- IPLA licensing
 - execution-based licensing
 - Getting Started Sub-Capacity Pricing [6](#)
 - reference-based licensing [6](#)
 - z/OS-based licensing [6](#)
- IPLA migration grace period, *See* Multi-Version Measurement
- IPLA MQ-Based virtual product [112](#)
- IPLA PL/I-Based virtual product [112](#)
- IPLA Product id field
 - Product Summary Information section [111](#)
- IPLA Product Name field
 - Product Summary Information section [111](#)
- IPLA z/OS-Based virtual product [112](#)
- IPLA_ALL statement [40](#), [56](#)
- ISV, *See* independent software vendor
- ISV_SHOW_CONTAINERS statement [40](#), [56](#)
- ISVLIB DD statement [61](#)
- IWM063I, message issued by z/OS [442](#)
- IWPTRACE DD statement [66](#)

J

- Java workloads [16](#)
- JCL statements
 - DELETE EXEC [33](#)
 - EXEC [33](#)

JCL statements (*continued*)

JOB [32](#)

JOB statement [32](#)

Justification for low data collection field

Missing CPC Data Details section [115](#)

Missing LPAR Data Details section [114](#)

L

levels, SCRT current [439](#)

License Management Support (LMS) application

choosing functions [428](#)

confirmation [434](#)

customer support center [427](#)

data submission [428](#)

getting access [427](#)

overview [427](#)

submitting reports [434](#)

system list [430](#)

uploading reports [429](#)

viewing archived reports [437](#)

viewing reports [429](#)

web interface for submitting sub-capacity reports [427](#)

web site (URL) [427](#)

License Manager Support (LMS)

uploading data in a single file [65](#)

limitations of SCRT

in current version [439](#)

product features in different LPARs [440](#)

removed from current version

dedicated CPs [441](#)

running z/OS guests under z/VM [441](#)

support for 128 native and guest systems [440](#)

Wait Completion=Yes [441](#)

Single Version Charging [440](#)

SMF type 89 records [440](#)

z/OS guests [440](#)

LMS

customer support center [427](#)

getting access [427](#)

overview [427](#)

subscribing to informational email [437](#)

web interface for submitting sub-capacity reports [427](#)

web site (URL) [427](#)

LPAR group capping, examples of, See capping enforcement

by SCRT, examples

LPAR names and values fields

Product Max Contributors section [122](#)

LPAR/Group contribution to estimated peak field

Missing LPAR Data Details section [114](#)

LPARs

assigned to a group [125](#)

changing times in [15](#)

license capacity example [15](#)

utilization capacity [8](#)

M

machine replacements [442](#)

Machine Serial Number field

customer information section [106](#)

machine software model changes [441](#)

machine summary section of multiplex report [337](#), [348](#), [394](#)

Machine Type/Model field

customer information section [106](#)

maximum concurrent LPAR utilization [4](#)

maximum utilization

for native systems [119](#)

message IWM063I (on z/OS) [442](#)

messages

internal SCRT restrictions [494](#)

SCRT89 record utility

SCRTRU001 [495](#)

SCRTRU002 [495](#)

SCRTRU003 [495](#)

SCRTRU004 [495](#)

SCRTRU005 [495](#)

z/OS systems

***** [449](#)

SCRTTOOL001 [449](#)

SCRTTOOL002 [449](#)

SCRTTOOL003 [449](#)

SCRTTOOL004 [449](#)

SCRTTOOL005 [449](#)

SCRTTOOL006 [450](#)

SCRTTOOL007 [450](#)

SCRTTOOL009 [450](#)

SCRTTOOL012 [450](#)

SCRTTOOL015 [451](#)

SCRTTOOL017 [451](#)

SCRTTOOL025 [451](#)

SCRTTOOL026 [451](#)

SCRTTOOL030 [452](#)

SCRTTOOL031 [452](#)

SCRTTOOL032 [452](#)

SCRTTOOL033 [452](#)

SCRTTOOL036 [452](#)

SCRTTOOL045 [453](#)

SCRTTOOL047 [453](#)

SCRTTOOL048 [454](#)

SCRTTOOL050 [454](#)

SCRTTOOL054 [454](#)

SCRTTOOL055 [455](#)

SCRTTOOL065 [456](#)

SCRTTOOL068 [456](#)

SCRTTOOL069 [456](#)

SCRTTOOL070 [456](#)

SCRTTOOL072 [457](#)

SCRTTOOL073 [457](#)

SCRTTOOL074 [457](#)

SCRTTOOL075 [457](#)

SCRTTOOL076 [457](#)

SCRTTOOL077 [457](#)

SCRTTOOL080 [458](#)

SCRTTOOL091 [458](#)

SCRTTOOL092 [458](#)

SCRTTOOL093 [458](#)

SCRTTOOL094 [458](#)

SCRTTOOL095 [458](#)

SCRTTOOL099 [458](#)

SCRTTOOL100 [459](#)

SCRTTOOL101 [459](#)

SCRTTOOL102 [459](#)

SCRTTOOL103 [459](#)

SCRTTOOL104 [460](#)

SCRTTOOL105 [460](#)

messages (*continued*)
z/OS systems (*continued*)

SCRTTOOL106 [460](#)
SCRTTOOL107 [460](#)
SCRTTOOL109 [460](#)
SCRTTOOL110 [460](#)
SCRTTOOL119 [461](#)
SCRTTOOL120 [461](#)
SCRTTOOL121 [461](#)
SCRTTOOL123 [461](#)
SCRTTOOL124 [461](#)
SCRTTOOL125 [461](#)
SCRTTOOL127 [462](#)
SCRTTOOL129 [462](#)
SCRTTOOL130 [462](#)
SCRTTOOL131 [462](#)
SCRTTOOL134 [462](#)
SCRTTOOL135 [462](#)
SCRTTOOL136 [463](#)
SCRTTOOL137 [463](#)
SCRTTOOL138 [463](#)
SCRTTOOL139 [463](#)
SCRTTOOL140 [464](#)
SCRTTOOL141 [464](#)
SCRTTOOL142 [464](#)
SCRTTOOL143 [464](#)
SCRTTOOL148 [465](#)
SCRTTOOL149 [465](#)
SCRTTOOL150 [465](#)
SCRTTOOL153 [465](#)
SCRTTOOL154 [465](#)
SCRTTOOL155 [465](#)
SCRTTOOL158 [466](#)
SCRTTOOL161 [466](#)
SCRTTOOL171 [466](#)
SCRTTOOL173 [467](#)
SCRTTOOL174 [467](#)
SCRTTOOL175 [467](#)
SCRTTOOL176 [467](#)
SCRTTOOL177 [468](#)
SCRTTOOL178 [468](#)
SCRTTOOL179 [468](#)
SCRTTOOL180 [469](#)
SCRTTOOL200 [469](#)
SCRTTOOL201 [469](#)
SCRTTOOL202 [469](#)
SCRTTOOL203 [469](#)
SCRTTOOL204 [470](#)
SCRTTOOL205 [470](#)
SCRTTOOL210 [470](#)
SCRTTOOL211 [470](#)
SCRTTOOL212 [470](#)
SCRTTOOL220 [470](#)
SCRTTOOL221 [471](#)
SCRTTOOL230 [471](#)
SCRTTOOL240 [471](#)
SCRTTOOL241 [471](#)
SCRTTOOL242 [471](#)
SCRTTOOL243 [471](#)
SCRTTOOL244 [472](#)
SCRTTOOL245 [472](#)
SCRTTOOL250 [472](#)
SCRTTOOL251 [472](#)
SCRTTOOL252 [473](#)

messages (*continued*)
z/OS systems (*continued*)

