

z/OS
3.2

File System Messages and Codes



Note

Before using this information and the product it supports, read the information in [“Notices” on page 361.](#)

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 1996, 2025.**

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

Contents

About this document.....	v
How to provide feedback to IBM.....	vii
Summary of changes.....	ix
Summary of message changes for z/OS 3.2.....	xi
Summary of message changes for z/OS 3.1.....	xiii
Chapter 1. Introduction to zFS messages and codes.....	1
Chapter 2. IOEZnnnnnt: zFS messages.....	5
Chapter 3. IOEZHnnnnnt: zFS Health Checker messages.....	221
Appendix A. Reason codes.....	241
Appendix B. Accessibility.....	359
Notices.....	361
Trademarks.....	365
Index.....	367

About this document

Detailed explanations and recovery actions are provided for the messages issued by z/OS File System (zFS). zFS is one of the file systems that is used by z/OS UNIX System Services (z/OS UNIX). See [z/OS File System Administration](#) for more information about zFS.

Prior to V2R4, this document was titled *z/OS Distributed File Service Messages and Codes*. Because V2R3 was the last release to support the Distributed File System / Server Message Block (DFS/SMB) functionality, the title is now *z/OS File System Messages and Codes* but the DFS/SMB messages are still included in the documentation.

Who should use this document

This document is intended for operators, administrators, programmers, and help desk representatives who require explanations and recoveries for the messages issued by zFS. A knowledge of TCP/IP communications and the z/OS UNIX operating system can help you use this guide more effectively.

z/OS information

This information explains how z/OS references information in other documents and on the web.

When possible, this information uses cross-document links that go directly to the topic in reference using shortened versions of the document title. For complete titles and order numbers of the documents for all products that are part of z/OS, see *z/OS Information Roadmap*.

Using this documentation

Two methods of installing z/OS are provided with your z/OS license: ServerPac z/OSMF portable software instance (ServerPac) and CBPDO. For each of these, z/OS Planning for Installation describes what IBM does for you, what you receive from IBM, and what actions you need to take. It is important that you are familiar with the information that comes with those two installation methods. The information is needed when you install z/OS UNIX, along with the other elements and features.

z/OSMF Software Management produces the jobs that are used to install your z/OSMF portable software instance (ServerPac). For CBPDO users, the [z/OS Program Directory](#) describes how to use the SMP/E RECEIVE, APPLY, and ACCEPT commands to install your order. Both describe the installation verification procedures (IVPs) that you perform to ensure that your installation is proceeding successfully. They also contain some customization information.

Information about zFS configuration on other IBM systems can be found in the configuration guide for those systems.

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](#).

Summary of changes

This information includes terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Note: IBM z/OS policy for the integration of service information into the z/OS product documentation library is documented on the z/OS Internet Library under [IBM z/OS Product Documentation Update Policy](http://www.ibm.com/docs/en/zos/latest?topic=zos-product-documentation-update-policy) (www.ibm.com/docs/en/zos/latest?topic=zos-product-documentation-update-policy).

Summary of message changes for z/OS 3.2

The following messages are new, changed, or no longer issued for z/OS File System Messages and Codes in z/OS 3.2.

New

The following messages are new.

IOEZ01040A
IOEZ01042E
IOEZ01060E
IOEZ01061E
IOEZ01076E
IOEZ01077I
IOEZ01078E
IOEZ01079E
IOEZ01080I
IOEZ01081I
IOEZ01082I
IOEZ01083I
IOEZ01084I
IOEZ01085E
IOEZ01086E
IOEZ01087I
IOEZ01088I
IOEZ01089E

Changed

The following messages are changed.

IOEZH0083I
IOEZH0084I

Deleted

The following messages are no longer issued.

None.

Summary of message changes for z/OS 3.1

The following messages are new, changed, or no longer issued for z/OS File System Messages and Codes in z/OS 3.1.

New

The following messages are new.

IOEZH0076I
IOEZH0077I
IOEZH0078I
IOEZH0079I
IOEZH0080I
IOEZH0083I
IOEZH0084I
IOEZH0088I
IOEZ01061I

Changed

The following messages are changed.

None.

Deleted

The following messages are no longer issued.

IOEN000I–IOEN00515I
IOEP00001A–IOEP12403I
IOEX18000A–IOEX18234K

Chapter 1. Introduction to zFS messages and codes

Detailed explanations and recovery actions are provided to supplement the messages that are issued by z/OS File System (zFS) and the zFS health checks. Processing options that can be used to control zFS are documented in *z/OS File System Administration*.

In z/OS V2R3, the statement of direction was updated to indicate that V2R3 would be the last release to support the DFS/SMB (Distributed File System/Server Message Block) functionality. IBM had previously announced that the Network File System (NFS) was the strategic file sharing protocol for the z/OS platform. The related messages and return codes are still included in the documentation.

Where to find the most current message information

In the z/OS library, as messages are added to the system, they are added to the documents. Similarly, when messages are changed on the system, they are changed in the documents. However, when a message is deleted from the system because they are no longer issued, the message is **not** deleted from the documents.

To find the most current description of a message, see the following resources:

[z/OS Internet library \(www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary\)](http://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary)
[IBM Documentation \(www.ibm.com/docs/en/zos\)](http://www.ibm.com/docs/en/zos)

Understanding zFS messages

Messages are divided into three categories: action, error, and informational. Those of the form IOEZnnnnnt are issued by zFS, while messages prefixed with IOEZHnnnnnt are issued from the zFS health checker.

