

z/OS
3.2

Tailored Fit Pricing for IBM Z



Note

Before using this information and the product it supports, read the information in [“Notices” on page 149](#).

This edition applies to IBM® z/OS® 3.2 (5655-ZOS) and to all subsequent releases and modifications until otherwise indicated in new editions.

Last updated: 2025-09-30

© **Copyright International Business Machines Corporation 2017, 2025.**

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

Contents

Figures.....	v
Tables.....	ix
About the Tailored Fit Pricing Content Solution.....	xi
How to provide feedback to IBM.....	xiii
What's new in Tailored Fit Pricing for IBM Z?.....	xv
Chapter 1. What is Tailored Fit Pricing for IBM Z?.....	1
Deploying workload-specific solutions.....	2
Getting started.....	3
Software Consumption Solution.....	3
Hardware Consumption Solution.....	3
Enterprise Capacity Solution.....	3
Application Development and Test Solution.....	4
New Application Solution.....	4
Obtain the solution ID from LMS.....	5
Scenario: New Application Solution.....	5
Chapter 2. Requirements for Tailored Fit Pricing for IBM Z.....	9
Chapter 3. Setting up a colocated Tailored Fit Pricing DevTest or New	
 Application solution.....	11
WLM definitions for a colocated Tailored Fit Pricing DevTest or New Application Solution.....	11
Some WLM basics for Tailored Fit Pricing.....	12
Using IBM Cloud Provisioning for Tailored Fit Pricing.....	17
Using z/OSMF to define WLM Tailored Fit Pricing definitions.....	18
Using ISPF to define WLM Tailored Fit Pricing definitions.....	18
Using the Definition Menu.....	18
Working with tenant resource groups.....	19
Working with tenant report classes.....	20
WLM application messages.....	20
Using system command RESET to change the service class of work when associated with a tenant	
resource group.....	21
Chapter 4. Using SCRT for Tailored Fit Pricing.....	23
SCRT commands for the Software Consumption Solution.....	23
SCRT commands for the Hardware Consumption Solution.....	23
SCRT commands for a dedicated DevTest or New Application solution.....	23
SCRT commands for a colocated DevTest or New Application solution.....	24
Dedicated LPARs.....	24
Using the SOLUT parameter.....	24
Reporting for Tailored Fit Pricing.....	24
Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution.....	25
Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution.....	27
Workload Classification Supplement.....	30

Reporting for Tailored Fit Pricing for IBM Z.....	33
SPECIAL DD statement.....	55
Using SCRT on Windows and Linux systems.....	59
Using the SCRT graphical user interface.....	59
Using the SCRT command line interface.....	76
Chapter 5. Using workload reporting services to collect Container Services	
performance data.....	79
When to use the workload reporting services.....	79
Interpreting report class data.....	79
Using the IWM4QTNT service.....	80
IWM4QTNT — Query tenant resource group consumption.....	80
Mapping a service definition.....	86
Querying the active classification rules.....	87
Structure of the XML service definition (DTD).....	87
Chapter 6. MVS System Management Facilities (SMF) record type changes for	
Tailored Fit Pricing.....	93
Record Type 70 (X'46') — RMF Processor Activity.....	93
Record mapping.....	93
Subtype 1 — CPU, PR/SM, and ICF activity.....	95
Record Type 72 (X'48') — Workload Activity, Storage Data, and Serialization Delay.....	103
Subtype 3 — Workload activity.....	104
Record Type 79 (X'4F') — RMF Monitor II activity.....	107
Subtype 1 — Address Space State Data.....	107
Subtype 2 — Address Space Resource Data.....	112
Subtype 5 — Address Space SRM Data.....	117
Record type 89 (X'59') — Usage Data.....	118
Record environment.....	119
Record mapping.....	120
Record Type 99 (X'63') — System Resource Manager Decisions.....	133
Record type 1156 (X'484') — Module fetch.....	138
Record mapping.....	139
Chapter 7. Using RMF with colocated DevTest and New Application Solutions.....	143
CPU Activity - SMF record type 70-1.....	144
Appendix A. Accessibility.....	147
Notices.....	149
Terms and conditions for product documentation.....	150
IBM Online Privacy Statement.....	151
Policy for unsupported hardware.....	151
Minimum supported hardware.....	151
Trademarks.....	152
Index.....	153

Figures

1. Tailored Fit Pricing lets you place solutions as required by business needs.....	2
2. Tailored Fit Pricing: Interaction of solution, SMF, RMF, and SCRT.....	6
3. Tailored Fit Pricing data recorded on behalf of a single tenant resource group.....	7
4. Definition Menu panel.....	19
5. Create a Tenant Resource Group panel.....	20
6. Create a Tenant Report Class.....	20
7. Sample report for Software Consumption Solution: Section B5.....	27
8. Sample report for Software Consumption Solution: CPS1 header section.....	27
9. Sample report for Software Consumption Solution: CPS2 header section.....	27
10. Sample SCRT report: CPC MSU Rate Above Permanent Capacity Rating section (I5).....	28
11. Sample SCRT report: Detailed Interval Data section (V9).....	29
12. Example of CPSTRACE output (formatted).....	37
13. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Customer information and Tool Information sections.....	39
14. 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.....	40
15. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Active Containers section.....	41
16. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Billing Summary section.....	41
17. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Active Products section.....	42
18. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: MSU Consumption Data Collection section.....	43
19. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Interval Consumption section.....	43

20. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Product Metrics section.....	44
21. 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.....	45
22. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Product Max Contributors section.....	45
23. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Product Grid Snapshot section.....	46
24. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Product Grid Snapshot section.....	46
25. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Max Contributors section.....	47
26. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container MSU Consumption Contributors section.....	48
27. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Multiplex summary section.....	50
28. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Machine summary section.....	50
29. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Machine summary section for CPC M1C1.....	52
30. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Active Containers section.....	53
31. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Container Billing Summary and Container Active Products sections.....	53
32. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Customer verification section for CPC M1C1 - Machine information.....	54
33. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Product Max Contributors and Product Grid Snapshot sections.....	54
34. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Container Product Grid Snapshot section.....	55
35. SCRT for Windows and Linux: Customer parameter panel.....	61
36. SCRT for Windows and Linux: SMF or SCRT89 data selection panel.....	62
37. SCRT for Windows and Linux: Customer number panel.....	63
38. SCRT for Windows and Linux: Reporting period panel.....	64

39. SCRT for Windows and Linux: NO89 product selection panel — NO89 Product List tab.....	65
40. SCRT for Windows and Linux: NO89 product selection panel — Product Assignment tab.....	66
41. SCRT for Windows and Linux: Input CSV tab	67
42. SCRT for Windows and Linux: Processing Options tab	68
43. SCRT for Windows and Linux: Detected solution IDs panel.....	69
44. SCRT for Windows and Linux: Container settings assignment panel.....	70
45. SCRT for Windows and Linux: Exclude panel	71
46. SCRT for Windows and Linux: Exclude panel — Set Exclude options.....	72
47. SCRT for Windows and Linux: Exclude panel — Set Exclude options for MSU consumption.....	73
48. SCRT for Windows and Linux: Report preview panel	74
49. SCRT for Windows and Linux: Trace Options tab on the report preview panel	75
50. SCRT for Windows and Linux: Report results panel	76
51. Mixed-class-indication timestamp in relation to the time interval.....	80
52. RMF report class report showing a tenant report class.....	143
53. WLMGL Report - Service Policy Page showing tenant resource groups.....	143

Tables

1. SCRT SPECIAL control statements..... 55

2. SCRT command line options..... 77

3. Using IWM4QTNT with the workload reporting services..... 80

4. Return and Reason Codes for the IWM4QTNT Macro..... 84

5. Valid name spaces and corresponding functionality levels..... 90

6. CPU Activity - Conditions Based on SMF Record Type 70–1..... 144

About the Tailored Fit Pricing Content Solution

Purpose of this information This information is a collection of all of the information that you need to understand and exploit Tailored Fit Pricing for IBM Z®. Some of the information contained in this content solution also exists elsewhere in the z/OS library.

Who should read this information This information is intended for system programmers who are responsible for capacity planning, and who are responsible for MVS workload management. There is also information for people writing programs or subsystems that monitor performance.

Related information

For an interactive start and technical resources such as articles and workflows, see [Tailored Fit Pricing for IBM Z](http://www.ibm.com/support/z-content-solutions/tailored-fit-pricing/) (www.ibm.com/support/z-content-solutions/tailored-fit-pricing/).

To find the complete z/OS library, go to [IBM Documentation](http://www.ibm.com/docs/en/zos) (www.ibm.com/docs/en/zos).

How to provide feedback to IBM

We welcome any feedback that you have, including comments on the clarity, accuracy, or completeness of the information. For more information, see [How to send feedback to IBM](#).

What's new in Tailored Fit Pricing for IBM Z?

July 2025 enhancements

Tailored Fit Pricing for IBM Z is updated as follows:

- A new Tailored Fit Pricing for Hardware (TFP Hardware) solution, TFP Hardware for zIIP workloads, is added, which is available exclusively for IBM z17®.
- Information about the SCRT Workload Classification Supplement report is added.

May 2021 enhancements

Tailored Fit Pricing for IBM Z is updated as follows:

- The Hardware Consumption Solution is added.
- Enterprise Consumption Solution is renamed Software Consumption Solution.
- The option to specify a solution ID using the **SOLUT** system parameter, and a corresponding SMF89SolutionID field in System Management Facility (SMF) record type 89, are added for APAR OA60198.

Chapter 1. What is Tailored Fit Pricing for IBM Z?

Tailored Fit Pricing delivers simplicity, transparency, and predictability of pricing. It offers these comprehensive alternatives to the rolling 4-hour (R4HA) average:

- Software Consumption Solution (formerly known as Enterprise Consumption Solution), which is a cloud-like usage-based licensing model. Compute is measured on the basis of MSUs consumed, which removes the need for manual or automated capping. You can configure your systems to support optimal response times and service level agreements.
- Hardware Consumption Solution, which delivers flexible, cloud-like pricing that enables you to leverage additional always-on hardware capacity for a range of use cases.

There are 3 Tailored Fit Pricing for Hardware (TFP Hardware) solutions, the first 2 of which are specific to z/OS:

TFP Hardware for General Purpose CP workloads

Delivers a flexible, cloud-like pricing solution that helps you adapt to disruptive events, manage changing workload dynamics, and support organic workload growth. This solution also requires a Software Consumption Solution.

TFP Hardware for zIIP workloads

Enables you to cost-effectively design and deploy new workloads (such as AI inference, OpenShift, and zCX) at scale on z/OS with workload-optimized pricing, cloud-like flexibility, and no long-term commitment. Unlike TFP Hardware for General Purpose CP workloads, TFP Hardware for zIIP workloads does not require a Software Consumption Solution.

TFP Hardware for Linux® workloads

Delivers the same benefits as TFP Hardware for General Purpose CP workloads, but for LinuxONE and Linux on IBM Z.

The TFP Hardware solutions can be deployed together, but they do not require one another; each could also be deployed as a standalone solution. In other words, the TFP Hardware solutions are independent of, but compatible with, one another.

- Enterprise Capacity Solution, which is a full-capacity licensing model that offers maximum predictability and simplicity.

It also offers these solutions for specific workloads that can coexist alongside both traditional and the new enterprise models:

- Application Development and Test Solution (DevTest). which, by removing the need for aggressive cost controls around development and test, promotes a healthy development and test environment on z/OS
- New Application Solution, which allows new applications to be tightly integrated with colocated workloads, with the price predictability of a dedicated environment.

All of these solutions dramatically simplify pricing and deliver flexible deployment options that are tailored to reflect your environment. They allow you to embrace the best technical fit, greatly reducing the need to architect for software costs.

Technology in z/OS provides the framework to enable Tailored Fit Pricing. It provides the capability to meter and report on specific workloads in a similar manner, regardless of the solution deployed. Manual tagging and tracking or other increased monthly overhead is not required.

Requirements

For a summary of the conditions you must meet to implement a Tailored Fit Pricing solution, see [Chapter 2, “Requirements for Tailored Fit Pricing for IBM Z,” on page 9.](#)

Related resources

For an interactive start and technical resources such as articles and workflows, see [Tailored Fit Pricing for IBM Z \(www.ibm.com/support/z-content-solutions/tailored-fit-pricing/\)](https://www.ibm.com/support/z-content-solutions/tailored-fit-pricing/).

Deploying workload-specific solutions

When associated with a Tailored Fit Pricing for workload-specific solution such as an Application Development and Test Solution, SMF and the Sub-Capacity Reporting Tool (SCRT) use Tailored Fit Pricing to understand the resource consumption of the solution. This allows SCRT to remove the direct impact of these solutions from the rolling four-hour average of the LPARs where it runs. SCRT also provides an isolated view of the container environment, including the products in use.

With Tailored Fit Pricing for IBM Z, there is no need to isolate a solution to a separate LPAR for pricing purposes and you are not required to build spreadsheets manually to track specific workloads. Tailored Fit Pricing gives you the freedom to deploy a solution where it makes the most sense based on a technical evaluation.

After an initial set up is performed, the SMF data that SCRT uses contains all the information required for Tailored Fit Pricing. There is no need for any complex data collection and analysis, even for cases where the new solution is colocated with other workloads.

You can deploy the Tailored Fit Pricing workload-specific solutions solutions as follows.

- New Application Solution. Add new approved z/OS solutions either:
 - Colocated in an existing LPAR with other unrelated workloads without directly impacting R4HA
 - To a dedicated LPAR.

For a colocated solution, you can measure the MSU used for the Tailored Fit Pricing colocated solution independently of the other workloads in the LPAR. Optionally, the consumption by the solution may be capped using a function similar to WLM Resource Group capping. The system records data about the colocated solution for accounting SCRT reporting purposes. SCRT analyzes the recorded data and produces a report for consumption by fulfillment systems, removing the workload's MSU utilization from the R4HA.

- Application Development and Test Solution (DevTest), for z/OS-based development and test workloads. The DevTest environment is typically on dedicated LPARs.

Figure 1 on page 2 illustrates the possible placement of workload-specific solutions solutions.

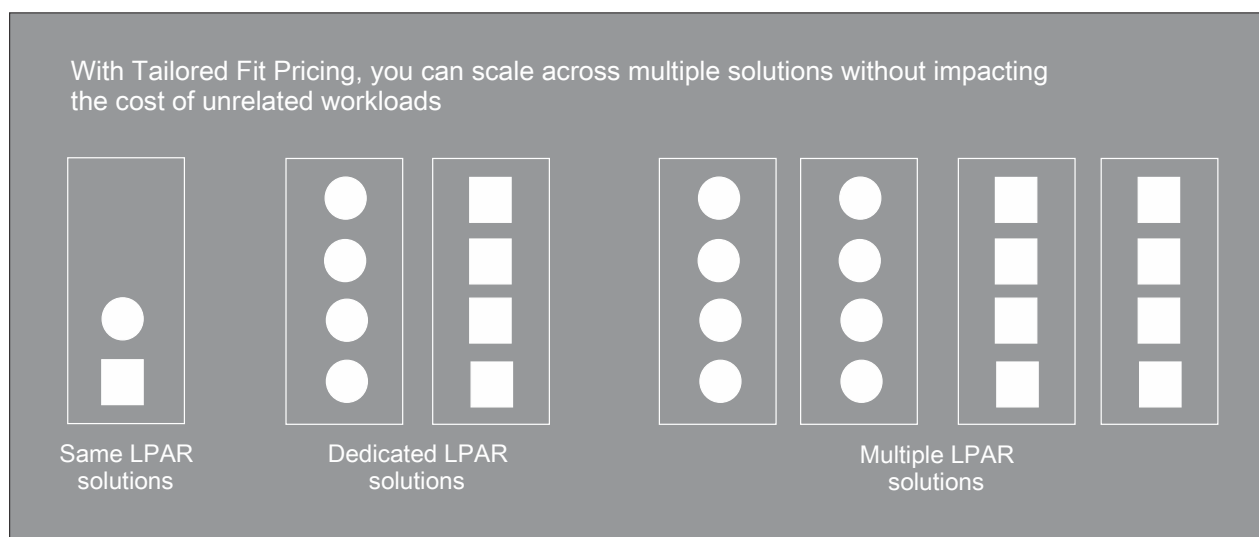


Figure 1. Tailored Fit Pricing lets you place solutions as required by business needs

Getting started

The setup required varies with the Tailored Fit Pricing solution.

Software Consumption Solution

The Software Consumption Solution is a cloud-like usage-based licensing model. Compute is measured on the basis of MSUs consumed, which removes the need for manual or automated capping.

Use the Software Consumption Solution in conjunction with the [“Application Development and Test Solution” on page 4](#) for a comprehensive end-to-end pricing solution for all stages of the application lifecycle.

Implement a Software Consumption Solution as follows:

1. Work with an IBM Sales representative to define the Software Consumption Solution. This includes separate solution IDs for your development and production environments.
2. Obtain the solution IDs from the License Management Support Web portal. For more information, see [“Obtain the solution ID from LMS” on page 5](#).
3. **Dedicated-LPAR solution:** Use either of the following methods to associate the solution IDs for the Software Consumption Solution and the Application Development and Test Solutions:
 - The SCRT **CONTAINER** command to specify only those LPARs that are part of the qualified solution. For more information, see [Chapter 4, “Using SCRT for Tailored Fit Pricing,” on page 23](#).
 - The **SOLUT** system parameter when IPLing a z/OS system as part of the qualified solution. **SOLUT** is available with the appropriate level of z/OS service. For more information, see [“Dedicated LPARs” on page 24](#).
4. Send the SCRT report to IBM. The solutions are reported separately.

Colocated-LPAR solution: Update your SCRT JCL to reflect the solution ID, which is generated as part of the agreement with IBM. For more information, see [“SCRT commands for the Software Consumption Solution” on page 23](#).

Hardware Consumption Solution

The Hardware Consumption Solution introduces cloud-like pricing for IBM Z hardware, with an always-on, subscription-based corridor of pay-for-use capacity, on top of purchased capacity, that can alleviate the impact of short, unpredictable spikes in business-critical workloads.

Implement a Hardware Consumption Solution as follows:

1. Work with an IBM Sales representative to define the Hardware Consumption Solution and the Software Consumption Solution that is required with it.

IBM uses a workload analysis to determine eligibility and fit for the Hardware Consumption Solution. For details on a Software Consumption Solution, see [“Software Consumption Solution” on page 3](#).
2. Update your SCRT JCL with control statements for interval rate data reporting. For more information, see [Chapter 4, “Using SCRT for Tailored Fit Pricing,” on page 23](#).
3. Send the SCRT report to IBM.

Enterprise Capacity Solution

The Enterprise Capacity Solution is a tailored full-capacity licensing model, offering predictable cost. Charges are connected to the size of the physical environment and are calculated based on the estimated mix of workloads running, with the flexibility to vary actual usage across workloads. Charges include increased capacity for development and test environments and reduced pricing for all types of workload growth. Charging based on the overall size of the physical environment removes the need for manual or

automated capping. You can configure your systems to support optimal response times and service level agreements.

To implement an Enterprise Capacity Solution, work with an IBM Sales representative to define a solution. The agreement includes an allowance for application development and test workloads.

Because this is a full-capacity solution, there is no reporting requirement.

Application Development and Test Solution

The Application Development and Test (DevTest) Solution, by removing the need for aggressive cost controls around development and test, promotes a healthy development and test environment on z/OS.

To implement a DevTest Solution, do the following.

1. Work with an IBM Sales representative to define a DevTest Solution. This includes separate solution IDs for your development and production environments. The agreement also includes the decision to use either an LPAR that is dedicated to the workload or an LPAR on which the workload is colocated with other workloads. The DevTest environment is typically on dedicated LPARs.
2. Obtain the solution IDs from the License Management Support Web portal. For more information, see [“Obtain the solution ID from LMS” on page 5](#).
3. **Dedicated-LPAR solution:** Use either of the following methods to associate the solution IDs with the LPAR or system:
 - The SCRT **CONTAINER** command to specify only those LPARs that are part of the qualified solution. For more information, see [Chapter 4, “Using SCRT for Tailored Fit Pricing,” on page 23](#).
 - The **SOLUT** system parameter when IPLing a z/OS system as part of the qualified solution. **SOLUT** is available with the appropriate level of z/OS service. For more information, see [“Dedicated LPARs” on page 24](#).

Colocated-LPAR solution: Create definitions in WLM and then update your SCRT JCL to reflect the solution ID, which is generated as part of the agreement with IBM. For more information, see [Chapter 3, “Setting up a colocated Tailored Fit Pricing DevTest or New Application solution,” on page 11](#).

For more information about updating the SCRT JCL, see [Chapter 4, “Using SCRT for Tailored Fit Pricing,” on page 23](#).

4. Send the SCRT report to IBM. The solutions are reported separately.

New Application Solution

The New Application Solution allows new applications to be tightly integrated with colocated workloads, with the price predictability of a dedicated environment.

To implement a New Application Solution, do the following.

1. Work with an IBM Sales representative to define a New Application Solution. This includes separate solution IDs for your development and production environments. The agreement also includes the decision to use either an LPAR that is dedicated to the workload or an LPAR on which the workload is colocated with other workloads.
2. Obtain the solution IDs from the License Management Support Web portal. For more information, see [“Obtain the solution ID from LMS” on page 5](#).
3. **Dedicated-LPAR solution:** Use either of the following methods to associate the solution IDs with the LPAR or system:
 - The SCRT **CONTAINER** command to specify only those LPARs that are part of the qualified solution. For more information, see [Chapter 4, “Using SCRT for Tailored Fit Pricing,” on page 23](#).
 - The **SOLUT** system parameter when IPLing a z/OS system as part of the qualified solution. **SOLUT** is available with the appropriate level of z/OS service. For more information, see [“Dedicated LPARs” on page 24](#).

Colocated-LPAR solution: Create definitions in WLM and then update your SCRT JCL to reflect the solution ID, which is generated as part of the agreement with IBM. For more information, see [Chapter 3, “Setting up a colocated Tailored Fit Pricing DevTest or New Application solution,”](#) on page 11.

For more information about updating the SCRT JCL, see [Chapter 4, “Using SCRT for Tailored Fit Pricing,”](#) on page 23.

4. Send the SCRT report to IBM. The solutions are reported separately.

Obtain the solution ID from LMS

Clients work with an IBM Sales representative to define a solution with an agreed upon price, which triggers the creation of a **Solution ID**. The IBM-provided Solution ID is a 64-character string identifying an approved workload. For any approved solution, the Solution ID ties the system environment to the solution.

The Solution ID is assigned by the IBM License Management Support (LMS) portal. Login to the License Management Support (LMS) (www.ibm.com/software/lms) to obtain the Solution ID. You also use the LMS portal to submit SCRT reports to IBM.

Scenario: New Application Solution

The following scenario provides a high-level illustration of how Tailored Fit Pricing helps the solution architect and system programmer, improving deployment decisions on IBM Z servers. It describes the actions to be taken as well as the roles of other elements, such as WLM, RMF, and SCRT, for a New Application Solution.

- [“Step 1: Work with IBM Sales to define a solution and get a solution ID”](#) on page 5
- [“Step 2: For a colocated solution, create Tailored Fit Pricing definitions in WLM”](#) on page 5
- [“Step 3: For a colocated solution, WLM, RMF and SMF track and record solution usage data”](#) on page 6
- [“Step 4: For either a colocated or dedicated solution, use SCRT to collect billing data and report to IBM”](#) on page 8.

Step 1: Work with IBM Sales to define a solution and get a solution ID

Work with an IBM Sales representative to define a New Application Solution with an estimated size of 100 MSUs. This agreement generates a solution ID that correlates the IBM sales contract with the actual solution, and purchased capacity is used for billing purposes. You can obtain the solution ID from the LMS portal.

Where do you go from here? Once you have worked with IBM Sales to design a solution, you'll know whether you are going to create a colocated solution, a dedicated solution, or a solution spanning multiple dedicated LPARs. Proceed as follows based on that decision:

- For a colocated solution, go to [“Step 2: For a colocated solution, create Tailored Fit Pricing definitions in WLM”](#) on page 5 and [“Step 3: For a colocated solution, WLM, RMF and SMF track and record solution usage data”](#) on page 6.
- For a solution on a dedicated LPAR or multiple dedicated LPARs, go directly to [“Step 4: For either a colocated or dedicated solution, use SCRT to collect billing data and report to IBM”](#) on page 8.

Step 2: For a colocated solution, create Tailored Fit Pricing definitions in WLM

For a colocated solution, the system programmer creates tenant resource groups and tenant report classes in the WLM service definition, using either the ISPF WLM administrative application or the z/OSMF WLM task, to identify the WLM-classified workloads that constitute the approved solution workload.

- The tenant resource group defines the address spaces and independent enclaves making up the solution and supports metering and optional capping of the workload's MSU utilization. When defining

the tenant resource group for Tailored Fit Pricing purposes, the system programmer must also include the 64-character solution ID obtained in [“Step 1: Work with IBM Sales to define a solution and get a solution ID”](#) on page 5 from the LMS web site. The solution ID is the link between the approved workload and SCRT’s evaluation, and exclusion from the rolling 4-hour average.

- The system programmer also defines one or more tenant report classes. Tenant report classes are similar to standard report classes. However, tenant report classes are assigned to a tenant resource group and thus provide the metering capability for the tenant resource group. You can configure the information that is reported on behalf of tenant resource groups by associating their tenant report classes with address spaces or enclaves using classification rules.

See the following for more information about setting up WLM:

- [“WLM definitions for a colocated Tailored Fit Pricing DevTest or New Application Solution”](#) on page 11
- [“Using z/OSMF to define WLM Tailored Fit Pricing definitions”](#) on page 18
- [“Using ISPF to define WLM Tailored Fit Pricing definitions”](#) on page 18

Step 3: For a colocated solution, WLM, RMF and SMF track and record solution usage data

Once the WLM service definition is activated with tenant resource group definitions and associated classification rules, WLM tracks the processor consumption on behalf of the tenant resource group-specified workload. RMF then writes each tenant resource group's consumption data to SMF Type 70, Subtype 1 records every recording interval. SCRT analyzes the recorded data and produces a report for consumption by IBM's fulfillment systems, removing the workload's MSU utilization from the R4HA.

The SMF data can also be used:

- By z/OS system programmers or performance analysts to monitor workload resource usage and performance
- For other metering purposes, such as chargeback related to the specific workload.

SMF records written by RMF contain a reference to the tenant resource group, which ties the processing back to the solution ID for SCRT processing.

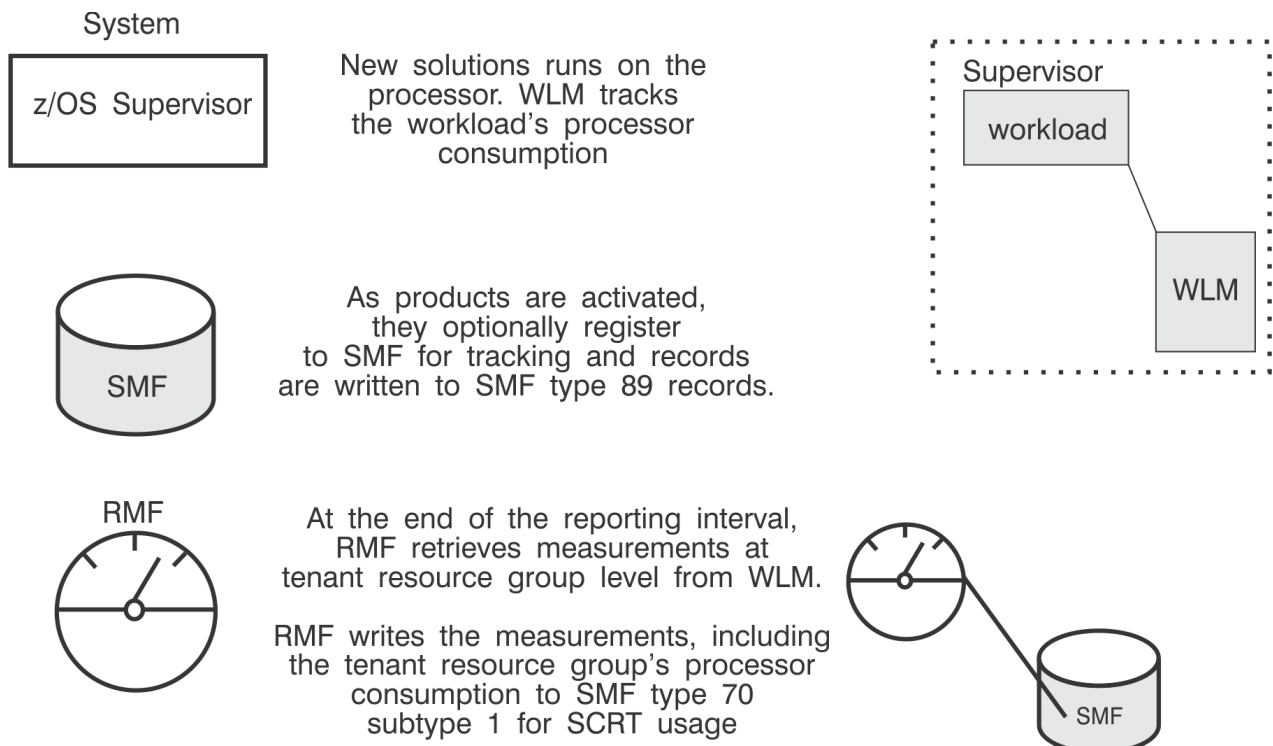


Figure 2. Tailored Fit Pricing: Interaction of solution, SMF, RMF, and SCRT

The system programmer can instruct WLM to limit (cap) processor consumption related to that workload, specified in the tenant resource group definitions. Similarly, the cap can be removed (or not specified), allowing processor consumption for that workload to continue unrestrained based on expected workload consumption, without the workload behavior modification that occurs when a CPU cap is reached.

Tailored Fit Pricing for IBM Z generally requires WLM work qualifiers that correspond to address space or independent enclaves. For specific details on qualification rules, see the terms and conditions of the Tailored Fit Pricing Solution.

Tailored Fit Pricing

Many IBM subsystems and sub-capacity products record their own CPU usage directly to SMF using IFAUSAGE or IFAEDREG, which is written as SMF Type 89 Subtypes 1 and 2 records, respectively. In both Type 70 and Type 89 records, data is written in tenant resource group segments, defining the relevant data for each tenant resource group. The example in [Figure 3 on page 7](#) shows the SMF data recorded on behalf of CICS regions and other information.

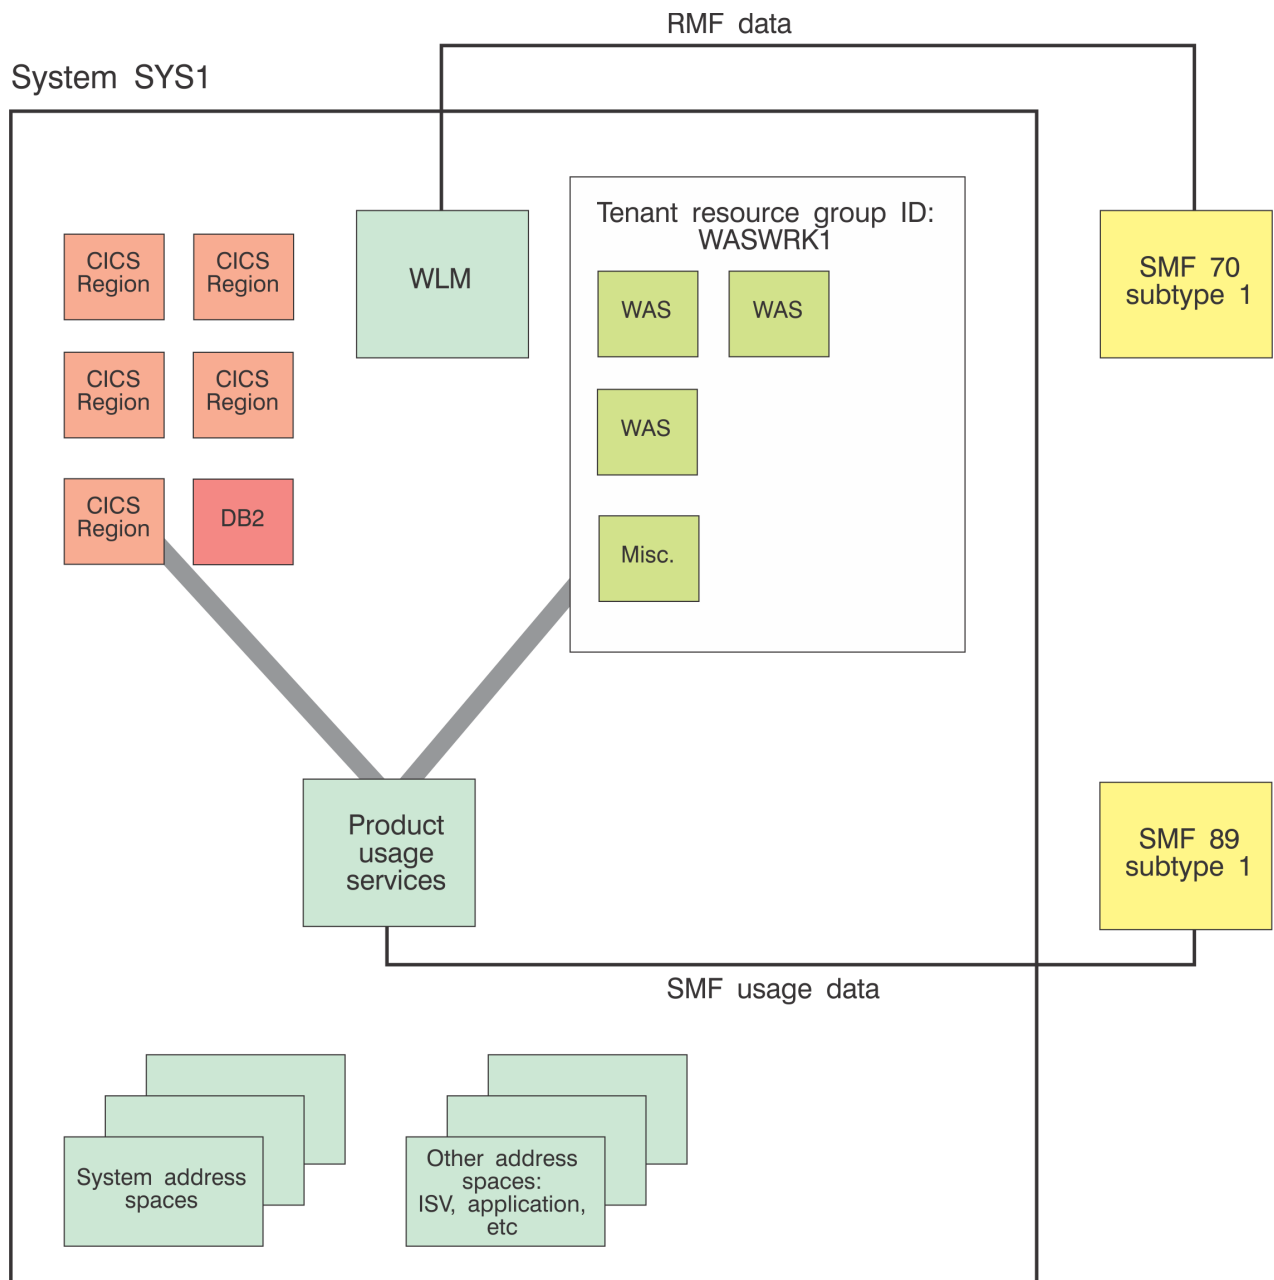


Figure 3. Tailored Fit Pricing data recorded on behalf of a single tenant resource group

For more information, see:

- [Chapter 7, “Using RMF with colocated DevTest and New Application Solutions,” on page 143](#)
- [Chapter 6, “MVS System Management Facilities \(SMF\) record type changes for Tailored Fit Pricing,” on page 93](#)

Step 4: For either a colocated or dedicated solution, use SCRT to collect billing data and report to IBM

To collect billing data and report back to IBM, update SCRT with the solution ID. (For a dedicated solution, you can specify the solution ID in parmlib instead.) When creating the next report, SCRT reads appropriate SMF records and reports MSU use for each container. The report includes the solution ID for all defined solutions and, as appropriate, data about the tenant resource groups and other identifying information for each product.