SCRTTOOL253 [473](#)
SCRTTOOL254 [473](#)
SCRTTOOL255 [473](#)
SCRTTOOL256 [473](#)
SCRTTOOL257 [474](#)
SCRTTOOL258 [474](#)
SCRTTOOL259 [474](#)
SCRTTOOL260 [474](#)
SCRTTOOL261 [474](#)
SCRTTOOL262 [474](#)
SCRTTOOL263 [475](#)
SCRTTOOL264 [475](#)
SCRTTOOL265 [475](#)
SCRTTOOL275 [475](#)
SCRTTOOL276 [475](#)
SCRTTOOL277 [475](#)
SCRTTOOL278 [476](#)
SCRTTOOL280 [476](#)
SCRTTOOL281 [476](#)
SCRTTOOL282 [476](#)
SCRTTOOL295 [476](#)
SCRTTOOL296 [477](#)
SCRTTOOL297 [477](#)
SCRTTOOL300 [477](#)
SCRTTOOL301 [477](#)
SCRTTOOL302 [477](#)
SCRTTOOL303 [477](#)
SCRTTOOL319 [478](#)
SCRTTOOL320 [478](#)
SCRTTOOL321 [478](#)
SCRTTOOL322 [478](#)
SCRTTOOL323 [478](#)
SCRTTOOL324 [479](#)
SCRTTOOL325 [479](#)
SCRTTOOL326 [479](#)
SCRTTOOL327 [479](#)
SCRTTOOL328 [479](#)
SCRTTOOL329 [479](#)
SCRTTOOL330 [480](#)
SCRTTOOL334 [480](#)
SCRTTOOL335 [480](#)
SCRTTOOL336 [480](#)
SCRTTOOL337 [480](#)
SCRTTOOL338 [480](#)
SCRTTOOL343 [481](#)
SCRTTOOL344 [481](#)
SCRTTOOL345 [481](#)
SCRTTOOL346 [481](#)
SCRTTOOL347 [481](#)
SCRTTOOL348 [482](#)
SCRTTOOL349 [482](#)
SCRTTOOL350 [482](#)
SCRTTOOL351 [482](#)
SCRTTOOL352 [482](#)
SCRTTOOL353 [482](#)
SCRTTOOL354 [483](#)
SCRTTOOL355 [483](#)
SCRTTOOL356 [483](#)
SCRTTOOL357 [483](#)
SCRTTOOL358 [483](#)
SCRTTOOL359 [483](#)
SCRTTOOL360 [484](#)

messages (*continued*)

z/OS systems (*continued*)

SCRTTOOL361 [484](#)
SCRTTOOL362 [484](#)
SCRTTOOL363 [484](#)
SCRTTOOL364 [484](#)
SCRTTOOL365 [485](#)
SCRTTOOL366 [485](#)
SCRTTOOL367 [485](#)
SCRTTOOL368 [485](#)
SCRTTOOL369 [485](#)
SCRTTOOL370 [486](#)
SCRTTOOL371 [486](#)
SCRTTOOL372 [486](#)
SCRTTOOL373 [486](#)
SCRTTOOL374 [486](#)
SCRTTOOL375 [486](#)
SCRTTOOL376 [487](#)
SCRTTOOL377 [487](#)
SCRTTOOL378 [487](#)
SCRTTOOL380 [487](#)
SCRTTOOL381 [487](#)
SCRTTOOL382 [487](#)
SCRTTOOL383 [487](#)
SCRTTOOL384 [487](#)
SCRTTOOL385 [488](#)
SCRTTOOL386 [488](#)
SCRTTOOL387 [488](#)
SCRTTOOL388 [488](#)
SCRTTOOL389 [488](#)
SCRTTOOL390 [488](#)
SCRTTOOL398 [488](#)
SCRTTOOL399 [488](#)
SCRTTOOL400 [489](#)
SCRTTOOL401 [489](#)
SCRTTOOL402 [489](#)
SCRTTOOL403 [489](#)
SCRTTOOL404 [489](#)
SCRTTOOL405 [489](#)
SCRTTOOL406 [489](#)
SCRTTOOL407 [490](#)
SCRTTOOL408 [490](#)
SCRTTOOL409 [490](#)
SCRTTOOL410 [490](#)
SCRTTOOL411 [490](#)
SCRTTOOL412 [490](#)
SCRTTOOL413 [490](#)
SCRTTOOL414 [491](#)
SCRTTOOL415 [491](#)
SCRTTOOL416 [491](#)
SCRTTOOL425 [491](#)
SCRTTOOL426 [491](#)
SCRTTOOL427 [491](#)
SCRTTOOL428 [491](#)
SCRTTOOL429 [492](#)
SCRTTOOL601 [492](#)
SCRTTOOL602 [492](#)
SCRTTOOL603 [492](#)
SCRTTOOL605 [492](#)
SCRTTOOL610 [492](#)
SCRTTOOL611 [493](#)
SCRTTOOL612 [493](#)
SCRTTOOL613 [493](#)
SCRTTOOL620 [494](#)

messages (*continued*)

z/OS systems (*continued*)

SCRTTOOL621 [494](#)
SCRTTOOL650 [494](#)
Midrange Workload License Charges (MWLC) [5](#)
Migration Pricing Option (MPO), *See* Multi-Version Measurement
Missing CPC Data Details section of sub-capacity report [114](#), [115](#), [384](#)
Missing LPAR Data Details section [112](#)
Missing LPAR Data Details section of sub-capacity report [113](#), [384](#)
mixed z/VM guest systems, calculating MSUs for [133](#)
MLC Product id field
 Product Summary Information section [110](#)
MLC Product Name field
 Product Summary Information section [110](#)
MLC_ALL statement [40](#), [56](#)
Mobile Workload Pricing (MWP)
 data collection requirements [202](#)
 explanation [213](#), [223](#)
 monthly reporting process [202](#)
 MWP defining programs [199](#)
 overview
 mobile workload classification [200](#)
 qualifications [201](#)
 sample
 configuration [207](#), [217](#)
 control statement input for INPUTCSV-based reporting [208](#)
 control statement input for WLM-based reporting [218](#)
 sub-capacity report [209](#), [219](#)
 sample for INPUTCSV-based reporting [207](#)
 sample for WLM classification-based reporting [217](#)
SCRT externals
 DIAGMSG=WLMMOB control statement [206](#)
 Generate_Detailed_Data control statement [203](#)
 IGNORE_WLM_MWP_DATA control statement [206](#)
 INPUTCSV DD statement [203](#)
 INPUTCSV input data format [204](#)
 Mobile Workload Pricing footnote [203](#), [206](#)
 Mobile/zCAP Pricing Detailed Data section [204](#), [207](#)
SCRT externals for CSV input [203](#)
SCRT externals for WLM input [205](#)
 workload management (WLM) [200](#)
Mobile/zCap Pricing Detailed Data section of sub-capacity report [216](#), [217](#), [226–228](#), [247–250](#), [268](#), [269](#), [278–280](#)
model changes, machine software [441](#)
model upgrades, machine software [441](#)
MSU capping enforcement by SCRT, examples, *See* capping enforcement by SCRT, examples
MSU Consumption Data Collection section of sub-capacity report [387](#)
MSU values for mixed z/VM guest systems [133](#)
MSUS
 calculating for z/OS guest systems [134](#)
Multi-Version Measurement (MVM)
 externals [367](#)
 overview [367](#)
 prerequisites [367](#)
 report differences [368](#)