The last part of the message number (*t*) is a single-character operator code that represents the type of recovery action to be taken by the operator in response to the message. These operator codes are associated with the different levels of severity, as listed in “Severity levels” on page 2. The recovery actions are represented by the following characters:

- A**
Indicates that the operator must take immediate action, for example, recovering from a system failure.
- D**
Indicates that the operator must choose an alternative action, for example, responding yes or no.
- E**
Indicates that the operator must take eventual action, for example, loading paper in a printer.
- I**
Indicates that the message is informational and no action is required, for example, when a command completes successfully.

What is first-level message text?

Together, the message number and the message text that immediately follows the message number comprise the first-level text. This is the message information that is returned to the operator console or user display.

What is second-level message text?

The explanation and responses that supplement the message number and message text comprise the second-level text.

Severity levels

Each message contains a severity field that indicates the level of severity associated with the message. These severity levels are described in the following list:

svc_c_sev_fatal

Indicates a non-recoverable error that probably requires manual intervention. Usually, permanent loss or damage has occurred that results in the program terminating immediately, such as database corruption. Messages of this severity have an operator code of A.

svc_c_sev_error

Indicates that an unexpected event has occurred that does not terminate the program and which can be corrected by human intervention. The program continues although certain functions or services might remain unavailable. Examples include performance degradation that results in a loss of function, such as a timeout, or a specific request or action that cannot be completed, such as trying to add an object to a directory system when the object already exists. Messages of this severity have an operator code of A or E.

svc_s_sev_user

Indicates that a usage error on a public API has occurred, such as a syntax error. Messages of this severity have an operator code of A or E.

svc_c_sev_warning

Indicates one of the following conditions; messages of this severity have an operator code of E or I.

- An error occurred that was automatically corrected by the program or system. An example of an error corrected by a program is when a configuration file is not found during configuration and a message is issued warning the user that certain internal defaults were used.
- A condition has been detected which may be an error depending on whether the effects of the condition are acceptable. For example, a directory is deleted and a warning message is issued that all files contained in the directory will also be deleted.
- A condition exists that, if left uncorrected, will eventually result in an error.

svc_c_sev_notice

Indicates major events, such as the start of a server, completion of server initialization, or an offline server. Messages of this severity have an operator code of E or I.

svc_sev_notice_verbose

Indicates events of special interest, such as statistical information, key data values, use of default settings, and version information. However, it does not indicate program flow or normal events. Messages of this severity have an operator code of I.

Setting slip traps to obtain diagnosis data

When you receive a zFS reason code and need additional diagnosis information, IBM Support Center might ask you set a slip trap to collect a dump when you recreate the problem.

As a general example, you can perform the following steps to obtain a dump on a specific zFS reason code:

1. Enter the following command to determine the ASID for the zFS Physical File System:

```
D A,ZFS
```

2. Look for the ASID of the ZFS address space in A=nnnn.
3. Use the following SLIP command to produce a dump when a specific reason code is issued:

```
SLIP SET,IF,A=SYNCSVCD,RANGE=(10?+8C?+F0?+1F4?),  
ASIDLST=(nnnn),DATA=(13R??+1b0,EQ,xxxxxxx),DSPNAME=('OMVS'.*),  
SDATA=(ALLNUC,PSA,CSA,LPA,TRT,SQA,RGN,SUM),j=jobname,END
```

nnnn

The zFS ASID

xxxxxxx

The 4-byte zFS reason code to trap (for example, EF17624E)

jobname

The optional job name that is expected to issue the error (for example, j=IBMUSER).

Chapter 2. IOEZnnnnnt: zFS messages

The following messages may result from zFS processing.

IOEZ00001E **zFS I/O error error occurred for aggregate *aggregate*.**

Explanation:

A physical I/O error occurred on aggregate *aggregate*. Additional messages are provided that give more information about the error. If the error occurred with file system metadata, the aggregate may be disabled for writing.

System action:

The program continues.

Administrator response:

Determine if there is a real hardware error. If there is, then take steps to correct the problem. Otherwise, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00002E **MMRE error id=*id* cis=*count* if=*inflags* of=*outflags* buf=*buffer* CI=*ci*.**

Explanation:

This message documents additional information about a physical I/O error that occurred on a zFS aggregate. This information is intended for IBM service personnel, and is used to provide more information about the details of the I/O error.

System action:

The program continues.

Administrator response:

Determine if there is a real hardware error. If there is, then take steps to correct the problem. Otherwise, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00003E **While opening minor device *minor_number*, could not open dataset *dataset_name*.**

Explanation:

The device driver cannot open a zFS aggregate.

System action:

The program continues.

Administrator response:

Verify that the named data set exists and is accessible.

Severity:

svc_c_sev_error

IOEZ00004I **Formatting to *control_interval_size* control interval *control_interval_number* for primary extent of *aggregate_name*.**

Explanation

The zFS device driver encountered an unformatted aggregate and is loading it.

In the message text:

aggregate_name

Name of aggregate.

control_interval_size

The control interval size.

control_interval_number

Number of control intervals.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00005I

Primary extent loaded successfully for *dataset_name*.

Explanation:

The device driver has finished loading the primary extent of the named aggregate.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00006E

Error *ErrorCode* occurred while loading dataset '*dataset_name*'.

Explanation:

An error occurred while the device driver was loading the named linear data set.

System action:

The program continues.

Administrator response:

Verify that the data set exists and that the application can write to it.

Severity:

svc_c_sev_error

IOEZ00007A

zFS kernel unable to set up to receive Operator commands.