When the system programmer sends the SCRT report to IBM, the IBM uses the solution ID to correlate the workload with the entitled Tailored Fit Pricing solution for that client, validates the workload and its entitled capacity, and handles software billing based on the report.

z/OS clients run the SCRT job monthly to prepare a report of CPU consumed on behalf of workloads run on the system.

For colocated solutions, the CPU resource reported for the address spaces and independent enclaves that are defined to Tailored Fit Pricing are subtracted by SCRT from the LPAR 4- hour rolling average reported for billing purposes. This allows Tailored Fit Pricing workloads to be deployed to existing z/OS systems, colocated with existing workloads, without impacting the monthly license charges that are associated with z/OS and other middleware running on the same system.

For more information, see [Chapter 4, “Using SCRT for Tailored Fit Pricing,” on page 23](#).

Chapter 2. Requirements for Tailored Fit Pricing for IBM Z

The following sections describe the requirements for Tailored Fit Pricing:

- [“Requirements for Tailored Fit Pricing” on page 9](#)
- [“Requirements for using IBM Cloud Provisioning for Tailored Fit Pricing” on page 10](#)
- [“Requirements for using z/OSMF for Tailored Fit Pricing” on page 10](#)

Requirements for Tailored Fit Pricing

This topic outlines the technical requirements that must be met to implement a Tailored Fit Pricing solution. Some requirements are described in more detail in other topics. For additional offering requirements, see the Tailored Fit Pricing announcement letter.

The requirements for Tailored Fit Pricing vary with the solution. The specific requirements for a solution must be met before IBM can accept and process sub-capacity reports in which Tailored Fit Pricing solutions are reported. These requirements are in addition to any requirements that are described in a Tailored Fit Pricing solution announcement letter.

- 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.Pricing.Infrastructure fix category.

- Hardware requirements vary based on the offering.
 - The Hardware Consumption Solution, TFP Hardware for General Purpose CP workloads, is available for customers with IBM z15® or later hardware, z/OS general purpose processors, and a Software Consumption Solution. There can be no dedicated CPs on any LPARs, and WAIT COMPLETION = NO must be specified.
 - The Hardware Consumption Solution, TFP Hardware for zIIP workloads, is available for customers with IBM z17 hardware and zIIP processors. No Software Consumption Solution is needed.
- Sub-Capacity Reporting Tool (SCRT) V28.1, or later. (SCRT requires that you use the latest available version. See [Using the Sub-Capacity Reporting Tool \(ibm.biz/SCRTUsersGuide\)](#).)

SCRT is run for each sub-capacity reporting period and the resulting sub-capacity report is submitted to IBM on a monthly basis. The Hardware Consumption Solution requires that at least one native (not running under z/VM®) z/OS system be active at all times during the report period with an RMF (or equivalent) interval of 1, 3, 5, or 15 minutes.

- When a unique solution ID is used to identify the Tailored Fit Pricing solution, you obtain the solution ID through the License Management Support (LMS) website once the solution has been approved.
 - For solutions deployed on an existing sub-capacity LPAR (collocated solutions), you must associate the solution with the tenant resource groups 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 tenant resource group.
 - For solutions deployed on separate LPARs, you must do one of the following:
 - Use the **CONTAINER** command to specify only those LPARs that are part of the solution. The **CONTAINER** command associates the LPARs with the solution ID.

- Specify the **SOLUT** system parameter to associate the z/OS system being IPLed as part of the solution with the solution ID. **SOLUT** is available with the appropriate level of z/OS service.

A solution name is also used to identify the Tailored Fit Pricing 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 that you submit your report, the solution name must match the name as defined in LMS for the container (with the associated solution ID).

Requirements for using IBM Cloud® Provisioning for Tailored Fit Pricing

To use Cloud Provisioning and Management for z/OS for Tailored Fit Pricing, you must have z/OS V2.3 or later, with the PTFs applied for APAR PI88944.

Requirements for using z/OSMF for Tailored Fit Pricing

- To get z/OSMF WLM support for Tailored Fit Pricing definitions, you must have z/OS V2.2 or later, with the PTFs applied for APAR P189361.
- To get z/OSMF RMF support for Tailored Fit Pricing, you must have z/OS V2.2 or later, with the PTFs applied for APAR P189935.

Chapter 3. Setting up a colocated Tailored Fit Pricing DevTest or New Application solution

A colocated DevTest or New Application solution lets you add new approved z/OS solutions in an existing LPAR with other unrelated workloads without directly impacting the R4HA.

Once the WLM service definition is activated with tenant resource group definitions and associated classification rules, WLM tracks the processor consumption on behalf of the tenant resource group-specified workload. RMF then writes each tenant resource group's consumption data to SMF Type 70, Subtype 1 records every recording interval. SCRT analyzes the recorded data and produces a report for consumption by IBM's fulfillment systems, removing the workload's MSU utilization from the R4HA. SCRT analyzes the recorded data and produces a report for consumption by fulfillment systems, removing the workload's MSU utilization from the R4HA.

Use the following sections to set up a colocated Tailored Fit Pricing solution:

- [“WLM definitions for a colocated Tailored Fit Pricing DevTest or New Application Solution” on page 11](#)
- [Chapter 4, “Using SCRT for Tailored Fit Pricing,” on page 23](#)

WLM definitions for a colocated Tailored Fit Pricing DevTest or New Application Solution

For a colocated Tailored Fit Pricing DevTest or New Application solution, the service level administrator must define tenant resource groups, tenant report classes, and classification rules as part of the WLM service definition.

The tenant resource groups and tenant report classes are comparable to resource groups and report classes. They allow for the metering and optional capping of workloads, along with the ability to map those workloads directly to Tailored Fit Pricing for IBM Z solutions.

- Tenant resource groups have sysplex scope, with support for up to 32 tenant resource groups across the sysplex. For each tenant resource group, the system programmer defines one or more tenant report classes, which can be used in classification rules to associate work with the tenant resource group.
- Tenant report classes are similar to standard report classes. You can configure the information that is reported on behalf of tenant resource groups by associating their tenant report classes with address spaces or enclaves using classification rules.

With this new function, the WLM Administrative Application level changes to 32, which is only available with z/OS V2R2 and above when APAR OA52312 is applied. The functionality level of the WLM service definition changes to 32 as soon as a tenant resource group is defined. A WLM service definition with functionality level 32 cannot be extracted, displayed, modified, installed or activated from a z/OS V2R2 or V2R3 system without APAR OA52312 applied. As soon as the functionality level of the WLM service definition changes to LEVEL032, all actions must be taken from a z/OS V2R2 or z/OS V2R3 system with this APAR applied.

Use *z/OS MVS Planning: Workload Management* for complete information on setting up a complete service definition for your colocated workload. In this document we just show the Tailored Fit Pricing highlights:

- [“Some WLM basics for Tailored Fit Pricing” on page 12](#)
- You can define the tenant resource groups and tenant report classes for a Tailored Fit Pricing solution as follows:
 - [“Using z/OSMF to define WLM Tailored Fit Pricing definitions” on page 18](#)
 - [“Using ISPF to define WLM Tailored Fit Pricing definitions” on page 18](#)
- [“WLM application messages” on page 20.](#)

- Once you have WLM definitions set-up for your solution, you can use the following information:
 - [Chapter 5, “Using workload reporting services to collect Container Services performance data,” on page 79](#)

Some WLM basics for Tailored Fit Pricing

This topic provides some basic information about WLM related to Tailored Fit Pricing.

Defining tenant report classes

Optionally, classification rules can assign incoming work to a tenant report class. Tenant report classes are similar to report classes and reported likewise. However, it is required that a tenant report class be assigned to a tenant resource group. A tenant resource group allows for additional metering capabilities or association to a solution.

You can define up to 2047 tenant report classes per service definition whereby the sum of report classes and tenant report classes may not exceed 2047.

Defining tenant report classes	
Name	Tenant report class name
Description	Description of the tenant report class
Tenant Resource Group	Name of the tenant resource group associated with the tenant report class

Name

Eight character identifier of the tenant report class. Each tenant report class must be unique within a service definition and may not have the same name as a report class.

Description

Up to 32 characters that describe the tenant report class.

Tenant Resource Group

Eight character identifier of the tenant resource group that is associated with the tenant report class.

When using tenant report classes in classification rules, note the following:

- A tenant report class cannot be specified on a classification rule with a Reporting Attribute of MOBILE, CATEGORYA, or CATEGORYB. Workload management can report on processor consumption either based on tenant resource groups or based on special reporting options, but not both at the same time.
- A classification rule cannot categorize work into a tenant report class and a service class which is associated with a resource group. Work cannot be subject to capping by means of resource groups and tenant resource groups at the same time.
- As with report classes, tenant report classes are homogeneous or heterogeneous. WLM workload reporting services provide less meaningful data for heterogeneous than for homogeneous tenant report classes. Thus, it is recommended to define separate tenant report classes for each service class and assign them all to the same tenant resource group. If your tenant report class might become heterogeneous, the WLM ISPF applications displays an appropriate warning message.
- The SYSTEM and SYSSTC service classes cannot be associated with a resource group. But, through classification with tenant report classes such work could become part of a tenant resource group that has a processor or memory limit defined:
 - The processor consumption of SYSTEM and SYSSTC work is counted towards the limit of the respective Tenant Resource Group. But, because it represents very important work which is latency sensitive, it will not be capped. Therefore, it is possible that other work may need to be capped more or that the Tenant Resource Group exceeds the defined limit.

- Address spaces classified into the SYSTEM or SYSSTC service classes will not be associated to a memory pool. Thus, the memory pool of the Tenant Resource Group can fully be used by other work associated with it.
- A memory limit overrules the storage critical attribute assigned in classification rules and any protective storage target managed through SRM. IBM recommends that you do not classify memory sensitive work into a memory pool.

Defining tenant resource groups

Tenant resource groups allow the metering and optional capping of workloads, along with the ability to map those workloads directly to Tailored Fit Pricing for IBM Z solutions. A tenant resource group is comparable to a resource group but accepts and processes an IBM provided 64-character Solution ID. While a resource group is associated with service classes, a tenant resource group is associated with tenant report classes. You must define a tenant resource group before you can define tenant report classes.

When you specify a maximum capacity or memory limit for the tenant resource group, WLM limits the amount of processor capacity or memory available to work, which is classified to the tenant report classes associated with the tenant resource group.

You can define up to 32 tenant resource groups per service definition.

Defining tenant resource groups	
Name	Tenant Resource Group Name
Description	Description of the tenant resource group.
Tenant ID	Tenant identifier.
Tenant Name	Descriptive name for the Tenant ID.
Solution ID	IBM provided 64-character solution ID.
Tenant Resource Group Type	Description of the tenant resource group type. This is the same as the resource group type.
Capacity Maximum	Specifies the maximum amount of processor capacity that work associated with the tenant resource group may use.
Include Specialty Processor Consumption	Specifies whether capacity maximum applies to the sum of general purpose processor consumption and specialty processor consumption.
Memory Limit	Maximum amount of real memory that address spaces associated with the tenant resource group through classification may use on the local system. The value has a system scope.

Name

Eight characters that identify the name of the tenant resource group. Each tenant resource group must be unique within a service definition and may not have the same name as a resource group.

Description

Up to 32 characters that describe the tenant resource group.

Tenant ID

Up to eight characters that identify a tenant.

Tenant Name

Up to 32 characters that provide a descriptive name for the Tenant ID.

Solution ID

The 64 character Solution ID as provided by IBM.

Tenant Resource Group Type

Optionally, a tenant resource group allows for the control of the maximum processor consumption. Refer to [“Defining resource groups”](#) on page 14.

Maximum Capacity

CPU service that this tenant resource group may use. *Maximum* applies to all tenant report classes associated with the tenant resource group. *Maximum* is enforced. There is no default maximum value. Tenant resource group capping will not be enforced while system recovery boost is active in a partition.

Include Specialty Processor Consumption

The attribute specifies whether capacity maximum applies not only to general purpose processors but also to specialty processors. The default is **no**, which ignores CPU consumption of specialty processors when managing the maximum capacity. If **yes** is specified, the total CPU consumption on general purpose and specialty processors is limited by the Maximum Capacity.

Memory Limit

Maximum amount of real memory that address spaces that are associated with the tenant resource group through classification may use on the local system. The attribute is specified in GB. The attribute value has system scope. Please refer to [“Defining resource groups”](#) on page 14 for a detailed description for setting a memory limit.

Tenant report classes representing transaction-oriented work, such as CICS or IMS transactions, can only be assigned to tenant resource groups without a maximum capacity defined. If you assign a tenant resource group with a maximum capacity, the WLM ISPF application displays an appropriate warning message. Although the tenant resource group is accepted, the capacity limit is ignored for the CICS and IMS transactions.

Defining resource groups

A *resource group* is an amount of processor capacity and/or memory. It is optional. Unless you have some special need to limit or protect processor capacity or memory for a group of work, you should skip defining resource groups and let workload management manage all of the processor and memory resource to meet performance goals. You use a resource group to:

- Limit the amount of processor capacity available to one or more service classes.
- Set a minimum for processor capacity for one or more service classes if the work is not achieving its goals.
- Define a minimum and maximum amount of processor capacity sysplex-wide, or on a system level.
- Specify whether capacity values of the resource groups apply to general purpose processors only or to general purpose and specialty processors.
- Limit the amount of memory capacity that is available to one or more service classes on a system level.

You can specify a minimum and maximum amount of processor capacity and a maximum amount of memory to a resource group. You can assign only one resource group to a service class. You can assign multiple service classes to the same resource group. You can define up to 32 resource groups per service definition.

Keep in mind your service class goals when you assign a service class to a resource group. Given the combination of the goals, the importance level, and the resource capacity, some goals might not be achievable when capacity is restricted.

Setting a maximum processing capacity

If work in a resource group is consuming more processor resources than the specified maximum processor capacity, the system caps the associated work accordingly to slow down the rate of processor resource consumption. The system might use several mechanisms to slow down the rate of processor resource consumption, including swapping the address spaces, changing their dispatching priority, and

capping the amount of processor service that can be consumed. Reporting information reflects that the service class might not be achieving its goals because of the resource group capping.

Setting a minimum processing capacity

By setting a minimum processing capacity, you create an overriding mechanism to circumvent the normal rules of importance. If the work in a resource group is not meeting its goals, then workload management attempts to provide the defined minimum amount of CPU resource to that work.

Setting a memory limit

By specifying a memory limit, you explicitly restrict the use of real memory of work that runs in address spaces that are associated with the resource group through classification. For a resource group with a memory limit, the system creates a memory pool. A memory pool does not reserve real memory for use by the pool, but rather tracks the aggregate usage in order to limit the total usage by address spaces connected to the pool.

An address space that is associated with the resource group through classification connects to the memory pool when the address space, or a new job, starts. In that case, all its frames are counted toward the pool limit. When a memory pool approaches its limit, the system takes actions such as initiating self-stealing to page out memory pool pages and thus free up memory pool frames. This protects the real memory allocation of other work that is running on the system.

An address space can be switched to another memory pool or back to system storage either by activating another service policy, or resetting it to another service class.

IBM recommends that you use memory pools only when it is required to limit the use of memory by workloads and for applications which provide guidance on how to operate them in a memory pool.

When you install and activate a service definition that deletes an existing resource group with a memory limit, the system defers deletion of the associated memory pool until all address spaces associated with the memory pool disconnect and end.

When a memory pool approaches its limit, address spaces starting up and connecting to the pool are deferred until enough frames are available through self-stealing from the pool.

Defining resource groups	
Name	Resource Group name
Description	Description of resource group
Resource Group Type	Description of resource group type
Capacity Maximum	Can be calculated in various ways, depending on which resource group is used, and is explained in the following.
Capacity Minimum	Can be calculated in various ways, depending on which resource group is used, and is explained in the following.
Include Specialty Processor Consumption	Specifies whether minimum and maximum capacity applies to the sum of general purpose processors and specialty processor consumption.
Memory Limit	Maximum amount of real memory the address spaces that are associated with the resource group through classification may use on the local system. The value has a system scope.

Name

Eight characters that identify the name of the resource group. Each resource group must be unique within a service definition.

Description

Up to 32 characters that describe the resource group.

Resource Group Type

Resource groups allow to define a guaranteed maximum and minimum CPU consumption for work on the sysplex and on each individual member of the sysplex. This allows to:

- Prioritize work on a system-level basis.
- Control the minimum and maximum resource consumption.

The following types of resource groups are valid:

Resource Group Type 1

The capacity is specified in unweighted CPU service units per second, the value must be between 0 and 99999999.

Minimum and maximum capacity applies sysplex-wide, that is, WLM ensures that the limits are met within the sysplex.

Resource Group Type 2

The capacity is specified as a percentage of the LPAR share in the general purpose processor pool, the value must be between 0 and 99999. To accommodate specialty processor capacity, values greater than 100 may be specified.

Minimum and maximum capacity has a system scope, that is, WLM ensures that the limits are met on each system within the sysplex.

Resource Group Type 3

The capacity is specified as a number of general purpose processors (CPs), a number of 100 represents the capacity of 1 CP. The number should be between 0 and 999999. To accommodate specialty processors that run at a different speed, a number greater than 100 must be specified to represent the capacity of one specialty processor.

Minimum and maximum capacity has a system scope, that is WLM ensures that the limits are met on each system within the sysplex.

Resource Group Type 4

The capacity is specified in accounted workload MSU that is based on captured time. Minimum and maximum capacity is processor consumption that is expressed in million service units per hour and applies sysplex-wide, that is, WLM ensures that the limits are met within the sysplex. Minimum and maximum must be a value between 0 and 999999.

Resource group type 4 is intended to simplify the specification of a limit expressed in MSU. This limit only applies to captured TCB and SRB times. System management time, also known as uncaptured time, is not included. Furthermore, this limit is managed on a short interval. Thus, it is no four-hour rolling average MSU consumption.

The "CPU capacity table" in *z/OS MVS Planning: Workload Management* shows the service units per second by CPU model. Also, refer to *Large Systems Performance Reference for IBM Z* to determine the MSU rating of your sysplex.

Capacity

Identifies the amount of available capacity you want workload management to allocate to the resource group. Capacity includes cycles in both TCB and SRB mode. Resource group minimum can equal resource group maximum.

Maximum

CPU service that this resource group might use. *Maximum* specified for this resource group applies to all service classes in that resource group combined. *Maximum* is enforced. There is no default maximum value. If specified, *Maximum* must be greater than 0.

Minimum

CPU service that should be available for this resource group when work in the group is missing its goals. The default is 0. If a resource group is not meeting its minimum capacity and work in that resource group is missing its goal, workload management attempts to give CPU resource to that work, even if the action causes more important work (outside the resource group) to miss its goal. If there is discretionary work in a resource group that is not meeting its minimum capacity, WLM attempts to give the discretionary work more CPU resource if that action does not cause other work to miss its goal.

The minimum capacity setting has no effect when work in a resource group is meeting its goals.

Memory Limit

Maximum amount of memory that address spaces that are associated with the resource group through classification might consume on the local system. The attribute is specified as absolute value in GB. The attribute value has system scope.

Include Specialty Processor Consumption

The attribute specifies whether capacity minimum and maximum apply not only to general purpose processors but also to specialty processors. The default is **no**, which ignores CPU consumption of specialty processors when managing the guaranteed minimum and maximum capacity. If **yes** is specified, the total CPU consumption on general purpose and specialty processors is applied.

Note:

1. You cannot assign a resource group to service classes representing transaction-oriented work, such as CICS or IMS transactions. The ISPF application notifies you with an error message if you attempt to do so. If you want to assign a minimum or a maximum processor capacity and a maximum amount of memory to CICS or IMS work, you can do so by assigning a resource group to their regions. For example, suppose that you have three service classes representing your CICS works: CICSTRN, CICSORS, and CICSTORS. CICSTRN represents all of your online CICS transactions, and has one period with a short response time goal. CICSORS and CICSTORS represent all of your CICS AOR and TOR regions, respectively, that process the online transactions. To assign a maximum processor capacity and a maximum amount of memory to your CICSTRN work, define a resource group, and assign it to the regions. So you assign the resource group to the CICSORS and CICSTORS service classes.
2. Similarly, resource groups with a memory limit cannot be applied to enclave service classes. However, because enclave service classes can be used anywhere, unlike CICS or IMS transaction service classes, the ISPF application does not notify you with an error message if you attempt to do so. As for CICS or IMS, a resource group with a memory limit must be assigned to the service class of the address spaces that join the enclaves.
3. A memory limit overrules the storage critical attribute assigned in classification rules and also any protective storage target managed through SRM.
4. Resource group processor capacity capping is implemented by marking the work units that belong to resource group non-dispatchable for some time slices and dispatchable for the remaining time slices (awake slices). Depending on the configuration, it might not be possible to enforce very low resource group limits. The granularity to which a resource group limit can be managed depends on how much service the work can consume in a system or across the Sysplex, respectively, during one awake time slice. Beginning with z/OS V2.1 the granularity of awake slices is 1/256 of the time.
5. When resource groups are managed based on the general purpose processor service (the attribute, Include Specialty Processor Consumption, specifies no) the dispatchability attribute is also honored by zIIPs.
6. Resource group capping will not be enforced while system recovery boost is active in a partition.

Using IBM Cloud Provisioning for Tailored Fit Pricing

You can use IBM Cloud Provisioning and Management for z/OS to set up collocated DevTest or New Application solutions for Tailored Fit Pricing. For cloud provisioning, use the Resource Management task

in the Cloud Provisioning category of z/OSMF to define domains (systems) and tenants (users). You can define a tenant as a container for Tailored Fit Pricing by specifying a Solution ID for the tenant. Then, any software that you provision for that tenant is treated as part of the solution. This simplifies setup required for collocated Tailored Fit Pricing, because the Resource Management task does the z/OSMF Workload Management setup for you -- including creating the tenant resource group and tenant report classes, and generating the classification rules.

For information about IBM Cloud Provisioning and Management for z/OS, see the following:

Cloud Provisioning (www.ibm.com/docs/en/zos/2.5.0?topic=help-cloud-provisioning)

Using z/OSMF to define WLM Tailored Fit Pricing definitions

For a colocated DevTest or New Application solution, the tenant resource group and tenant report classes for your Tailored Fit Pricing solution are part of the WLM service definition for the workload.

The workload management task in z/OS Management Facility (z/OSMF) provides a browser-based user interface that you can use to manage z/OS workload management (WLM) service definitions and provide guidelines for WLM to use when allocating resources. Specifically, you can define, modify, view, copy, import, export, and print WLM service definitions. You can also install a service definition into the WLM couple data set for the sysplex, activate a service policy, and view the status of WLM on each system in the sysplex.

Using ISPF to define WLM Tailored Fit Pricing definitions

Note that IBM recommends that you use z/OSMF to define WLM Tailored Fit Pricing definitions. If you do use the WLM ISPF Administrative Application to maintain your service definition, save the service definition using the XML format to ensure that it can be read by older levels of the WLM Administrative Application or the IWMINSTL sample job.

You can use ISPF to define the WLM Tailored Fit Pricing definitions. For a colocated solution, the tenant resource group and tenant report classes for your Tailored Fit Pricing solution are part of the WLM service definition for the workload. Use *z/OS MVS Planning: Workload Management* for the complete information on setting up a complete service definition for your workload using the ISPF administrative application - in this section, we'll just hit the Tailored Fit Pricing highlights. Make sure you use the z/OS V2R3 level or higher of *z/OS MVS Planning: Workload Management* (even if you are using Tailored Fit Pricing on z/OS V2R2).

Using the Definition Menu

The definition menu is the central place for entering your service definition. When you set up a service definition, you must enter a service definition name and optionally, a description on the Definition Menu.

Figure 4 on page 19 shows a sample Definition Menu with the service definition name and a description filled in.

```

File Utilities Notes Options Help
-----
Functionality LEVEL001          Definition Menu          WLM Appl LEVEL025
Command ==> -----

Definition data set . . : none

Definition name . . . . . _____ (Required)
Description . . . . . _____

Select one of the following options.
--- 1. Policies                               12. Tenant Resource Groups
    2. Workloads                             13. Tenant Report Classes
    3. Resource Groups
    4. Service Classes
    5. Classification Groups
    6. Classification Rules
    7. Report Classes
    8. Service Definition Options
    9. Application Environments
   10. Scheduling Environments
   11. Guest Platform Management Provider

```

Figure 4. Definition Menu panel

Changes related to this panel for Tailored Fit Pricing include:

Tenant resource groups(new)

A tenant resource group is comparable to a resource group but accepts and processes a 64-character Solution ID. The processor consumption of all work classified into tenant report classes assigned to the tenant resource group is provided for metering capabilities.

Tenant report classes (new)

A tenant report class is a report class that is assigned to a tenant resource group. When assigning work in classification rules to a tenant report class, the processor consumption is provided for metering capabilities of the tenant resource group.

Policy overrides

You can specify a policy override for a tenant resource group. You specify a policy override by selecting Policies from the Definition Menu and then specifying the appropriate action code.

Working with tenant resource groups

To define a tenant resource group, chose option **12** on the Definition Menu. Define a name, and optionally a description, a tenant ID and name, and a 64-character Solution ID. If you want to specify a capacity limit, define the type (1, 2, 3, or 4) and capacity maximum. If you want to include the consumption of specialty processors in the capacity maximum, specify YES in the corresponding field. For a memory limit, specify the maximum amount in GB.

```

Tenant-Resource-Group  Notes  Options  Help
-----
                                Create a Tenant Resource Group
Command ==> -----

Enter or change the following information:

Tenant Resource Group Name  _____ (required)
Description . . . . . -----
Tenant ID . . . . . -----
Tenant Name . . . . . -----
Solution ID . . . . . -----

-----
Define Capacity: --  1. In Service Units (Sysplex Scope)
                   2. As Percentage of the LPAR share (System Scope)
                   3. As a Number of CPs times 100 (System Scope)
                   4. In accounted workload MSU (Sysplex Scope)
Maximum Capacity . . . . . -----
Include Specialty Processor Consumption NO (YES or NO)
Memory Limit (System Scope) . . . . . -----

```

Figure 5. Create a Tenant Resource Group panel

Once you have created a tenant resource group, any other time you choose the tenant resource group option from the definition menu, the application displays a selection list. From here, you can modify your tenant resource group, as well as print, and browse it.

Working with tenant report classes

To define a tenant report class, chose option **13** on the Definition Menu. Define the name of the tenant report class, and optionally a description. You must assign a tenant resource group to the tenant report class. You can type ? in the tenant resource group name field for a list of tenant resource groups.

You can use tenant report classes in classification rules to categorize work.

```

Tenant-Report-Class  Notes  Options  Help
-----
                                Create a Tenant Report Class
Command ==> -----

Enter or change the following information:

Tenant Report Class Name . . . _____ (Required)
Description . . . . . -----
Tenant Resource Group Name . . _____ (Required; name or ?)

```

Figure 6. Create a Tenant Report Class

Once you have created a tenant report class, any other time you choose the tenant report class option from the definition menu, the application displays a selection list. From here, you can modify your tenant report class, as well as print, and browse it.

You can also type ? in the report class field on the Modify Rules for a Subsystem Type panel for a selection list of tenant report classes.

WLM application messages

IWMAM313 **No more than 2047 report classes and tenant report classes may be defined.**

System action

The requested operation is not performed.

Explanation

It is not possible to define more than 2047 report classes and tenant report classes.

Programmer response

Do not use more than 2047 report classes and tenant report classes.

Module	IWMAM540	No more than 32 tenant resource groups may be defined.
Workload management (WLM)		
IWMAM512	No more than 32 resource groups may be defined.	Explanation
Explanation		It is not possible to define more than 32 tenant resource groups.
It is not possible to define more than 32 resource groups.	System action	The requested operation is not performed.
System action		Programmer response
The requested operation is not performed.		Do not use more than 32 tenant resource groups.
Programmer response		Module
Do not use more than 32 resource groups.		Workload management (WLM)
Module		
Workload management (WLM)		

Using system command RESET to change the service class of work when associated with a tenant resource group

Use system command RESET | E *jobname* with the SRVCLASS keyword to change the service class of work currently in execution. Resetting to a new service class also resumes quiesced work. The SRVCLASS keyword explains how the system handles the case when the service class is associated with a tenant resource group.

For complete information about the RESET command, see [z/OS MVS System Commands](#).

Use this command only at the direction of the system programmer.

The syntax for this form of the RESET command is:

```
RESET | E jobname [,A=asid] ,
           {SRVCLASS=classname}
           {QUIESCE | Q}
           {RESUME}
```

jobname

The name of the job, time-sharing user, or started task whose performance characteristics you want to change. This command affects the current job step and all subsequent job steps in this execution.

A=*asid*

The hexadecimal address space identifier (ASID) of the job, time-sharing user, or started task you want to change. You can specify this keyword before or after the SVRCLASS, QUIESCE, or RESUME parameters.

This parameter is required if there is more than one job with the same job name.

SRVCLASS=*classname*

The name of the service class to be assigned to the job or address space. Resetting to a new service class also resumes quiesced work.

When you issue a RESET against a server (for example, an address space) to a new service class, the goals associated with that service class are ignored. However, the resource group associated with the new service class is honored. The one exception is the case where the goal for a server is honored when the transactions it is serving have been assigned a discretionary goal.

If the address space you want to change is associated with a tenant resource group, and the new service class is associated with a resource group, the address space is disconnected from the tenant resource group. If the new service class is not associated with a resource group, RESET does not change the connection of the address space with a tenant resource group.

There may be special circumstances under which you would wish to reset an address space with a SYSTEM or SYSSTC service class. See the “Defining Classification Rules” topic in *z/OS MVS Planning: Workload Management* for information about the use of the SYSTEM and SYSSTC service classes.

RESET SRVCLASS will remain in effect until one of the following events occurs:

- The job ends.
- The policy is switched to a new policy in which the target service class has been deleted.
- A RESET RESUME command is issued.

QUIESCE|Q

Requests that the target job or address space be quiesced; that is, given the lowest possible performance characteristics. QUIESCE swaps out swappable work, effectively shutting off that work. QUIESCE just lowers the performance of non-swappable work, leaving it swapped in.

RESET QUIESCE will remain in effect until one of the following events occurs:

- The job ends.
- A RESET RESUME command is issued.

RESUME

Specifies that a job or address space be reclassified. If the job or address space was quiesced by a previous RESET *jobname*, QUIESCE command, or if the job or address space was assigned to a different service class, RESUME causes the work to be reclassified according to the service policy in effect and resumes processing at the performance targets specified in the service policy.

The classification rules used are those in effect at the time the RESET command is issued.

Chapter 4. Using SCRT for Tailored Fit Pricing

The Sub-Capacity Reporting Tool (SCRT) allows you to generate Sub-Capacity reports or Multiplex reports.

The procedure for using SCRT for Tailored Fit Pricing is to update the SCRT JCL SPECIAL DD statement before you run the job that produces the SCRT report. You use

- For a Software Consumption Solution, a CONTAINER and UPDATE CONTAINER command
- For a Hardware Consumption Solution, INTERVAL_RATE_DATA and DETAIL_INTERVAL_RATE_DATA control statements
- For a dedicated LPAR DevTest or New Application solution, a CONTAINER and UPDATE CONTAINER command
- For a colocated LPAR DevTest or New Application solution, an UPDATE CONTAINER command.