- Multi-Version Measurement (MVM) (*continued*)
 - sample report [368](#)
 - zNALC [151](#)
- multiplex [343](#)
- Multiplex Peak MSU field
 - Missing LPAR Data Details section [114](#)
- multiplex report
 - Active Containers section [397](#)
 - Container Active Products section [397](#)
 - Container Billing Summary section [397](#)
 - customer verification section for CPC
 - machine information [398](#)
 - detail data sections [340](#), [358](#)
 - detail LPAR data section for CPC [340](#), [358](#), [360](#), [361](#)
 - detail LPAR usage data section for CPC [359](#), [361](#)
 - Detailed Interval Data section [363–365](#)
 - machine information section for CPC [340](#), [358](#), [360](#)
 - machine summary section [337](#), [348](#), [394](#)
 - machine summary section for CPC
 - Active Containers section [396](#)
 - machine information [338](#), [340](#), [351](#), [355](#), [358](#), [360](#), [396](#)
 - missing CPC data details [339](#), [354](#), [357](#)
 - Missing CPC data details [396](#)
 - missing LPAR data details [339](#), [353](#), [356](#)
 - Missing LPAR data details [396](#)
 - product summary information [338](#), [352](#), [355](#)
 - Product summary information [396](#)
 - special conditions [351](#)
 - missing CPC data details section [339](#), [354](#), [357](#)
 - missing LPAR data details section [339](#), [354](#), [357](#)
 - Missing LPAR Data Details section [112](#)
 - multiplex summary section [337](#), [347](#), [394](#)
 - product features section for CPC [353](#), [356](#)
 - Product Grid Snapshot section [398](#)
 - product grid snapshot section for CPC [341](#), [359](#), [360](#), [362](#)
 - Product Max Contributors section [398](#)
 - product max contributors section for CPC [341](#), [359](#), [361](#), [362](#)
 - product summary information section for CPC [338](#), [352](#), [355](#)
 - sample for z/OS and z/VSE systems [336](#)
 - system recovery boost activity section for CPC [353](#), [356](#)
- multiplex report for Country Multiplex Pricing
 - sample report
 - customer verification section [339](#), [357](#)
 - detail data sections [358](#), [360](#)
 - detailed interval data section [362](#)
 - machine summary section [337](#), [348](#)
 - multiplex summary section [337](#), [347](#)
- multiplex summary section of multiplex report [337](#), [347](#), [394](#)
- MVM, See Multi-Version Measurement

N

- native z/OS systems (single-version), sample sub-capacity report [105](#)
- native z/OS systems, only zNALC workloads, sample sub-capacity report [153](#)
- native z/OS systems, sample single-version configuration [104](#)

- native z/OS systems, traditional and zNALC workloads, sample sub-capacity report [145](#)
- NO89 control statements
 - customizing [320](#)
- NO89 DD statement [61](#)
- non-MSU reporting metrics
 - metric types
 - simultaneous instances [371](#)
 - sample sub-capacity report [372](#)
 - SCRT externals for [371](#)

O

- On/Off Capacity on Demand (OOCOD) testing [443](#)
- OS Product ID field
 - Missing LPAR Data Details section [114](#)
- OS Product Name field
 - Missing LPAR Data Details section [114](#)
- OUTPUT DD statement [64](#)
- overview of SCRT process [8](#)

P

- Parallel Sysplex License Charges (PSLC) [5](#)
- parent product [6](#)
- PARMS control statements
 - customizing [320](#)
- PARMS DD statement [34](#)
- PARMS parameters
 - optional [36](#), [37](#)
 - required [34](#)
- percentage of data collected
 - affected by reporting period [16](#)
- performance of the SCRT job [13](#)
- placeholder fields
 - conditional [420](#)
 - optional [420](#)
 - required [421](#)
- planning for sub-capacity pricing using SCRT [10](#)
- planning to install SCRT [11](#)
- power save mode, static
 - sample Detail LPAR Data Section [508](#)
- SCRT support for [507](#)
- z/OS support for
 - MAXEVENTINTRECS parameter [509](#)
 - required APARs [509](#)
- z/TPF support for [511](#)
- z/VSE support for [513](#)
- prerequisites, SCRT [11](#)
- pricing
 - and LPARs [11](#)
 - IBM Z IPLA [6](#)
 - Monthly License Charges (MLC) [5](#)
- pricing, full capacity
 - metrics [5](#)
- pricing, sub-capacity
 - metrics [5](#)
- problem scenarios, common [442](#)
- problems
 - CPC changed during reporting period [441](#)
 - missing purchase order number [36](#)
 - SCRT generates error report [447](#)
 - SCRT limitations [440](#)

- problems (*continued*)
 - unusual situations affecting product use [17](#)
- procedures
 - getting access to LMS [427](#)
 - viewing and editing the sub-capacity report
 - spreadsheet method [420](#)
- process, overview of SCRT [8](#)
- product features [440](#)
- Product Features section of sub-capacity report [115](#), [116](#)
- Product Grid Snapshot section of multiplex report [398](#)
- Product Grid Snapshot section of sub-capacity report [123](#), [124](#), [130](#), [141](#), [144](#), [152](#), [157](#), [161](#), [164](#), [179](#), [197](#), [214](#), [224](#), [245](#), [266](#), [276](#), [293](#), [297](#), [326](#), [330](#), [332](#), [336](#), [390](#)
- Product Grid Snapshot Section of sub-capacity report [149](#)
- Product ID field
 - Product Max Contributors section [121](#)
- Product Max Contributors section of multiplex report [398](#)
- Product Max Contributors section of sub-capacity report [121](#), [123](#), [129](#), [130](#), [141](#), [143](#), [148](#), [151](#), [156](#), [160](#), [163](#), [178](#), [196](#), [212](#), [223](#), [244](#), [265](#), [275](#), [293](#), [296](#), [326](#), [329](#), [332](#), [335](#), [374](#), [389](#)
- Product Name field
 - Product Max Contributors section [121](#)
- Product Summary Information section
 - Footnotes field [373](#)
 - non-MSU-based products [372](#)
 - Tool Units field [373](#)
- Product Summary Information section of ISV product report [416](#)
- Product Summary Information section of sub-capacity report [110](#), [112](#), [128](#), [129](#), [147](#), [150](#), [153](#), [155](#), [159](#), [163](#), [176](#), [191](#), [192](#), [210](#), [211](#), [220](#), [221](#), [241](#), [242](#), [262](#), [263](#), [272](#), [273](#), [292](#), [295](#), [324](#), [328](#), [331](#), [334](#), [369](#), [373](#), [384](#)
- PRODUCT_ID parameter, EXCLUDE control statement [46](#)
- products running in each LPAR [123](#)
- purchase order number field [36](#)
- Purchase Order Number field
 - customer information section [106](#)