Explanation:

The program was not able to setup to receive MODIFY commands from the operator.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_fatal

IOEZ00008E

***ProgName*: Unable to open debug parameter dataset '*ParmFile*'.**

Explanation:

The *ProgName* program was not able to open the zFS debug parameter data set. The default values will be used for zFS debugging parameters.

System action:

The program continues.

Administrator response:

Ensure that the *ParmFile* data set exists and that the *ProgName* program has read data set authority to it. If the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00009I

***ProgName*: Using default values for zFS debug parameters.**

Explanation:

The *ProgName* program is using default values for zFS debug parameters.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00010A *ProgName: Incorrect parameter record 'BadRecord' is ignored.*

Explanation:

The *ProgName* program found an incorrect record in the zFS debug parameters data set.

System action:

The program ends.

Administrator response:

Correct the record and start *ProgName* again.

Severity:

svc_c_sev_fatal

IOEZ00011A *ProgName: Operand must be numeric (BadArg).*

Explanation

The *ProgName* program found a non-numeric value specified for a numeric operand.

System action:

The program ends.

Administrator response:

Specify a numeric value and start *ProgName* again.

Severity:

svc_c_sev_fatal

IOEZ00012A *ProgName: BadArg exceeds maximum value MaxValue.*

Explanation:

The *ProgName* program found a numeric operand that was too large.

System action:

The program ends.

Administrator response:

Ensure that the operand *BadArg* is not larger than the *MaxValue* and restart *ProgName*.

Severity:

svc_c_sev_fatal

IOEZ00013A *ProgName: Operand must be string (BadArg).*

Explanation:

The *ProgName* program found a missing string operand.

System action:

The program ends.

Administrator response:

Ensure that the *BadArg* operand is specified and restart *ProgName*.

Severity:

svc_c_sev_fatal

IOEZ00014A *ProgName: BadArg exceeds max string length Maxlen.*

Explanation:

The *ProgName* program found a string operand that is too long.

System action:

The program ends.

Administrator response:

Ensure that the *BadArg* operand string is not greater than *Maxlen* and restart *ProgName*.

Severity:

svc_c_sev_fatal

IOEZ00015A

ProgName: Partition start list 'BadArg' not delimited by parentheses.

Explanation:

The *ProgName* program found a syntax error in operand *BadArg* in the zFS debug parameter data set.

System action:

The program ends.

Administrator response:

Ensure that the *BadArg* operand string is not greater than *Maxlen* and restart *ProgName*.

Severity:

svc_c_sev_fatal

IOEZ00016I

ProgName: Modify complete for Parameter Value.

Explanation:

The *ProgName* program has successfully set the value of the program parameter *Parameter* to *Value*.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00017A

ProgName: Incorrect parameter debug record BadRecord in dataset dataset_name.

Explanation:

The *ProgName* program found a syntax error in ioedebg record *BadRecord*.

System action:

The program ends.

Administrator response:

Correct the error in the ioedebg record and restart *ProgName*.

Severity:

svc_c_sev_fatal

IOEZ00018A

ProgName: Incorrect parameter BadParameter.

Explanation:

The *ProgName* program found an incorrect parameter *BadParameter* in the zFS debug parameters data set.

System action:

The program ends.

Administrator response:

Specify a valid parameter and restart *ProgName*.

Severity:

svc_c_sev_fatal

IOEZ00019E***ProgName: Parameter Parameter is not valid for modify command CommandName.*****Explanation**

The specified parameter *Parameter* is not valid for the *CommandName* command.

System action

The program continues.

Administrator response

Additional information about the syntax of the *CommandName* command is found in the [z/OS File System Administration](#).

Severity

svc_c_sev_error

IOEZ00020I***progrname: prodname featurename Version ver.rel.mod Service Level slv.
Created on date.*****Explanation**

This message is issued when the level of zFS is queried using the QUERY command. The system returns the product name, feature name, version, release, modlevel, service level and creation date of the zFS program.

System action

The program continues.

Severity

svc_c_sev_notice

IOEZ00023E***ProgName: modify command is not valid - Parm.*****Explanation:**

The syntax of the MODIFY command is not correct.

System action:

The program continues.

Administrator response:

Additional information about the syntax of the *Parm* command is found in the [z/OS File System Administration](#).

Severity:

svc_c_sev_error

IOEZ00024E***ProgName: MODIFY command - Parm failed.*****Explanation:**

The MODIFY command *Parm* failed.

System action:

The program continues.

Administrator response:

Additional information about the syntax of the *Parm* command is found in [z/OS File System Administration](#). Verify that the command was entered with valid parameters. Also, ensure that the command was issued in an environment where it is supported. If the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00025I***ProgName: MODIFY command - Parm completed successfully.*****Explanation:**The MODIFY command *Parm* was successful.**System action:**

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00032I***ProgramName: in-memory trace table has been reset.*****Explanation:**

The in-memory trace table has been reset by the RESET MODIFY command.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00033E***ProgramName: could not open trace output dataset.*****Explanation:**

The in-memory trace table could not be printed because the output trace data set could not be opened.

System action:

The program continues.

Administrator response:

Ensure that the trace_dsn parameter specifies a valid data set name that can be opened for output.

Severity:

svc_c_sev_error

IOEZ00034I***ProgramName: printing contents of trace to dataset <DatasetName>.*****Explanation:**The in-memory trace table is being written to data set *DatasetName*.**System action:**

The program continues.

Severity:

svc_c_sev_notice

IOEZ00035I***No trace table to print.*****Explanation:**

There was no in-memory trace table found to be printed.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00036I***Printing contents of table at address Address name: Name.*****Explanation:**The contents of in-memory trace table name *Name* is being printed.**System action:**

The program continues.