There is no SCRT reporting requirement for an Enterprise Capacity Solution.

As an alternative, for a solution with a dedicated LPAR, you can specify the **SOLUT** system parameter in the IEASYSxx member of parmlib when IPLing a z/OS system as part of the qualified solution. The **SOLUT** parameter associates the system with the solution ID. **SOLUT** is available with the appropriate level of z/OS service. For more information, see [“Dedicated LPARs”](#) on page 24.

SCRT uses the Tailored Fit Pricing data to understand the resource consumption, and it automatically removes that container’s MSU consumption from the LPAR’s peak rolling 4-hour average (R4HA).

In this information, we cover the highlights for Tailored Fit Pricing. For complete information about SCRT, see *Using the Sub-Capacity Reporting Tool*, available from [Getting Started with the Sub-Capacity Reporting Tool](#) (www.ibm.com/downloads/cas/OVPO2DJW).

SCRT commands for the Software Consumption Solution

For a Software Consumption Solution, use these commands:

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

```
UPDATE CONTAINER,  
ID=solutionID,  
SET_NAME="solutionName"
```

Use one CONTAINER command for each machine that the solution runs on and one UPDATE CONTAINER command for each solution ID.

SCRT commands for the Hardware Consumption Solution

For a Hardware Consumption Solution, include the following control statements in the SPECIAL DD of the SCRT job:

```
INTERVAL_RATE_DATA  
DETAIL_INTERVAL_RATE_DATA
```

SCRT commands for a dedicated DevTest or New Application solution

For a dedicated solution, use these commands:

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

```
UPDATE CONTAINER,  
ID=solutionID,  
SET_NAME="solutionName"
```

SCRT commands for a colocated DevTest or New Application solution

For a colocated solution, use this command:

```
UPDATE CONTAINER,  
ID=solutionID,  
SET_NAME="solutionName"
```

Dedicated LPARs

For a solution with a dedicated LPAR, the method used to associated the solution ID with the system or LPAR affects reporting.

- 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.

Using the SOLUT parameter

The **SOLUT** system parameter in the IEASYSxx member of parmlib associates a system with the solution ID when you IPL a z/OS system as part of a dedicated-LPAR solution. **SOLUT** requires the appropriate level of z/OS service. The details of the **SOLUT** parameter follow.

SOLUT

```
SOLUT=solution-id
```

This parameter specifies a 64-character solution ID (*solution-id*) that is provided to you by IBM.

The solution ID associates the system with a Tailored Fit Pricing solution. All system activity that is not otherwise associated with a different solution ID by way of WLM tenant resource groups will be associated with the specified solution ID. The solution ID is reported in SMF type 89 subtype 1 and subtype 2 records.

Value range: An IBM-provided, 64-character Tailored Fit Pricing solution ID.

Default: None. (The system is not associated with a Tailored Fit Pricing solution.)

Associated parmlib member: None.

Reporting for Tailored Fit Pricing

The reporting requirements vary with the solution. See:

- [“Reporting for Tailored Fit Pricing for IBM Z - Software Consumption Solution”](#) on page 25 for information about a Software Consumption Solution
- [“Reporting for Tailored Fit Pricing for IBM Z - Hardware Consumption Solution”](#) on page 27 for information about a Hardware Consumption Solution
- [“Workload Classification Supplement”](#) on page 30 for information about CPC utilization by workload class.
- [“Reporting for Tailored Fit Pricing for IBM Z”](#) on page 33 for information about an Application Development and Test and New Application Solution.

The Enterprise Capacity Solution has no reporting requirement.

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.

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.
- 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=WLM MOB
- 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 7 on page 27 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-S0000000000
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 7. 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 8 on page 27 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 8. 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 9 on page 27 shows the CPS2 header section for the Production1 container.

```

==CPS2=====
Solution ID          D911111-N6EE848-5A09B977D0-NTZZZZZ-1DDE-4F4B-821D-3D8B30-301310
Solution Name        DevTest1
Peak Four Hour Rolling Average 4139

```

Figure 9. 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.

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.)

SCRT externals for Hardware Consumption reporting

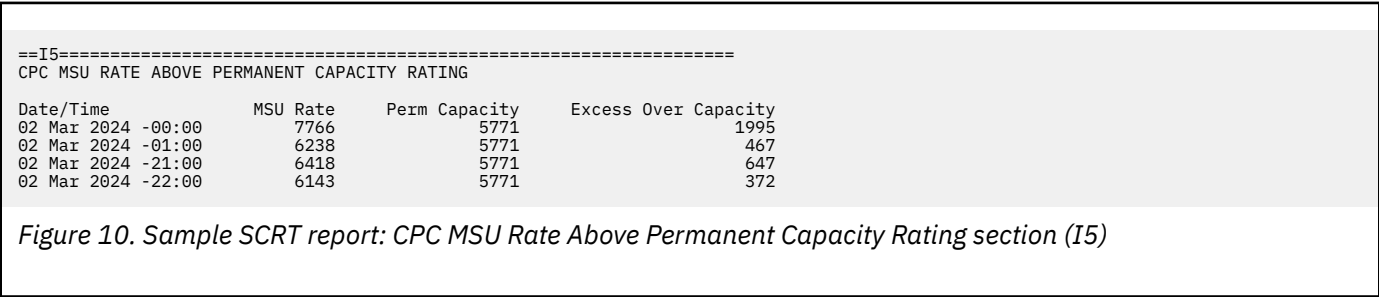
SCRT provides the following externals in support of Hardware Consumption reporting:

- 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 28.
- A new Detailed Interval Data section (V9) in SCRT reports. For details about this section, see “Sample report for Hardware Consumption reporting” on page 28.

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 10 on page 28 shows an example of 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 11 on page 29 shows an example of section V9.

==V9=====																
DETAILED INTERVAL DATA																
	Date/Time	IntervalStart	IntervalEnd	IntLen	TypeMod	Rating	Hour R4HA	Hour Consumed	MSU	MSU Rate	Perm Model	Perm Capacity	Temp Model	Temp Capacity	Var Model	Var 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	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	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	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	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	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	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	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	7898

Figure 11. 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.

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 MxCx header (see the description of the machine identifier for CMP reporting), 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 Release          30.1.x
Name of Person Submitting Report  Joe Contact
E-Mail Address of Report Submitter nobody@example.com
Phone Number of Report Submitter 555-555-5555

Customer Name              Joe Customer
Run Date/Time              01 Apr 2025 - 15:37
Reporting Period           2 Jan, 2025 - 1 Feb, 2025 inclusive (31 days)
Number of processors reported 1

Machine Identifier         M1C1
Machine Serial Number      51-77777
Machine Type               8561
```

==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%	...

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).

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 [csscrtspcialdd.dita#csscrtspcialdd/wtf](https://www.ibm.com/support/ctgdocsscp/eng/SSCSCRTSPECIALDD.dita#csscrtspcialdd/wtf).

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 [csscrtspecialdd.dita#csscrtspecialdd/igcon](#).

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 [csscrtspecialdd.dita#csscrtspecialdd/upsol](#).

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 [csscrtspecialdd.dita#csscrtspecialdd/igcon](#).

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 12 on page 37 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 12. 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.

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-NTZZZZZ-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.

Figure 13 on page 39 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 13. 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 14 on page 40 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-Z0S	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 14. 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.

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 15 on page 41 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 15. 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 16 on page 41](#).

```

==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 16. 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 17 on page 42 shows the Container Active Products section for Tailored Fit Pricing solution CPS1.

```

==E7=====
CONTAINER ACTIVE PRODUCTS

MLC Product Name                MLC Product ID
z/OS V2                        5650-Z0S
CICS TS for z/OS V5           5655-Y04

IPLA Product Name              IPLA Product ID
File Export for z/OS V1       5697-I12

```

Figure 17. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Active Products section

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 42 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 18. 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 42, 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 19. 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 20 on page 44](#).

```

=====
==J7=====
CONTAINER PRODUCT METRICS
Product Name                Product ID   Product Units   Product Metric
IBM z/OS Connect EE V2     5655-CEE    40              Simultaneous Instances

```

Figure 20. Sample sub-capacity report for a collocated solution for Tailored Fit Pricing for IBM Z: Container Product Metrics section

In Figure 20 on page 44, 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.

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

Figure 21 on page 45 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 21. 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.

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.

Figure 22 on page 45 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 22. 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.

Figure 23 on page 46 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 23. 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 24 on page 46 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 TGCICS21	SYS2 TGCICS21	SYS4 TGCICS21	SYS5 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 24. 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 25 on page 47 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 25. 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 25 on page 47, 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 26 on page 48 shows an example of the Container MSU Consumption Contributors section for a container named CPS1.

```

=====
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 26. 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 26 on page 48, 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 27 on page 50](#) 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 27. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Multiplex summary section

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 28 on page 50 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
DB2 for z/OS	(All)	84	27 Nov 2024 - 21:00
DB2 11 for z/OS	5615-DB2	84	27 Nov 2024 - 21:00
DB2 10 for z/OS	5605-DB2	84	27 Nov 2024 - 21:00
CICS TS for z/OS V4	5655-S97	84	27 Nov 2024 - 21:00
WebSphere MQ for z/OS V7	5655-R36	84	27 Nov 2024 - 21:00
IPLA Product Name	Number	MSU	Time
WebSphere Transformation Extender for z/OS V8	5655-R95	2	27 Nov 2024 - 20:00
IPLA z/OS-Based	(All)	84	27 Nov 2024 - 21:00
Container Identifier	Container Name	MSU	Time
CPS1	DevTest Solution	56	27 Nov 2024 - 21:00
Container Identifier	Container Name	TOTAL MSU Consumption	
CPS1	DevTest Solution	5234	5234

Figure 28. 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.

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 29 on page 52](#).

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	LPAR1
z/OS V2	5650-ZOS	200	50	27 Nov 2024 - 18:00	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 29. 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 30. 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 31 on page 53](#).

```

==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 31. 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 41](#).

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.

The detail data sections for customer verification follow the machine information section (M7) shown in [Figure 32 on page 54](#).

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 32. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Customer verification section for CPC M1C1 - Machine information

Figure 33 on page 54 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 33. Sample multiplex report for a standalone Tailored Fit Pricing solution with Country Multiplex Pricing: Product Max Contributors and Product Grid Snapshot sections

Figure 34 on page 55 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 34. 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.

SPECIAL DD statement

You can use the optional SPECIAL DD statement to specify one or more of the SCRT control statements listed in [Table 1 on page 55](#).

Table 1. SCRT SPECIAL control statements

Control statement name	Usage
CONTAINER	Associate an LPAR with a Tailored Fit Pricing solution ID.
IGNORE CONTAINER	Ignore a Tailored Fit Pricing solution for a specified time interval.
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.
UPDATE SOLUTION_ID	Replace a solution ID associated with a container tenant resource group with a new or corrected solution ID.

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.

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.

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.

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.

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.

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.

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.
2. Run the SCRT for Windows and Linux graphical user interface (GUI), as described in [“Running the SCRT graphical user interface”](#) on page 60.

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 76.

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.

Figure 35 on page 61 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.)

IBM Sub-Capacity Reporting Tool

File Help

Sub-Capacity Reporting Tool (SCRT)
This wizard will guide you through the report creation process.

Customer Name: Customer

Customer Number: 2000000000

Contact Name: Mr. Customer

Contact Email: customer@example.com

Contact Phone: 555-555-5555

Cancel Back Next

Figure 35. SCRT for Windows and Linux: Customer parameter panel

Fill in the parameter values as described in "PARMS DD statement" in the SCRT user's guide, *Using the Sub-Capacity Reporting Tool*, available from the [Getting Started with the Sub-Capacity Reporting Tool](http://www.ibm.com/downloads/cas/OVPO2DJW) (www.ibm.com/downloads/cas/OVPO2DJW). 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 36 on page 62 shows an example of the SMF or SCRT89 data selection panel.

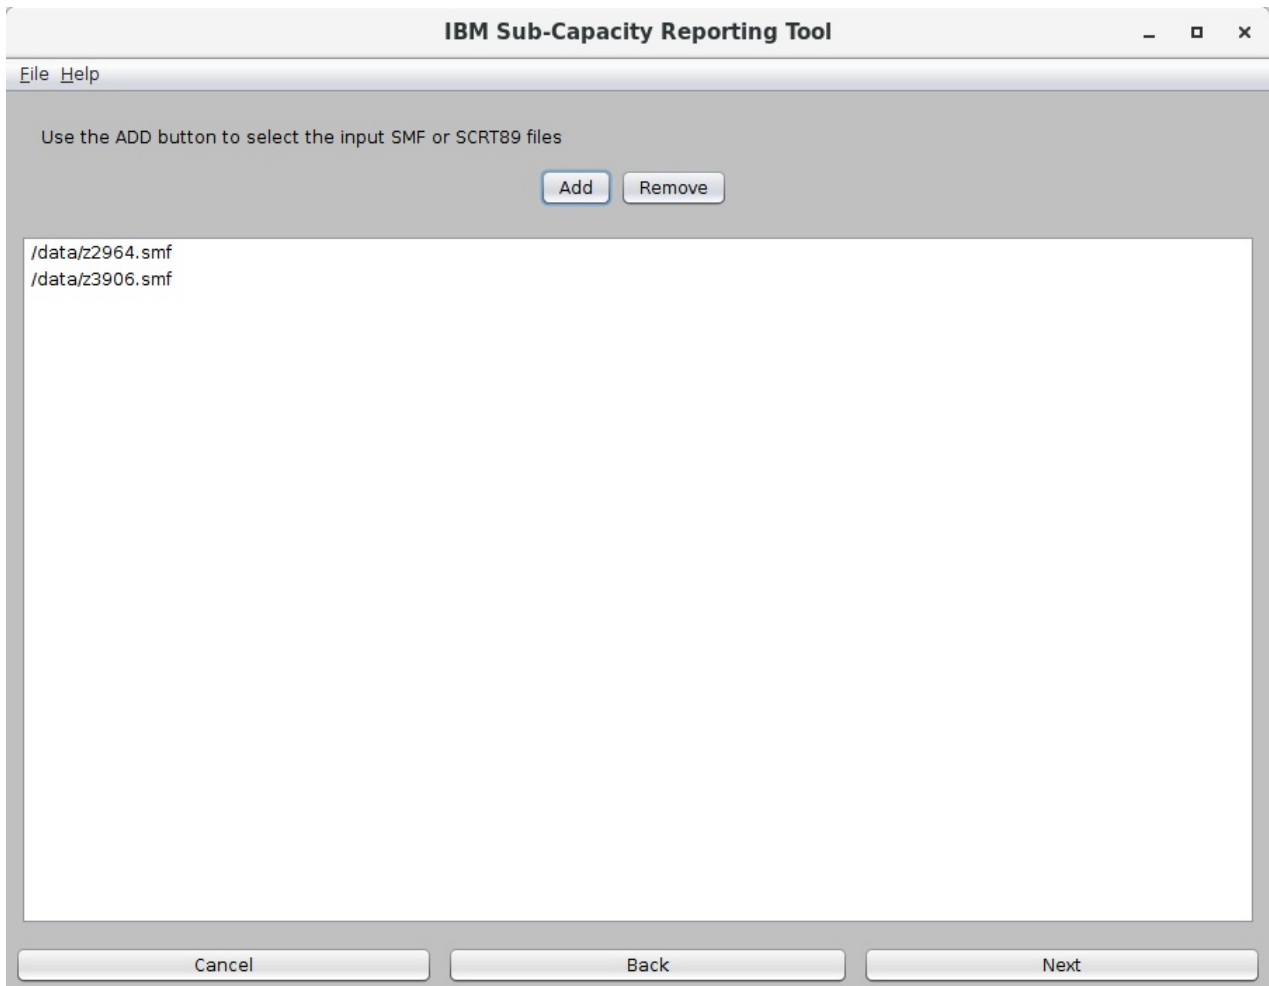


Figure 36. SCRT for Windows and Linux: SMF or SCRT89 data selection panel

Important: SCRT expects the SMF or SCRT89 data to retain a specific data format.

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 37 on page 63 shows an example of the customer number panel.




Figure 37. 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 38 on page 64](#) shows an example of the reporting period panel.




Figure 38. 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 39 on page 65 shows an example of the NO89 product selection panel with the **NO89 Product List** tab selected.

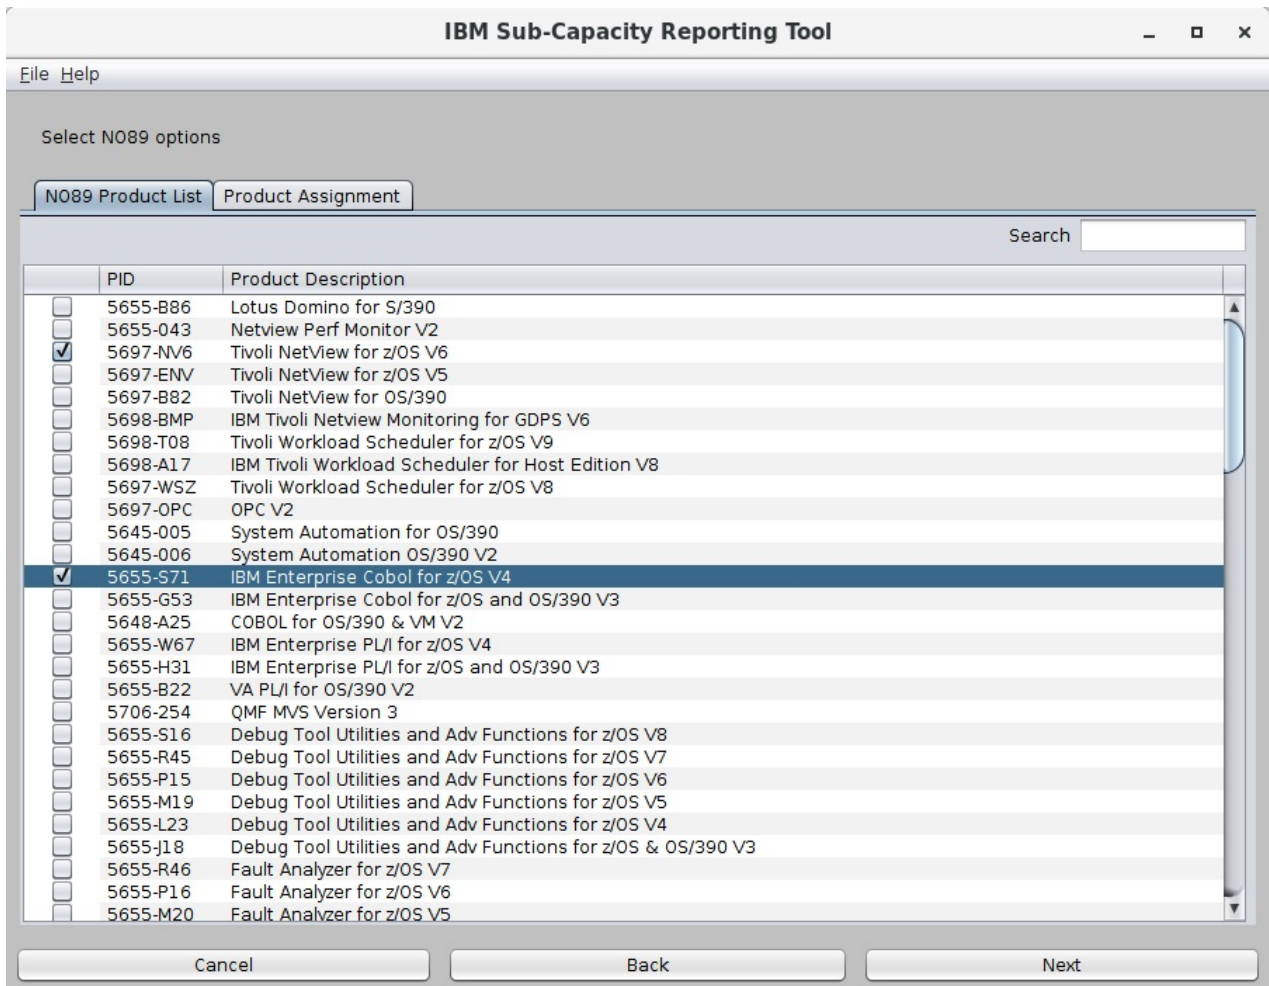


Figure 39. SCRT for Windows and Linux: NO89 product selection panel — NO89 Product List tab

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 40 on page 66 shows an example of the NO89 product selection panel with the **Product Assignment** tab selected.

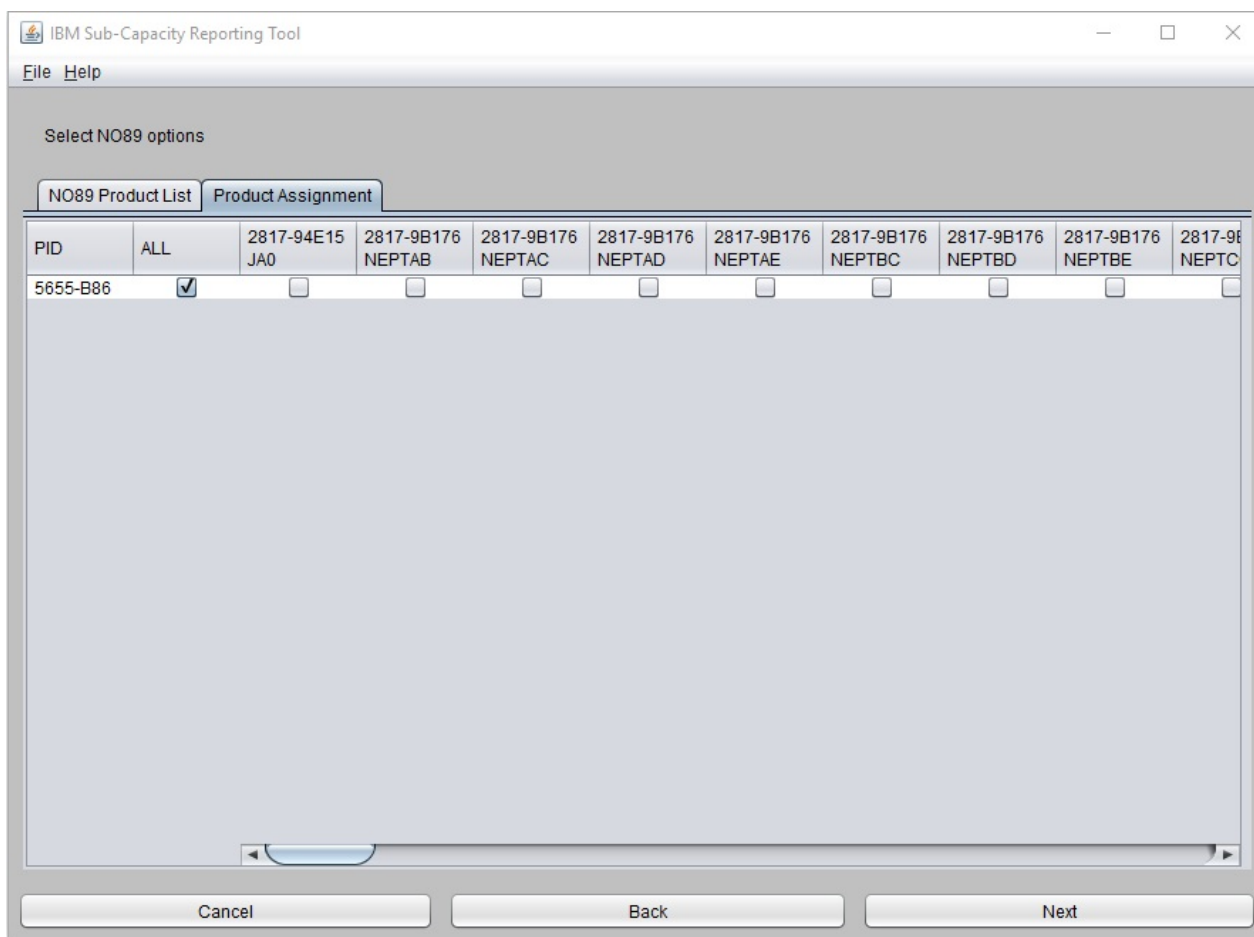


Figure 40. SCRT for Windows and Linux: NO89 product selection panel — Product Assignment tab

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.

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 41 on page 67 shows an example of the **Input CSV** tab.




Figure 41. 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 42 on page 68](#) shows an example of the **Processing Options** tab.




Figure 42. 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 cssrtspecialdd.dita#cssrtspecialdd 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.

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 43 on page 69](#) shows an example of the detected solution IDs panel.

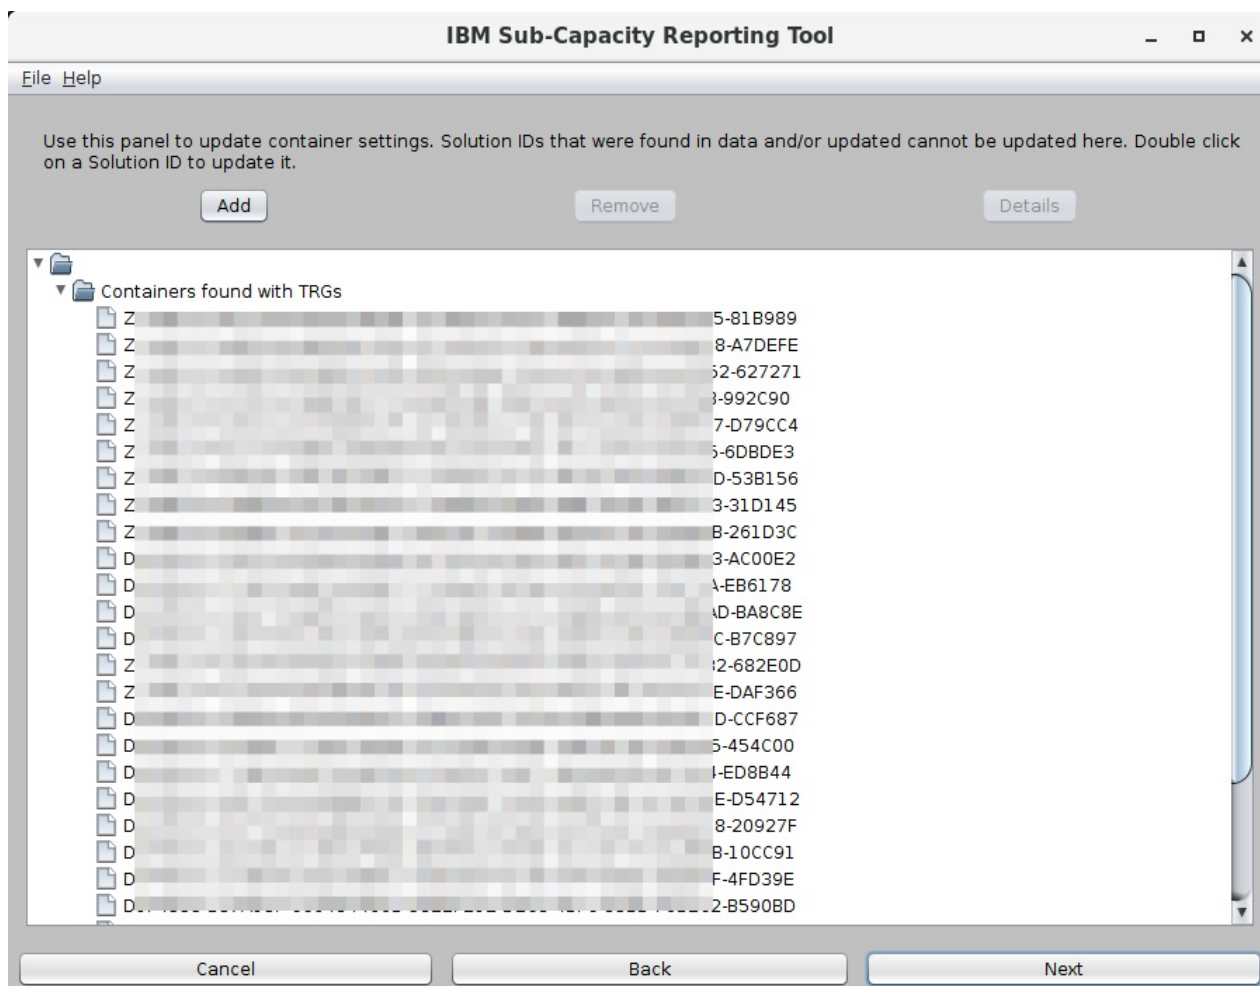


Figure 44. SCRT for Windows and Linux: Container settings assignment panel

- 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.

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.

Figure 45 on page 71 shows an example of the Exclude panel.

IBM Sub-Capacity Reporting Tool

File Help

Please update list of exclude options.

CPCs Images Products MSU Consumption

Search

Add	CPC
<input type="checkbox"/>	8561-22E78

Remove Details

CPC	Time	Hour	ID
-----	------	------	----

Cancel Back Next

Figure 45. 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 46 on page 72](#).

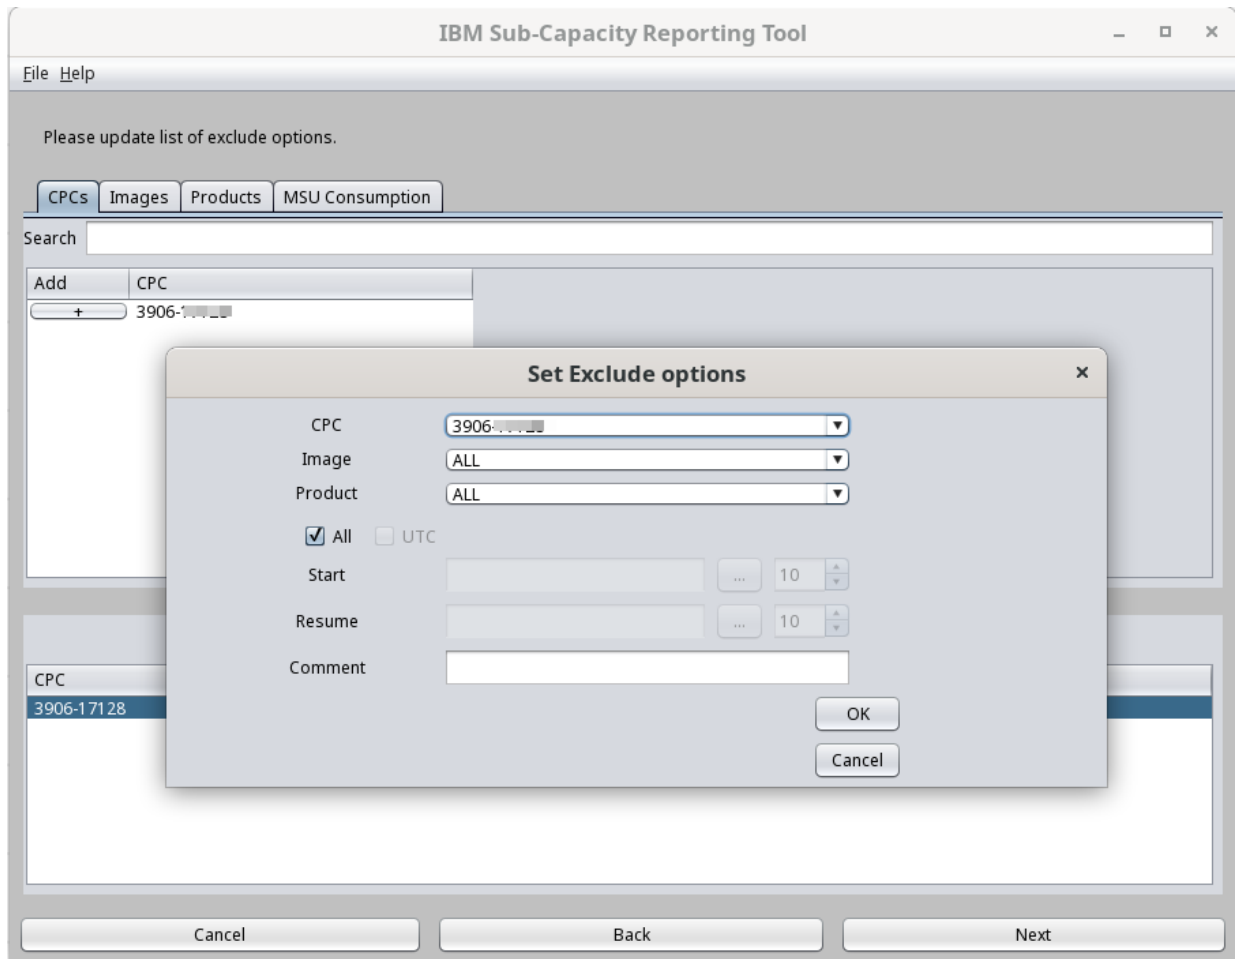


Figure 46. SCRT for Windows and Linux: Exclude panel – Set Exclude options

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 47 on page 73](#).
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.