Q

- QPP indicator, in MLC Product Name field [110](#)
- Quality Partnership Program (QPP) [110](#)

R

- receiving the sub-capacity report (IBM) [425](#)
- Recoup Utility, z/TPF [281](#), [289](#)
- reference guide, software pricing [8](#)
- reference-based IPLA licensing [6](#)
- replacements, machine [442](#)
- REPORT_TIME control statement
 - Special_Local_Time option [57](#)
 - UTC_Time option [57](#)
- REPORT_TIME statement [40](#), [56](#)
- report, SCRT [3](#)
- report, sub-capacity [3](#)
- reporting period
 - determining [16](#)
 - displayed in sub-capacity report for native z/OS systems [107](#)
- Reporting Period field
 - Tool Information section [107](#)

- reports from different time zones [119](#)
- requirements
 - sub-capacity pricing for z/VSE systems [306](#)
- requirements for z/OS [21](#)
- requirements, SCRT [11](#)
- Resolved field
 - Missing CPC Data Details section [115](#)
 - Missing LPAR Data Details section [114](#)
- Resource Measurement Facility (RMF) [12](#)
- RESUME parameter, EXCLUDE control statement [48](#)
- return codes [448](#)
- rolling average, 4-hour [16](#)

S

- sample
 - NO89 DD statement [64](#)
 - time zone processing [119](#)
- sample configuration
 - Country Multiplex Pricing [345](#)
 - Getting Started Sub-Capacity Pricing [173](#)
 - IBM Z Collocated Application Pricing [237](#)
 - INPUTCSV-based IBM Z Workload Pricing for Cloud [260](#)
 - INPUTCSV-based Mobile Workload Pricing [207](#)
 - Integrated Workload Pricing [188](#)
 - native z/OS systems [104](#)
 - WLM-based IBM Z Workload Pricing for Cloud [269](#)
 - WLM-based Mobile Workload Pricing [217](#)
 - z/OS V1 and z/OS V2 systems [127](#)
- sample multiplex report
 - z/OS and z/VSE systems [336](#)
- sample sub-capacity report
 - Country Multiplex Pricing [337](#), [339](#), [347](#), [348](#), [357](#), [358](#), [360](#), [362](#)
 - Getting Started Sub-Capacity Pricing [174](#)
 - IBM Z Collocated Application Pricing [239](#)
 - IBM Z Workload Pricing for Cloud, INPUTCSV-based [261](#)
 - IBM Z Workload Pricing for Cloud, WLM-based [271](#)
 - Integrated Workload Pricing [189](#)
 - Mobile Workload Pricing, INPUTCSV-based [209](#)
 - Mobile Workload Pricing, WLM-based [219](#)
 - native z/OS systems, only zNALC workloads [153](#)
 - native z/OS systems, traditional and zNALC workloads [145](#)
 - single-version z/OS native systems [105](#)
 - sub-capacity report for z/OS V1.6 guest systems [139](#)
 - sub-capacity report for z/OS V1.7 and later guest systems [141](#)
 - z/OS and z/TPF systems [294](#)
 - z/OS and z/VSE systems [333](#)
 - z/OS V1 and z/OS V2 native systems [128](#)
 - z/OS V1 and z/OS V2 traditional and zNALC workloads [149](#)
 - z/OS V1.6 guests, traditional and zNALC workloads, different LPARs [154](#)
 - z/OS V1.6 guests, traditional and zNALC workloads, same LPARs [157](#)
 - z/OS V1.7 guests, traditional and zNALC workloads [161](#)
 - z/TPF systems [290](#)
 - z/VSE systems only [323](#)
 - z/VSE V4 and V5 systems [327](#)
 - z/VSE V6, V5, and V4 systems [330](#)
- SCRT
 - prerequisites [11](#)

- SCRT (*continued*)
 - requirements [11](#)
 - skill requirements [xv](#)
 - subscribing to informational email via LMS [437](#)
- SCRT control statements, customizing [320](#)
- SCRT current levels [439](#)
- SCRT fixes [439](#)
- SCRT limitations
 - SMF type 89 records [440](#)
- SCRT PARMs parameters
 - optional [36](#), [37](#)
 - required [34](#)
- SCRT process
 - for z/VSE systems [305](#)
- SCRT process overview [8](#)
- SCRT report [3](#)
- SCRT Softcap Enforcement section of sub-capacity report [116](#)
- SCRT89
 - extracting records [14](#)
 - records from different time zones [14](#)
- SCRT89 processing
 - on Windows or Linux [320](#)
 - on z/OS
 - steps [317](#)
 - preparing for, on z/OS [316](#)
- SCRT89 record
 - files required [309](#)
 - input for SCRT [309](#)
- SCRT89 record utility
 - installing [287](#)
 - using [288](#)
- SCRT89 records
 - backing up
 - rule [322](#)
 - backup considerations [322](#)
 - collecting from z/TPF [283](#)
 - collecting from z/VSE [309](#)
 - for sub-capacity reporting [314](#)
 - producing [309](#)
 - retaining [14](#)
- Select Application License Charges (SALC) [6](#)
- sending the sub-capacity report to IBM
 - via the web [427](#)
- sending TSAD to IBM [514](#)
- simultaneous instances metric [371](#)
- Single Version Charging [440](#)
- Single Version Charging (SVC) [343](#)
- single-byte character set (SBCS) [35](#), [38](#), [39](#), [419](#), [421](#), [431](#)
- SKCMT [311](#)
- SKCMTINI [313](#)
- SKCMTREP [313](#)
- skeleton
 - customizing [313](#)
 - SKCMT [311](#)
 - SKCMTINI [313](#)
 - SKCMTREP [313](#)
 - SKUBERG [313](#)
- skills
 - requirements for SCRT [xv](#)
- SKUBERG [313](#)
- SMF DD statement [65](#)
- SMF records