Severity:

svc_c_sev_notice

IOEZ00037I Trace table is already being printed.

Explanation:

The contents of in-memory trace table is already being printed.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00038E Incorrect record length *Reclen* found while printing trace table.

Explanation:

A record was found to have an incorrect length while printing the trace table.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00039E Mismatched record length *SReclen*, *ERclen* found while printing trace table.

Explanation:

A record was found to have a starting record length that did not match the ending record length in the trace table.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00040E An incorrect record at offset *offset* was found while printing trace table.

Explanation

An incorrect record was found when printing the in-memory trace table.

In the message text:

offset

The offset to the incorrect record in the trace table.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00041I No start record, trace wrapped, total records *Records*, *Bytes* bytes to format.

Explanation:

An in-memory trace table wrapped. The whole table is being printed. Important data may be lost.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00042I **Start record found, total records *Records*, *Bytes* bytes to format.**

Explanation:

An in-memory trace is being printed. The trace table has not wrapped since the last printing or program start.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00043I ***ProgramName*: print of in-memory trace table has completed.**

Explanation:

An in-memory trace table print has completed.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00044I ***Aggregate Name* attached successfully.**

Explanation:

Aggregate Name was successfully attached by zFS.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00045I **Could not attach *aggregate Name* because its already attached.**

Explanation:

Aggregate Name could not be attached because it is already attached.

System action:

The program continues.

Severity:

svc_c_sev_warning

IOEZ00046E **Error *ErrorCode* received while attaching *aggregate Name*.**

Explanation:

Aggregate Name could not be attached. See [Return codes \(errno\)](#) in [z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program continues.

Administrator response:

Ensure that the aggregate is accessible to zFS in the desired attach mode (read-only or read-write). Also, ensure that you are not mounting or attaching an aggregate read-write when it is already mounted or attached read-write on another (non-sysplex coordinated) system. If the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00048I **Detaching aggregate *Name*.****Explanation:**

Aggregate *Name* is being detached.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00050I ***ProgramName*: Stop command received.****Explanation:**

The zFS kernel has received an operator stop request. All aggregates will be detached and the zFS kernel will terminate.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00051I **An error occurred in program IOEFSKN.****Explanation:**

This message is issued when an internal error has occurred in program IOEFSKN. The zFS kernel (IOEFSKN) is being internally restarted. zFS is attempting to recover from the internal error. Some programs may see I/O errors.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_notice

IOEZ00052I ***programe*: Initializing *prodname* *featurename* Version *ver.rel.mod*
Service Level *slv*. Created on *date*.****Explanation:**

This message is issued when the named zFS program starts. It identifies the product name, feature name, version, release, modlevel, service level and creation date of the zFS program.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00053E **Error setting up termination thread, code = *ErrorCode*.****Explanation:**

An error was encountered attempting to define the routine used to handle the **stop** command from the operator or z/OS UNIX kernel. The result is that a modify omvs,stopdfs=zfs command will not function properly. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00054E **Error starting console thread, code = *ErrorCode*.****Explanation:**

An error was encountered attempting to start a thread to accept operator MODIFYcommand. **F** **zFS** commands will be disabled as a result. See [Return codes \(errno\)](#) in [z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00055I **zFS kernel: initialization complete.****Explanation:**

The zFS kernel has completed initialization. File systems can now be mounted.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00057I **zFS kernel program IOEFSKN is ending.****Explanation:**

IOEFSKN is terminating. It will be restarted by the zFS control program IOEFSCM. Some programs may see I/O errors.

System action:

The program (IOEFSCM) continues

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_notice

IOEZ00062A **zFS kernel: could not create admin thread, code *ErrorCode*.****Explanation:**

A zFS kernel could not create a thread to process administration commands. See [Return codes \(errno\)](#) in [z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program ends.

Severity:

svc_c_sev_fatal

Administrator Response:

Determine if enough storage is specified for the zFS kernel. If increasing storage does not help the problem, contact your service representative.

IOEZ00064I **General Registers RRegNumber: *Reg1 Reg2 Reg3 Reg4*.****Explanation:**

A zFS kernel was driven for recovery for a z/OS UNIX process. This message shows the registers at time of error.

System action:

The program continues.

Severity:

svc_c_sev_warning

Administrator Response:

This message will be followed by additional messages indicating more about the problem. Recovery may not be a problem.

IOEZ00068E **zFS file system *Name* exceeds *Threshold%* full (*blocks1*/*blocks2*) (WARNING).****Explanation:**

A zFS file system used space has exceeded the administrator defined threshold specified on the `fsfull` option. The numbers in parentheses are the number of 1K blocks used in the file system and the number of 1K blocks in the total file system, respectively.

System action:

The program continues.

Severity:

svc_c_sev_warning

Administrator Response:

If appropriate, increase the file system quota.

IOEZ00069I **zFS file system *Name* is now below *Threshold%* full (*blocks1*/*blocks2*).****Explanation:**

A zFS file system used space has fallen below the administrator defined threshold specified on the `fsfull` option. The numbers in parentheses are the number of 1K blocks used in the aggregate and the number of 1K blocks in the total aggregate, respectively.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00070E **Incorrect value *Perms* for permissions.****Explanation:**

An incorrect value *Perms* for the `-perms` field of the `ioeagfmt` command was specified. Either the permissions specified are not valid or you do not have UID=0 (nor READ authority to the SUPERUSER.FILESYS.PFCTL profile in the UNIXPRIV class) and you are attempting to specify a value for permissions that is not the default.

System action:

The program ends.