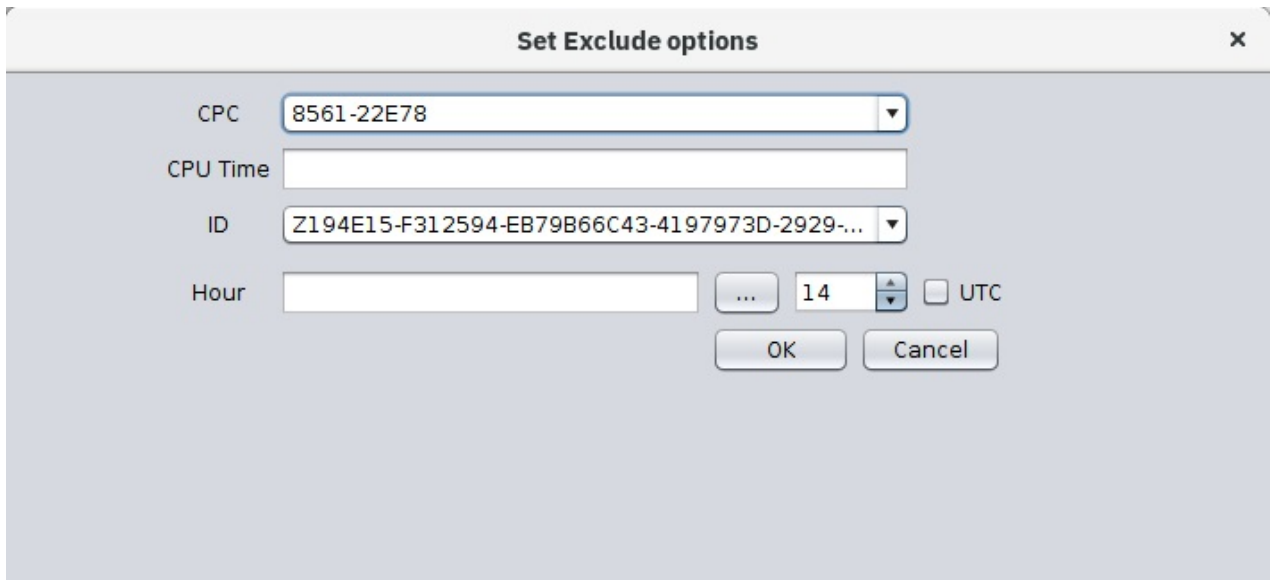
A screenshot of a 'Set Exclude options' dialog box. The dialog has a title bar with a close button (X). Inside, there are four input fields: 'CPC' with a dropdown menu showing '8561-22E78', 'CPU Time' with an empty text box, 'ID' with a dropdown menu showing 'Z194E15-F312594-EB79B66C43-4197973D-2929-...', and 'Hour' with an empty text box. To the right of the 'Hour' field is a button with three dots, followed by a spinner box showing '14', and a checkbox labeled 'UTC'. At the bottom right are 'OK' and 'Cancel' buttons.

Figure 47. 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 48 on page 74](#) 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 48. 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 49 on page 75](#).




Figure 49. 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 50 on page 76](#) shows an example of the exclude panel.




Figure 50. 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.

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 2 on page 77 describes the command line options.

Table 2. SCRT command line options			
Long form	Short form	Required	Description
--container	<i>none</i>	no	For Tailored Fit Pricing for IBM Z, associates an LPAR with a container solution ID. For details, see ./containerPricingcliopts.dita#containerPricingcliopts/cont.
--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 ./containerPricingcliopts.dita#containerPricingcliopts/ignore.
--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 ./containerPricingcliopts.dita#containerPricingcliopts/upcont.
--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 ./containerPricingcliopts.dita#containerPricingcliopts/upsol.

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=Z1nnnnnn-nnnnnnnn-nnnnnnnnnn-nnnnnnnn-nnnn-nnnn-nnnn-nnnnnn-nnnnnn
```

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=D911111-N4ED9AD-DB3D215CAF-NTZZZZZZ-07F5-4EDF-AF34-E062EB-0ADCD7,  
SET_NAME="Container-Name"
```

- Specifying a single NO89 product:

```
--update-container ID=D911111-N4ED9AD-DB3D215CAF-NTZZZZZZ-07F5-4EDF-AF34-E062EB-0ADCD7,  
SET_NAME="Container-Name",NO89_PRODUCT_ID=nnnn-nnn
```

- Specifying multiple NO89 products:

```
--update-container ID=D911111-N4ED9AD-DB3D215CAF-NTZZZZZZ-07F5-4EDF-AF34-E062EB-0ADCD7,  
SET_NAME="Container-Name",NO89_PRODUCT_ID=(nnnn-nnn,mmmm-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-mmmmmmmmm-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.

Chapter 5. Using workload reporting services to collect Container Services performance data

The workload reporting services are intended for monitoring or reporting products to collect performance data. They replace some of the existing methods of collecting data, and provide a means to collect data for the goal-oriented processing with the service policy.

When to use the workload reporting services

Because the data collection is cumulative, performance monitors can collect information based on their own reporting intervals. But the collection is stopped and re-started when a significant change occurs in workload management, such as when a new policy is activated, or a system failure occurs. Performance monitors should “bookend” their intervals when such a significant change occurs in workload management. For each of these events, event notification facility (ENF) signals are issued. SRM samples the states of address spaces, and an ENF signal is issued when a new set of samples is available. The performance monitor can use the ENF signals to guide its reporting interval. For example, when an ENF code for a new policy activation is issued, the performance monitor can end its last reporting interval, and start the next reporting interval for the newly activated service policy.

To enable the performance monitor to know when a workload reporting change is taking place, such as when a policy is activated, there is a ENF system event code. ENF event code 41 and its related qualifiers indicate when changes are taking place related to service policies. ENF code 41 guides the performance monitor's reporting intervals, and helps it to issue the services at the appropriate times.

Interpreting report class data

Through classification rules WLM allows you to associate transactions with report classes or tenant report classes for reporting purposes. Since service IWMRCOLL processes tenant report classes like report classes, the term report class always refers to both report classes and tenant report classes in the following.

Report classes can be used to report on a subset of transactions running in a single service class but also combine transactions running in different service classes within one report class. In the first case a report class is called homogeneous, and in the second case it is called heterogeneous.

A report class period is homogeneous if there is only one service class found contributing to this report class period in a given report interval. To allow a reporting product to determine whether a report class period is homogeneous in its reporting interval, WLM offers two indicators returned by IWMRCOLL:

mixed-class-indication timestamp

This timestamp indicates when workload data associated with a different service class last contributed data to a report class period that was currently collecting data from another service class (see [Figure 51 on page 80](#)).

service class index

This index indicates the last service class whose data was collected in the report class period.

[Figure 51 on page 80](#) illustrates the concept of the mixed-class-indication timestamp in relation to the time interval in which a caller collects workload data.




Figure 51. Mixed-class-indication timestamp in relation to the time interval

A caller invokes IWMRCOLL twice in order to get interval data, first at time t_0 to start the interval and second at time t_1 to end collecting data. With the second invocation at t_1 , the caller gets back the mixed-class-indication timestamp. If the returned timestamp is $\geq t_0$, as it is for caller 1, it means that transaction data from a different service class started contributing data to the same report class period. The report class is heterogeneous. If the returned timestamp is outside the interval (smaller than t_0), as it is for caller 2, it means that the report class remained homogenous during the calling interval.

Being able to see that a period is homogeneous allows the reporting product to format response time distribution buckets and work manager delay data for this period while it would not report this data for a heterogeneous period. If the timestamp indicates that the report class is heterogeneous, it is recommended to collapse the periods which means to report the data as if the report class had only one period.

Using the IWM4QTNT service

The IWM4QTNT service provides CPU service consumption of tenant resource groups defined in the WLM service definition. Long-term average service on general purpose processors used by the tenant resource groups is provided as well as the consumption on specialty engines.

Table 3 on page 80 shows a sample sequence of how you can use IWM4QTNT to obtain CPU service consumption of tenant resource groups on a single system.

Table 3. Using IWM4QTNT with the workload reporting services

Action	Reason
Issue REQSRMST SYSEVENT	To get information about IWM4QTNT availability
Set up a reporting interval	To prepare for subsequent IWM4QTNT requests
Issue IWM4QTNT specifying ANSLEN and QUERYLEN	To obtain length of storage needed. IWM4QTNT returns ANSTOKN required for subsequent calls to IWM4QTNT
Issue GETMAIN	To get storage needed to hold information returned by IWM4QTNT
Issue IWM4QTNT ANSTOKN=token	To get CPU service consumption data per tenant resource group mapped by IWMWQTAA

IWM4QTNT – Query tenant resource group consumption

IWM4QTNT is the interface reporting products should use to obtain CPU service consumption of tenant resource groups defined in the WLM service definition. Long-term average service on general purpose processors used by the tenant resource groups is provided as well as the consumption on specialty engines.

To help the caller keep track of changes in workload management, this service returns a token, ANSTOKN. ANSTOKN is a required input on all subsequent calls to IWM4QTNT. When a change occurs in workload management, for example, when a new policy is activated, IWM4QTNT returns a new token value. The caller's code should check the reason codes to see if the ANSTOKN has changed since the last call to IWM4QTNT. If the token has changed, the performance monitor should reset its reporting interval. If the token has not changed, the performance monitor can continue with its existing reporting interval.

There are also some ENF event codes to keep track of changes in workload management. For information about the ENF codes, see "ENF event codes and meanings" in *z/OS MVS Programming: Authorized Assembler Services Guide*

The caller must provide an area of storage in the ANSAREA parameter and the length of that area in the ANSLLEN parameter for IWM4QTNT to return the CPU consumption. IWM4QTNT return the actual length of the information in the QUERYLEN parameter. The answer area is mapped by the IWMWQTAA data area.

If the storage area provided is insufficient, no data is returned by IWM4QTNT but an appropriate return and reason code is issued and the required amount of storage is returned in the QUERYLEN parameter. If a user does not know the size of the answer area required by the service, he should issue IWM4QTNT with ANSLLEN set to zero. The length of the answer area will be placed in QUERYLEN.

Environment

The requirements for the caller are:

Minimum authorization:	Supervisor state or program key mask (PKM) allowing keys 0-7
Dispatchable unit mode:	Task
Cross memory mode:	Any PASN, any HASN, any SASN
AMODE:	64-bit
ASC mode:	Primary or access register. If in access register ASC mode, specify SYSSTATE ASCENV = AR before invoking IWM4QTNT.
Interrupt status:	Enabled for I/O and external interrupts
Locks:	No locks held.
Control parameters:	Control parameters must be in the primary address space. The caller of IWM4QTNT must provide storage for an answer area mapped by IWMWQTAA. This answer area may reside in the caller's primary address space, or in a dataspace accessible via the current unit of work's dispatchable unit access list (DUAl).

Programming requirements

You must include the CVT and the IWMYCON mapping macros in the program.

Restrictions

The caller cannot have an EUT FRR established.

Input register information

Before issuing the IWM4QTNT macro, the caller does not have to place any information into any register unless using it in register notation for a particular parameter, using it as a base register, or using it to provide the ALET of the storage area.

Output register information

When control returns to the caller, the GPRs contain:

Register

Contents

0

Reason code if GR15 return code is non-zero

1

Used as work registers by the system

2-13

Unchanged

14

Used as work registers by the system

15

Return code

When control returns to the caller, the ARs contain:

Register

Contents

0-1

Used as work registers by the system

2-13

Unchanged

14-15

Used as work registers by the system

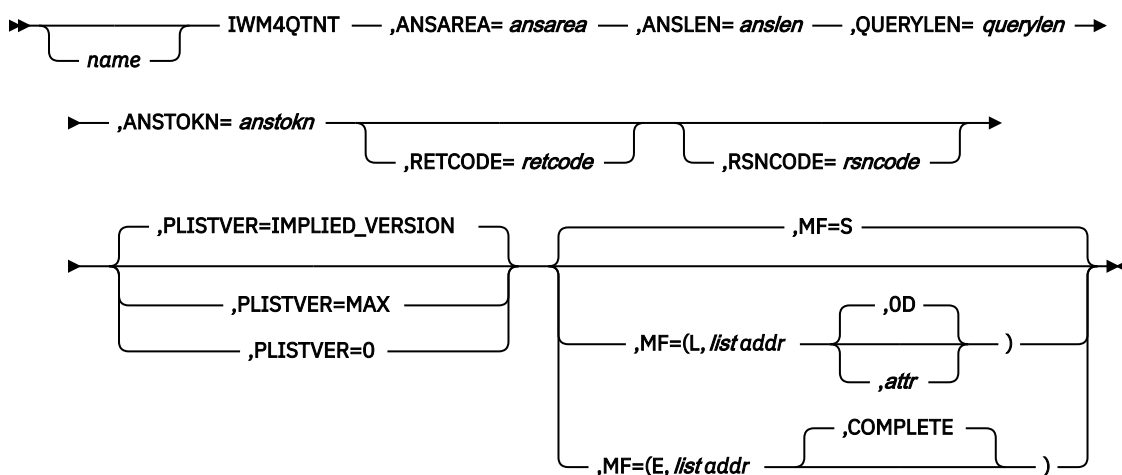
Some callers depend on register contents remaining the same before and after issuing a service. If the system changes the contents of registers on which the caller depends, the caller must save them before issuing the service, and restore them after the system returns control.

Performance implications

None.

Syntax

The syntax of the IWM4QTNT macro is as follows:



Parameters

The parameters are explained as follows:

name

An optional symbol, starting in column 1, that is the name on the IWM4QTNT macro invocation. The name must conform to the rules for an ordinary assembler language symbol.

,ANSAREA=ansarea

A required output parameter that contains the address of a storage area to hold the information returned by IWM4QTNT. The area is mapped by the IWMWQTAA mapping macro.

To code: Specify the RS-type address, or address in register (2)-(12) of a character field.

,ANSLEN=anslen

A required input parameter that contains the length of the storage area (answer area) you are providing on ANSAREA.

To code: Specify the RS-type address or address in register (2)-(12) of a fullword field.

,ANSTOKN=anstoken

A required input/output parameter that contains a token value. On your first call to IWM4QTNT, you specify ANSTOKEN as an output parameter. IWM4QTNT provides a token value that is required for subsequent calls to IWM4QTNT.

To code: Specify the RS-type address, or address in register (2)-(12) of an 8-character field.

,MF=S

,MF=(L,list addr)

,MF=(L,list addr,attr)

,MF=(L,list addr,OD)

,MF=(E,list addr)

,MF=(E,list addr,COMPLETE)

An optional input parameter that specifies the macro form.

Use MF=S to specify the standard form of the macro, which builds an inline parameter list and generates the macro invocation to transfer control to the service. MF=S is the default.

Use MF=L to specify the list form of the macro. Use the list form together with the execute form of the macro for applications that require reentrant code. The list form defines an area of storage that the execute form uses to store the parameters. Only the PLISTVER parameter may be coded with the list form of the macro.

Use MF=E to specify the execute form of the macro. Use the execute form together with the list form of the macro for applications that require reentrant code. The execute form of the macro stores the parameters into the storage area defined by the list form, and generates the macro invocation to transfer control to the service.

,list addr

The name of a storage area to contain the parameters. For MF=S and MF=E, this can be an RS-type address or an address in register (1)-(12).

,attr

An optional 1- to 60-character input string that you use to force boundary alignment of the parameter list. Use a value of 0F to force the parameter list to a word boundary, or 0D to force the parameter list to a doubleword boundary. If you do not code *attr*, the system provides a value of 0D.

,COMPLETE

Specifies that the system is to check for required parameters and supply defaults for omitted optional parameters.

,PLISTVER=IMPLIED_VERSION

,PLISTVER=MAX

,PLISTVER=0

An optional input parameter that specifies the version of the macro. PLISTVER determines which parameter list the system generates. PLISTVER is an optional input parameter on all forms of the

macro, including the list form. When using PLISTVER, specify it on all macro forms used for a request and with the same value on all of the macro forms. The values are:

- **IMPLIED_VERSION**, which is the lowest version that allows all parameters specified on the request to be processed. If you omit the PLISTVER parameter, IMPLIED_VERSION is the default.
- **MAX**, if you want the parameter list to be the largest size currently possible. This size might grow from release to release and affect the amount of storage that your program needs.

If you can tolerate the size change, IBM recommends that you always specify PLISTVER=MAX on the list form of the macro. Specifying MAX ensures that the list-form parameter list is always long enough to hold all the parameters you might specify on the execute form, when both are assembled with the same level of the system. In this way, MAX ensures that the parameter list does not overwrite nearby storage.

- **0**, if you use the currently available parameters.

To code: Specify one of the following:

- IMPLIED_VERSION
- MAX
- A decimal value of 0

,QUERYLEN=querylen

A required output parameter that contains the length of the storage area required by IWM4QTNT to contain all the performance data for active tenants while the ANSTOKN is valid.

To code: Specify the RS-type address, or address in register (2)-(12) of a fullword field.

,RETCODE=retcode

An optional output parameter into which the return code is to be copied from GPR 15. If you specify 15, GPR15, REG15, or R15 (with or without parentheses), the value will be left in GPR 15.

To code: Specify the RS-type address of a fullword field, or register (2)-(12), or (15), (GPR15), (REG15), or (R15).

,RSNCODE=rsncode

An optional output parameter into which the reason code is to be copied from GPR 0. If you specify 0, 00, GPR0, GPR00, REG0, REG00, or R0 (with or without parentheses), the value will be left in GPR 0.

To code: Specify the RS-type address of a fullword field, or register (0) or (2)-(12), (00), (GPR0), (GPR00), REG0), (REG00), or (R0).

ABEND codes

None.

Return codes and reason codes

The following table identifies the hexadecimal return and reason codes and the equate symbol associated with each reason code. IBM support personnel may request the entire reason code, including the **xxxx** value.

Table 4. Return and Reason Codes for the IWM4QTNT Macro		
Return code	Reason code	Equate symbol, Meaning, and action
0	—	Equate Symbol: IwmRetCodeOk Meaning: Successful completion. All requested data returned. Action: None required.
4	—	Equate Symbol: IwmRetCodeWarning Meaning: Successful completion, unusual conditions noted.

Table 4. Return and Reason Codes for the IWM4QTNT Macro (continued)

Return code	Reason code	Equate symbol, Meaning, and action
4	xxxx040A	<p>Equate Symbol: IwmRsnCodeOutputAreaTooSmall</p> <p>Meaning: The output area supplied is too small to receive all the available information. The correct answer area length is returned in the QUERYLEN field.</p> <p>Action: None required. If necessary, reinvoke the service with an output area of sufficient size to receive all information.</p>
4	xxxx040F	<p>Equate Symbol: IwmRsnCodeStateInvDataRet</p> <p>Meaning: The token value specified on the ANSTOKN keyword is associated with a WLM state that is no longer valid. The new system state is represented by the token returned in the ANSTOKN field. The answer area provided is large enough to contain the available data. However, the new answer area length is returned in the QUERYLEN field.</p> <p>Action: Reinvoke the service with the token passed with the ANSTOKN keyword.</p>
8	—	<p>Equate Symbol: IwmRetCodeInvocError</p> <p>Meaning: Invalid invocation environment or parameters.</p>
8	xxxx0801	<p>Equate Symbol: IwmRsnCodeSrbMode</p> <p>Meaning: The caller is in SRB mode.</p> <p>Action: Avoid requesting this function while in SRB mode.</p>
8	xxxx0803	<p>Equate Symbol: IwmRsnCodeDisabled</p> <p>Meaning: The caller is disabled.</p> <p>Action: Avoid requesting this function while disabled.</p>
8	xxxx0804	<p>Equate Symbol: IwmRsnCodeLocked</p> <p>Meaning: The caller is locked.</p> <p>Action: Avoid requesting this function while locked.</p>
8	xxxx0808	<p>Equate Symbol: IwmRsnBadPl</p> <p>Meaning: Error accessing parameter list.</p> <p>Action: Check for possible storage overlay.</p>
8	xxxx0810	<p>Equate Symbol: IwmRsnCodeEutFrr</p> <p>Meaning: The caller has an EUT FRR set.</p> <p>Action: Avoid requesting this function with an EUT FRR set.</p>
8	xxxx0830	<p>Equate Symbol: IwmRsnCodeBadAlet</p> <p>Meaning: The caller specified an invalid alet for the storage pointed to by the ANSAREA keyword.</p> <p>Action: Check for possible storage overlay of the parameter list or variable.</p>
8	xxxx0832	<p>Equate Symbol: IwmRsnCodeStateInvNoDatRet</p> <p>Meaning: The token value specified on the ANSTOKN keyword is associated with a WLM state that is no longer valid. A new token has been returned. The storage provided is not large enough to contain all of the data available because of the state change. No data was returned. The length of the new answer area required is returned in the QUERYLEN field.</p> <p>Action: Reinvoke the service with an output area of sufficient size to receive all information and the token passed with the ANSTOKN keyword.</p>

Table 4. Return and Reason Codes for the IWM4QTNT Macro (continued)

Return code	Reason code	Equate symbol, Meaning, and action
C	—	Equate Symbol: IwmRetCodeEnvError Meaning: Environmental error
C	xxxx0C01	Equate Symbol: IwmRsnCodeNoStg Meaning: Storage is not available for the request. Action: There is a storage shortage. The function may work successfully at a later time.
C	xxxx0C0A	Equate Symbol: IwmRsnCodeSuspended Meaning: Data collection is suspended as a result of a component error. No data can be returned for this IWM4QTNT invocation, future invocations may be successful. Action: Reinvoke this service.
10	—	Equate Symbol: IwmRetCodeCompError Meaning: Component error Action: No action required. The function may work successfully if invoked again.

Example

For tenant consumption information from a system, specify:

```
IWM4QTNT  ANSAREA=(R6), ANSLN=(R8),
           QUERYLEN=QLEN, ANSTOKN=ATOKN,
           RETCODE=RC, RSNCODE=RSN, MF=(E, MFQTNT)
```

Mapping a service definition

The service definition is installed and extracted from the WLM couple data set either in XML format, or as a data area mapped by the IWMSERVD mapping macro. The XML structure is defined by the DTD described in “Structure of the XML service definition (DTD)” on page 87. The IWMSERVD macro points to the following sections:

IWMSVDEF

Maps the following service definition information:

- Service policies
- Workloads
- Service classes
- Report classes and tenant report classes
- Resource groups and tenant resource groups

IWMSVDCR

Maps the service definition classification rule information.

IWMSVNPA

Maps the service definition notepad area.

IWMSVAEA

Maps the service definition application environments.

IWMSVSEA

Maps the service definition scheduling environments.

Querying the active classification rules

IWMCQRY lets a caller query the classification rules associated with the active service policy. The classification rules determine how incoming work is assigned a service class and/or report class or tenant report class. The data returned by this macro is mapped by IWMSVDCR. For a description of the IWMSVDCR macro, see *z/OS MVS Data Areas, Vol 3*.

Optionally, a caller can request the active service policy identifier by specifying the **POLICY_ID** parameter. This is the active policy containing the classification rules returned by this macro. The caller can then compare the service policy ID with the policy identification information returned by the IWMPQRY macro to ensure they are the same.

Some data sections in the IWMSVDCR data area may not be available through IWMCQRY. For example, the time stamps indicating when a classification GROUP was last updated and by whom may not be available. For a complete list of fields that are not available refer to IWMSVDCR as described in *z/OS MVS Data Areas, Vol 3*. The complete classification rules associated with a service policy are returned by the IWMDEXTR macro and mapped by IWMSVDCR.

A caller can use the classification rules together with the active service policy to determine the goals associated with incoming work. The service class goals are in the active service policy mapping returned by the IWMPQRY service.

The information returned is not serialized upon return to the caller, and so may be out of date if a service definition was modified, installed, and a new policy activated.

Structure of the XML service definition (DTD)

This section describes the following:

- The structure of the XML output of IWMDEXTR
- The layout of the XML service definition (DTD) that can be passed to IWMDINST

To obtain XML output, specify the TYPE=XML parameter for the IWMDEXTR service.

To install such an XML service definition with the IWMDINST service, also specify the TYPE=XML parameter.

The following DTD defines the structure of an XML service definition:

```
<!ELEMENT ServiceDefinition ( Name, Description?, CreationDate?, CreationUser?,
ModificationDate?, ModificationUser?, Level, ReplId?, ProdId?, Notes,
ResourceGroups, TenantResourceGroups?, Workloads, ServicePolicies, ReportClasses, TenantReportClasses?,
ClassificationGroups, Classifications, ServiceParameter,
ApplicationEnvironments?, Resources?, SchedulingEnvironments?, GMPSettings?,
Extensions? ) >

<!ATTLIST ServiceDefinition
    xmlns          CDATA #IMPLIED
    codepage        CDATA #IMPLIED >

<!ELEMENT Name ( #PCDATA ) >
<!ELEMENT Description ( #PCDATA ) >
<!ELEMENT CreationDate ( #PCDATA ) >
<!ELEMENT CreationUser ( #PCDATA ) >
<!ELEMENT ModificationDate ( #PCDATA ) >
<!ELEMENT ModificationUser ( #PCDATA ) >
<!ELEMENT Level ( #PCDATA ) >
<!ELEMENT ReplId ( #PCDATA ) >
<!ELEMENT ProdId ( #PCDATA ) >

<!ELEMENT Notes ( Note* ) >
<!ELEMENT Note ( #PCDATA ) >

<!ELEMENT ResourceGroups ( ResourceGroup* ) >
<!ELEMENT ResourceGroup ( Name, Description?, CreationDate, CreationUser,
ModificationDate, ModificationUser, Type?, CapacityMinimum?,
CapacityMaximum?, MemoryLimit?, IncludeSpecialtyProcessorConsumption? ) >

<!ELEMENT TenantResourceGroups ( TenantResourceGroup* ) >
<!ELEMENT TenantResourceGroup (
Name, Description?, SolutionId?, TenantId?, TenantName?,
CreationDate, CreationUser,
ModificationDate, ModificationUser, Type?, CapacityMinimum?,
CapacityMaximum?, MemoryLimit?, IncludeSpecialtyProcessorConsumption? ) >
```

```

<!ELEMENT Type ( #PCDATA ) >

<!ELEMENT SolutionId ( #PCDATA ) >
<!ELEMENT TenantId ( #PCDATA ) >
<!ELEMENT TenantName ( #PCDATA ) >

<!ELEMENT CapacityMaximum ( #PCDATA ) >
<!ELEMENT CapacityMinimum ( #PCDATA ) >
<!ELEMENT MemoryLimit ( #PCDATA ) >

<!ELEMENT IncludeSpecialtyProcessorConsumption ( #PCDATA ) >

<!ELEMENT Workloads ( Workload* ) >
<!ELEMENT Workload ( Name, Description?, CreationDate, CreationUser,
    ModificationDate, ModificationUser, ServiceClasses ) >

<!ELEMENT ServiceClasses ( ServiceClass* ) >
<!ELEMENT ServiceClass ( Name, Description?, CreationDate, CreationUser,
    ModificationDate, ModificationUser, CPUCritical?, IOPriorityGroup?,
    HonorPriority?, ResourceGroupName?, Goal ) >

<!ELEMENT ResourceGroupName ( #PCDATA ) >

<!ELEMENT Goal ( (AverageResponseTime | PercentileResponseTime | Velocity)*,
    Discretionary? ) >

<!ELEMENT AverageResponseTime ( Importance, Duration?, ResponseTime ) >

<!ELEMENT Importance ( #PCDATA ) >
<!ELEMENT Duration ( #PCDATA ) >
<!ELEMENT ResponseTime ( #PCDATA ) >

<!ELEMENT PercentileResponseTime ( Importance, Duration?, ResponseTime,
    Percentile ) >

<!ELEMENT Percentile ( #PCDATA ) >

<!ELEMENT Velocity ( Importance, Duration?, Level ) >

<!ELEMENT Discretionary EMPTY >

<!ELEMENT ServicePolicies ( ServicePolicy* ) >
<!ELEMENT ServicePolicy ( Name, Description?, CreationDate, CreationUser,
    ModificationDate, ModificationUser, ServiceClassOverrides,
    ResourceGroupOverrides, TenantResourceGroupOverrides? ) >

<!ELEMENT ServiceClassOverrides ( ServiceClassOverride* ) >
<!ELEMENT ServiceClassOverride ( ServiceClassName, CPUCritical?,
    IOPriorityGroup?, HonorPriority?, ResourceGroupName?, Goal ) >

<!ELEMENT ServiceClassName ( #PCDATA ) >
<!ELEMENT CPUCritical ( #PCDATA ) >
<!ELEMENT IOPriorityGroup ( #PCDATA ) >
<!ELEMENT HonorPriority ( #PCDATA ) >

<!ELEMENT ResourceGroupOverrides ( ResourceGroupOverride* ) >
<!ELEMENT ResourceGroupOverride ( ResourceGroupName, Type?,
    CapacityMinimum?, CapacityMaximum?, MemoryLimit??,
    IncludeSpecialtyProcessorConsumption? ) >

<!ELEMENT TenantResourceGroupOverrides ( TenantResourceGroupOverride* ) >
<!ELEMENT TenantResourceGroupOverride ( TenantResourceGroupName, Type?,
    CapacityMinimum?, CapacityMaximum?, MemoryLimit?,
    IncludeSpecialtyProcessorConsumption? ) >

<!ELEMENT ReportClasses ( ReportClass* ) >
<!ELEMENT ReportClass ( Name, Description?, CreationDate,
    CreationUser, ModificationDate, ModificationUser ) >

<!ELEMENT TenantReportClasses ( TenantReportClass* ) >
<!ELEMENT TenantReportClass ( Name, TenantResourceGroupName,
    Description?, CreationDate,
    CreationUser, ModificationDate, ModificationUser ) >

<!ELEMENT TenantResourceGroupName ( #PCDATA ) >

<!ELEMENT ClassificationGroups ( ClassificationGroup* ) >
<!ELEMENT ClassificationGroup ( Name, Description?, CreationDate,
    CreationUser, ModificationDate, ModificationUser, QualifierType,
    QualifierNames ) >

<!ELEMENT QualifierType ( #PCDATA ) >

<!ELEMENT QualifierNames ( QualifierName* ) >
<!ELEMENT QualifierName ( Name, Description?, Start? ) >

<!ELEMENT Classifications ( Classification* ) >
<!ELEMENT Classification ( SubsystemType, Description?, CreationDate,
    CreationUser, ModificationDate, ModificationUser,
    DefaultServiceClassName?, DefaultReportClassName?, EWLClassification?,
    ClassificationRules? ) >

<!ELEMENT SubsystemType ( #PCDATA ) >

```



```

<!ELEMENT DefaultServiceClassName ( #PCDATA ) >
<!ELEMENT DefaultReportClassName ( #PCDATA ) >

<!ELEMENT ClassificationRules ( ClassificationRule* ) >
<!ELEMENT ClassificationRule ( Description?, QualifierType, QualifierValue,
    Start?, ServiceClassName?, ReportClassName?, StorageCritical?, RegionGoal?,
    ReportingAttribute?, Boost?, ClassificationRule* ) >

<!ELEMENT QualifierValue ( #PCDATA ) >
<!ELEMENT Start ( #PCDATA ) >
<!ELEMENT ReportClassName ( #PCDATA ) >
<!ELEMENT RegionGoal ( #PCDATA ) >
<!ELEMENT StorageCritical ( #PCDATA ) >
<!ELEMENT Boost ( #PCDATA ) >
<!ELEMENT ReportingAttribute ( #PCDATA ) >
<!ELEMENT ServiceParameter ( ServiceCoefficients, ServiceOptions? ) >

<!ELEMENT ServiceCoefficients ( CPU, IOC, MSO, SRB )? >

<!ELEMENT CPU ( #PCDATA ) >
<!ELEMENT IOC ( #PCDATA ) >
<!ELEMENT MSO ( #PCDATA ) >
<!ELEMENT SRB ( #PCDATA ) >

<!ELEMENT EWLClassification ( #PCDATA ) >

<!ELEMENT ServiceOptions ( IOPriorityManagement, DynamicAliasManagement?,
    IOPriorityGroupsEnabled?, DeactivateDiscretionaryGoalManagement? ) >

<!ELEMENT IOPriorityManagement ( #PCDATA ) >
<!ELEMENT DynamicAliasManagement ( #PCDATA ) >
<!ELEMENT IOPriorityGroupsEnabled ( #PCDATA ) >
<!ELEMENT DeactivatedDiscretionaryGoalManagement ( #PCDATA ) >

<!ELEMENT ApplicationEnvironments ( ApplicationEnvironment* ) >
<!ELEMENT ApplicationEnvironment ( Name, Description?, SubsystemType, Limit,
    ProcedureName?, StartParameter? ) >

<!ELEMENT StartParameter ( #PCDATA ) >
<!ELEMENT Limit ( #PCDATA ) >
<!ELEMENT ProcedureName ( #PCDATA ) >

<!ELEMENT Resources ( Resource* ) >
<!ELEMENT Resource ( Name, Description? ) >

<!ELEMENT SchedulingEnvironments ( SchedulingEnvironment* ) >
<!ELEMENT SchedulingEnvironment ( Name, Description?, ResourceNames ) >

<!ELEMENT ResourceNames ( ResourceName* ) >
<!ELEMENT ResourceName ( Name, RequiredState ) >

<!ELEMENT RequiredState ( #PCDATA ) >

<!ELEMENT GMPSettings ( Activation, ExcludedHostSystems? ) >
<!ELEMENT Activation ( #PCDATA ) >
<!ELEMENT ExcludedHostSystems ( ExcludedHostSystem* ) >
<!ELEMENT ExcludedHostSystem ( Name ) >

<!ELEMENT Extensions ( ServiceDefinitionExtensions?,
    ResourceGroupExtensions?, ResourceGroupAttributeExtensions?,
    WorkloadExtensions?, ServiceClassExtensions?,
    ServiceClassAttributeExtensions?, ServicePolicyExtensions?,
    ReportClassExtensions?, ClassificationExtensions?,
    ApplicationEnvironmentExtensions?, ResourceExtensions?,
    SchedulingEnvironmentHeaderExtensions?,
    SchedulingEnvironmentExtensions?,
    SchedulingEnvironmentResourceExtensions? ) >

<!ELEMENT ServiceDefinitionExtensions ( ServiceDefinitionExtension* ) >
<!ELEMENT ServiceDefinitionExtension ( VendorId?, RelatedObject,
    ExtensionData? ) >

<!ELEMENT VendorId ( #PCDATA ) >
<!ELEMENT RelatedObject ( #PCDATA ) >
<!ELEMENT ExtensionData ( #PCDATA ) >