- SMF records (*continued*)
 - collecting [13](#)
 - do not delete [18](#)
 - duplicate [13](#)
 - retaining [14](#)
- SMF type 70 subtype 1 and type 89 records [4](#), [12](#)
- SMF/SCRT89 Input Data Statistics section of sub-capacity report [119](#), [140](#), [143](#), [148](#), [156](#), [160](#), [163](#), [178](#), [196](#), [211](#), [221](#), [242](#), [264](#), [273](#), [293](#), [296](#), [325](#), [329](#), [332](#), [335](#), [389](#)
- software asset manager, sub-capacity analysis by [420](#)
- Software Consumption Solution , See Tailored Fit Pricing for IBM Z - Software Consumption Solution
- software model changes, machine [441](#)
- software pricing reference guide [8](#)
- Solution Edition offering [38](#), [45](#)
- Special conditions section of sub-capacity report [220](#), [272](#)
- Special Conditions section of sub-capacity report
 - example [108](#), [110](#)
 - Exclude statements [108](#)
 - for Integrated Workload Pricing [110](#)
 - INCLUDE statements [109](#)
- SPECIAL control statements
 - customizing [321](#)
- SPECIAL DD statement
 - ASSIGN ZWPC [40](#), [41](#)
 - CONTAINER [40](#), [41](#)
 - CONTAINER_INTERVAL_CONSUMPTION [40](#), [43](#)
 - COUNTRY_MULTIPLEX_PRICING [40](#), [43](#)
 - DETAIL_INTERVAL_RATE_DATA [40](#), [43](#)
 - DIAGMSG [40](#), [43](#)
 - ENTERPRISE_TAILORED_FIT_PRICING [40](#), [44](#)
 - EXCLUDE [40](#), [44](#)
 - EXCLUDE control statement
 - COMMENT parameter [49](#)
 - CPC parameter [45](#)
 - IMAGE_ID parameter [46](#)
 - PRODUCT_ID parameter [46](#)
 - RESUME parameter [48](#)
 - START parameter [47](#)
 - using UTC time [47](#)
 - EXCLUDE MSU_CONSUMPTION [51](#)
 - GENERATE_DETAILED_DATA [40](#), [52](#)
 - IGNORE CONTAINER [40](#), [53](#)
 - IGNORE_WLM_MWP_Data [54](#)
 - IGNORE_WLM_MWP_DATA [40](#)
 - INCLUDE [40](#), [54](#)
 - INTEGRATED_WORKLOAD_PRICING [40](#), [55](#)
 - INTERVAL_RATE_DATA [40](#), [56](#)
 - IPLA_ALL [40](#), [56](#)
 - ISV_SHOW_CONTAINERS [40](#), [56](#)
 - MLC_ALL [40](#), [56](#)
 - REPORT_TIME [40](#), [56](#)
 - REPORT_TIME control statement
 - Special_Local_Time option [57](#)
 - UTC_Time option [57](#)
 - SUPPRESS_REPORT_SECTION [40](#), [58](#)
 - UPDATE CONTAINER [40](#), [59](#)
 - UPDATE SOLUTION_ID [41](#), [59](#)
 - z/TPF_Transformation_Engines_ Offering [60](#)
 - z/TPF_TRANSFORMATION_ENGINES_OFFERING [41](#)
- Special_Local_Time option, REPORT_TIME control statement [57](#)
- standard sub-capacity pricing. See traditional sub-capacity pricing

- START parameter, EXCLUDE control statement [47](#)
- static power save mode
 - sample Detail LPAR Data Section [508](#)
 - SCRT support for [507](#)
 - z/OS support for
 - MAXEVENTINTRECS parameter [509](#)
 - required APARs [509](#)
 - z/TPF support for [511](#)
 - z/VSE support for [513](#)
- STDENV DD statement [34](#)
- Sub-Capacity Capping Enforced by SCRT section of sub-capacity report [124](#), [131](#), [141](#), [144](#), [149](#), [152](#), [157](#), [161](#), [164](#), [178](#), [179](#), [196](#), [197](#), [297](#)
- sub-capacity pricing
 - based on LPAR utilization capacity [8](#)
 - for System z New Application License Charges (zNALC) [145](#)
 - for z/TPF systems [281](#)
 - for z/VM guest systems [133](#)
 - IPLA [6](#)
 - metrics
 - Advanced Entry Workload License Charges (AEWLC) [5](#)
 - Advanced Workload License Charges (AWLC) [5](#)
 - Entry Workload License Charges (EWLC) [5](#)
 - IBM Z New Application License Charges (zNALC) [7](#)
 - Integrated Workload Pricing (IWP) [5](#)
 - Midrange Workload License Charges (MWLC) [5](#)
 - Select Application License Charges (SALC) [6](#)
 - System z New Application License Charges (zNALC) [5](#)
 - Workload License Charges (WLC) [5](#)
 - prerequisites [11](#)
 - terms and conditions [514](#)
 - types supported [7](#)
 - using SCRT to plan for [10](#)
- sub-capacity report
 - Active Containers section [385](#)
 - auditing [17](#)
 - billing-related section [105](#)
 - Capacity Changes section [116](#)
 - checklist for submission [423](#)
 - Container Active Products section [386](#)
 - Container Billing Summary section [385](#)
 - Container Interval Consumption section [387](#)
 - Container Max Contributors section [391](#), [399](#)
 - Container MSU Consumption Contributors section [392](#)
 - Container Product Grid Snapshot section [390](#), [399](#)
 - Container Product Metrics section [388](#)
 - CPC MSU Rate Above Permanent Capacity Rating section [405](#)
 - customer information section
 - Customer Comments field [107](#)
 - Customer Name field [106](#)
 - Customer Number field [106](#)
 - Machine Rated Capacity field [106](#)
 - Machine Serial Number field [106](#)
 - Machine Type/Model field [106](#)
 - Purchase Order Number field [106](#)
 - submitter email address field [105](#)
 - submitter name field [105](#)
 - submitter phone number field [106](#)
 - customer verification section [118](#), [142](#)
 - date to submit [423](#)