Severity:

svc_c_sev_error

Administrator Response:

Correct the value and re-execute the `ioeagfmt` command. Note that the value for `-perms` cannot be greater than the hexadecimal value `xFFF`.

IOEZ00077I **HFS-compatibility aggregate *aggregate_name* has been successfully created.****Explanation:**

An HFS-compatibility aggregate was successfully created by the `ioeagfmt` program.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00078E **zFS aggregate *Name* exceeds *Threshold%* full (*blocks1*/*blocks2*) (WARNING).****Explanation:**

A zFS aggregate has exceeded the space utilization threshold that is specified on the `aggrfull` (or `fsfull`) option. The numbers in parentheses are the number of 8 K blocks used in the aggregate and the number of 8 K blocks in the total aggregate, respectively.

System action:

The program continues.

Severity:

`svc_c_sev_warning`

Administrator Response:

If appropriate, increase the aggregate size.

IOEZ00079I

zFS aggregate *Name* is now below *Threshold*% full (*blocks1*/*blocks2*).

Explanation:

A zFS aggregate has now below the space utilization threshold specified on the `aggrfull` (or `fsfull`) option. The numbers in parentheses are the number of 8 K blocks used in the aggregate and the number of 8 K blocks in the total aggregate, respectively.

System action:

The program continues.

Severity:

`svc_c_sev_notice`

IOEZ00080A

Signal *signal* received, dump in progress.

Explanation:

A zFS program is abending, and a dump is in progress.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

`svc_c_sev_fatal`

IOEZ00081A

**zFS: PSW from Machine State: *psw1 psw2* Abend Code from CIB: *abend*
Reason Code from CIB: *reason* Load Module Name: *module*.**

Explanation:

Abend information at time of dump. See [z/OS MVS System Codes](#) for additional information about abend codes.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

`svc_c_sev_fatal`

IOEZ00082A

Immediate end of processing requested: *message_string*.

Explanation:

The zFS program is stopping due to an unrecoverable error.

System action:

The program ends abnormally.

Administrator response:

Contact your service representative.

Severity:

`svc_c_sev_fatal`

IOEZ00083A **Assertion Failed:Assertion Line:LineNo File:FileName.****Explanation:**

An internal error has occurred in the zFS kernel. A dump will be taken.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_fatal

IOEZ00084E **Syntax error string *String* on line(s) *StartLine:EndLine* in configuration dataset *FileName*.****Explanation**

A syntax error was found while parsing the zFS parameter data set. The line in error is shown as an aid to the administrator.

In the message text:

String

The text in error.

StartLine

The starting line number of the statement in error.

EndLine

The ending line number of the statement in error.

FileName

The configuration filename.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Correct the line in error and try the operation again.

IOEZ00085E **Error *ErrorCode* received opening configuration dataset *Dsname*.****Explanation:**

An error was received while attempting to open the configuration data set. The error code from the DFSMS OPEN macro instruction. See [*z/OS DFSMS Macro Instructions for Data Sets*](#) for additional information about OPEN return codes.

System action:

The program continues.

Administrator response:

Ensure that the zFS parameter data set is accessible to the program and try the operation again.

Severity:

svc_c_sev_error

IOEZ00087I **The parameter, *Parameter1*, requires *String* non-null argument value.****Explanation:**

The parameter requires one or more values, but none were specified.

System action:

The program ends.

User response:

Severity:
svc_c_sev_warning

Explanation:
The value specified for the parameter was not valid.

System action:
The program ends.

User response:
Enter the command with a valid value for the parameter.

Severity:
svc_c_sev_warning

Explanation:
Some commands require either a UID of 0 or READ access to the SUPERUSER.FILESYS.PFSCCTL profile in the z/OS UNIXPRIV class. Some commands require UPDATE access to the data set profile. To find the privilege required for zFS commands, see *z/OS File System Administration*.

System action:
The program continues.

Severity:
svc_c_sev_error

Explanation:
An incorrect value *Owner* for the `-owner` field of the **zfsadm format**, **ioeofsutl format**, or **ioeagfmt** command was specified. Either the owner specified does not exist or you do not have UID=0 (nor READ authority to the SUPERUSER.FILESYS.PFCTL profile in the UNIXPRIV class) and you are attempting to specify a value for an owner that is not the default.

System action:
The program ends.

Administrator response:
Correct the value and re-execute the **zfsadm format, ioefsutl format, or ioeagfmt** command.

Severity:
svc_c_sev_error

Explanation:
An incorrect value *Group* for the `-group` field of the **zfsadm format, ioefsutl format, or ioeagfmt** command was specified. Either the owner specified does not exist or you do not have UID=0 (nor READ authority to the SUPERUSER.FILESYS.PFCTL profile in the UNIXPRIV class) and you are attempting to specify a value for an owner that is not the default.

System action:
The program ends.

Administrator response:
Correct the value and re-execute the **zfsadm format, ioefsutl format, or ioeagfmt** command.

Severity:
svc c sev error

18 z/OS: z/OS File System Messages and Codes

Explanation:

An incorrect value *Perms* for the *-perms* field of the **zfsadm format**, **ioefsutl format**, or **ioeagfmt** command was specified.

System action:

The program ends.

Administrator response:

Correct the value and re-execute the **zfsadm format**, **ioefsutl format**, or **ioeagfmt** command.

Severity:

svc_c_sev_error

IOEZ00096E Incorrect value *Size* for *size*.**Explanation:**

An incorrect value *Size* for the *-size* field of the **zfsadm** command was specified.

System action:

The program ends.

Administrator response:

Correct the value and re-execute the **zfsadm** command.