<!ELEMENT ResourceGroupExtensions ( ResourceGroupExtension* ) >
<!ELEMENT ResourceGroupExtension ( VendorId?, RelatedObject, ServicePolicyName?,
    ExtensionData? ) >

<!ELEMENT ServicePolicyName ( #PCDATA ) >

<!ELEMENT ResourceGroupAttributeExtensions ( ResourceGroupAttributeExtension* ) >
<!ELEMENT ResourceGroupAttributeExtension ( VendorId?, RelatedObject,
    ServicePolicyName?, ExtensionData? ) >

<!ELEMENT WorkloadExtensions ( WorkloadExtension* ) >
<!ELEMENT WorkloadExtension ( VendorId?, RelatedObject, ServicePolicyName?,
    ExtensionData? ) >

<!ELEMENT ServiceClassExtensions ( ServiceClassExtension* ) >
<!ELEMENT ServiceClassExtension ( VendorId?, RelatedObject, ServicePolicyName?,
    ExtensionData? ) >

```

```

<!ELEMENT ServiceClassAttributeExtensions ( ServiceClassAttributeExtension* ) >
<!ELEMENT ServiceClassAttributeExtension (VendorId?, RelatedObject,
ServicePolicyName?, ExtensionData?) >

<!ELEMENT ServicePolicyExtensions ( ServicePolicyExtension* ) >
<!ELEMENT ServicePolicyExtension (VendorId?, RelatedObject, ServicePolicyName?,
ExtensionData?) >

<!ELEMENT ReportClassExtensions ( ReportClassExtension* ) >
<!ELEMENT ReportClassExtension (VendorId?, RelatedObject, ServicePolicyName?,
ExtensionData?) >

<!ELEMENT ClassificationExtensions ( ClassificationExtension* ) >
<!ELEMENT ClassificationExtension (VendorId?, RelatedObject, ExtensionData?) >

<!ELEMENT ApplicationEnvironmentExtensions ( ApplicationEnvironmentExtension* ) >
<!ELEMENT ApplicationEnvironmentExtension (VendorId?, RelatedObject,
ExtensionData?) >

<!ELEMENT ResourceExtensions ( ResourceExtension* ) >
<!ELEMENT ResourceExtension (VendorId?, RelatedObjectName?, ExtensionData?) >

<!ELEMENT SchedulingEnvironmentHeaderExtensions
( SchedulingEnvironmentHeaderExtension* ) >
<!ELEMENT SchedulingEnvironmentHeaderExtension
(VendorId?, RelatedObject, ExtensionData?) >

<!ELEMENT SchedulingEnvironmentExtensions
( SchedulingEnvironmentExtension* ) >
<!ELEMENT SchedulingEnvironmentExtension
(VendorId?, RelatedObject, ExtensionData?) >