- sub-capacity report (*continued*)
 - Detail LPAR Data section
 - 2nd Highest field [120](#)
 - Date/Time field [120](#)
 - Highest field [120](#)
 - Hour Count field [120](#)
 - Detail LPAR Usage Data section
 - Total MSU Consumed field [121](#)
 - Detailed Interval Data section [406](#)
 - environments [447](#)
 - Getting Started Sub-Capacity Pricing Assigned Hours section [214](#), [224](#), [245](#), [266](#), [276](#)
 - Getting Started Sub-Capacity Pricing Assigned section [179](#), [180](#), [214](#), [224](#), [246](#), [266](#), [276](#)
 - Getting Started Sub-Capacity Pricing Equal to Traditional Sub-Capacity section [179](#), [180](#)
 - Getting Started Sub-Capacity Pricing Not Assigned section [180](#)
 - Group Capacity LPARs section [124](#), [125](#), [149](#), [181](#), [197](#)
 - handling by IBM [425](#)
 - handling unusual situations [17](#)
 - Missing CPC Data Details section
 - Detected field [115](#)
 - first column [115](#)
 - Justification for low data collection field [115](#)
 - Resolved field [115](#)
 - Missing LPAR Data Details section
 - Date/Time field [114](#)
 - Detected field [113](#)
 - Est Peak MSU field [114](#)
 - first column [113](#)
 - Justification for low data collection field [114](#)
 - LPAR/Group contribution to estimated peak field [114](#)
 - Multiplex Peak MSU field [114](#)
 - OS Product ID field [114](#)
 - OS Product Name field [114](#)
 - Resolved field [114](#)
 - Mobile/zCap Pricing Detailed Data section [216](#), [217](#), [226–228](#), [247–250](#), [268](#), [269](#), [278–280](#)
 - MSU Consumption Data Collection section [387](#)
 - preparing to submit to IBM [423](#)
 - Product Features section [115](#), [116](#)
 - Product Grid Snapshot section [123](#), [124](#), [130](#), [141](#), [144](#), [149](#), [152](#), [157](#), [161](#), [164](#), [179](#), [197](#), [214](#), [224](#), [245](#), [266](#), [276](#), [293](#), [297](#), [326](#), [330](#), [332](#), [336](#), [390](#)
 - Product Max Contributors section
 - Date/Time field [122](#)
 - group names and values fields [122](#)
 - Highest field [122](#)
 - LPAR names and values fields [122](#)
 - Product ID field [121](#)
 - Product Name field [121](#)
 - Product Summary Information section
 - Customer Comments field [111](#)
 - customer MSUs field [191](#)
 - Customer MSUs field [111](#), [176](#)
 - footnotes field [191](#), [210](#), [220](#), [241](#), [263](#), [272](#)
 - Footnotes field [177](#)
 - IPLA Product ID field [111](#)
 - IPLA Product Name field [111](#)
 - MLC Product ID field [110](#)
 - MLC Product Name field [110](#)
 - tool MSUs field [191](#), [210](#), [220](#), [241](#), [263](#), [272](#)

sub-capacity report (*continued*)

- Product Summary Information section (*continued*)
 - Tool MSUs field [111](#), [176](#)
 - sample for Country Multiplex Pricing [347](#)
 - sample for Getting Started Sub-Capacity Pricing [174](#)
 - sample for IBM Z Collocated Application Pricing [239](#)
 - sample for INPUTCSV-based IBM Z Workload Pricing for Cloud [261](#)
 - sample for INPUTCSV-based Mobile Workload Pricing [209](#)
 - sample for Integrated Workload Pricing [189](#)
 - sample for native z/OS, only zNALC workloads [153](#)
 - sample for native z/OS, traditional and zNALC workloads [145](#)
 - sample for single-version z/OS native systems [105](#)
 - sample for WLM-based IBM Z Workload Pricing for Cloud [271](#)
 - sample for WLM-based Mobile Workload Pricing [219](#)
 - sample for z/OS and z/TPF systems [294](#)
 - sample for z/OS and z/VSE systems [333](#)
 - sample for z/OS V1 and z/OS V2 native systems [128](#)
 - sample for z/OS V1 and z/OS V2 traditional and zNALC workloads [149](#)
 - sample for z/OS V1.6 guest systems [139](#)
 - sample for z/OS V1.6 guests, traditional and zNALC workloads, different LPARs [154](#)
 - sample for z/OS V1.6 guests, traditional and zNALC workloads, same LPARs [157](#)
 - sample for z/OS V1.7 and later guest systems [141](#)
 - sample for z/OS V1.7 guests, traditional and zNALC workloads [161](#)
 - sample for z/TPF systems [290](#)
 - sample for z/VSE systems only [323](#)
 - sample for z/VSE V4 and V5 systems [327](#)
 - sample for z/VSE V6, V5, and V4 systems [330](#)
 - SCRT Softcap Enforcement section [116](#)
 - SMF/SCRT89 Input Data Statistics section [119](#), [140](#), [143](#), [148](#), [156](#), [160](#), [163](#), [178](#), [196](#), [211](#), [221](#), [242](#), [264](#), [273](#), [293](#), [296](#), [325](#), [329](#), [332](#), [335](#), [389](#)
 - Special Conditions section
 - example [108](#), [110](#)
 - Exclude statements [108](#)
 - for Integrated Workload Pricing [110](#)
 - INCLUDE statements [109](#)
 - Sub-Capacity Capping Enforced by SCRT section [124](#), [131](#), [141](#), [144](#), [149](#), [152](#), [157](#), [161](#), [164](#), [178](#), [179](#), [196](#), [197](#), [297](#)
 - Sub-hour Data Collection Details section [125](#)
 - submitting to IBM [427](#)
 - submitting via the web [427](#)
 - System Recovery Boost Usage Details section [126](#)
 - System Recovery Boost Usage section [117](#)
 - Tool Information section
 - Reporting Period field [107](#)
 - Tool Release field [107](#)
 - viewing as a spreadsheet file [420](#)
 - zIIP Usage Data section [117](#)
- sub-capacity report for Country Multiplex Pricing
 - sample configuration [345](#)
- sub-capacity report for Getting Started Sub-Capacity Pricing
 - sample configuration [173](#)
 - sample report [174](#)
- sub-capacity report for guest systems
 - sub-capacity report for guest systems (*continued*)
 - billing-related section [142](#)
 - customer verification section [142](#)
 - sub-capacity report for IBM Z Collocated Application Pricing
 - sample configuration [237](#)
 - sample report [239](#)
 - sub-capacity report for INPUTCSV-based IBM Z Workload Pricing for Cloud
 - sample configuration [260](#)
 - sample report [261](#)
 - sub-capacity report for INPUTCSV-based Mobile Workload Pricing
 - sample configuration [207](#)
 - sample report [209](#)
 - sub-capacity report for Integrated Workload Pricing
 - sample configuration [188](#)
 - sample report [189](#)
 - sub-capacity report for native z/OS systems
 - Capacity Changes section [116](#)
 - customer verification section [118](#)
 - Detail LPAR Data section [119](#)
 - Detail LPAR Usage Data section [120](#)
 - Group Capacity LPARs section [125](#)
 - Missing CPC Data Details section [114](#)
 - Missing LPAR Data Details section [112](#)
 - Product Features section [115](#)
 - Product Grid Snapshot section [123](#)
 - Product Max Contributors section [121](#)
 - sample defined capacity value used section [123](#)
 - sample Detail LPAR Data section [120](#)
 - sample Detail LPAR Usage Data section [121](#)
 - sample Group Capacity LPARs section [124](#)
 - sample Product Grid Snapshot section [124](#)
 - sample single-version configuration [104](#)
 - sample Special Conditions section [108](#), [109](#)
 - SCRT Softcap Enforcement section [116](#)
 - Sub-Capacity Capping Enforced by SCRT section [124](#)
 - Sub-hour Data Collection Details section [125](#)
 - System Recovery Boost Usage Details section [126](#)
 - System Recovery Boost Usage section [117](#)
 - zIIP Usage Data section [117](#)
 - sub-capacity report for native z/OS, only zNALC workloads
 - sample report [153](#)
 - sub-capacity report for native z/OS, traditional and zNALC workloads
 - sample report [145](#)
 - sub-capacity report for single-version z/OS native systems
 - sample report [105](#)
 - sub-capacity report for WLM-based IBM Z Workload Pricing for Cloud
 - sample configuration [269](#)
 - sample report [271](#)
 - sub-capacity report for WLM-based Mobile Workload Pricing
 - sample configuration [217](#)
 - sample report [219](#)
 - sub-capacity report for z/OS and z/TPF systems
 - sample report [294](#)
 - sub-capacity report for z/OS V1 and z/OS V2 native systems
 - sample report [128](#)