Severity:

svc_c_sev_error

IOEZ00100E File system *Name* was not found.**Explanation**

A **zfsadm** command was executed for a file system that was not found to be attached to the zFS kernel.

System action

The request fails.

Administrator response:

Retry the operation with a valid file system name.

Severity:

svc_c_sev_error

IOEZ00105I File system *Name* deleted successfully.**Explanation**

A file system was successfully deleted.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00106I A total of *Count* aggregates are attached.**Explanation:**

This message shows the total number of aggregates attached.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00109E

Could not retrieve parameter dataset name from zFS kernel.

Explanation:

A **zfsadm** command could not communicate with the zFS kernel to retrieve the parameter data set name. One possible reason is the zFS kernel is not started.

System action:

The program ends.

Administrator response:

Ensure that the zFS kernel is started before running the **zfsadm** command.

Severity:

svc_c_sev_error

IOEZ00110E

Could not read the zFS parameter dataset.

Explanation:

A **zfsadm** command could not read the specified zFS debug parameters data set.

System action:

The program continues.

Administrator response:

Ensure that the debug parameters data set is accessible to the **zfsadm** command.

Severity:

svc_c_sev_error

IOEZ00112E

Must specify at least -all or -aggregate options.

Explanation:

A **zfsadm** command did not specify a required parameter.

System action:

The program ends.

Administrator response:

Retry the command specifying valid options.

Severity:

svc_c_sev_error

IOEZ00117I

Aggregate Name attached successfully.

Explanation:

An aggregate was attached successfully.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00118I

Aggregate Name is already attached.

Explanation:

An aggregate could not be attached because it is already attached.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00119E

Aggregate Name could not be attached, error code=ErrorCode reason code=ReasonCode.

Explanation:

An aggregate could not be attached. The error code is shown with messages that will further explain the error. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, “Reason codes,”](#) on page 241 for a description of the reason code.

System action:

The program continues.

Administrator response:

Use the error code and reason code to determine the cause of the error. Retry the command. If the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00120E

Syntax error with aggregate full parameter - Parameter.

Explanation:

An aggregate could not be attached because the aggrfull processing option is improperly specified.

System action:

The program continues.

Administrator response:

Correct the aggregate full parameter and retry the command.

Severity:

svc_c_sev_error

IOEZ00122I

Aggregate Name detached successfully.

Explanation:

An aggregate was successfully detached by the **zfsadm** command.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00123I

Could not detach aggregate Name because it was not attached.

Explanation:

An aggregate could not be detached because the zFS kernel did not have it attached when the command was executed.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00124E

Error detaching aggregate Name, error code=ErrorCode reason code=ReasonCode.

Explanation:

An aggregate could not be detached due to an unexpected error. Possibly the aggregate was busy with another command, or file systems in the aggregate are mounted. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, “Reason codes,”](#) on page 241 for a description of the reason code.

System action:

The program ends.

Administrator response:

If there are no file systems on the aggregate that are mounted, and there are no other commands running against the aggregate, and the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00127I **No file systems found for aggregate *Name*.**
Explanation:

A **zfsadm lsfs** command was issued to query file system information. No file systems were found in the aggregate.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00129I **Total of *Total* file systems found for aggregate *Name*.**
Explanation:

A **zfsadm** command was issued to list file system information.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00131E **File system name *filesys_name* too long; must be less than *bad_number* chars.**
Explanation:

The specified file system name length is too long.

System action:

The request fails.

Administrator response:

Correct the file system name length to be fewer than the maximum number of characters listed in the error message.

Severity:

svc_c_sev_error

IOEZ00132E **Incorrect file system name *filesys_name*; should not be a number.**
Explanation:

The specified file system must be a name, not a number.

System action:

The program continues.

Administrator response:

Correct the file system to be a character name. Then try the request again.

Severity:

svc_c_sev_error

IOEZ00133E

program_name: Number (bad_number) is too large. Specify a smaller number.

Explanation:

The specified number is too large.

System action:

The program continues.

Administrator response:

Specify a smaller number. Then try the request again.

Severity:

svc_c_sev_error

IOEZ00134E

Incorrect file system name *filesys_name*; cannot end in .bak.

Explanation:

Specify a file system name that does not end in .bak.

System action:

The program continues.

Administrator response:

Correct the file system name. Then try the request again.

Severity:

svc_c_sev_error

IOEZ00135E

Aggregate Name was not found.

Explanation:

A **zfsadm** command was executed for an aggregate that was not found to be attached to the zFS kernel.

System action:

The request fails.

Administrator response:

Retry the operation with a valid aggregate name.

Severity:

svc_c_sev_error

IOEZ00136E

File system *filesys_id_hi*,*filesys_id_low* is busy: *error_text*.

Explanation:

The file system is busy and cannot be accessed.

System action:

The program continues.

Administrator response:

Correct the error listed in the error text. Then try the request again. If the problem continues, contact the service representative.

Severity:

svc_c_sev_error

IOEZ00138E

Open error on file(*dsname*), errno(*errno*).

Explanation:

An error occurred opening the messages data set pointed to by `msg_input_dsn` in the configuration data set.

System action:

The program continues. Messages will be from in core message table.

Administrator response:

Check that the msg_input_dsn variable in the configuration data set and verify that the messages data set exists.

Severity:

svc_c_sev_error

IOEZ00139E read error on file(dsname), errno(errno).

Explanation:

An error occurred reading the messages data set pointed to by msg_input_dsn in the configuration data set.

System action:

The program continues. Messages will be from the core message table.

Administrator response:

Check the msg_input_dsn variable in the configuration data set and verify that the messages data set is valid.