<!ELEMENT SchedulingEnvironmentResourceExtensions
( SchedulingEnvironmentResourceExtension* ) >
<!ELEMENT SchedulingEnvironmentResourceExtension
(VendorId?, RelatedObject, ExtensionData?) >

```

Table 5 on page 90 lists the valid name spaces and the corresponding functionality levels:

Table 5. Valid name spaces and corresponding functionality levels	
Name space	Level
http://www.ibm.com/xmlns/prod/zwlm/1993/09/ServiceDefinition.xsd	001
http://www.ibm.com/xmlns/prod/zwlm/1994/09/ServiceDefinition.xsd	002
http://www.ibm.com/xmlns/prod/zwlm/1997/03/ServiceDefinition.xsd	003
http://www.ibm.com/xmlns/prod/zwlm/1997/09/ServiceDefinition.xsd	004
http://www.ibm.com/xmlns/prod/zwlm/1998/09/ServiceDefinition.xsd	006
http://www.ibm.com/xmlns/prod/zwlm/1999/03/ServiceDefinition.xsd	007
http://www.ibm.com/xmlns/prod/zwlm/1999/09/ServiceDefinition.xsd	008
http://www.ibm.com/xmlns/prod/zwlm/2000/09/ServiceDefinition.xsd	011
http://www.ibm.com/xmlns/prod/zwlm/2001/09/ServiceDefinition.xsd	013
http://www.ibm.com/xmlns/prod/zwlm/2005/12/ServiceDefinition.xsd	017
http://www.ibm.com/xmlns/prod/zwlm/2006/09/ServiceDefinition.xsd	019
http://www.ibm.com/xmlns/prod/zwlm/2008/09/ServiceDefinition.xsd	021
http://www.ibm.com/xmlns/prod/zwlm/2009/09/ServiceDefinition.xsd	023
http://www.ibm.com/xmlns/prod/zwlm/2010/09/ServiceDefinition.xsd	025
http://www.ibm.com/xmlns/prod/zwlm/2012/09/ServiceDefinition.xsd	029
http://www.ibm.com/xmlns/prod/zwlm/2015/12/ServiceDefinition.xsd	030
http://www.ibm.com/xmlns/prod/zwlm/2016/12/ServiceDefinition.xsd	031
http://wlm.ibm.com/xmlns/prod/zwlm/2017/12/ServiceDefinition.xsd	032

<i>Table 5. Valid name spaces and corresponding functionality levels (continued)</i>	
Name space	Level
http://www.ibm.com/xmlns/prod/zwlm/2017/09/ServiceDefinition.xsd	035
http://www.ibm.com/xmlns/prod/zwlm/2022/05/ServiceDefinition.xsd	038
http://www.ibm.com/xmlns/prod/zwlm/2022/09/ServiceDefinition.xsd	040

Chapter 6. MVS System Management Facilities (SMF)

record type changes for Tailored Fit Pricing

Tailored Fit Pricing updates or adds the following SMF record types:

- Record type 70 (X'46') , subtype 1
- Record type 72 (X'48'), subtype 3
- Record type 79 (X'4F'), subtypes 1, 2, and 5
- Record type 89 (X'59'), subtype 1
- Record type 99 (X'63'), subtypes 1 and 2
- Record type 1156 (X'484'), subtype 1

For complete information about SMF records, see [z/OS MVS System Management Facilities \(SMF\)](#).

Record Type 70 (X'46') – RMF Processor Activity

Tailored Fit Pricing updates SMF record type 70, subtype 1. In this section we hit just the sections updated for Tailored Fit Pricing. For complete information about SMF record type 70, see [z/OS MVS System Management Facilities \(SMF\)](#).

Record type 70 is written for each measurement interval and when the session terminates. As with all SMF records produced by RMF, it contains a header section followed by the RMF product section.

Macro to symbolically address record type 70: The SMF record mapping macro for all records produced by RMF is ERBSMFR. Its format is ERBSMFR (n1,n2,...) where n1,n2, ... are the SMF record types you want to map. Note that the parentheses are required only when two or more record types are specified. The mapping macro resides in SYS1.MACLIB.

For information on using RMF, see [z/OS Resource Measurement Facility User's Guide](#). For information on Monitor I and II, see [z/OS RMF Report Analysis](#).

Record mapping

Header/self-defining section

This section contains the common SMF record headers fields and the triplet fields (offset/length/number), if applicable, that locate the other sections on the record.

Offsets	Name	Length	Format	Description
0	0 SMF70LEN	2	binary	Record length. This field and the next field (total of four bytes) form the RDW (record descriptor word).
2	2 SMF70SEG	2	binary	Segment descriptor (see record length field).

Offsets	Name	Length	Format	Description
4	4 SMF70FLG	1	binary	System indicator: Bit Meaning when set 0 New record format 1 Subtypes used 2 Reserved. 3–6 Version indicators* 7 System is running in PR/SM mode.
5	5 SMF70RTY	1	binary	Record type 70 (X'46').
6	6 SMF70TME	4	binary	Time since midnight, in hundredths of a second, that the record was moved into the SMF buffer.
10	A SMF70DTE	4	packed	Date when the record was moved into the SMF buffer, in the form 0cyyddF.
14	E SMF70SID	4	EBCDIC	System identification (from the SMFPRMxx SID parameter).
18	12 SMF70SSI	4	EBCDIC	Subsystem identification (RMF).
22	16 SMF70STY	2	binary	Record subtype.
24	18 SMF70TRN	2	binary	Number of triplets in this record. A triplet is a set of three SMF fields (offset/length/number values) that defines a section of the record. The offset is the offset from the RDW.
26	1A	2		Reserved.
28	1C SMF70PRS	4	binary	Offset to RMF product section from the RDW.
32	20 SMF70PRL	2	binary	Length of RMF product section.
34	22 SMF70PRN	2	binary	Number of RMF product sections.
Individual header extension for subtype 1:				
36	24 SMF70CCS	4	binary	Offset to CPU control section from RDW.
40	28 SMF70CCL	2	binary	Length of CPU control section.
42	2A SMF70CCN	2	binary	Number of CPU control section.
44	2C SMF70CPS	4	binary	Offset to CPU data section from RDW.
48	30 SMF70CPL	2	binary	Length of CPU data section.
50	32 SMF70CPN	2	binary	Number of CPU data sections in this record.
52	34 SMF70ASS	4	binary	Offset to ASID data area section from RDW.
56	38 SMF70ASL	2	binary	Length of ASID data area section.
58	3A SMF70ASN	2	binary	Number of ASID data area sections.
60	3C SMF70BCS	4	binary	Offset to PR/SM partition data section from RDW.
64	40 SMF70BCL	2	binary	Length of PR/SM partition data section.
66	42 SMF70BCN	2	binary	Number of PR/SM partition data sections.
68	44 SMF70BVS	4	binary	Offset to PR/SM logical processor data section from RDW.
72	48 SMF70BVL	2	binary	Length of PR/SM logical processor data section.
74	4A SMF70BVN	2	binary	Number of PR/SM logical processor data sections.
76	4C SMF70CNS	4	binary	Offset to CPU-identification name sections.
80	50 SMF70CNL	2	binary	Length of CPU-identification name section.

Offsets	Name	Length	Format	Description
82	52 SMF70CNN	2	binary	Number of CPU-identification name sections.
84	54 SMF70COS	4	binary	Offset to logical core data section from RDW.
88	58 SMF70COL	2	binary	Length of logical core data section.
90	5A SMF70CON	2	binary	Number of logical core data sections.
92	5C SMF70TNS	4	binary	Offset to Tenant Resource Group data section from RDW.
96	60 SMF70TNL	2	binary	Length of Tenant Resource Group data section.
98	62 SMF70TNN	2	binary	Number of Tenant Resource Group data sections.
100	64 SMF70WCS	4	binary	Offset to workload classification data section from RDW.
104	68 SMF70WCL	2	binary	Length of workload classification data section.
106	6A SMF70WCN	2	binary	Number of workload classification data sections.
Individual header extension for subtype 2:				
36	24 SMF7023S	4	binary	Offset to Cryptographic CCA Coprocessor data section.
40	28 SMF7023L	2	binary	Length of Cryptographic CCA Coprocessor data section.
42	2A SMF7023N	2	binary	Number of Cryptographic CCA Coprocessor data sections.
44	2C SMF7024S	4	binary	Offset to Cryptographic Accelerator data section.
48	30 SMF7024L	2	binary	Length of Cryptographic Accelerator data section.
50	32 SMF7024N	2	binary	Number of Cryptographic Accelerator data sections.
52	34 SMF702CS	4	binary	Offset to ICSF Services data section.
56	38 SMF702CL	2	binary	Length of ICSF Services data section.
58	3A SMF702CN	2	binary	Number of ICSF Services data sections.
60	3C SMF7025S	4	binary	Offset to Cryptographic PKCS11 Coprocessor data section.
64	40 SMF7025L	2	binary	Length of Cryptographic PKCS11 Coprocessor data section.
66	42 SMF7025N	2	binary	Number of Cryptographic PKCS11 Coprocessor data sections.

Subtype 1 – CPU, PR/SM, and ICF activity

CPU control section

There is one section per record.

Offsets	Name	Length	Format	Description
0	0 SMF70MOD	2	binary	CPU processor family.
2	2 SMF70VER	1	binary	CPU version number — meaning varies with model number.
3	3 SMF70BNP	1	binary	Number of physical processors assigned for use by PR/SM.

Offsets	Name	Length	Format	Description
4	4 SMF70INB	1	binary	<p>PR/SM indicator bits.</p> <p>Bit</p> <p>Meaning when set</p> <p>0 PR/SM diagnose X'204' failure.</p> <p>1 Number of physical processors has changed.</p> <p>2 Dispatch interval time has been changed.</p> <p>3 An additional partition, that is not included in the count of configured partitions, is presented with a name of "PHYSICAL". This partition includes all of the uncaptured time that was used by the LPAR management time support feature but could not be attributed to a specific logical partition.</p> <p>4 PR/SM - Diagnose X'204' extended data is supported.</p> <p>5 Simplified Diagnose X'204' data provided for system running as z/VM guest. CPU consumption by z/VM itself provided with partition data section for logical partition named PHYSICAL.</p> <p>6 Power information is available.</p> <p>7 Failure obtaining power information.</p>
5	5 SMF70STF	1	binary	<p>Flags.</p> <p>Bit</p> <p>Meaning when set</p> <p>0 The STSI facility is available for the CPC.</p> <p>1 Physical CPU adjustment factor has been changed.</p> <p>2 Service units available to MVS image have been changed.</p> <p>3 SMF70LAC is provided for systems running in LPAR mode, as a z/VM guest, or under an alternate virtual machine. The value no longer includes CPU wait times.</p> <p>4 SMF70MDL is the model-capacity identifier and SMF70HWM is the physical model. If this bit is OFF, SMF70MDL represents both model-capacity and physical model.</p> <p>5 OPT parameter BLWLTRPCT changed.</p> <p>6 OPT parameter BLWLINTHD changed.</p> <p>7 Field SMF70GAU is valid.</p>
6	6 SMF70GTS	2	binary	<p>Dispatch accumulated interval time in milliseconds. A zero value indicates that the dispatch interval was dynamically determined.</p>

Offsets	Name	Length	Format	Description																
8	8 SMF70MDL	16	EBCDIC	CPC model identifier indicating the total capacity of the CPC, including all types of active capacity. See bit 4 of SMF70STF.																
24	18 SMF70DSA	2	binary	Number of Diagnose samples.																
26	1A SMF70IFA	2	binary	Number of zAAPs online at the end of the interval.																
28	1C SMF70CPA	4	binary	Physical CPU adjustment factor based on alternate CPU capability. This value is replaced by SMF70CPA_actual and SMF70CPA_scaling_factor.																
32	20 SMF70WLA	4	binary	Processor capacity available to MVS image measured in MSUs (millions of service units) per hour. The value takes into account whether or not the image has a defined capacity limit. (For systems running as VM guest, this is the VM capacity).																
36	24 SMF70LAC	4	binary	Long-term average of CPU service (millions of service units). Scope of the value depends on bit 3 of SMF70STF.																
40	28 SMF70HOF	8	binary	Hypervisor date/time offset in STCK format (aka Sysplex timer offset).																
48	30 SMF70HWM	16	EBCDIC	CPC physical model identifier. Valid if bit 4 of SMF70STF is set.																
64	40 SMF70SUP	2	binary	Number of zIIPs online at the end of the interval.																
66	42 SMF70GJT	8	EBCDIC	Time in STCK format when the partition that wrote this record has joined or left a capacity group (last change of group name). Also set at IPL time, when the partition is not a member of a capacity group.																
74	4A SMF70POM	4	EBCDIC	EBCDIC plant code that identifies the plant of manufacture for the configuration. The plant code is left-justified with trailing blank characters if necessary.																
78	4E SMF70CSC	16	EBCDIC	EBCDIC sequence code of the configuration. The sequence code is right-justified with leading EBCDIC zeroes if necessary.																
94	5E SMF70HHF	1	binary	Additional flags. <table><tr><th>Bit</th><th>Meaning when set</th></tr><tr><td>0</td><td>HiperDispatch mode supported</td></tr><tr><td>1</td><td>HiperDispatch mode is active</td></tr><tr><td>2</td><td>HiperDispatch mode changed during interval</td></tr><tr><td>3</td><td>Failure returned by HISMT service. Values in Logical Core data section and values provided in SMF70MCF, SMF70MCFS, SMF70MCFI, SMF70CF, SMF70CFS, SMF70CFI, SMF70ATD, SMF70ATDS, and SMF70ATDI are invalid.</td></tr><tr><td>4</td><td>Absolute MSU capping is active for this partition.</td></tr><tr><td>5</td><td>SMF70OS_PRTCT is valid.</td></tr><tr><td>6-7</td><td>Reserved.</td></tr></table>	Bit	Meaning when set	0	HiperDispatch mode supported	1	HiperDispatch mode is active	2	HiperDispatch mode changed during interval	3	Failure returned by HISMT service. Values in Logical Core data section and values provided in SMF70MCF, SMF70MCFS, SMF70MCFI, SMF70CF, SMF70CFS, SMF70CFI, SMF70ATD, SMF70ATDS, and SMF70ATDI are invalid.	4	Absolute MSU capping is active for this partition.	5	SMF70OS_PRTCT is valid.	6-7	Reserved.
Bit	Meaning when set																			
0	HiperDispatch mode supported																			
1	HiperDispatch mode is active																			
2	HiperDispatch mode changed during interval																			
3	Failure returned by HISMT service. Values in Logical Core data section and values provided in SMF70MCF, SMF70MCFS, SMF70MCFI, SMF70CF, SMF70CFS, SMF70CFI, SMF70ATD, SMF70ATDS, and SMF70ATDI are invalid.																			
4	Absolute MSU capping is active for this partition.																			
5	SMF70OS_PRTCT is valid.																			
6-7	Reserved.																			
95	5F SMF70CR	1	binary	ZEP field 0.																

Offsets	Name	Length	Format	Description
96 60	SMF70PMI	4	binary	Accumulated number of blocked dispatchable units per second that may get promoted in their dispatch priority. To get the average promote event rate, divide SMF70PMI by SMF70SAM.
100 64	SMF70PMU	4	binary	Number of blocked dispatchable units being promoted during the interval.
104 68	SMF70PMW	4	binary	Accumulated number of address spaces and enclaves being blocked during the interval. To get the average number of waiters for promote, divide SMF70PMW by SMF70SAM.
108 6C	SMF70PMP	4	binary	Maximum number of address spaces and enclaves found being blocked during the interval.
112 70	SMF70PMT	2	binary	1/1000s of the CPU capacity for promote slices (OPT parameter BLWLTRPCT).
114 72	SMF70PML	2	binary	Swapped-in starvation threshold. When an address space or enclave has not received CPU service within this time interval although it has ready-to-run work, it is considered being blocked (OPT parameter BLWLINTHD).
116 74	SMF70MPC	16	EBCDIC	CPC model identifier indicating the permanent capacity of the CPC, without the temporarily increased capacity, the temporarily available replacement capacity, and the corridor capacity. The identifier is left justified with trailing blanks if necessary. This field is zero, if not supported by the hardware.
132 84	SMF70MTC	16	EBCDIC	CPC model identifier indicating the temporary capacity of the CPC, which is the total of permanent capacity and temporarily increased capacity, without the temporarily available replacement capacity and the corridor capacity. The identifier is left justified with trailing blanks if necessary. This field is zero, if not supported by the hardware.
148 94	SMF70MCR	4	binary	CPC model capacity rating associated with the model as identified by SMF70MDL. This field is zero, if not supported by the hardware.
152 98	SMF70MPR	4	binary	CPC permanent model capacity rating associated with the model as identified by SMF70MPC. This field is zero, if not supported by the hardware.
156 9C	SMF70MTR	4	binary	CPC temporary model capacity rating associated with the model as identified by SMF70MTC. This field is zero, if not supported by the hardware.
160 A0	SMF70ZEP	4	binary	ZEP field 1.
164 A4	SMF70ZER	8	binary	ZEP field 2.
172 AC	SMF70ZEE	8	binary	ZEP field 3.
180 B4	SMF70ZEC	8	binary	ZEP field 4.
188 BC	SMF70NRM	4	binary	Normalization factor for zIIP. Multiply zIIP time by this value and divide by 256 to get the equivalent time on a CP.
192 C0	SMF70GAU	4	binary	Long-term average of CPU service in millions of service units which would be allowed by the limit of the capacity group but is not used by its members. If the value is negative, the group is capped. Valid if bit 7 of SMF70STF is set.
196 C4	SMF70ZEI	8	binary	ZEP field 5.

Offsets	Name	Length	Format	Description
204	CC SMF70NCR	4	binary	Nominal model-capacity rating in MSU/hour. When non-zero, this value is associated with the nominal model capacity as identified in field SMF70MDL. When field SMF70CAI contains a value of 100, this value equals the value in field SMF70MCR.
208	D0 SMF70NPR	4	binary	Nominal permanent model-capacity rating in MSU/hour. When non-zero, this value is associated with the nominal permanent model capacity as identified in field SMF70MPC. When field SMF70CAI contains a value of 100, this value equals the value in field SMF70MPR.
212	D4 SMF70NTR	4	binary	Nominal temporary model-capacity rating in MSU/hour. When non-zero, this value is associated with the nominal temporary model capacity as identified in field SMF70MTC. When field SMF70CAI contains a value of 100, this value equals the value in field SMF70MTR.
216	D8 SMF70CAI	1	binary	Capacity-adjustment indication. When zero, the indication is not reported. When in the range from 1 to 99, some amount of reduction is indicated. When 100, the machine is operating at its normal capacity. Temporary capacity changes that affect machine performance (for example, CBU or OOCOD) are not included.
217	D9 SMF70CCR	1	binary	Capacity-change reason. Valid if SMF70CAI is non-zero. When 0, no capacity change took place. When 1, the capacity change is due to the setting of a manual control. When greater than 1, the capacity change is due to an internal machine condition or due to an external machine exception.
218	DA SMF70MCP	2	binary	Maximum CPU ID available for this IPL.
220	DC SMF70ICP	2	binary	Highest CPU ID installed at IPL time.
222	DE SMF70CCP	2	binary	Highest CPU ID currently installed. This number can increase upon dynamic CPU addition.
224	E0 SMF70CPA_actual	4	binary	Physical CPU adjustment factor based on Model Capacity Rating (will be used for converting processor time to service units). This value together with SMF70CPA_scaling_factor replaces SMF70CPA.
228	E4 SMF70CPA_scaling_factor	4	binary	Scaling factor for SMF70CPA_actual.
232	E8 SMF70MCF	4	binary	Multithreading maximum capacity numerator for general purpose processors. Divide this value by 1024 to get the multithreading maximum capacity factor for all general purpose processors that were configured ONLINE for the complete interval.
236	EC SMF70MCFS	4	binary	Multithreading maximum capacity numerator for zIIP. Divide this value by 1024 to get the multithreading maximum capacity factor for all zIIPs that were configured ONLINE for the complete interval. A zero value is reported if no zIIP is currently installed.

Offsets	Name	Length	Format	Description
240	F0 SMF70MCFI	4	binary	Multithreading maximum capacity numerator for zAAP. Divide this value by 1024 to get the multithreading maximum capacity factor for all zAAPs that were configured ONLINE for the complete interval. A zero value is reported if no zAAP is currently installed.
244	F4 SMF70CF	4	binary	Multithreading capacity numerator for general purpose processors. Divide this value by 1024 to get the multithreading capacity factor for all general purpose processors that were configured ONLINE for the complete interval.
248	F8 SMF70CFS	4	binary	Multithreading capacity numerator for zIIP. Divide this value by 1024 to get the multithreading capacity factor for all zIIPs that were configured ONLINE for the complete interval. A zero value is reported if no zIIP is currently installed.
252	FC SMF70CFI	4	binary	Multithreading capacity numerator for zAAP. Divide this value by 1024 to get the multithreading capacity factor for all zAAPs that were configured ONLINE for the complete interval. A zero value is reported if no zAAP is currently installed.
256	100 SMF70ATD	4	binary	Average Thread Density numerator for general purpose processors . Divide this value by 1024 to get the average number of active threads for all general purpose processors that were dispatched to physical hardware and configured ONLINE for the complete interval.
260	104 SMF70ATDS	4	binary	Average Thread Density numerator for zIIP. Divide this value by 1024 to get the average number of active threads for all zIIPs that were dispatched to physical hardware and configured ONLINE for the complete interval. A zero value is reported if no zIIP is currently installed.
264	108 SMF70ATDI	4	binary	Average Thread Density numerator for zAAP. Divide this value by 1024 to get the average number of active threads for all zAAPs that were dispatched to physical hardware and configured ONLINE for the complete interval. A zero value is reported if no zAAP is currently installed.
268	10C SMF70LACM	4	binary	Long-term average of CPU service (millions of service units) consumed by transactions classified with reporting attribute MOBILE. If an address space or enclave is part of a tenant resource group, it will not contribute to SMF70LACM.
272	110 SMF70LACA	4	binary	Long-term average of CPU service (millions of service units) consumed by transactions classified with reporting attribute CATEGORYA. If an address space or enclave is part of a tenant resource group, it will not contribute to SMF70LACA.
276	114 SMF70LACB	4	binary	Long-term average of CPU service (millions of service units) consumed by transactions classified with reporting attribute CATEGORYB. If an address space or enclave is part of a tenant resource group, it will not contribute to SMF70LACB.
280	118 SMF70ADJ	4	binary	Logical adjustment factor for CPU rate.
284	11C SMF70LACCR	4	binary	Long-term average of CPU service (millions of service units) consumed by DFSMS data set encryption. Valid only for IBM z14® and later CPCs.

Offsets	Name	Length	Format	Description
288 120	SMF70MaxPU	2	binary	When non-zero, this field indicates how many processor cores are physically available in this particular machine. When the value is 0, it is not defined for this model.
290 122	SMF70OS_PRTCT	1	binary	When non-zero, the OSPROTECT system parameter with a value other than SYSTEM is in effect. X'01' indicates OSPROTECT=1. For machines after IBM z14, may be 0 with OSPROTECT=1.
291 123	*	1	binary	Reserved.
292 124	SMF70MDL_CBP	16	EBCDIC	Reserved for future use.
308 134	SMF70MCR_CBP	4	binary	Reserved for future use.
312 138	SMF70NCR_CBP	4	binary	Reserved for future use.
316 13C	SMF70LAC_CBP	4	binary	Reserved for future use.
320 140	SMF70CPA_actual_CBP	4	binary	Reserved for future use.
324 144	SMF70_IPL_TIME	8	binary	IPL time of partition, in TOD format.
332 14C	SMF70_TRG_M_CNT	4	binary	Number of times sampling of tenant resource group memory consumption happened.
336 150	SMF70CRW	4	binary	Reserved.
340 154	SMF70CPC_TYPE	4	binary	CPC Type.
344 158	SMF70MDL_VAR	16	EBCDIC	CPC model identifier indicating the variable capacity of the CPC. The identifier is left-justified with trailing blanks, if necessary. This field is zero if not supported by the hardware.
360 168	SMF70MVCR	4	binary	Model variable capacity rating. When non-zero, this value is associated with the model capacity as identified in the SMF70MDL_VAR field.
364 16C	SMF70NVCR	4	binary	Nominal model variable capacity rating. When non-zero, this value is associated with the nominal model capacity as identified in the SMF70MDL_VAR field. When the SMF70CAI field contains a value of 100, this value equals the value in the SMF70MVCR field.
368 170	SMF70ZSU_on_zIIP	8	binary	Unweighted zIIP-eligible service units spent on zIIP for the entire system.
376 178	SMF70ZSU_on_CP	8	binary	Unweighted zIIP-eligible service units spent on CP for the entire system.
384 180	SMF70JSU_on_zIIP	8	binary	Unweighted zIIP-eligible Java service units spent on zIIP for the entire system.
392 188	SMF70JSU_on_CP	8	binary	Unweighted zIIP-eligible Java service units spent on CP for the entire system.
400 190	SMF70CPE_LO	2	binary	Low threshold value of OPT parameter CPENABLE.
402 192	SMF70CPE_HI	2	binary	High threshold value of OPT parameter CPENABLE.
404 194	*	4	*	Reserved.
408 198	SMF70MDL_REP	16	EBCDIC	CPC model identifier indicating the replacement capacity of the CPC, which is the total of permanent capacity, temporarily increased capacity, and temporarily available replacement capacity, without the corridor capacity. The identifier is left-justified with trailing blanks, if necessary. This field is zero if not supported by the hardware.

Offsets	Name	Length	Format	Description
424 1A8	SMF70MR CR	4	binary	Model replacement capacity rating. When non-zero, this value is associated with the model capacity as identified in the SMF70MDL_REP field.
428 1AC	SMF70NR CR	4	binary	Nominal model-replacement-capacity rating. When non-zero, this value is associated with the nominal model capacity as identified in the SMF70MDL_REP field. When the SMF70CAI field contains a value of 100, this value equals the value in the SMF70MR CR field.
432 1B0	SMF70_CPU Power	8	binary	Accumulated microwatts for all CPU resources allocated to the LPAR during the interval. Divide by SMF70_NumPowerSamples to retrieve the average power consumption of the interval.
440 1B8	SMF70_Storage Power	8	binary	Accumulated microwatts for all storage resources allocated to the LPAR during the interval. Divide by SMF70_NumPowerSamples to retrieve the average power consumption of the interval.
448 1C0	SMF70_IO Power	8	binary	Accumulated microwatts for all I/O resources allocated to the LPAR during the interval. Divide by SMF70_NumPowerSamples to retrieve the average power consumption of the interval.
456 1C8	SMF70_CPC Total Power	8	binary	Accumulated microwatts for all electrical and mechanical components in the CPC. Divide by SMF70_NumPowerSamples to retrieve the average power consumption of the interval.
464 1D0	SMF70_CPC Unass Res Power	8	binary	Accumulated microwatts for all types of resources in the standby or reserved state. Divide by SMF70_NumPowerSamples to retrieve the average power consumption of the interval.
472 1D8	SMF70_CPC Infra Power	8	binary	Accumulated microwatts for all subsystems in the CPC which do not provide CPU, storage, or I/O resources to logical partitions. These include service elements, cooling systems, power distribution, and network switches, among others. Divide by SMF70_NumPowerSamples to retrieve the average power consumption of the interval.
480 1E0	SMF70_NumPowerSamples	2	binary	Number of power samples for the interval.
482 1E2	*	6	*	Reserved.
488 1E8	SMF70_PowerPartitionName	8	EBCDIC	The name of the LPAR to which the LPAR-specific power fields apply.

Tenant Resource Group data section

This section contains general tenant resource group and tenant information as well as processor consumption measurements for the interval.

Offsets	Name	Length	Format	Description
0	0 SMF70_TRG_NAME	8	EBCDIC	Tenant resource group name.
8	8 SMF70_TRG_DESC	32	EBCDIC	Tenant resource group description.
40	28 SMF70_TRG_TNTID	8	EBCDIC	Tenant identifier.
48	30 SMF70_TRG_TNTNAME	32	EBCDIC	Tenant name.
80	50 SMF70_TRG_SBID	64	EBCDIC	Solution ID.
144	90 SMF70_TRG_SUCP	8	binary	Service units on CPs consumed by tenant resource group.
152	98 SMF70_TRG_SUIFA	8	binary	Service units on zAAPs consumed by tenant resource group.
160	A0 SMF70_TRG_SUSUP	8	binary	Service units on zIIPs consumed by tenant resource group.

Offsets	Name	Length	Format	Description
168	A8 SMF70_TRG_LAC	4	binary	Long-term average service on general purpose processors in millions of service units per hour consumed by tenant resource group.
172	AC SMF70_TRG_LAC_CBP	4	binary	Reserved for future use.
176	B0 SMF70_TRG_FLAGS	2	binary	Reserved for future use.
178	B2	2		Reserved.
180	B4 SMF70_TRG_MEM	8	binary	Memory consumption of tenant resource group in units of 4K frames.

Workload classification data section

There is one workload classification data section for each processor type. The section is described by triplet SMF70WCS and includes the workload classification samples array.

Offsets	Name	Length	Format	Description
0	0 SMF70WC_CPU_Type	1	binary	CPU type as defined by SMF70TYP.
1	1 *	3	*	Reserved.
4	4 SMF70WC_Array_Offset	4	binary	Offset to workload classification samples array (relative to the start of this data section).
8	8 SMF70WC_Element_Size	2	binary	Size of an element in the samples array.
10	A SMF70WC_Classes_Num	2	binary	Number of workload classes in the samples array (equal to its dimension).

Workload classification samples array

An array (with dimension equal to the number of workload classes) is included in each workload classification data section and is pointed to by SMF70WC_Array_Offset.

Offsets	Name	Length	Format	Description
0	0 SMF70WC_Samples	8	binary	Count of samples attributed to the indexed workload class.

Record Type 72 (X'48') – Workload Activity, Storage Data, and Serialization Delay

Tailored Fit Pricing updates SMF record type 72, subtype 3. In this section we hit just the sections updated for Tailored Fit Pricing. For complete information about SMF record type 72, see [z/OS MVS System Management Facilities \(SMF\)](#).

Subtype 3: Workload Activity — is written for each service class and active report class in the active service policy. A report class becomes active as soon as work has been assigned to that report class.

Workload Manager control section

Identifies the policy, workload and service/report class name and contains workload data.

Resource Group data section

Contains information about the resource group to which the service class or tenant report class belongs.

Macro to Symbolically Address Record Type 72: The SMF record mapping macro for all records produced by RMF is ERBSMFR. Its format is ERBSMFR (*n1,n2,...*) where *n1,n2, ...* are the SMF record types you want to map. Note that the parentheses are required only when two or more record types are specified. The mapping macro resides in SYS1.MACLIB.

For information on using RMF, see *z/OS Resource Measurement Facility User's Guide*. For more information on Monitor I, II, and III, see *z/OS RMF Report Analysis*. For information on the MVS workload manager, see *z/OS MVS Planning: Workload Management* and *z/OS MVS Programming: Workload Management Services*.

Subtype 3 – Workload activity

Workload manager control section

Offsets	Name	Length	Format	Description
0	0 R723MSCF	1	binary	<p>Service/Report class flags.</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Indicator for a report class</p> <p>1 Workload activity data not available</p> <p>2 Policy data not available</p> <p>3 Execution velocity includes I/O delays</p> <p>4 Indicator for CPU protection</p> <p>5 Indicator for storage protection</p> <p>6 Indicator for dynamic alias tuning</p> <p>7 Indicator for I/O priority group HIGH</p>
1	1 R723MFLG	1	binary	<p>Flags.</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Indicator for zAAP crossover.</p> <p>1 Indicator for zAAP honor priority.</p> <p>2 Indicator for zIIP honor priority.</p> <p>3 Failure returned by HISMT service. Multithreading maximum capacity numerator values are invalid.</p> <p>4 Indicator that service class is not eligible for honor priority processing. When on, specialty engine eligible work in this service class will not be offloaded to CPs for help processing.</p> <p>5 Indicator for a tenant report class</p> <p>6 Service class and tenant report class periods that are associated with a resource group and have assigned a discretionary goal are excluded from workload management.</p> <p>7 WLM batch initiator management is AI-infused.</p>
2	2 R723MFL2	1	binary	<p>Additional flags:</p> <p>Bit</p> <p>Meaning when set</p> <p>0 At least one zAAP is online in the sysplex (or local system when not joined to a sysplex).</p> <p>1–7 Reserved.</p>

Offsets	Name	Length	Format	Description
3	3 *	1	*	Reserved.
4	4 R723MNSP	8	EBCDIC	Policy name.
12	C R723MDSP	32	EBCDIC	Policy description.
44	2C R723MTPA	8	binary	Local time/date of policy activation (STCK format).
52	34 R723MCPU	4	binary	CPU service coefficient * 10,000.
56	38 R723MIOC	4	binary	I/O service coefficient. Always zero.
60	3C R723MMSO	4	binary	Storage service coefficient. Always zero.
64	40 R723MSRB	4	binary	SRB service coefficient * 10,000
68	44 R723MTVL	4	binary	WLM sample interval (in milliseconds).
72	48 R723MTV#	4	binary	Number of times when WLM sampling code ran.
76	4C R723MOPT	2	EBCDIC	Suffix of the IEAOPTxx parmlib member.
78	4E	2		Reserved.
80	50 R723MWNM	8	EBCDIC	Workload name.
88	58 R723MWDE	32	EBCDIC	Workload description.
120	78 R723MCNM	8	EBCDIC	Service/Report class name.
128	80 R723MCDE	32	EBCDIC	Service/Report class description.
160	A0 R723MCPG	2	binary	Number of periods belonging to this service or report class.
162	A2 R723MSUB	1	binary	Number of entries in the work/resource manager state section belonging to a subsystem.
163	A3 *	3		Reserved.
166	A6 R723MERF	6	EBCDIC	Enqueue residency CPU service factor.
172	AC R723MADJ	4	binary	Adjustment factor for CPU rate.
176	B0 R723MIDN	8	EBCDIC	Service definition name.
184	B8 R723MIDD	32	EBCDIC	Service definition description.
216	D8 R723MTDI	8	binary	Local time/date the service definition was installed (STCK format).
224	E0 R723MIDU	8	EBCDIC	Userid that installed the service definition.
232	E8 R723CLSC	8	EBCDIC	Service class that last contributed to this report class. Blank if this is a service class.
240	F0 R723NFFI	4	binary	Normalization factor for zAAP. Multiply zAAP service times or service units with this value and divide by 256 to calculate the CP equivalent value.
244	F4 R723NFFS	4	binary	Normalization factor for zIIP. Multiply zIIP service units with this value and divide by 256 to calculate the CP equivalent value.
248	F8 R723NADJ	4	binary	Nominal adjustment factor for CPU rate.
252	FC R723CECA	4	binary	CEC adjustment factor.
256	100 R723MCF	4	binary	Multithreading maximum capacity numerator for general purpose processors. Divide this value by 1024 to get the MT maximum capacity factor for all general purpose processors that were configured ONLINE for the complete interval.
260	104 R723MCFs	4	binary	Multithreading maximum capacity numerator for zIIP. Divide this value by 1024 to get the multithreading maximum capacity factor for all zIIPs that were configured ONLINE for the complete interval. A zero value is reported if no zIIP is currently installed.

Offsets	Name	Length	Format	Description
264 108	R723MCFI	4	binary	Multithreading maximum capacity numerator for zAAP. Divide this value by 1024 to get the multithreading maximum capacity factor for all zAAPs that were configured ONLINE for the complete interval. A zero value is reported if no zAAP is currently installed.
268 10C	R723CPA_actual	4	binary	Physical CPU adjustment factor based on Model Capacity Rating.
272 110	R723CPA_scaling_factor	4	binary	Scaling factor for R723CPA_actual.
276 114	R723CPA_actual_zCBP	4	binary	Reserved for future use.

Resource Group Data Section

Offsets	Name	Length	Format	Description
0	0 R723GGNM	8	EBCDIC	Resource group name.
8	8 R723GGDE	32	EBCDIC	Resource group description.
40	28 R723GGLT	1	binary	<p>Resource group flags.</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Maximum capacity was specified</p> <p>1 Minimum capacity was specified</p> <p>2 Specification of R723GGMN and R723GGMX is in percentage of the LPAR share rather than in service units. In addition, the scope of the resource group is system-wide rather than sysplex-wide.</p> <p>3 Specification of R723GGMN and R723GGMX is in percentage of a single processor capacity rather than in service units. In addition, the scope of the resource group is system-wide rather than sysplex-wide.</p> <p>4 Memory limit was specified.</p> <p>5 Specification of R723GGMN and R723GGMX is in MSU/h rather than in service units.</p> <p>6 Specialty processor consumption is included into the WLM capping algorithms, i.e. R723GGMN and R723GGMX limit the combined general purpose and specialty processor consumption.</p> <p>7 Reserved.</p>
41	29 R723GGTF	1	binary	<p>Tenant Resource Group Flags.</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Indicator for a tenant resource group</p> <p>1-7 Reserved</p>
42	2A	2		Reserved.
44	2C R723GGMN	4	binary	If bit 1 of R723GGLT is ON, minimum capacity of the resource group. If bit 2, bit 3, and bit 5 of R723GGLT are OFF, this value is in unweighted CPU service units per second. In addition, the scope of the resource group is sysplex-wide. If bit 2, bit 3, or bit 5 of R723GGLT is ON, see the description of R723GGLT.

Offsets	Name	Length	Format	Description
48	30 R723GGMX	4	binary	If bit 0 of R723GGTL is ON, maximum capacity of the resource group. If bit 2, bit 3, and bit 5 of R723GGTL are OFF, this value is in unweighted CPU service units per second. In addition, the scope of the resource group is sysplex-wide. If bit 2, bit 3, or bit 5 of R723GGTL is ON, see the description of R723GGTL.
52	34 R723GGML	4	binary	If bit 4 of R723GGTL is ON, memory limit (in GB) of the resource group. The scope of the resource group is system-wide.
56	38 R723GGTI	8	EBCDIC	Tenant identifier. Only valid if bit 0 of R723GGTF is ON.
64	40 R723GGTN	32	EBCDIC	Tenant name. Only valid if bit 0 of R723GGTF is ON.
96	60 R723GGKY	64	EBCDIC	Solution ID. Only valid if bit 0 of R723GGTF is ON.

Record Type 79 (X'4F') – RMF Monitor II activity

Tailored Fit Pricing updates SMF record type 79, subtypes 1, 2, and 5. In this section we hit just the sections updated for Tailored Fit Pricing. For complete information about SMF record type 70, see [z/OS MVS System Management Facilities \(SMF\)](#).

Reference information:

- For information on using RMF, see *z/OS Resource Measurement Facility User's Guide*.
- For information on Monitor I and II, see *z/OS RMF Report Analysis*.
- For more information on performance groups, see *z/OS MVS Initialization and Tuning Guide*.

Record type 79 is written during a Monitor II background session when feedback is requested as SMF records. It is written at each measurement interval and when the session is terminated. It contains a section that is identical for all Monitor II reports and a subtype section that is unique for each report. The subtypes are:

Subtype 1

Contains information that describes address space state data (and address space state data by job name) for each address space identifier included.

Subtype 2

Contains information that describes address space resource data (and address space resource data by job name) activity. The length depends on the number of devices.

Subtype 5

Contains information that describes address space SRM data (and address space SRM data by job name).

Macro to symbolically address record type 79

The SMF record mapping macro for all records produced by RMF is ERBSMFR. Its format is ERBSMFR (*n1,n2,...*) where *n1,n2, ...* are the SMF record types you want to map. Note that the parentheses are required only when two or more record types are specified. The mapping macro resides in SYS1.MACLIB.

Subtype 1 – Address Space State Data

ASD and ASDJ data section

Offsets	Name	Length	Format	Description
0	0 R791ASID	2	binary	Address space identifier.
2	2 R791JBN	8	EBCDIC	Name of job.
10	A R791DMN	2	binary	Reserved.

Offsets	Name	Length	Format	Description
12	C R791NPG	2	binary	Reserved.
14	E R791PGP	2	binary	Reserved.
16	10 R791TTOD	4	binary	Real time into transaction (milliseconds).
20	14 R791CL	2	EBCDIC	<p>Current location. (Set to IN when all other indicators are off.)</p> <p>Contents</p> <p>Meaning</p> <p>DL Out queue/delayed</p> <p>IN In storage</p> <p>LO Logically swapped out</p> <p>NS Non-swappable</p> <p>PR Privileged</p> <p>OT Swapped out and ready</p> <p>WL Wait queue/long wait</p> <p>WM Wait queue/MSO</p> <p>WO Wait queue/reasons other than WM, WL, or WT</p> <p>WT Wait queue/terminal wait</p> <p>>> Transitioning out</p> <p><< Transitioning in.</p>
22	16 R791TAS	2	binary	<p>Type of user</p> <p>Contents</p> <p>Meaning</p> <p>0 Batch</p> <p>1 Started task</p> <p>2 Mount task</p> <p>3 TSO/E</p> <p>4 ASCH</p> <p>5 OMVS address space.</p>

Offsets	Name	Length	Format	Description
24	18 R791SRC	2	EBCDIC	Reason for last swap-out Contents Meaning TO Terminal output TI Terminal input LW Long wait XS Auxiliary storage shortage RS Central storage shortage DW Detected wait MP Memory Pool shortage NQ CAP enqueue EX CAP exchange US CAP uni-swap TS Transition swap IC Improve central storage usage IP Improve system paging rate MR Make room for a user who has been swapped out too long AW APPC WAIT (swapped out, because waiting for APPC services) IW OMVS input wait OW OMVS output wait SR In-real swap 00 Unknown.
26	1A R791DP	2	binary	Dispatcher priority.
28	1C	6		Reserved.
34	22 R791SWC	2	binary	Transaction swap count.
36	24 R791SWMR	2	binary	SRM work load recommendation value.
38	26	4		Reserved.
42	2A R791WMS	4	binary	SRM service for the current transaction since the last swap-in.
46	2E R791TCPU	4	binary	CPU time (TCB + SRB) for current job step, in milliseconds.
50	32	4		Reserved.
54	36 R791ESCT	4	binary	Number of pages on expanded storage frames.
58	3A	2		Reserved.
60	3C R791PIN	4	binary	Page-in count.

Offsets	Name	Length	Format	Description
64	40 R791TRTM	4	binary	Transaction residency time, in milliseconds.
68	44 R791FLG	1	binary	Bit Meaning when set 0 Cross memory address space 1 Data in R791CTAR is valid 2 Data in R791VAL is valid 3 Reserved. 4 If ON: this address space is a server address space If OFF: goal specified for this address space is being honored by WLM 5 Address space has been quiesced by a RESET command 6 Address space matched a classification rule in the active policy which prevents managing the region based on the response time goals of its served transactions 7 Server has temporal affinity to clients.
69	45 R791FLG2	1	binary	Additional bits. Bit Meaning when set 0 Service class assigned by classification, or RESET SRVCLASS was designated CPU-critical in the active policy. 1 Address space matched a classification rule in the active policy which was designated storage-critical. 2 Address space is serving transactions which belong to a service class that was designated storage-critical in the active policy's classification rules, or is running in SYSTEM/SYSSTC. 3 CPU protection was assigned either to the address space or to transaction service classes being served by the space, or implicitly. 4 Storage protection was assigned either to the address space or to transaction service classes being served by the space. 5 The dispatching priority of the address space is currently promoted due to a chronic resource contention. 6 Address space is a CICS TOR that matched a classification rule in the active policy which allows managing the region based on the region goals but also ensures that completed transactions are reported and used for management of the CICS AORs. 7 Honor priority ineligibility was assigned either to the address space or to transaction service classes being served by the space.
70	46 R791FMCT	4	binary	Number of central storage frames.
74	4A R791WSS	4	binary	Working set at last swap in.