- sub-capacity report for z/OS V1 and z/OS V2 systems
 - sample configuration [127](#)
- sub-capacity report for z/OS V1 and z/OS V2 traditional and zNALC workloads
 - sample report [149](#)
- sub-capacity report for z/OS V1.6 and earlier guests
 - interpreting [141](#)
- sub-capacity report for z/OS V1.6 guest systems
 - sample report [139](#)
- sub-capacity report for z/OS V1.6 guests, traditional and zNALC workloads, different LPARs
 - sample report [154](#)
- sub-capacity report for z/OS V1.6 guests, traditional and zNALC workloads, same LPARs
 - sample report [157](#)
- sub-capacity report for z/OS V1.7 and later guest systems
 - sample report [141](#)
- sub-capacity report for z/OS V1.7 guests, traditional and zNALC workloads
 - sample report [161](#)
- sub-capacity report for z/TPF systems
 - sample report [290](#)
- sub-capacity report for z/VSE
 - sample report or z/OS and z/VSE systems [333](#)
 - sample report or z/VSE systems only [323](#)
 - sample report or z/VSE V4 and V5 systems [327](#)
 - sample report or z/VSE V6, V5, and V4 systems [330](#)
- sub-capacity reporting
 - Country Multiplex Pricing [343](#)
 - for System z New Application License Charges (zNALC) [145](#)
 - for z/TPF systems [281](#)
 - for z/VM guest systems [133](#)
 - for z/VSE systems
 - Country Multiplex Pricing [336](#)
 - Getting Started Sub-Capacity Pricing [165](#)
 - IBM Z Collocated Application Pricing [229](#)
 - IBM Z Workload Pricing for Cloud [251](#)
 - Integrated Workload Pricing [183](#)
 - Mobile Workload Pricing [199](#)
 - multiple z/OS versions [127](#)
 - native z/OS systems [103](#)
 - products with non-MSU reporting metrics
 - metric types [371](#)
 - sample sub-capacity report [372](#)
 - SCRT externals for [371](#)
 - Tailored Fit Pricing for IBM Z [375](#)
 - Tailored Fit Pricing for IBM Z - Software Consumption Solution [401](#)
 - z/TPF Transformation Engines Offering [299](#)
- Sub-Capacity Reporting Tool (SCRT)
 - customizing JCL [25](#)
 - for Windows and Linux
 - installing in console mode [71](#)
 - obtaining [70](#)
 - requirements [69](#)
 - using the CLI [90](#)
 - using the GUI [73](#)
 - installing on z/OS [23](#)
 - limitations [439](#)

- Sub-Capacity Reporting Tool (SCRT) (*continued*)
 - messages [67](#)
 - obtaining for z/OS [21](#)
 - overview [3](#)
 - prerequisites [11](#)
 - requirements for z/OS [21](#)
 - return codes [448](#)
 - running on z/OS [24](#)
 - using on z/OS systems [21](#)
- Sub-hour Data Collection Details section of sub-capacity report [125](#)
- submission date [423](#)
- submitter email address field
 - customer information section [105](#)
- submitter name field
 - customer information section [105](#)
- submitter phone number field
 - customer information section [106](#)
- submitting sub-capacity reports
 - date required [423](#)
 - preparing [423](#)
 - via the web [427](#)
- summary of changes
 - for SCRT 29.1.0 [521](#)
 - for SCRT 30.1 xix
 - for V27.1.0 [526](#)
 - for V27.2.0 [525](#)
 - for V28.1.0 [523](#)
 - for V28.2.0 [522](#)
- support, obtaining [439](#)
- SUPPRESS_REPORT_SECTION statement [40](#), [58](#)
- SVC, See Single Version Charging
- SYSID [453](#)
- Sysplex Timer offset [14](#), [15](#)
- system ID [453](#)
- system identities, reported by SCRT
 - z/OS guest systems [138](#)
 - z/OS systems [113](#), [115](#)
 - z/TPF systems [282](#)
 - z/VSE systems [308](#)
- system identities, z/OS native systems [103](#)
- System Management Facility (SMF)
 - extracting records [13](#)
 - input for SCRT [4](#), [12](#)
 - records from different time zones [14](#)
 - type 70 subtype 1 and type 89 records [4](#), [12](#)
 - using [12](#)
- system message IWM063I, z/OS [442](#)
- System Recovery Boost Usage Details section of sub-capacity report [126](#)
- System Recovery Boost Usage section of sub-capacity report [117](#)
- System z Application Assist Processor (zAAP) [111](#)
- System z Integrated Information Processor (zIIP) [111](#)
- System z New Application License Charges (zNALC)
 - sample sub-capacity report
 - native z/OS, only z/NALC workloads [153](#)
 - native z/OS, traditional and z/NALC workloads [145](#)
 - z/OS V1.6 guests, traditional and z/NALC workloads, different LPARs [154](#)
 - z/OS V1.6 guests, traditional and z/NALC workloads, same LPARs [157](#)
 - z/OS V1.7 guests, traditional and z/NALC workloads [161](#)

System z New Application License Charges (zNALC) (continued) **U**

understanding [145](#)

z/OS guest systems [153](#)

T

Tailored Fit Pricing for IBM Z

interaction with traditional sub-capacity [377](#)

ISV reporting [381](#)

overview [375](#)

requirements [375](#)

sample report

control statement input [402](#)

sample sub-capacity report for collocated solution

billing-related section [382](#)

customer verification section [388](#)

sample sub-capacity report for standalone solution [393](#)

SCRT example for a collocated solution [381](#)

SCRT example for a standalone solution [392](#)

SCRT externals for [378](#)

SCRT sample configuration for collocated solution [381](#)

SCRT sample configuration for standalone solution [392](#)

SCRT sample control statement input for collocated solution [381](#)

SCRT sample control statement input for standalone solution [393](#)

Tailored Fit Pricing for IBM Z - Hardware Consumption Solution [405](#)

Tailored Fit Pricing for IBM Z - Software Consumption Solution

overview [401](#)

requirements [401](#)

sample report

configuration [402](#)

SCRT externals for [401](#)

technical support, obtaining [439](#)

terms and conditions for sub-capacity pricing

sending TSAD data [514](#)

time changes in LPARs

Sysplex Timer offset [15](#)

virtual TOD clock (VTOD) [15](#)

time zones [14](#), [119](#)