Severity:

svc_c_sev_error

IOEZ00140I Using default message table which is in English.

Explanation:

An error occurred opening or reading the messages data set pointed to by msg_input_dsn in the configuration data set. Verify that the message data set name is valid, exists, and is not damaged.

System action:

The program continues. Messages will be from in core message table.

Administrator response:

Check the msg_input_dsn variable in the configuration data set and verify that the messages data set is valid.

Severity:

svc_c_sev_warning

IOEZ00141E Insufficient storage available.

Explanation:

An error occurred attempting to allocate storage.

System action:

The program continues. Messaging may not occur.

Administrator response:

Check that the task was started with sufficient storage.

Severity:

svc_c_sev_error

IOEZ00142E Message control block is NULL, Cannot print message ID(xmsgid).

Explanation:

The messaging support main control block is corrupted. The message number that was being displayed is reported.

System action:

The program continues. Messaging may not occur.

Administrator response:

Task may need to be restarted to correct the messaging control block.

Severity:

svc_c_sev_error

IOEZ00143E Unable to retrieve message text for message ID(xmsgid).

Explanation:

Correct the msg_output_dsn parameter and be sure it points to a data set that either exists, or zFS has write access to it.

Severity:

svc_c_sev_error

IOEZ00163I **Aggregate Name successfully quiesced.**

Explanation:

An aggregate was successfully quiesced by the **zfsadm** command.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00164I **Could not quiesce aggregate Name because it was not attached.**

Explanation:

An aggregate could not be quiesced because the zFS kernel did not have it attached when the command was executed.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00165E **Error ErrorCode reason code ReasonCode received quiescing aggregate Name.**

Explanation:

An aggregate could not be quiesced due to an unexpected error, or possibly the aggregate was busy with another command or operation. See the section on return codes in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, "Reason codes," on page 241](#) for a description of the reason code.

System action:

The program ends.

Severity:

svc_c_sev_error

Administrator Response:

If there are no other commands or operations running against the aggregate, and the problem continues, contact your service representative.

IOEZ00166I **Aggregate Name successfully unquiesced.**

Explanation:

An aggregate was successfully unquiesced by the **zfsadm** command.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00167I **Could not unquiesce aggregate Name because it was not attached.**

Explanation:

An aggregate could not be unquiesced because the zFS kernel did not have it attached when the command was executed.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00168E

Error *ErrorCode* reason code *ReasonCode* received unquiescing aggregate *Name*.

Explanation:

An aggregate could not be quiesced due to an unexpected error, or possibly the aggregate was busy with another command or operation. See Return codes (errno) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, “Reason codes,” on page 241](#) for a description of the reason code.

System action:

The program ends.

Administrator response:

If there are no other commands or operations running against the aggregate, and the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00170E

Open of file system *FilesystemName* failed with rc *ReturnCode*.

Explanation:

In order to quiesce file system *FilesystemName*, it needs to be opened. This operation failed with a return code of *ReturnCode*. Therefore, the aggregate is not quiesced. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program continues.

Administrator response

Use the return code to determine why the open failed. Then, correct the problem and reissue the quiesce command.

Severity:

svc_c_sev_error

IOEZ00171I

Close of file system *FilesystemName* failed with rc *ReturnCode*.

Explanation:

In order to unquiesce file system *FilesystemName*, it needs to be closed. This operation failed with a return code of *ReturnCode*. The unquiescing of the aggregate will continue until all file systems have been attempted to be closed. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program continues.

Administrator response

Use the return code to determine why the close failed, and correct any problem that may render the file system unusable.

Severity:

svc_c_sev_notice

IOEZ00173I **Aggregate Name successfully grown.**

Explanation:

An aggregate was successfully increased in size by the **zfsadm** command.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00175E **Error growing aggregate *Name*, error code=*ErrorCode* reason code=*ReasonCode*.**

Explanation:

An aggregate could not be increased in size due to an unexpected error, or possibly the aggregate was busy with another command or operation. You may have specified a new size smaller than the current size, a new size of 0 with a zero secondary allocation on the data set, or there may not be sufficient space on the volumes. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, "Reason codes," on page 241](#) for a description of the reason code.

System action:

The program ends.

Administrator response:

Correct the size specification and ensure that there is space on the volume. If there are no other commands or operations running against the aggregate, and the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00178I **The configuration data set *Filename* could not be opened.**

Explanation

Since zFS could not open the configuration data set, it will continue without one and will use defaults for all configuration options.

In the message text:

Filename

The name of the configuration data set.

System action:

The program continues.

Administrator response:

Look for related error messages in the system log to determine why the open failed. Ensure that the zFS configuration data set is accessible to the program and retry the operation.

Severity:

svc_c_sev_notice

IOEZ00179I **There is no configuration dataset. Default configuration values will be used.**

Explanation:

This message notifies the user that default configuration values will be used.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00181E **Could not open trace output dataset.**

Explanation:

A zFS program could not print the in-memory trace table.

System action:

The program continues.

Administrator response:

Ensure that the data set that should contain the trace exists, is writable, and is not allocated to another job or user.

Severity:

svc_c_sev_error

IOEZ00182E **Could not write to trace output dataset.**

Explanation

A zFS program could not print the in-memory trace table.

System action:

The program continues.

Administrator response

Ensure that the data set that should contain the trace exists, is writable, and is not allocated to another job or user.

Severity:

svc_c_sev_error

IOEZ00183E **File system *FileSystem* is busy.**

Explanation:

A **zfsadm** command was issued against a file system that is currently the target of another **zfsadm** operation, or the file system is mounted and the **zfsadm** command is not valid for a mounted file system.