78	4E R791TWSS	4	binary	RSM target working set size.

Offsets	Name	Length	Format	Description
82	52 R791ESHP	4	binary	Number of hiperspace expanded storage pages used by job.
86	56 R791ESVI	4	binary	Number of VIO expanded storage pages used by job.
90	5A R791HIN	4	binary	Number of ESO hiperspace page-ins by block.
94	5E R791HRMS	4	binary	Number of ESO hiperspace read misses by job (a read miss is an attempt to read a frame that is not in expanded storage).
98	62 R791BPIN	4	binary	Number of blocked pages brought in from DASD.
102	66 R791PINE	4	binary	Number of pages brought in from expanded storage.
106	6A R791BPNE	4	binary	Number of blocked pages brought in from expanded storage.
110	6E R791CTAR	4	binary	Central storage target number of frames.
114	72 R791VAL	4	binary	Recommendation value for working-set-managed address spaces.
118	76 R791SCL	8	EBCDIC	Service class name.
126	7E R791SCP	2	binary	Service class period.
128	80 R791WKLD	8	EBCDIC	Workload name.
136	88 R791RGRP	8	EBCDIC	Resource group name.
144	90 R791SPI	4	binary	Number of page-ins from auxiliary storage for shared page groups.
148	94 R791CMNI	4	binary	Number of common pages for current transaction.
152	98 R791PNV	4	binary	Number of non-VIO pages for current transaction.
156	9C R791PVIO	4	binary	Number of VIO pages for current transaction.
160	A0 R791EXCT	4	binary	EXCP count for this step.
164	A4 R791TCPC	4	binary	Total CPU time consumed in this address space, in milliseconds.
168	A8 R791ASST	4	binary	CPU time consumed by preemptible-class SRBs running on behalf of this address space, in milliseconds.
172	AC R791PHTM	4	binary	CPU time consumed by preemptible-class SRBs running in this address space, in milliseconds.
176	B0 R791RCL	8	EBCDIC	Report class name.
184	B8 R791MLIM	8	binary	Address space memory limit, in megabytes.
192	C0 R791TIFA	4	binary	CPU time in milliseconds consumed on zAAPs.
196	C4 R791TCP	4	binary	CPU time in milliseconds consumed on standard CPs. Only valid if zAAPs or zIIPs are in the configuration.
200	C8 R791TIFC	4	binary	CPU time in milliseconds consumed on standard CPs by work that was eligible for zAAP.
204	CC R791NFFI	4	binary	Normalization factor for zAAP time. Used to convert between real zAAP times and "normalized" zAAP times, that is, the equivalent time on a standard CP. Multiply R791TIFA by this value and divide by 256 to calculate the normalized zAAP time.
208	D0 R791TSUP	4	binary	CPU time in milliseconds consumed on zIIPs
212	D4 R791TSUC	4	binary	CPU time in milliseconds consumed on standard CPs by work that was eligible for zIIP.
216	D8 R791NFFS	4	binary	Normalization factor for zIIP time. Used to convert between real zIIP times and "normalized" zIIP times, that is, the equivalent time on a standard CP. Multiply R791TSUP by this value and divide by 256 to calculate the normalized zIIP time.
220	DC R791EXCW	8	binary	EXCP count (double word).
228	E4 R791PHTA	4	binary	zAAP-only equivalent of R791PHTM. This is normalized time.
232	E8 R791PHTI	4	binary	zIIP-only equivalent of R791PHTM. This is normalized time.

Offsets	Name	Length	Format	Description
236	EC R791FLG3	1	binary	<p>Additional flags.</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Service class was assigned by classification, or RESET SRVCLASS belongs to I/O priority group HIGH in the active policy.</p> <p>1 I/O priority group HIGH was assigned either to the address space or to transaction service classes served by the address space.</p> <p>2 R791RGRP is the name of a tenant resource group and R791RCL is the name of a tenant report class.</p> <p>3 General purpose and specialty processor consumption is considered by WLM capping algorithms for this address space.</p> <p>4 Reserved.</p> <p>5 If On: Boost attribute can be defined on WLM classification rules for started tasks.</p> <p>6 Address space matched a classification rule in the active policy which enables for recovery process boost. Only valid if bit 5 of R791FLG3 is set.</p> <p>7 CPU protection was implicitly assigned.</p>
237	ED	3		Reserved.

Subtype 2 – Address Space Resource Data

ARD and ARDJ data section

Offsets	Name	Length	Format	Description
0	0 R792ASID	2	binary	Address space identifier.
2	2 R792JBN	8	EBCDIC	Name of job.
10	A R792DMN	2	binary	Reserved.
12	C R792NPG	2	binary	Reserved.

Offsets	Name	Length	Format	Description
14	E R792CL	2	EBCDIC	Current location Contents Meaning DL Out queue/delayed IN In storage LO Logically swapped out NS Non-swappable PR Privileged OT Swapped out and ready WL Wait queue/long wait WM Wait queue/MSO WO Wait queue/reasons other than WM, WL, or WT WT Wait queue/terminal wait >> Transitioning out << Transitioning in.
16	10 R792TAS	2	binary	Type of user Contents Meaning 0 Batch 1 Started task 2 Mount task 3 TSO/E 4 ASCH 5 OMVS address space.
18	12 R792TRC	2	binary	Transaction count.
20	14 R792TTOD	4	binary	Transaction elapsed time, in milliseconds.
24	18 R792PRFX	4	binary	Number of private fixed frames.
28	1C	2		Reserved.
30	1E R792SVAR	4	binary	SRM service absorption rate for step.
34	22 R792TCPU	4	binary	Total TCB time for step, in milliseconds.
38	26 R792PSS1	4	binary	High order word - CPU page seconds, in milliseconds. One page in storage for one second is one page second.
42	2A R792PSS2	4	binary	Low order word - step product of frame, in milliseconds. One page in storage for one second is one page second.
46	2E R792EJST	4	binary	Total processor time (TCB+SRB), in milliseconds.

Offsets	Name	Length	Format	Description																		
50	32 R792TSRM	4	binary	Total SRM service for job or session.																		
54	36 R792RTM	4	binary	Resident time for step, in milliseconds.																		
58	3A R792EXCP	2	binary	EXCP count for this step.																		
60	3C R792CMNI	4	binary	Number of common pages for current transaction.																		
64	40 R792PNV	4	binary	Number of non-VIO pages for current transaction.																		
68	44 R792PVIO	4	binary	Number of VIO pages for current transaction.																		
72	48 R792FXBL	4	binary	Number of fixed frames below 16 megabytes.																		
76	4C R792PSWP	4	binary	Number of pages swapped in and out for current transaction.																		
80	50 R792LPAI	4	binary	Number of link pack area (LPA) pages paged in for current transaction.																		
84	54 R792CSAI	4	binary	Number of CSA pages paged in for current transaction.																		
88	58 R792LSQA	4	binary	Number of fixed local system queue area (LSQA) fixed frames.																		
92	5C R792NLQF	4	binary	Number of non-local system queue area (LSQA) fixed frames.																		
96	5E R792TDEV	4	binary	Total device connect time in milliseconds.																		
100	64	2		Reserved.																		
102	66 R792PIN	4	binary	Page-in count.																		
106	6A R792TRTM	4	binary	Transaction residency time.																		
110	6E R792FLG	1	binary	<div>Flags</div> <table><thead><tr><th>Bit</th><th>Meaning when set</th></tr></thead><tbody><tr><td>0</td><td>Cross-memory address space</td></tr><tr><td>1</td><td>Incorrect RSM data obtained for address space</td></tr><tr><td>2</td><td>Reserved</td></tr><tr><td>3</td><td>Reserved</td></tr><tr><td>4</td><td>If ON: this address space is a server address space If OFF: goal specified for this address space is being honored by WLM</td></tr><tr><td>5</td><td>Address space has been quiesced by a RESET command</td></tr><tr><td>6</td><td>Address space matched a classification rule in the active policy which prevents managing the region based on the response time goals of its served transactions</td></tr><tr><td>7</td><td>Server has temporal affinity to clients.</td></tr></tbody></table>	Bit	Meaning when set	0	Cross-memory address space	1	Incorrect RSM data obtained for address space	2	Reserved	3	Reserved	4	If ON: this address space is a server address space If OFF: goal specified for this address space is being honored by WLM	5	Address space has been quiesced by a RESET command	6	Address space matched a classification rule in the active policy which prevents managing the region based on the response time goals of its served transactions	7	Server has temporal affinity to clients.
Bit	Meaning when set																					
0	Cross-memory address space																					
1	Incorrect RSM data obtained for address space																					
2	Reserved																					
3	Reserved																					
4	If ON: this address space is a server address space If OFF: goal specified for this address space is being honored by WLM																					
5	Address space has been quiesced by a RESET command																					
6	Address space matched a classification rule in the active policy which prevents managing the region based on the response time goals of its served transactions																					
7	Server has temporal affinity to clients.																					

Offsets	Name	Length	Format	Description
111	6F R792FLG2	1	binary	<p>Additional bits.</p> <p>Bit Meaning when set</p> <p>0 Service class assigned by classification, or RESET SRVCLASS was designated CPU-critical in the active policy.</p> <p>1 Address space matched a classification rule in the active policy which was designated storage-critical.</p> <p>2 Address space is serving transactions which belong to a service class that was designated storage-critical in the active policy's classification rules, or is running in SYSTEM/SYSSTC.</p> <p>3 CPU protection was assigned either to the address space or to transaction service classes being served by the space, or implicitly.</p> <p>4 Storage protection was assigned either to the address space or to transaction service classes being served by the space.</p> <p>5 The dispatching priority of the address space is currently promoted due to a chronic resource contention.</p> <p>6 Address space is a CICS TOR that matched a classification rule in the active policy which allows managing the region based on the region goals but also ensures that completed transactions are reported and used for management of the CICS AORs.</p> <p>7 Honor priority ineligibility was assigned either to the address space or to transaction service classes being served by the space.</p>
112	70 R792LSQR	4	binary	Local system queue area (LSQA) pages in central storage.
116	74 R792LSQE	4	binary	Local system queue area (LSQA) pages in expanded storage.
120	78 R792ARS	4	binary	Average number of real frames for step.
124	7C R792TWSS	4	binary	SRM target working set size for this job.
128	80 R792PHSP	4	binary	Number of hiperspace pages for the current transaction.
132	84 R792EXCT	4	binary	EXCP count for this step.
136	88 R792SCL	8	EBCDIC	Service class name.
144	90 R792SCP	2	binary	Service class period.
146	92 R792WKLD	8	EBCDIC	Workload name.
154	9A R792RGRP	8	EBCDIC	Resource group name.
162	A2	2		Reserved.
164	A4 R792TCPC	4	binary	Total CPU time consumed in this address space, in milliseconds.
168	A8 R792ASST	4	binary	CPU time consumed by preemptible-class SRBs running on behalf of this address space, in milliseconds.
172	AC R792PHTM	4	binary	CPU time consumed by preemptible-class SRBs running in this address space, in milliseconds.
176	B0 R792FXAB	4	binary	Number of fixed frames between 16M and 2G (z/Architecture® mode).

Offsets	Name	Length	Format	Description																		
180	B4 R792TIFA	4	binary	CPU time in milliseconds consumed on zAAPs.																		
184	B8 R792TCP	4	binary	CPU time in milliseconds consumed on standard CPs. Only valid if zAAPs or zIIPs are in the configuration.																		
188	BC R792TIFC	4	binary	CPU time in milliseconds consumed on standard CPs by work that was eligible for zAAP.																		
192	C0 R792NFFI	4	binary	Normalization factor for zAAP time. Used to convert between real zAAP times and "normalized" zAAP times, that is, the equivalent time on a standard CP. Multiply R792TIFA by this value and divide by 256 to calculate the normalized zAAP time.																		
196	C4 R792TSUP	4	binary	CPU time in milliseconds consumed on zIIPs.																		
200	C8 R792TSUC	4	binary	CPU time in milliseconds consumed on standard CPs by work that was eligible for zIIP.																		
204	CC R792NFFS	4	binary	Normalization factor for zIIP time. Used to convert between real zIIP times and "normalized" zIIP times, that is, the equivalent time on a standard CP. Multiply R792TSUP by this value and divide by 256 to calculate the normalized zIIP time.																		
208	D0 R792EXCW	8	binary	EXCP count (double word).																		
216	D8 R792PHTA	4	binary	zAAP-only equivalent of R792PHTM. This is normalized time.																		
220	DC R792PHTI	4	binary	zIIP-only equivalent of R792PHTM. This is normalized time.																		
224	E0 R792FLG3	1	binary	<div>Additional flags.</div> <table><tr><th>Bit</th><th>Meaning when set</th></tr><tr><td>0</td><td>Service class was assigned by classification, or RESET SRVCLASS belongs to I/O priority group HIGH in the active policy.</td></tr><tr><td>1</td><td>I/O priority group HIGH was assigned either to the address space or to transaction service classes served by the address space.</td></tr><tr><td>2</td><td>R792RGRP is the name of a tenant resource group.</td></tr><tr><td>3</td><td>General purpose and specialty processor consumption is considered by WLM capping algorithms for this address space.</td></tr><tr><td>4</td><td>Reserved.</td></tr><tr><td>5</td><td>If On: Boost attribute can be defined on WLM classification rules for started tasks.</td></tr><tr><td>6</td><td>Address space matched a classification rule in the active policy which enables for recovery process boost. Only valid if bit 5 of R792FLG3 is set.</td></tr><tr><td>7</td><td>CPU protection was implicitly assigned.</td></tr></table>	Bit	Meaning when set	0	Service class was assigned by classification, or RESET SRVCLASS belongs to I/O priority group HIGH in the active policy.	1	I/O priority group HIGH was assigned either to the address space or to transaction service classes served by the address space.	2	R792RGRP is the name of a tenant resource group.	3	General purpose and specialty processor consumption is considered by WLM capping algorithms for this address space.	4	Reserved.	5	If On: Boost attribute can be defined on WLM classification rules for started tasks.	6	Address space matched a classification rule in the active policy which enables for recovery process boost. Only valid if bit 5 of R792FLG3 is set.	7	CPU protection was implicitly assigned.
Bit	Meaning when set																					
0	Service class was assigned by classification, or RESET SRVCLASS belongs to I/O priority group HIGH in the active policy.																					
1	I/O priority group HIGH was assigned either to the address space or to transaction service classes served by the address space.																					
2	R792RGRP is the name of a tenant resource group.																					
3	General purpose and specialty processor consumption is considered by WLM capping algorithms for this address space.																					
4	Reserved.																					
5	If On: Boost attribute can be defined on WLM classification rules for started tasks.																					
6	Address space matched a classification rule in the active policy which enables for recovery process boost. Only valid if bit 5 of R792FLG3 is set.																					
7	CPU protection was implicitly assigned.																					
225	E1	3		Reserved.																		

Subtype 5 – Address Space SRM Data

ASRM and ASRMJ Data Section

Offsets	Name	Length	Format	Description
0	0 R795ASID	2	binary	Address space identifier.
2	2 R795JBN	8	EBCDIC	Name of job.
10	A R795DMN	2	binary	Reserved.
12	C R795NPG	2	binary	Reserved.
14	E R795PGP	2	binary	Reserved.
16	10 R795TTOD	4	binary	Real time into transaction.
20	14 R795CL	2	EBCDIC	<p>Current location (set to IN when all other indicators are off)</p> <p>Contents</p> <p>Meaning</p> <p>DL Out queue/delayed</p> <p>IN In storage</p> <p>LO Logically swapped out</p> <p>NS Non-swappable</p> <p>PR Privileged</p> <p>OT Swapped out and ready</p> <p>WL Wait queue/long wait</p> <p>WM Wait queue/MSO</p> <p>WO Wait queue/reasons other than WM, WL, or WT</p> <p>WT Wait queue/terminal wait</p> <p>>> Transitioning out</p> <p><< Transitioning in.</p>
22	16 R795TAS	2	binary	<p>Type of user</p> <p>Contents</p> <p>Meaning</p> <p>0 Batch</p> <p>1 Started task</p> <p>2 Mount task</p> <p>3 TSO/E</p> <p>4 ASCH</p> <p>5 OMVS address space.</p>
24	18 R795TROD	4	binary	Transaction resident time.

Offsets	Name	Length	Format	Description
28	1C R795TCNT	2	binary	Transaction count.
30	1E R795SWC	2	binary	Transaction swap count.
32	20 R795CPUS	4	binary	Total processor service units for transaction (zeros when ASID is out of storage).
36	24 R795MSOS	4	binary	Total main storage origin (MSO) service units for transaction (zeros when ASID is out of storage).
40	28 R795IOCS	4	binary	Total IOC service units for transaction (zeros when ASID is out of storage).
44	2C R795WMS	4	binary	Total service units for transaction (zeros when ASID is out of storage).
48	30 R795TOTL	4	binary	Total service units for job or TSO/E session (zeros when ASID is out of storage).
52	34 R795TOT	4	binary	Total service units for transaction since last swap-in.
56	38 R795SRBS	4	binary	Total SRB service units for transaction (zeros when ASID is out of storage).
60	3C R795FLG	1	binary	Flags. Bit Meaning when set 0-2 Reserved 3 Reserved 4 If ON: this address space is a server address space If OFF: goal specified for this address space is being honored by WLM 5 Address space has been quiesced by a RESET command 6 R795RGRP is the name of a tenant resource group. 7 Reserved.
61	3D R795SCL	8	EBCDIC	Service class name.
69	45 R795SCP	2	binary	Service class period.
71	47 R795WKLD	8	EBCDIC	Workload name.
79	4F R795RGRP	8	EBCDIC	Resource group name.

Record type 89 (X'59') – Usage Data

The type 89 record provides information about product usage on a particular MVS system. The usage reporting program analyzes the data collected in the type 89 record. For more information see [z/OS MVS Product Management](#). The record is generated on a scheduled interval (1 hour maximum).

Record type 89 has two subtypes:

- Subtype 1 – Usage data

Contains, for the scheduled interval, summary usage data for all products across the system that have registered to request usage recording. These products must issue the IFAUSAGE macro to specify:

- Registration information.
- Level and scope of data collection (task or address space level).
- Start and end of collection period.

- Subtype 2 — State Data

Contains, for the scheduled interval, summary state data for all products across the system that have registered to indicate that they are running. These products issue the MVS Register service to indicate that they are running. MVS uses information a product supplies to determine if the product is enabled and to maintain a list of active products.

The installation controls the scheduling of the type 89 record by checking the INTERVAL value specified for the SMF address space. Because SMF is a started task, this is the INTERVAL value for SUBSYS=STC in the SMFPRMxx member. If the INTERVAL value is less than or equal to one hour, then that value is used as the reporting interval for type 89 records. If that value is greater than one hour, or if no INTERVAL value is specified, then one hour is used as the reporting interval for type 89 recording.

There are two sets of interval START and STOP times in the record:

- Usage data interval START/STOP.
- Reporting interval START/STOP.

The usage data interval represents the hourly buckets that the usage reporting program records product usage in. This interval is synchronized to the top of the hour.

The reporting interval represents the increment when the type 89 records are generated and is also synchronized to the top of the hour. For example, if you specified an interval value of 30 minutes, type 89 records would be generated at 9:00, 9:30, 10:00... If you are collecting usage data at the task level, you may want to synchronize interval processing to the top of the hour in your SMFPRMxx member because task level data collection is scheduled by interval processing.

SMF type 89 records are generated on the interval as requested; if no products are registered, then a type 89 record is generated with a product count of 0.

Product intersection time data sections are generated when a product registers at the ADDRSP level, and then invokes a product that registers at the TASK level. When no product intersections occur, the product intersection count in SMF89CNN is zero.

SMFPRMxx parameters are described in [*z/OS MVS Initialization and Tuning Reference*](#). The usage reporting program is described in [*z/OS MVS Product Management*](#). The Register service is described in *MVS Programming: Registration Services*.

Record environment

The following conditions exist for the generation of each of the subtypes of this record:

Macro

SMFEWTM, BRANCH=YES (record exit: IEFU84)

Mode

Subtype	Mode
---------	------

1	SRB
---	-----

2	SRB
---	-----

Storage Residency

31-bit

SUBSYS

STC

Record mapping

Header/self-defining section

This section contains the common SMF record header fields and, if applicable, the triplet fields (offset/length/number) that locate the other sections on the record.

Offsets	Name	Length	Format	Description
0	0 SMFOLEN	2	binary	Record length. This field and the next field (total of four bytes) form the RDW (record descriptor word).
2	2 SMFOSEG	2	binary	Segment descriptor (see record length field).
4	4 SMFOFLG	1	binary	System indicator: Bit Meaning when set 0-2 Reserved 3-6 Version indicators* 7 Reserved.
5	5 SMFORTH	1	binary	Record type 0 (X'00')
6	6 SMFOTME	4	binary	Time since midnight, in hundredths of a second, that the record was moved into the SMF buffer.
10	A SMFODTE	4	packed	Date when the record was moved into the SMF buffer, in the form 0cyydddF.
14	E SMFOSID	4	EBCDIC	System identification (from the SID parameter).
18	12 SMFOJWT	4	binary	Limit, in minutes, of continuous wait for the job (taken from JWT parameter). Continuous wait time is defined as time spent waiting while the application program is in control. For example, for data sets allocated dynamically (while the application program is running, for example) either or both of the following count toward a job's continuous wait time: <ul style="list-style-type: none">• The time required to recall a data set from HSM Migration Levels 1 or 2• The time required to mount a tape If a data set was allocated statically (for a DD statement, for example) these activities will not be counted towards the job's continuous wait time.
22	16 SMFOBUF	4	binary	This field contains meaningless information.
26	1A SMFOVST	4	binary	Number of 1K bytes in virtual storage.

Offsets	Name	Length	Format	Description
30	1E SMFOOPT	1	binary	<p>SMF options:</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Reserved</p> <p>1 Reserved</p> <p>2 Reserved.</p> <p>3 Data set accounting. Record types selected. This bit is on when one of the following record types is selected: 14, 15, 17, 18, 62, 63, 64, 67 or 68.</p> <p>4 Volume accounting. Record types 10 or 69 selected.</p> <p>5 Reserved.</p> <p>6 Type 17 records will be written for temporary data sets (REC(ALL)).</p> <p>7 Reserved.</p>
31	1F SMFORST	4	binary	<p>Number of 1K bytes in central storage.</p> <p>This field is only accurate below 4 TB. For larger systems, this field contains X'FFFFFFFF'. The SMFORS4K field also represents the system size and is not limited to 4 TB. See SMFORS4K and SMFORS4K_Not_Valid for additional instructions to determine if SMFORS4K contains valid information.</p>
35	23 SMFORSV	1		Reserved.
36	24 SMFOOSL	8	EBCDIC	MVS product name.
44	2C SMF0SYN	8	EBCDIC	System name (from the SYSNAME parameter in the IEASYSxx parmlib member).
52	34 SMF0SYP	8	EBCDIC	Sysplex name (from the SYSPLEX parameter in the COUPLExx parmlib member).
60	3C SMF0TZ	4	binary	<p>Difference in time between local time and Greenwich mean time in binary units of 1.048576 seconds. The value of SMF0TZ is copied from the CVTTZ field. For more information about the CVTTZ field, see the CVT mapping macro in <i>z/OS MVS Data Areas</i> in the <i>z/OS Internet library</i> (www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary).</p>
64	40 SMF0MSWT	4	binary	Started task wait time limit (SMFPRMxx SWT(<i>hhmm</i>) value) converted to minutes.
68	44 SMF0MTWT	4	binary	TSO wait time limit (SMFPRMxx TWT(<i>hhmm</i>) value) converted to minutes.
72	48 SMF0TBUF	2	binary	The number of megabytes specified by the SMFTBUFF IEASYSxx parmlib option. If the SMFTBUFF parameter value is incorrectly specified or SMFTBUFF is not specified, this value will be zero, and the default buffer size of 5 megabytes was used during IPL SMF initialization processing.
74	4A SMFO_End_V1	0	n/a	End of version 1.

Offsets	Name	Length	Format	Description
74	4A SMFORS4K	8	binary	<p>Number of 4K frame units in central storage.</p> <p>If the record length in SMFOLEN is SMFO_Len_V2, this field always contains valid information. To determine whether this field contains valid information when the record length in SMFOLEN is SMFO_Len_V3 or larger:</p> <ul style="list-style-type: none"> • If this field contains a non-zero value, it contains valid information. Also, the SMFORS4K_Not_Valid bit will be off. • If this field contains a zero value, it does not contain valid information. The SMFORS4K_Not_Valid bit will be on. <p>It is sufficient to check either the SMFORS4K field for a non-zero value or the SMFORS4K_Not_Valid bit to determine whether the field contains valid information. This field always contains valid information as of z/OS 2.5.</p>
82	52 SMFO_End_V2	0	n/a	End of version 2.
82	52 SMFO_Num_CrypCtrs	2	binary	Number of crypto counters supported by this system.
84	54 SMFO_Num_NNPICtrs	2	binary	Number of NNPI counters supported by this system.
86	56 SMFO_Flags	1	binary	<p>Flags</p> <p>Bit</p> <p>Meaning when set</p> <p>0 (SMFORS4K_Not_Valid) Indicates that SMFORS4K does not contain a valid value. When this bit is on, SMFORS4K will contain zeros. See the description of SMFORS4K for more information.</p> <p>1 - 7 Reserved.</p>
87	57 SMFO_End_V3	0	n/a	End of version 3.

Version information

When processing a type 0 record, the value in SMFOLEN can be used to determine the version of the record. Equates are provided for the lengths of each version:

Version 1

This version contains the fields through SMF0TBUF. SMFOLEN contains the value SMFO_Len_V1 (X'4A').

Version 2

This version contains the fields through SMFORS4K. SMFOLEN contains the value SMFO_Len_V2 (X'52').

Version 3

This version contains the fields through SMFO_Flags. SMFOLEN contains the value SMFO_Len_V3 (X'57').

Self-defining section

Offsets	Name	Length	Format	Description
0	SMF992COF	4	binary	Offset to class data sections from beginning of the record.
4	4 SMF992CLN	2	binary	Length of class data sections.
6	6 SMF992CON	2	binary	Number of class data sections. There is one per service class.
8	8 SMF992CPOF	4	binary	Offset to period data section from beginning of record (including RDW).
12	C SMF992CPLN	2	binary	Length of period data section.
14	E SMF992CPON	2	binary	Number of period data sections.

Offsets	Name	Length	Format	Description
16 10	SMF992ECOF	4	binary	Offset to EWLM class section from beginning of record (including RDW).
20 14	SMF992ECLN	2	binary	Length of EWLM class section.
22 16	SMF992ECON	2	binary	Number of EWLM class section.

Record product section

This section provides information about the type 89 record, the system, and the recording interval.

Triplet information

This section is located in the record using the following triplet fields, which are located in the "self-defining" section:

Offset

SMF89PRO

Length

SMF89PRL

Number

SMF89PRN - This field is always "1" because each type 89 record that is generated has one record product section.

Offsets	Name	Length	Format	Description
0 0	SMF89PNM	8	EBCDIC	Record product name - "SMF".
8 8	SMF89RVN	4	binary	Record version number - "1".
12 0C	SMF89OSL	8	EBCDIC	MVS system level (For example, SP4.3.0).
20 14	SMF89IST	4	binary	Reporting interval START Time (local, hundredths of a second from midnight). This field and SMF89IET define the recording interval. This is different from the usage data interval that is used to collect data into hourly buckets.
24 18	SMF89ISD	4	packed	Reporting interval START Date in the form 0cyydddF.
28 1C	SMF89IET	4	binary	Reporting interval END Time (local, hundredths of a second from midnight). This field and SMF89IST define the recording interval. This is different from the usage data interval that is used to collect data into hourly buckets.
32 20	SMF89IED	4	packed	Reporting interval END Date in the form 0cyydddF.
36 24	SMF89PFL	1	binary	Bit Meaning when set 0 Reserved. 1 Indicates that LICENSE=zNALC was specified in IEASYSxx. (SMF89ZNA) 2 - 7 Reserved.
37 25		3	binary	Reserved
40 28	SMF89HOF	8	binary	Hypervisor date/time offset in STCK format. When present, this field contains the sysplex timer offset value.
48 30	SMF89DTO	8	binary	Local data/time offset, copied from CVTLDTO.
56 38	SMF89_CoreMode_CP	2	binary	The number of CPUs that are active on a CP core.
58 3A	SMF89_CoreMode_zAAP	2	binary	The number of CPUs that are active on a zAAP core.

Offsets	Name	Length	Format	Description
60 3C	SMF89_CoreMode_zIIP	2	binary	The number of CPUs that are active on a zIIP core.

System ID section

This section provides information about the system (both hardware and software) at the time the usage data was collected.

Triplet information

This section is located on the record using the following triplet fields, which are located in the "self-defining" section:

Offset

SMF89SIO

Length

SMF89SIL

Number

SMF89SIN - This field is always 1 because each type 89 record that is generated has one system ID section.

Offsets	Name	Length	Format	Description
0 0	SMF89SYN	8	EBCDIC	MVS system name (SYSNAME from IEASYSxx).
8 8	SMF89UST	4	binary	Usage data interval START time (local, hundredths of a second from midnight). This is usually an hour value (such as 01:00:00.00) except in the case of the first record during an IPL (which reports the "IPL" time). This is different from the recording interval that is used to report on the generation of the usage records. This field and SMF89UET define the hour "bucket" that the usage data reflects. This field is only filled in for the SMF89 subtype 1 records.
12 0C	SMF89USD	4	packed	Usage data interval START Date in the form 0cyydddF. This field is only filled in for the SMF89 subtype 1 records.
16 10	SMF89UET	4	binary	Usage data interval END time (local, hundredths of a second from midnight). This is usually an hour value (such as 01:00:00.00). This field and SMF89UST define the hour "bucket" that the usage data reflects. This field is only filled in for the SMF89 subtype 1 records.
20 14	SMF89UED	4	packed	Usage data interval END date in the form 0cyydddF. This field is only filled in for the SMF89 subtype 1 records.
24 18 *		4	binary	Reserved.
28 1C *		4	binary	Reserved.
32 20	SMF89CMN	2	packed	CPU model number.
34 22	SMF89CVN	1	binary	CPU version number.

Offsets	Name	Length	Format	Description
35 23	SMF89LPI	1	binary	<p>LPAR indicators:</p> <p>Bit</p> <p>Meaning when set</p> <p>0</p> <p>The one digit LPAR ID contained in SMF89LP2 (bit 4) is valid (SMF89LPV).</p> <p>1</p> <p>The two digit LPAR ID contained in field SMF89LP3 is valid (SMF89LPM).</p> <p>2–3</p> <p>Reserved.</p> <p>4–7</p> <p>The one digit LPAR ID (X'0-F') (SMF89LP2).</p> <p>Notes:</p> <ol style="list-style-type: none"> 1. For a one digit LPAR ID (X'0-F'), both SMF89LPV (bit 0) and SMF89LPM (bit 1) is on and both SMF89LP2 (bits 4–7) and field SMF89LP3 contain the LPAR ID. 2. For the two digit LPAR ID (greater than X'F'), SMF89LPM (bit 1) is on, and SMF89LP3 contains the LPAR ID.
36 24	SMF89SER	3	packed	CPU serial number.
39 27	SMF89LP3	1	binary	LPAR ID.
40 28	SMF89RPP	4	binary	CPU relative processing power indicator.
44 2C	SMF89SPN	8	EBCDIC	Sysplex name (from the SYSPLEX parameter in the COUPLExx parmlib member).
52 34	SMF89CPT	6	EBCDIC	CPC type number (blanks if data is not available).
58 3A	SMF89CPM	3	EBCDIC	CPC model number (blanks if data is not available).
61 3D	SMF89CPS	12	EBCDIC	CPC sequence number (blanks if data is not available).
73 49	SMF89SIF	1	binary	<p>Bit</p> <p>Meaning when set</p> <p>0</p> <p>Field SMF89LPN is valid. (SMF89LNV)</p> <p>1</p> <p>This is the last record for this usage interval. (SMF89LCR)</p> <p>2–7</p> <p>Reserved.</p>
74 4A *		2	binary	Reserved.
76 4C	SMF89MNF	16	EBCDIC	V1-CPC manufacturer.
92 5C	SMF89TID	4	EBCDIC	V1-CPC type.
96 60	SMF89MDL	16	EBCDIC	V1-CPC model.
112 70	SMF89SQC	16	EBCDIC	V1-CPC sequence code.
128 80	SMF89POM	4	EBCDIC	V1-CPC plant of manufacturer.
132 84	SMF89CPC	4	binary	CPU capability.
136 88	SMF89CCC	2	binary	Configured CPU count.
138 8A	SMF89SCC	2	binary	Standby CPU count.
140 8C	SMF89MAF	30	binary	Array of multiprocessing CPU capability adjustment factors. This array contains information for only the first 15 general processors. Obtain additional processor information from RMF records, or issuing the STSI instruction.

Offsets	Name	Length	Format	Description
170	AA SMF89LPN	8	EBCDIC	LPAR name when SYSIB 2.2.2 is valid, when returned by the STSI instruction (such as when running under z/VM). Bit SMF89LPV is on when the field is valid. Avoid looking at this field unless SMF89LNV is on.
178	B2 SMF89_Capacity_Change_Cnt	2	binary	The number of processor capacity changes that occurred since the previous interval or event interval. This number is greater than 1 when the number of processor capacity changes exceeded the number specified in the MAXEVENTINTRECS parmlib option.
180	B4 SMF89_RCTPCPUA_Actual	4	binary	Physical CPU adjustment factor (this is the adjustment factor for converting CPU time to equivalent service in basic-mode with all processors online). Based on model capacity rating.
184	B8 SMF89_RCTPCPUA_Nominal	4	binary	Physical CPU adjustment factor (this is the adjustment factor for converting CPU time to equivalent service in basic-mode with all processors online). Based on nominal model capacity rating.
188	BC SMF89_RCTPCPUA_scaling_factor	4	binary	Scaling factor for SMF89_RCTPCPUA_Actual and SMF89_RCTPCPUA_Nominal.
192	C0 SMF89_Capacity_Adjustment_Ind	1	binary	When: 0 The indication is not reported. 1–99 Some amount of reduction is indicated. 100 The machine is operating in normal capacity. The Primary CPU and all secondary-type CPU are similarly affected.
193	C1 SMF89_Capacity_Change_Rsn	1	binary	Indicates the reason that is associated with the present value contained in SMF89_Capacity_Adjustment_Ind. The bit values of this field correspond to those described in RMCTZ_Capacity_Adjustment_Indication of the IRARMCTZ mapping macro. (See z/OS MVS Data Areas.)
194	C2 SMF89_Capacity_Flags	1	binary	Processor capacity flags. Bit Meaning when set 0 (SMF89_Event_Driven_Interval_Rec) When on, indicates that the current record was generated as a result of an event, rather than as a result of a standard interval expiration based on time. 1 (SMF89_Capacity_Data_err) When on, indicates that an error occurred while collecting the processor capacity data, therefore the following fields are unreliable: SMF89_RCTPCPUA_Actual SMF89_RCTPCPUA_Nominal SMF89_RCTPCPUA_scaling_factor SMF89_Capacity_Adjustment_Ind SMF89_Capacity_Change_Rsn 2 (SMF89_PCD_Rsvd_Exists) When on, indicates records generated on systems running z/OS V1R7 through z/OS V1R9. When off, indicates records generated on systems running z/OS V1R10 and later.
195	C3 *	1	binary	Reserved.
196	C4 SMF89ZNF	4	binary	zAAP normalization factor for zAAP service time.
200	C8 SMF89SNF	4	binary	zIIP Normalization factor for zIIP service time.

Offsets	Name	Length	Format	Description
204	CC SMF89SEQ	2	binary	Record sequence number when multiple records are written for the same interval.
206	CE SMF89SolutionID	64	EBCDIC	The Tailored Fit Pricing solution ID from the SOLUT system parameter; otherwise, binary zeros if the SOLUT parameter was not specified.
270	10E SMF89SI_RPB_Duration	8	binary	Total duration for the life of the IPL of recovery process boosts, in STCK format. Updated when RP boost is extended or ends. Valid when SMF89SIL is at least SMF89_Len_V3.

Subtype 1 – Usage data section

This section contains the product information (specified on the IFAUSAGE REGISTER request) and the usage data that has been collected for the interval specified by the start and end times (SMF89UST and SMF89UET) for that product.

There is one usage data section for each unique product identification (specified by owner, name, version, qualifier) that is actively registered for any part of that specified interval. The data reported is accumulated for ALL address spaces that had any interaction with the product.

Triplet information

This section is located on the record using the following triplet fields, which are located in the "self-defining" section:

Offset

SMF89UDO

Length

SMF89UDL

Number

SMF89UDN

Offsets	Name	Length	Format	Description
0	0 SMF89UPO	16	EBCDIC	Product owner or vendor name (specified on the PRODOWNER option of the IFAUSAGE macro).
16	10 SMF89UPN	16	EBCDIC	Product name (specified on the PRODNAME option of the IFAUSAGE macro).
32	20 SMF89UPV	8	EBCDIC	Product version (specified on the PRODVERS option of the IFAUSAGE macro).
40	28 SMF89UPQ	8	EBCDIC	Product qualifier (specified on the PRODQUAL option of the IFAUSAGE macro).
48	30 SMF89UPI	8	EBCDIC	Product ID number (specified on the PRODID option of the IFAUSAGE macro).
56	38 SMF89UCT	8	long floating point	Product TCB time (in hundredths of a second).
64	40 SMF89USR	8	long floating point	Product SRB time (in hundredths of a second).

Offsets	Name	Length	Format	Description
72 48	SMF89UFG	1	binary	<p>Usage entry flags</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Unauthorized register</p> <p>1 Ineligible for measured usage</p> <p>2 Unauthorized register with SAF-authorized UNAUTHSERV=LEVEL1 requested</p> <p>3 - 7 Unused</p>
73 49	SMF89UF2	1	binary	<p>Bit</p> <p>Meaning when set</p> <p>0 This product has product intersection time sections (SMF89HCS).</p> <p>1 All ADDRSP registrations of this product are the first registration in the respective address space (SMF89AFS).</p> <p>2 An ADDRSP registration of this product registered or deregister while active TASK level registrations were present in one or more address spaces (SMF89TSH).</p> <p>3 All ADDRSP level registrations were at a service level that support product intersection time (SMF89PLV).</p> <p>4 This intersection has Tenant Resource Group Sections. Check SMF89TCO for the offset to the first TRG Intersection section. Then check each of those in the section for an exact match of the product information to find the corresponding section (SMF89HTR)</p> <p>5-7 Reserved.</p>
74 4A	SMF89_BoostInfo	1	binary	<p>Boost information</p> <p>Bit</p> <p>Meaning when set</p> <p>0 zIIP boost was active at some point within the interval.</p> <p>1 Speed boost was active at some point within the interval.</p> <p>5-7 Boost class: 001: IPL 010: Shutdown 011: Recovery process</p> <p>Note: The boost class value is valid only when one or more boosts is active; that is, a boost active bit is also on.</p>

Offsets	Name	Length	Format	Description
75	4B SMF89URT	1	binary	<p>Data format of value in SMF89URD (specified on the FORMAT option of the IFAUSAGE macro FUNCTIONDATA request).</p> <p>Bit</p> <p>Meaning when set</p> <p>0 No data specified</p> <p>1 CPU time, in long floating point (in hundredths of a second)</p> <p>2 Binary (64-bit)</p> <p>3 Long floating point</p> <p>4 - 7 Reserved</p>
76	4C SMF89URD	8	various	Product specific resource data (specified by the data option of the IFAUSAGE macro FUNCTIONDATA request). SMF89URT identifies the format of the data in this field.
84	54 SMF89UZT	8	long floating point	Product offload engine time (hundredth of a second).
92	5C SMF89CountAsTrad	4	binary	Count of active address spaces in traditional (non-TRG) sub-capacity workload environment.
96	60 SMF89CountAsTrg	4	binary	Count of active address spaces in TRG workload environment.

Subtype 1 – Product intersection data section

This section contains information about intersections that occur between products registered with the IFAUSAGE service. Intersections are generated when a product registered at the ADDRSP level invokes a program that registers at the TASK level for a task in the current address space. The ADDRSP scope product is known as the containing product and the TASK scope product is known as the intersecting product.

There is one product intersection data section for each intersection detected. The data reported is accumulated for ALL address spaces that had any intersection between two products.

Triplet information

This section is located on the record using the following triplet fields, which are located in the "self-defining" section:

Offset

SMF89CNO

Length

SMF89CNL

Number

SMF89UDN

Offsets	Name	Length	Format	Description
0	0 SMF89CPO	16	EBCDIC	Product owner or vendor name (obtained PRODOWNER option of IFAUSAGE macro) of the containing product.
16	10 SMF89CPN	16	EBCDIC	Product Name (obtained PRODNAME option of IFAUSAGE macro) of the containing product.
32	20 SMF89CPV	8	EBCDIC	Product Version (obtained PRODVERS option of IFAUSAGE macro) of the containing product.
40	28 SMF89CPQ	8	EBCDIC	Product Qualifier (obtained PRODQUAL option of IFAUSAGE macro) of the containing product.

Offsets	Name	Length	Format	Description
48	30 SMF89CPI	8	EBCDIC	Product ID (obtained PRODID option of IFAUSAGE macro) of the containing product.
56	38 SMF89IPO	16	EBCDIC	Product Owner or Vendor Name (obtained PRODOWNER option of IFAUSAGE macro) of the intersecting product.
72	48 SMF89IPN	16	EBCDIC	Product Name (obtained PRODNAME option of IFAUSAGE macro) of the intersecting product.
88	58 SMF89IPV	8	EBCDIC	Product Version (obtained PRODVERS option of IFAUSAGE macro) of the intersecting product.
96	60 SMF89IPQ	8	EBCDIC	Product Qualifier (obtained PRODQUAL option of IFAUSAGE macro) of the intersecting product.
104	68 SMF89IPI	8	EBCDIC	Product ID (obtained PRODID option of IFAUSAGE macro) of the intersecting product.
112	70 SMF89CFG	1	binary	Bit Meaning when set 0 The container product was registered unauthorized (SMF89CUC). 1 The intersecting product was registered unauthorized (SMF89CUP). 2 Some time for the intersection was a result of a SCOPE(FUNCTION) registered product (SMF89CFC). 3 Some time for the intersection was a result of a SCOPE(ALL) registered product (SMF89CTC). 4 Intersection time might be complete for this product. Note: Not all products use (SMF89CGO). 5 This intersection has Tenant Resource Group Sections. Check SMF89TCO for the offset to the first TRG Intersection section. Then check each of those in the section for an exact match of the product information to find the corresponding section (SMF89CHTR). 6-7 Reserved.
113	71	7		Reserved
120	78 SMF89CCT	8	long floating point	Product Intersect TCB Time (in hundredths of a second)