Tool Information section of sub-capacity report [107](#), [147](#), [155](#), [159](#), [162](#), [175](#), [190](#), [209](#), [219](#), [240](#), [262](#), [271](#), [291](#), [295](#), [324](#), [328](#), [331](#), [334](#), [383](#)

tool MSUs field

Product Summary Information section [191](#), [210](#), [220](#), [241](#), [263](#), [272](#)

Tool MSUs field

Product Summary Information section [111](#), [176](#)

Tool Release field

Tool Information section [107](#)

Total MSU Consumed field

Detail LPAR Usage Data section [121](#)

total utilization

for native systems [120](#)

traditional sub-capacity pricing [8](#)

Transmit System Availability Data (TSAD) [514](#)

troubleshooting [439](#)

type 70 subtype 1 and type 89 SMF records [4](#), [12](#)

unusual situations [17](#)

UPDATE CONTAINER control statement [59](#)

UPDATE CONTAINER statement [40](#)

UPDATE SOLUTION_ID control statement [59](#)

UPDATE SOLUTION_ID statement [41](#)

upgrades, machine software model [441](#)

uploading a new sub-capacity report [429](#)

UTC (Coordinated Universal Time) [14](#)

UTC time, used on EXCLUDE control statement [47](#)

UTC_Time option, REPORT_TIME control statement [57](#)

utilization capacity [8](#)

utilization, 4-hour rolling average [16](#)

utilization, maximum

for native systems [119](#)

utilization, total

for native systems [120](#)

V

Variable WLC (VWLC) [5](#)

viewing and editing the sub-capacity report

as a spreadsheet [420](#)

considerations for [419](#)

placeholder fields [420](#)

viewing the sub-capacity report

on the Internet [429](#)

virtual products [112](#)

virtual TOD clock (VTOD) [15](#)

VM guests

JCL input for [67](#)

VMGUESTS DD statement [67](#)

VSE[®], 21st Century Software [3](#), [307](#)

VWLC [5](#)

W

Wait Complete=Yes [441](#)

web interface for submitting sub-capacity reports [427](#)

WebSphere Transformation Extender V8 (5655-R95),

considerations for

NO89 considerations [61](#)

Workload Classification Supplement [409](#)

Workload License Charges (WLC) [5](#)

Z

z/OS and z/TPF systems, sample sub-capacity report [294](#)

z/OS and z/VSE systems

sample multiplex report [336](#)

z/OS and z/VSE systems, sample sub-capacity report [333](#)

z/OS native systems (single-version), sample sub-capacity report [105](#)

z/OS native systems, identities for [103](#)

z/OS native systems, only zNALC workloads, sample sub-capacity report [153](#)

z/OS native systems, sample single-version configuration [104](#)

z/OS native systems, traditional and zNALC workloads, sample sub-capacity report [145](#)

z/OS system message IWM063I [442](#)

z/OS systems

- z/OS systems (*continued*)
 - using SCRT on [21](#)
- z/OS under z/VM
 - JCL input for [67](#)
 - SCRT limitations [440](#)
- z/OS Unreduced virtual product [112](#)
- z/OS V1 and z/OS V2 native systems, sample sub-capacity report [128](#)
- z/OS V1 and z/OS V2 systems, sample configuration [127](#)
- z/OS V1 and z/OS V2 traditional and zNALC workloads, sample sub-capacity report [149](#)
- z/OS V1.6 (and earlier) guest systems
 - interpreting the sub-capacity report [141](#)
- z/OS V1.6 guest systems, sample sub-capacity report [139](#)
- z/OS V1.6 guests, traditional and zNALC workloads, different LPARS, sample sub-capacity report [154](#)
- z/OS V1.6 guests, traditional and zNALC workloads, same LPARS, sample sub-capacity report [157](#)
- z/OS V1.7 and later guest systems, sample sub-capacity report [141](#)
- z/OS V1.7 guests, traditional and zNALC workloads, sample sub-capacity report [161](#)
- z/OS-based IPLA licensing [6](#)
- z/TPF
 - LPAR naming requirements [281](#)
 - sub-capacity reporting using SCRT89 records
 - data collection and processing [284](#)
 - multiple sub-capacity operating systems [284](#)
 - z/TPF and non-z/TPF in same LPAR [283](#)
 - sub-capacity reporting without SCRT89 records [282](#)
 - unusual situations involving [17](#)
- z/TPF Recoup Utility [281](#), [289](#)
- z/TPF systems, sample sub-capacity report [290](#)
- z/TPF Transformation Engines Offering
 - externals
 - control statement [301](#)
 - Customer Information section [301](#)
 - Product Grid Snapshot [303](#)
 - Product Max Contributors section [302](#)
 - Product Summary Information section [302](#)
 - Special Conditions section [301](#)
 - qualifications for [299](#)
 - restriction [300](#)
- z/TPF_Transformation_Engines_ Offering statement [60](#)
- z/TPF_TRANSFORMATION_ENGINES_OFFERING [41](#)
- z/VM guest systems
 - JCL for [67](#)
 - MSU values for mixed configurations [133](#)
 - sample sub-capacity report for z/OS V1.6 guests [139](#)
 - sample sub-capacity report for z/OS V1.7 and later guests [141](#)
- z/VSE
 - CMT files [310](#)
 - GETVIS [310](#)
 - NO89 control statements
 - customizing [320](#)
 - NO89 file [309](#)
 - NO89 values
 - example [321](#)
 - PARMS control statements
 - customizing [320](#)
 - preparing for CMT [313](#)
 - processing SCRT89 records on Windows or Linux [320](#)
- z/VSE (*continued*)
 - processing SCRT89 records on
 - z/OS
 - overview [314](#)
 - skeleton [310](#)
 - SPECIAL control statements
 - customizing [321](#)
 - sub-capacity pricing requirements [306](#)
 - transferring files to workstation [318](#)
 - transferring files to z/OS [318](#)
 - transferring SCRT89 records
 - steps [317](#), [318](#)
 - z/VSE systems only, sample sub-capacity report [323](#)
 - z/VSE systems, sample sub-capacity report [327](#), [330](#)
 - zAAP, IBM [16](#), [111](#)
 - zIIP Usage Data section of sub-capacity report [117](#)
 - zIIP, IBM [16](#), [111](#)
 - zNALC
 - sample sub-capacity report
 - native z/OS, only z/NALC workloads [153](#)
 - native z/OS, traditional and z/NALC workloads [145](#)
 - z/OS V1.6 guests, traditional and z/NALC workloads, different LPARs [154](#)
 - z/OS V1.6 guests, traditional and z/NALC workloads, same LPARs [157](#)
 - z/OS V1.7 guests, traditional and z/NALC workloads [161](#)
 - understanding [145](#)
 - z/OS guest systems [153](#)
 - zSeries Entry License Charges (zELC) [5](#)
 - zVSE
 - multiple versions on CPC [308](#)
 - SCRT89 record [309](#)



SC23-6845-22