128	80 SMF89CZT	8	long floating point	Product Intersect Offload Engine Time (in hundredths of a second)

Subtype 1 – Tenant resource group section

This section contains the product information (specified on the IFAUSAGE REGISTER request), tenant resource group name, and the usage data that has been collected for the interval specified by the start and end times (SMF89UST and SMF89UET) for that product.

There is one Tenant resource group Data Section for each unique product identification (specified by owner, name, version, qualifier) that is actively registered for any part of that specified interval. The data reported is accumulated for ALL address spaces that had any interaction with the product.

Triplet information

This section is located on the record using the following triplet fields, which are located in the "self-defining" section:

Offset

SMF89TRO

Length

SMF89TRL

Number

SMF89TRN

Offsets	Name	Length	Format	Description
0	0 SMF89TPO	16	EBCDIC	Product owner or vendor name (specified on the PRODOWNER option of the IFAUSAGE macro).
16	10 SMF89TPN	16	EBCDIC	Product name (specified on the PRODNAME option of the IFAUSAGE macro).
32	20 SMF89TPV	8	EBCDIC	Product version (specified on the PRODVERS option of the IFAUSAGE macro).
40	28 SMF89TPQ	8	EBCDIC	Product qualifier (specified on the PRODQUAL option of the IFAUSAGE macro).
48	30 SMF89TPI	8	EBCDIC	Product ID number (specified on the PRODID option of the IFAUSAGE macro).
56	38 SMF89TRG_Name	8	EBCDIC	Tenant resource group.
64	40 SMF89TCT	8	long floating point	Tenant resource group TCB time (hundredths of a second - floating point)
72	48 SMF89TSR	8	long floating point	TRG SRB Time (in hundredths of a second - floating point)
80	50 SMF89TZT	8	long floating point	TRG offload engine time (hundredth of a second).
88	58 SMF89TRGData	8	binary	Product-specific resource data (specified by the DATA option of the FUNCTIONDATA request of the IFAUSAGE macro).
96	60 SMF89TRGDataType	1	binary	<p>Data format of value in SMF89TRGData (specified by the FORMAT option of the FUNCTIONDATA request of the IFAUSAGE macro).</p> <p>Value</p> <p>Meaning when set</p> <p>0 No data specified</p> <p>1 CPU time, in long floating point (in hundredths of a second)</p> <p>2 Binary (64-bit)</p> <p>3 Long floating point</p> <p>4 - 7 Reserved</p>

Subtype 1 – Intersection data for tenant resource groups

This section contains information about intersections that occur between products registered with the IFAUSAGE service while running in a tenant resource group. Intersections are generated when a product registered at the ADDRSP level invokes a program that registers at the TASK level for a task in the current address space. The ADDRSP scope product is known as the containing product and the TASK scope product is known as the intersecting product.

There is one product intersection data section for each intersection detected. The data reported is accumulated for ALL address spaces that had any intersection between two products.

Triplet information

This section is located on the record using the following triplet fields, which are located in the "self-defining" section:

Offset

SMF89TCO

Length

SMF89TCL

Number

SMF89TCN

Offsets	Name	Length	Format	Description
0	0 SMF89TCPO	16	EBCDIC	Product owner or vendor name (specified on the PRODOWNER option of the IFAUSAGE macro).
16	10 SMF89TCPN	16	EBCDIC	Product name (specified on the PRODNAME option of the IFAUSAGE macro).
32	20 SMF89TCPV	8	EBCDIC	Product version (specified on the PRODVERS option of the IFAUSAGE macro).
40	28 SMF89TCPQ	8	EBCDIC	Product qualifier (specified on the PRODQUAL option of the IFAUSAGE macro).
48	30 SMF89TCPI	8	EBCDIC	Product ID number (specified on the PRODID option of the IFAUSAGE macro).
56	38 SMF89TIPO	16	EBCDIC	Intersecting Product Owner or Vendor Name (obtained PRODOWNER option of IFAUSAGE macro).
72	48 SMF89TIPN	16	EBCDIC	Intersecting Product Name (obtained PRODNAME option of IFAUSAGE macro).
88	58 SMF89TIPV	8	EBCDIC	Intersecting Product Version (obtained PRODVERS option of IFAUSAGE macro).
96	60 SMF89TIPQ	8	EBCDIC	Intersecting Product Qualifier (obtained PRODQUAL option of IFAUSAGE macro).
104	68 SMF89TIPI	8	EBCDIC	Intersecting Product ID (obtained PRODID option of IFAUSAGE macro).
112	70 SMF89T_TRG_Name	8	EBCDIC	Tenant resource group.
120	78 SMF89TCFG	1	binary	Usage Entry Flags Bit Meaning when set 0 UNAUTHORIZED REGISTER Requested on container product (SMF89TCUC). 1 UNAUTHORIZED REGISTER Requested on intersecting product (SMF89TCUP).
121	79	7		Reserved.
128	80 SMF89TCCT	8	long floating point	Product Intersect TCB time (hundredths of a second - floating point).
136	88 SMF89TCZT	8	long floating point	Product Intersection Offload Engine Time (hundredths of a second - floating point).

Subtype 2 – State tenant resource group data section

This section contains the product information (specified on the MVS register service or in the IFAPRDxx parmlib member), the tenant resource group name, and instance data that has been collected for the interval at the time when the record was collected for that product.

There is one state data section for each unique product identification registered (specified by owner, name, feature, version, release, and modification level) for any part of the interval.

Triplet information

This section is located on the record using the following triplet fields, which are located in the "self-defining" section:

Offset

SMF89TRO

Length

SMF89TRL

Number

SMF89TRN

Offsets	Name	Length	Format	Description
0	0 SMF89T2TRGProdOwner	16	EBCDIC	Product Owner or Vendor Name (from prodowner parameter of IFAEDREG or OWNER option of PRODUCT statement of IFAPRDxx)
16	10 SMF89T2TRGProdName	16	EBCDIC	Product name (from prodname parameter of IFAEDREG or NAME option of PRODUCT statement of IFAPRDxx)
32	20 SMF89T2TRGFeatureName	16	EBCDIC	Feature name (from featurename parameter of IFAEDREG or FEATURENAME option of PRODUCT statement of IFAPRDxx)
48	30 SMF89T2TRGProdVers	2	EBCDIC	Product version (from prodvers parameter of IFAEDREG or VERSION option of PRODUCT statement of IFAPRDxx)
50	32 SMF89T2TRGProdRel	2	EBCDIC	Product release (from prodrel parameter of IFAEDREG or RELEASE option of PRODUCT statement of IFAPRDxx)
52	34 SMF89T2TRGProdMod	2	EBCDIC	Product modification level (from prodmod parameter of IFAEDREG or MOD option of PRODUCT statement of IFAPRDxx)
54	36 SMF89T2TRGProdID	8	EBCDIC	Product ID (from prodowner parameter of IFAEDREG or OWNER option of PRODUCT statement of IFAPRDxx)
62	3E	2		Reserved
64	40 SMF89T2TrgName	8	EBCDIC	Tenant Resource Group Name
72	48 SMF89T2TrgNumInstances	4	binary	Current number of registration of this product currently in this TRG

Record Type 99 (X'63') – System Resource Manager Decisions

Tailored Fit Pricing updates SMF record type 99. In this section we hit just the sections updated for Tailored Fit Pricing. For complete information about SMF record type 99, see [z/OS MVS System Management Facilities \(SMF\)](#).

For information about how to use type 99, see [z/OS MVS Programming: Workload Management Services](#).

Record Type 99 (63) – System Resource Manager Decisions

This record type is written by the SRM component. The records contain:

- Performance data for each service class period
- Trace codes representing the SRM actions
- The data which SRM used to decide which actions to take
- The controls SRM is using to manage work.

Tailored Fit Pricing updates the following subtypes:

Subtype 1

Contains system level data, the trace of SRM actions, and data about resource groups. The SRM actions are recorded in trace codes. All trace codes are described in *z/OS MVS Programming: Workload Management Services*. A subtype 1 record is written every policy interval.

Subtype 2

Contains data for service classes. A subtype 2 record is written every policy interval for each service class if any period in the service class had recent activity.

Subtype 1 - Resource Group Entry Section

Offsets	Name	Length	Format	Description
0	SMF99_RGNAME	8	EBCDIC	Resource group name
8	8 SMF99_MIN_SR	4	binary	Minimum service rate for the resource group in unweighted CPU service units per second. When there is no minimum defined, this field is 0.
12	C SMF99_MAX_SR	4	binary	Maximum service rate for the resource group in unweighted CPU service units per second. When there is no maximum defined, this field is X'7FFFFFFF'.
16	10 SMF99_ACT_SR	4	binary	Service rate received in the last policy adjustment interval on the local system in unweighted CPU service units per second.
20	14 SMF99_SPAS	4	binary	Service per non-capped slice in unweighted CPU service units per second.
24	18 SMF99_SLICES	2	binary	The number of cap slices in which work in this resource group was capped.
26	1A SMF99_RHELPCNT0	2	binary	A count of the systems that can help special system address spaces (work at importance 0). The count can include any systems in the sysplex running in goal mode other than the local system.
28	1C SMF99_RHELPCNT1	2	binary	A count of the systems that can help work at importance 1. The count can include any systems in the sysplex running in goal mode other than the local system.
30	1E SMF99_RHELPCNT2	2	binary	A count of the systems that can help work at importance 2. The count can include any systems in the sysplex running in goal mode other than the local system.
32	20 SMF99_RHELPCNT3	2	binary	A count of the systems that can help work at importance 3. The count can include any systems in the sysplex running in goal mode other than the local system.
34	22 SMF99_RHELPCNT4	2	binary	A count of the systems that can help work at importance 4. The count can include any systems in the sysplex running in goal mode other than the local system.
36	24 SMF99_RHELPCNT5	2	binary	A count of the systems that can help work at importance 5. The count can include any systems in the sysplex running in goal mode other than the local system.
38	26 SMF99_RHELPCNT6	2	binary	A count of the systems that can help discretionary work (work at importance 6). The count can include any systems in the sysplex running in goal mode other than the local system.

Offsets	Name	Length	Format	Description
40	28 SMF99_LHELP_FLGS	1	binary	<p>Flag indicating whether the local system can help work at each importance level. 1 indicates it can help, 0 indicates it cannot help.</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Reserved.</p> <p>1 Local system can help work at importance 0.</p> <p>2 Local system can help work at importance 1.</p> <p>3 Local system can help work at importance 2.</p> <p>4 Local system can help work at importance 3.</p> <p>5 Local system can help work at importance 4.</p> <p>6 Local system can help work at importance 5.</p> <p>7 Local system can help work at importance 6.</p>
41	29 SMF99_RG_FLAGS	1	binary	<p>Resource group flags</p> <p>Bit</p> <p>Meaning when set</p> <p>0 Indicates that the resource group is dynamic</p> <p>1 Indicates that the resource group capacity is specified in percentage of the total LPAR capacity</p> <p>2 Indicates that the resource group capacity is specified in percentage of a single processor capacity</p> <p>3 Indicates that the resource group capacity is specified in MSU/h</p> <p>4 Indicates that the resource group capacity is specified in raw service units</p> <p>5 Indicates that specialty processor consumption is included in the group consumption</p> <p>6 Indicates that the resource group is a tenant resource group</p> <p>7 Reserved.</p>
42	2A *	2	EBCDIC	Reserved.
44	2C SMF_RG_PERC_MIN	4	binary	Percentage min value, if min/max are specified in percentages, or MSU min value if min/max are specified in MSU.
48	30 SMF_RG_PERC_MAX	4	binary	Percentage max value, if min/max are specified in percentages, or MSU max value if min/max are specified in MSU.
56	38 SMF99_RG_MEM_LIMIT	4	binary	Maximum memory limit, in GB.
105	69 *	3	EBCDIC	Reserved.
108	6C SMF99_RG_LACS	4	binary	Tenant resource group long-term average of CPU service in MSUs per hour. Only valid if Bit 6 of SMF99_RG_FLAGS is ON

Offsets	Name	Length	Format	Description
112	70 SMF99_RG_SUsIFA	8	binary	Tenant resource group's aggregated IFA service units. Only valid if Bit 6 of SMF99_RG_FLAGS is ON
120	78 SMF99_RG_SUsSUP	8	binary	Tenant resource group's aggregated SUP service units. Only valid if Bit 6 of SMF99_RG_FLAGS is ON
128	80 *	4	binary	Internal use by IBM
132	84 SMF99_RG_MEMSMPCNT	4	binary	Number of times storage frames were sampled in SMF99_RG_FRAMECNT
136	88 SMF99_RG_FRAMECNT	8	binary	Tenant resource group's aggregated amount of storage frames (4K + 1M + 2G, normalized to 4K). Only valid if tenant resource group.

Subtype 2 - Address space expanded storage access policy section

Offsets	Name	Length	Format	Description
0	SMF99_AS_ESP_ANAM	8	EBCDIC	Address space name.
8	8 SMF99_AS_ESP_AP	1	binary	Expanded storage access policy for demand pages. Value Meaning 1 protected 2 least recently used (LRU) 3 space available
9	9 SMF99_AS_ESP_VP	1	binary	Expanded storage access policy for VIO pages. Value Meaning 1 protected 2 least recently used (LRU) 3 space available
10	A SMF99_AS_ESP_HP	1	binary	Expanded storage access policy for hiperspace pages. Value Meaning 1 protected 2 least recently used (LRU) 3 space available
11	B SMF99_AS_ESP_ASID	2	binary	Address space ID.

Offsets	Name	Length	Format	Description
13	D SMF99_AS_ESP_FLAGS	1	binary	<p>Flags.</p> <p>Bit</p> <p>Meaning when set</p> <p>0</p> <p>Storage is protected at this instant.</p> <p>1</p> <p>Storage protection assigned to space by classification rule.</p> <p>2</p> <p>Address space is currently managed to region's goal rather than transaction server's goal.</p> <p>3</p> <p>Address space is non swappable.</p> <p>4</p> <p>Address space is currently managed to both region's and transaction server's goal.</p> <p>5</p> <p>When on, specialty engine work in this address space is ineligible for "Honor Priority Processing", i.e., it will not be offloaded to CPs for help processing.</p> <p>6</p> <p>Address space matched a classification rule in the active policy which enables for recovery process boost.</p> <p>7</p> <p>Reserved.</p>
14	E *	2	binary	Reserved.
16	10 SMF99_AS_ESP_CS_FMCT	4	binary	Number of central storage frames the address spaces owns.
20	14 SMF99_AS_ESP_ES_FMCT	4	binary	Number of expanded storage frames the address spaces own.
24	18 SMF99_AS_ESP_PPS_TAR	4	binary	Address space protective process storage target. See subtype 5 for other targets. This is the only target non-monitor address spaces can have.
28	1C SMF99_AS_FULL_PREEMPTION	6	EBCDIC	Full Preemption Counts.
28	1C SMF99_AS_CPSRP_SAMP	2	binary	One sample per IRACPSRP invocation.
30	1E SMP99_AS_CPSRP_CUR_FP_SAMP	2	binary	Amount of IRACPSRP samples running with full preemption.
32	20 SMP99_AS_CPSRP_PREV_FP_SAMP	2	binary	Previous value of FULL_PRE1.
34	22 SMF99_AS_HealthInd	1	binary	Health indicator.
35	23 *	1	binary	Reserved to align to WORD boundary.
36	24 SMF99_AS_TOTAL_SERVICE	4	binary	Total service units for the address space - OUCBWMS.
40	28 SMF99_AS_CPU_SERVICE	4	binary	Total CPU service units for the address space - OUCBCPU.
44	2C SMF99_AS_SRB_SERVICE	4	binary	Total SRB service units for the address space - OUCBSRB.
48	30 SMF99_AS_MSO_SERVICE	4	binary	Total MSO service units for the address space - OUCBMSO.
52	34 SMF99_AS_TRN_SERVICE	4	binary	Accumulated transaction service for the address space - OUCBTRS.
56	38 SMF99_AS_IO_SERVICE	4	binary	Total IO service units for the address space - OUCBIOC.

Offsets	Name	Length	Format	Description
60	3C SMF99_AS_DISP_COUNT	2	binary	Dispatchable count: the number of times that this address space has been found in subroutine CPUTLCK to be dispatchable yet no CPU time has accumulated for it - OUXBDSCN.
62	3E *	2	binary	Reserved to align to DWORD boundary.
64	40 SMF99_AS_IFA_SERVICE	8	binary	Total IFA service units for the address space – Oucbx_Time_On_Pro(pro_ifa) descaled.
72	48 SMF99_AS_IFACP_SERVICE	8	binary	Total IFA service units spent on CP for the address space - Oucbx_Time_Pro_On_CP(pro_ifa) descaled.
80	50 SMF99_AS_SUP_SERVICE	8	binary	Total SUP service units for the address space – Oucbx_Time_On_Pro(pro_sup) descaled.
88	58 SMF99_AS_SUPCP_SERVICE	8	binary	Total SUP service units spent on CP for the address space - Oucbx_Time_Pro_On_CP(pro_sup) descaled.
96	60 SMF99_AS_PB_SERVICE	8	binary	Transaction service units on standard CP reported for PBs running in this address space - OucbxPBCP.
104	68 SMF99_AS_PB_OFFLOAD_SERVICE	8	binary	Transaction service units on offload engines reported for PBs running in this address space - OucbxPBOffload.
112	70 SMF99_AS_PB_OFFLOADONCP_SERVICE	8	binary	Transaction service units on standard CP that were offload eligible reported for PBs running in this address space – OucbxPBOffloadOnCP.
120	78 SMF99_AS_ENCLAVE_TIME	4	binary	Accumulate tx active time of completed enclaves owned by this space - OUCBETIM.
124	7C SMF99_AS_ENCLAVE_CPU_SERVICE	4	binary	Accumulated CPU service of completed enclaves owned by this space – OUCBECPU.
128	80 SMF99_AS_ENCLAVE_IFA_TIME	8	binary	Total IFA time for the enclaves owned by the address space – OucbxEncTimeOnPro(pro_ifs).
136	88 SMF99_AS_ENCLAVE_IFACP_TIME	8	binary	Total IFA time spent on CP for the enclaves owned by the address space - OucbxEncTimeProOnCP(pro_ifs).
144	90 SMF99_AS_ENCLAVE_SUP_TIME	8	binary	Total SUP time for the enclaves owned by the address space – OucbxEncTimeOnPro(pro_sup).
152	98 SMF99_AS_ENCLAVE_SUPCP_TIME	8	binary	Total SUP time spent on CP for the enclaves owned by the address space - OucbxEncTimeProOnCP(pro_sup).
160	A0 SMF99_AS_BA_BrkLoc	18	character	Location element for each processor type which describes the breakup environment of this address space, or 0.
178	B2 SMF99_AS_BA_Loc	18	character	Location element for each processor type of this address space, or 0.
196	C4 SMF99_AS_BA_Flag_Area	3	binary	HiperDispatch status flags.
199	C7 *	1	binary	Reserved.
200	C8 SMF_AS_BA_Mem_Score	8	character	Memory score.
208	D0 SMF_AS_MP_Actual_Pool	8	character	Actual memory pool of the address space.
216	D8 SMF99_AS_MP_Requested_Pool	8	character	User requested memory pool, can differ from actual pool due to RSM constraints.
224	E0 SMF99_AS_TRC	8	character	Tenant report class of address space.
232	E8 SMF99_AS_TRG	8	character	Tenant resource group of address space.
240	F0 *	16	binary	For IBM use only.

Record type 1156 (X'484') – Module fetch

The SMF type 1156 record provides information about module fetch processing.

The type 1156 subtype-specific records are mapped using the following mapping macros:

- IFASMFH, which maps the SMF standard and extended header area of the record
- CSVS1156, which maps the data provided for SMF record type 1156.

Both IFASMFH and CSVS1156 reside in SYS1.MACLIB.

Subtypes

The SMF type 1156 record has the following subtypes:

Subtype 1

Information pertaining to a module load with an associated workload class.

Record environment

The following conditions exist for the generation of this record:

Macro

SMFEWTM, BRANCH=YES (record exit: IEFU84)

Storage residency

31-bit

Record mapping

The SMF type 1156 record is mapped as follows:

Header/self-defining section

This section contains the common SMF extended SMF record header version 1 fields and the triplet fields (offset/length/number), if applicable, that locate the other sections on the record.

Offsets	Name	Length	Format	Description
0 0	SMF1156_H_SHD	24	binary	Start of SMF standard header.
0 0	SMF1156_H_LEN	2	binary	RDW record length.
2 2	SMF1156_H_SEG	2	binary	RDW record segment descriptor.
4 4	SMF1156_H_Flag	1	binary	Flag byte. Bit Meaning when set 0 Reserved. 1 Record has a subtype. 2 Extended header is present. 3-7 Reserved.
5 5	SMF1156_H_RTY	4	binary	Record type - always 126.
6 6	SMF1156_H_Time	4	binary	Record write time, in hundredths of seconds since midnight.
10 A	SMF1156_H_Date	4	binary	Record write date, in packed decimal format, 0cyydddF.
14 E	SMF1156_H_SID	4	EBCDIC	System identification.
18 12	SMF1156_H_WID	4	EBCDIC	Subsystem identification.
22 16	SMF1156_H_STP	2	binary	Record subtype.
24 18	SMF1156_H_EV1	32	binary	Length of V1 extension.

Offsets	Name	Length	Format	Description
24 18	SMF1156_H_ELN	2	binary	Length of extension.
26 1A	SMF1156_H_VER	1	binary	SMF header version.
27 1B	SMF1156_H_FG1	1	binary	Additional flag byte. Bit Meaning when set 0 The IEFU86 exit was called for this record.
28 1C	SMF1156_H_STCKE	16	binary	Time when the record was written, in STCKE format.
44 2C	SMF1153_H_TZO	8	binary	Time zone offset, taken from CVTLDTO.
52 34	SMF1156_H_ETY	2	binary	Extended record type - always 1156.
54 36	SMF1156T1_NumTriplets	2	binary	Number of triplets that follow, initially 1.
58 3A	*	6	binary	Reserved.
64 40	SMF1156T1_Triplet0	8		Data section triplet.
64 40	SMF1156T1_DOffset	4	binary	Offset to data section, from the start of SMF1156T1 to the SMF1156T1D section (mapped by SMF1156T1D).
68 44	SMF1156T1_DLength	2		Length of data section.
70 46	SMF1156T1_DNum	2		Number of data sections. Always 1.
72 48	*	0	*	End of header/self-defining section.

Subtype 1 — Module load with associated workload class

An SMF type 1156 subtype 1 record is produced whenever a module is loaded with an associated workload class.

Triplet information: The subtype 1 specific area can be located in the SMF 1156 subtype 1 record by using the following triplet fields in the “Header/self-defining section” on [page 139](#). These fields are defined in the CSVS1156 mapping macro.

Offset

SMF1156T1_DOffset

Length

SMF1156T1_DLength

Number

SMF1156T1_DNum

The SMF type 1156 subtype 1 record is mapped as follows.

Data section

Information pertaining to a module load with an associated workload class. Locate this section by adding to the address of the SMF record the value in SMF1156T1_DOffset.

Offsets	Name	Length	Format	Description
0 0	SMF1156T1D	48	binary	Data section.
0 0	SMF1156T1D_Version	1	binary	Version of data section.

Offsets	Name	Length	Format	Description
1 1	SMF1156T1D_Flags	1	binary	<p>Flag byte.</p> <p>Bit (Name) Meaning when set</p> <p>0–1 Result. Exactly one of these 2 bits will be on:</p> <p>0 (SMF1156T1D_Success) Module fetch was successful.</p> <p>1 (SMF1156T1D_NotSuccess) Module fetch was not successful; there was a mismatch related to workload classification. For instance, the UUID is not valid.</p> <p>2–3 By. Exactly one of these 2 bits will be on:</p> <p>2 (SMF1156T1D_ByVolDSN) This was a fetch from a data set, so the SMF1156T1DMVD section is present.</p> <p>3 (SMF1156T1D_ByPath) This was a fetch by path from the file system, so the SMF1156T1DP section is present.</p> <p>4 (SMF1156T1D_Pathname_Truncated) Applies only when ByPath and indicates that the path name was longer than the maximum indicated by the SMF1156T1DP_PathName_MaxLen constant. The last MaxLen characters are placed in the SMF record.</p> <p>5–7 Reserved.</p>
2 2	*	1	*	Reserved.
3 3	SMF1156T1D_WC	1	binary	Workload classification.
4 4	SMF1156T1D_MVD_or_P_Offset	2	binary	Offset from start of record to the SMF1156T1DMVD or SMF1156T1DP section, whichever applies.
6 6	SMF1156T1D_NotSuccess_Reason	2	binary	For IBM use only, the reason when SMF1156T1D_NotSuccess is on.
8 8	SMF1156T1D_JobnameS	8	EBCDIC	Job name for the START, MOUNT, or LOGON command that is associated with this address space. If the first byte is X'00', no command is associated. (Source: ASSBJBNS)
16 10	SMF1156T1D_JobnameI	8	EBCDIC	Job name for the initiated program that is associated with this address space. If the first byte is X'00', no job name is associated. (Source: ASSBJBNI)
24 18	SMF1156T1D_JobID	8	EBCDIC	The job ID, if available; otherwise all zeros.
32 20	SMF1156T1D_UUID	16	binary	When not all zeros, the UUID associated with the module. When all zeros, the UUID is not available.
48 30	SMF1156T1D_Base_End	0	*	End of base.

Module/volume/data set name section

The module/volume/data set name section is present when SMF1156T1D_ByVolDSN is on and is located by adding SMF1156T1D_MVD_or_P_Offset to the address of the start of the SMF record.

Offsets	Name	Length	Format	Description
0 0	SMF1156T1DMVD	58	binary	Module/volume/data set name section.
0 0	SMF1156T1DMVD_Modname	8	EBCDIC	Module name.
8 8	SMF1156T1DMVD_VolDsn	50	EBCDIC	Volume and data set name.

Offsets	Name	Length	Format	Description
8 8	SMF1156T1DMVD_Vol	6	EBCDIC	When the first character is not X'00', the volume from which the module was fetched, padded on the right with EBCDIC blanks if the volume name is shorter than 6 characters.
14 E	SMF1156T1DMVD_Dsname	44	EBCDIC	When the first character is not X'00', the data set from which the module was fetched, padded on the right with EBCDIC blanks if the data set name is shorter than 44 characters.
58 3A *		0	*	End of SMF1156T1DMVD.

Path section

The path section is present when SMF1156T1D_ByPath is on and is located by adding SMF1156T1D_MVD_or_P_Offset to the address of the start of the SMF record.

Offsets	Name	Length	Format	Description
0 0	SMF1156T1DP	*	binary	Path section.
0 0	SMF1156T1DP_Header	16	binary	Path header.
0 0	SMF1156T1DP_PathToken	12	binary	A path uniqueness token provided by z/OS UNIX System Services.
12 C	SMF1156T1DP_LenPath	2	binary	The length of the path name in the next field (SMF1156T1DP_PathName). Only the number of characters indicated by this length field are set in the SMF1156T1DP_PathName field.
14 E	SMF1156T1DP_TotalLenPath	2	binary	When the SMF1156T1D_Pathname_Truncated bit is off, this matches SMF1156T1DP_LenPath. Otherwise, it is the length known to the system (but no more than 64 characters are provided within the SMF record).
16 10	SMF1156T1DP_PathName	*	EBCDIC	The path name (which could be a relative path name) provided as input to identify the file. If the length of the path name does not exceed 64, the whole path name is in the SMF record. If the length of the path name exceeds 64, the last 64 characters are used and the SMF1156T1D_Pathname_Truncated bit is set.
58 3A *		0	*	End of SMF1156T1DP

Constants:

SMF1156T1DP_PathName_MaxLen = 64

Chapter 7. Using RMF with colocated DevTest and New Application Solutions

IBM z/OS Resource Measurement Facility (RMF) is IBM's product for z/OS performance measurement and management. RMF also allows you to tune and configure your system according to your business needs.

The z/OS Data Gatherer interprets WLM data related to CPU consumption for all work running on the system, and records that data to SMF Type 70, 72, and 79 records. This allows the information to be accessed later for RMF reporting and comparisons, client chargeback purposes, and for use by SCRT to build a report of data for sub-capacity billing purposes.

In order to use RMF for Tailored Fit Pricing, you must have z/OS V2.2 or later, with the PTFs applied for APAR OA52694.

For Tailored Fit Pricing, the RMF Postprocessor Workload Activity includes WLM information for colocated DevTest and New Application about tenant resource groups and tenant report classes, as shown in the following figures.

Enhanced Report Class Report showing a tenant report class:

REPORT CLASS(ES)									
POLICY=BASEPOL		TENANT=TENANT01		REPORT CLASS=TRC00001		RESOURCE GROUP=TRGROUP1		REPORT CLASS(ES)	
DESCRIPTION =Tenant report class 1									
-TRANSACTIONS-									
AVG	4.00	ACTUAL	0	TOTAL	N/A	N/A	N/A	AVG ENC	0.00
MPL	4.00	EXECUTION	0	MOBILE	N/A	N/A	N/A	REM ENC	0.00
ENDED	0	QUEUED	0	CATEGORYA	N/A	N/A	N/A	MS ENC	0.00
END/S	0.00	R/S AFFIN	0	CATEGORYB	N/A	N/A			
#SWAPS	0	INELIGIBLE	0						
EXCTD	0	CONVERSION	0						
		STD DEV	0						

Figure 52. RMF report class report showing a tenant report class

W O R K L O A D A C T I V I T Y													PAGE 2			
z/OS 3.1		SYSPLEX UTCLPXC RPT VERSION 3.1 RMF			DATE 09/30/2021 TIME 03.00.00		INTERVAL 14.59.999			MODE = GOAL						
POLICY ACTIVATION DATE/TIME 09/14/2019 09.00.04																
- SERVICE POLICY PAGE -																
SERVICE DEFINITION: CMBSVDEF WLM BASEPOL COMBAT for WBG							-SERVICE		DEFINITION		COEFFICIENTS-		-NORM FACTORS-			
INSTALL DATE: 09/07/2019 13.47.34 INSTALLED BY: RDOWNEY							IOC		CPU		SRB MSO		AAP IIP			
POLICY: BASEPOL WLM BASEPOL COMBAT for WBG							0.0		1.0		1.0 0.0000		1.0000 1.0000			
DISCRETIONARY GOAL MANAGEMENT: YES																
DYNAMIC ALIAS MANAGEMENT: YES																
I/O PRIORITY MANAGEMENT: YES																
SYSTEMS																
---ID---		OPT	SU/SEC	CAP%	--TIME--	INTERVAL	--BOOST--	---ID---	OPT	SU/SEC	CAP%	--TIME--	INTERVAL	--BOOST--		
CB8E		CB	59259.3	100	03.00.00	00.14.59	Inactive	CB89		CB	68965.5	100	03.00.00	00.14.59	Inactive	
RESOURCE GROUPS																
--NAME--		TYPE		-----DESCRIPTION-----			-SYSTEM-		---CPU		CONSUMPTION---		-----CPU CAPACITY-----		---MEMORY---	
									#CPS		MSU SU/SEC		MIN MAX DEFINED AS		USAGE LIMIT	
HWTRG1		TRG		Hardware Container TRG #1					0.22		27 14K					
							CB8E		0.11		14 7K					
							CB89		0.11		13 8K				98M	
				-----REPORT CLASSES			HWTRC1		0.17		21 11K				102M	
							HWTRC3		0.05		6 3K					
HWTRG2		TRG		Hardware Container TRG #2					0.21		25 13K					
							CB8E		0.11		13 6K				111M	
							CB89		0.10		12 7K				131M	
				-----REPORT CLASSES			HWTRC2		0.21		25 13K					
RGPTYP1		RG		V1r12 LI1370 testing, type 1					0.00		0 0		5 100		SU/SEC	
							CB8E		0.00		0 0				0	
							CB89		0.00		0 0				0	
				-----SERVICE CLASSES			RGRP1		0.00		0 0					
RGPTYP3		RG		v1r12 li1370 testing, type 3					0.00		0 0		0.01 1.00		NUMBER OF CPs	
							CB8E		0.00		0 0				0	
							CB89		0.00		0 0				0	
									0.00		0 0					
							RGRP3		0.00		0 0					
							RGRP3B		0.00		0 0					
							RGRP3C		0.00		0 0					

Figure 53. WLMGL Report - Service Policy Page showing tenant resource groups

For complete information on RMF, see:

- *z/OS RMF User's Guide*

CPU Activity - SMF record type 70-1

z/OS RMF Reporter User's Guide describes how to use RMF overview conditions. This section covers only those conditions that can be used to generate Overview reports based on data in SMF type 70 subtype 1 records.

One of the following qualifiers is possible:

cluster

Name of the sysplex or cluster

coreid

A processor identifier (one or two hexadecimal digits) that either identifies a logical core (when LOADxx PROCVIEW CORE is in effect) or a logical processor (when LOADxx PROCVIEW CPU is in effect).

If the qualifier is omitted, the values represent the average of all logical processors or cores.

cpuid

A processor identifier which must be in the format *cpuid*[.*threadid*]

cpuid is a processor identifier (one or two hexadecimal digits) that either identifies a logical core (when LOADxx PROCVIEW CORE is in effect) or a logical processor (when LOADxx PROCVIEW CPU is in effect).

threadid is an optional thread identifier (0 or 1) that identifies a thread that is executing on the logical core designated by *cpuid*. It is ignored when LOADxx PROCVIEW CPU is in effect. If LOADxx PROCVIEW CORE is in effect and *threadid* is omitted, the values represent the average of all threads executing on the logical core.

Examples: 0A, 3F.0, A.1

If the qualifier is omitted, the values represent the average of all logical processors or cores.

lpar

Logical partition name

group

Group of logical partitions managed towards a common group capacity limit

trg

Tenant resource group name

Table 6. CPU Activity - Conditions Based on SMF Record Type 70-1				
Condition	Condition Name	Qualifier	Source	Algorithm
Long-term average of CPU service (millions of service units) consumed by a tenant resource group	TRGLACS	trg	SMF70_TRG_LAC	Value or comparison
Service units on general purpose processors consumed by a tenant resource group per second	TRGCP	trg	SMF70_TRG_SUCP SMF70INT	TRG_SUCP *1000 / INT
Service units on zAAPs consumed by a tenant resource group per second	TRGAAP	trg	SMF70_TRG_SUIFA SMF70INT	TRG_SUIFA *1000 / INT
Service units on zIIPs consumed by a tenant resource group per second	TRGIIP	trg	SMF70_TRG_SUSUP SMF70INT	TRG_SUSUP *1000 / INT
General purpose processor consumption in terms of 1/100 of a CP	TRGCPN	trg	SMF70_TRG_SUCP SMF70ADJ SMF70INT	(TRG_SUCP *ADJ) / (INT*160000)
zAAP processor consumption in terms of 1/100 of a CP	TRGAAPN	trg	SMF70_TRG_SUIFA SMF70ADJ SMF70INT	(TRG_SUIFA *ADJ) / (INT*160000)

Table 6. CPU Activity - Conditions Based on SMF Record Type 70–1 (continued)				
Condition	Condition Name	Qualifier	Source	Algorithm
zIIP processor consumption in terms of 1/100 of a CP	TRGIIPN	trg	SMF70_TRG_SUSUP SMF70ADJ SMF70INT	$(TRG_SUSUP * ADJ) / (INT * 160000)$

Appendix A. Accessibility

Accessible publications for this product are offered through [IBM Documentation for z/OS \(www.ibm.com/docs/en/zos\)](http://www.ibm.com/docs/en/zos).

If you experience difficulty with the accessibility of any z/OS documentation see [How to Send Feedback to IBM](#) to leave documentation feedback.

Notices

This information was developed for products and services that are 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
United States of America*

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*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 states 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.

This information could include missing, incorrect, or broken hyperlinks. Hyperlinks are maintained in only the HTML plug-in output for IBM Documentation. Use of hyperlinks in other output formats of this information is at your own risk.

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 Corporation
Site Counsel
2455 South Road*

Poughkeepsie, NY 12601-5400
USA

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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 the names and addresses used by an actual business enterprise 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.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or

reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's name, email address, phone number, or other personally identifiable information for purposes of enhanced user usability and single sign-on configuration. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at ibm.com/privacy and IBM's Online Privacy Statement at ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at ibm.com/software/info/product-privacy.

Policy for unsupported hardware

Various z/OS elements, such as DFSMSdfp, JES2, 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 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).

Microsoft, Windows, Windows NT, and the Windows logo 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.

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.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Index

Special Characters

--container command line option [77](#)
--ignore-container command line option [78](#)
--update-container command line option [78](#)
--update-solution-id command line option [78](#)

A

accessibility
 contact IBM [147](#)
Active Containers section of multiplex report [53](#)
Active Containers section of sub-capacity report [41](#)
address space resource data
 RMF - record type 79-2 [112](#)
address space SRM data
 RMF - record type 79-5 [117](#)
address space state data
 RMF - record type 79-1 [107](#)
application
 definition menu [18](#)
assistive technologies [147](#)

C

capacity, minimum and maximum [13](#), [14](#)
central processor unit (CPU) activity
 Postprocessor
 overview/exception control statements
 [144](#)
 SMF record type 70-1 [144](#)
classification rules
 querying [87](#)
CLI [76](#)
command line interface (CLI) [76](#)
command line options
 --container [77](#)
 --ignore-container [78](#)
 --update-container [78](#)
 --update-solution-id [78](#)
contact
 z/OS [147](#)
Container Active Products section of multiplex report [53](#)
Container Active Products section of sub-capacity report [42](#)
Container Billing Summary section of multiplex report [53](#)
Container Billing Summary section of sub-capacity report [41](#)
CONTAINER control statement [55](#)
Container Interval Consumption section of sub-capacity report [43](#)
Container Max Contributors section of sub-capacity report [47](#), [55](#)
Container MSU Consumption Contributors section of sub-capacity report [48](#)
Container Product Grid Snapshot section of sub-capacity report [46](#), [55](#)
Container Product Metrics section of sub-capacity report [44](#)
CONTAINER statement [55](#)

CPC MSU Rate Above Permanent Capacity Rating section of sub-capacity report [28](#)
customer information section of sub-capacity report [39](#)

D

DD statements
 SPECIAL
 CONTAINER [55](#)
 IGNORE CONTAINER [55](#), [57](#)
 UPDATE CONTAINER [55](#), [58](#)
 UPDATE SOLUTION_ID [55](#), [59](#)
Detail LPAR Data section of sub-capacity report [45](#)
Detailed Interval Data section of sub-capacity report [29](#)
DTD - XML service definition [87](#)

E

ENF event code [41](#) [79](#)
EXCEPT (exceptional value) control statement
 Postprocessor
 CPU activity (type 70-1 SMF record) [144](#)
exception report
 Postprocessor
 exceptions list [144](#)

F

functionality levels, XML [90](#)

G

graphical user interface (GUI) [59](#)
GUI [59](#)

H

Hardware Consumption Solution reporting [27](#)
heterogeneous tenant report class [12](#)
homogeneous tenant report class [12](#)

I

IGNORE CONTAINER statement [55](#), [57](#)
interval rate data, reporting of [27](#)
IWM4QTNT macro [86](#)
IWM4TAF macro [80](#)
IWMAM313 message [20](#)
IWMAM512 message [21](#)
IWMAM540 message [21](#)

K

keyboard
 navigation [147](#)

keyboard (*continued*)
PF keys [147](#)
shortcut keys [147](#)

L

layout - XML service definition [87](#)

M

machine summary section of multiplex report [50](#)
maximum capacity [13](#), [14](#)
memory
 limit [13](#), [14](#)
memory limit [13](#), [14](#)
minimum capacity [13](#), [14](#)
Missing CPC Data Details section of sub-capacity report [40](#)
Missing LPAR Data Details section of sub-capacity report [40](#)
MSU Consumption Data Collection section of sub-capacity report [43](#)
multiplex report
 Active Containers section [53](#)
 Container Active Products section [53](#)
 Container Billing Summary section [53](#)
 customer verification section for CPC
 machine information [54](#)
 machine summary section [50](#)
 machine summary section for CPC
 Active Containers section [52](#)
 machine information [52](#)
 Missing CPC data details [52](#)
 Missing LPAR data details [52](#)
 Product summary information [52](#)
 multiplex summary section [50](#)
 Product Grid Snapshot section [54](#)
 Product Max Contributors section [54](#)
multiplex summary section of multiplex report [50](#)

N

name spaces, XML [90](#)
navigation
 keyboard [147](#)

O

overview report
 Postprocessor
 conditions list [144](#)
OVW (overview condition) control statement
 Postprocessor
 CPU activity (type 70–1 SMF record) [144](#)

P

Postprocessor
 CPU activity (type 70–1 SMF record) [144](#)
Product Grid Snapshot section of multiplex report [54](#)
Product Grid Snapshot section of sub-capacity report [46](#)
Product Max Contributors section of multiplex report [54](#)
Product Max Contributors section of sub-capacity report [45](#)
Product Summary Information section of sub-capacity report [40](#)

R

resource group
 defining [14](#)
 limitations [14](#)
 maximum capacity [14](#)
 minimum capacity [14](#)
RMF (Resource Measurement Facility)
 Type 723 - workload activity (goal mode) [104](#)
 Type 79 - Monitor II snapshot data
 Type 79-1 - address space state data [107](#)
 Type 79-2 - address space resource data [112](#)
 Type 79-5 - address space SRM data [117](#)

S

SCRT commands
 Hardware Consumption Solution [23](#)
shortcut keys [147](#)
SMF record
 type 70–1 (CPU activity)
 overview/exception control statements
 [144](#)
SMF/SCRT89 Input Data Statistics section of sub-capacity report [45](#)
Software Consumption Solution , See Tailored Fit Pricing for IBM Z - Software Consumption Solution
SOLUT parameter in IEASYSxx [24](#)
SPECIAL DD statement
 CONTAINER [55](#)
 IGNORE CONTAINER [55](#), [57](#)
 UPDATE CONTAINER [55](#), [58](#)
 UPDATE SOLUTION_ID [55](#), [59](#)
sub-capacity report
 Active Containers section [41](#)
 Container Active Products section [42](#)
 Container Billing Summary section [41](#)
 Container Interval Consumption section [43](#)
 Container Max Contributors section [47](#), [55](#)
 Container MSU Consumption Contributors section [48](#)
 Container Product Grid Snapshot section [46](#), [55](#)
 Container Product Metrics section [44](#)
 CPC MSU Rate Above Permanent Capacity Rating section [28](#)
 customer information section [39](#)
 Detail LPAR Data section [45](#)
 Detailed Interval Data section [29](#)
 Missing CPC Data Details section [40](#)
 Missing LPAR Data Details section [40](#)
 MSU Consumption Data Collection section [43](#)
 Product Grid Snapshot section [46](#)
 Product Max Contributors section [45](#)
 Product Summary Information section [40](#)
 SMF/SCRT89 Input Data Statistics section [45](#)
 Tool Information section [39](#)
sub-capacity reporting
 Tailored Fit Pricing for IBM Z [33](#)
 Tailored Fit Pricing for IBM Z - Software Consumption Solution [25](#)
Sub-Capacity Reporting Tool (SCRT)
 for Windows and Linux
 using the CLI [76](#)
 using the GUI [59](#)

T

- Tailored Fit Pricing
 - requirements [9](#)
- Tailored Fit Pricing for IBM Z
 - interaction with traditional sub-capacity [34](#)
 - overview [33](#)
 - sample report
 - control statement input [26](#)
 - sample sub-capacity report for collocated solution
 - billing-related section [39](#)
 - customer verification section [44](#)
 - sample sub-capacity report for standalone solution [49](#)
 - SCRT example for a collocated solution [38](#)
 - SCRT example for a standalone solution [48](#)
 - SCRT externals for [35](#)
 - SCRT sample configuration for collocated solution [38](#)
 - SCRT sample configuration for standalone solution [48](#)
 - SCRT sample control statement input for collocated solution [38](#)
 - SCRT sample control statement input for standalone solution [49](#)
- Tailored Fit Pricing for IBM Z - Hardware Consumption Solution [27](#)
- Tailored Fit Pricing for IBM Z - Software Consumption Solution
 - overview [25](#)
 - sample report
 - configuration [26](#)
 - SCRT externals for [25](#)
- tenant report class
 - assigning [12](#)
 - definition [12](#)
 - heterogeneous [12](#)
 - inhomogeneous [12](#)
- tenant resource group
 - defining [13](#)
 - limitations [13](#)
 - maximum capacity [13](#)
 - minimum capacity [13](#)
- Tool Information section of sub-capacity report [39](#)
- trademarks [152](#)

U

- UPDATE CONTAINER control statement [58](#)
- UPDATE CONTAINER statement [55](#)
- UPDATE SOLUTION_ID control statement [59](#)
- UPDATE SOLUTION_ID statement [55](#)
- usage
 - data set
 - DD DATA [118](#)
 - DD* [118](#)
 - record type [89](#) [118](#)
- user interface
 - ISPF [147](#)
 - TSO/E [147](#)

W

- workload activity (goal mode)
 - RMF - record type 72-3 [104](#)
- Workload Classification Supplement [30](#)

- workload reporting services
 - when to use [79](#)

X

- XML name spaces and functionality levels [90](#)
- XML output of IWMDEXTR [87](#)
- XML service definition (DTD) [87](#)



Product Number: 5655-ZOS