System action:

The program ends.

Administrator response:

The administrator can determine if the file system is mounted by issuing the **zfsadm lsfs** command or the **df** command. If the file system is busy due to another **zfsadm** command, then retry this command when the previous command completes.

Severity:

svc_c_sev_error

IOEZ00184E **Duplicate file system *FileSystem* found on aggregate *AggrName*, already attached on aggregate *PrevAggrName*.**

Explanation:

An aggregate was being attached and zFS detected that the aggregate has a file system defined that matches the name of a file system on another previously attached aggregate. The file system reported as a duplicate is not attached and is unavailable for processing. The original file system on the previously attached aggregate is available for processing.

System action:

The program continues.

Administrator response:

The administrator can correct the situation by renaming the file system on the aggregate that is currently attached, and then by re-attaching the aggregate that had the duplicate name.

Severity:

svc_c_sev_error

IOEZ00185E **DASD volser *Volser* is offline, aggregate *AggrName* cannot be attached.**

Explanation:

An aggregate was being attached and zFS detected that the aggregate is on a DASD volume that is offline.

System action:

The program continues.

Administrator response:

To improve performance and availability, ensure that the specified DASD volume is online and issue a remount to read-only and then remount back to read/write. Issue the remounts during a period of low activity to try to avoid disruption of the remount to read-only. You can also choose to do nothing and continue to run with z/OS UNIX function shipping from this system.

Severity:

svc_c_sev_error

IOEZ00186E **Error *ErrorCode* reason *ReasonCode* encountered obtaining status on volser *Volser* for aggregate *AggrName*, macro *MacroName*.**

Explanation:

An aggregate was being attached and zFS could not obtain information via the *MacroName* macro on the DASD volser shown in the message. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, "Reason codes," on page 241](#) for a description of the reason code.

System action:

The program continues.

Administrator response:

Check the z/OS documentation on the *MacroName* macro to determine the source of the error. If the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00187I **Access Registers AR *RegNumber* *Reg1* *Reg2* *Reg3* *Reg4*.**

Explanation:

A zFS kernel was driven for recovery for a z/OS UNIX process. This message shows the access registers at time of error.

System action:

The program continues.

Administrator response:

This message will be followed by additional messages indicating more about the problem, oftentimes recovery is not a problem.

Severity:

svc_c_sev_warning

IOEZ00188A **zFS kernel out of storage, total of *TotalBytesM* bytes used, size attempted was *Size*, class *class*, type *type*.**

Explanation:

The zFS kernel ran out of storage when attempting to obtain *Size* bytes. *TotalBytes* is the number of megabytes that has been obtained by the zFS kernel. *Class* and *Type* is information for the IBM service personnel.

System action:

The program ends.

Administrator response:

Try restarting zFS with smaller cache sizes. This message is followed by output showing the zFS kernel storage map. If the problem continues, contact your service representative.

Severity:

svc_c_sev_fatal

IOEZ00190E File system *FileSystem* is mounted, operation cannot be performed.

Explanation:

A **zfsadm** command determined that the file system was mounted. The file system cannot be deleted when it is still mounted.

System action:

The program ends.

Administrator response:

Unmount the file system before deleting it.

Severity:

svc_c_sev_error

IOEZ00191E Error deleting file system *FileSystem*, error code *ErrorCode* reason code *ReasonCode*.

Explanation:

An unexpected error occurred while attempting to delete the file system. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, "Reason codes,"](#) on page 241 for a description of the reason code.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00199E Aggregate *Name* cannot be detached because it is quiesced.

Explanation:

An aggregate could not be detached because it is quiesced.

System action:

The program ends.

Administrator response:

If there are operations in progress against the aggregate such as a grow or a backup of the aggregate, then this message is expected. Retry the detach after the grow or backup is complete. Additionally, the **zfsadm unquiesce** command will unquiesce the aggregate and allow it to be detached.

Severity:

svc_c_sev_error

IOEZ00200E Error *ErrorCode* occurred while closing dataset '*dataset_name*'.

Explanation:

An error occurred while the device driver was closing the named linear data set. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program continues.

Administrator response:

Verify that data set exists and the application can write to it.

Severity:

svc_c_sev_error

IOEZ00201E	Aggregate Name cannot be detached because it contains a file system which is mounted.
-------------------	--

Explanation:

An aggregate could not be detached because it contains a mounted file system.

System action:

The program ends.

Administrator response:

Unmount any file systems and retry the command. The **zfsadm lsfs** command can be used to determine which file systems are mounted.

Severity:

svc_c_sev_error

IOEZ00202E	Aggregate Name cannot be grown because it is attached read-only.
-------------------	---

Explanation:

An aggregate could not be grown because it is attached in read-only mode.

System action:

The program ends.

Administrator response:

Detach the aggregate and reattach the aggregate in read-write mode and rerun the command.

Severity:

svc_c_sev_error

IOEZ00207E	Aggregate Aggrname was not found.
-------------------	--

Explanation:

An attempt was made to attach or format an aggregate, but no VSAM linear data set could be found with the given name.

System action:

The program ends.

Administrator response:

Reissue the command specifying a valid zFS aggregate name.

Severity:

svc_c_sev_error

IOEZ00208E	The aggregate name Aggrname is too long. The maximum length for an aggregate name is length.
-------------------	---

Explanation:

The aggregate name specified is too long. Aggregate names cannot exceed the maximum length specified.

System action:

The program ends.

Administrator response:

Reissue the command specifying a valid zFS aggregate name.

Severity:

svc_c_sev_error

IOEZ00209E	Error ErrorCode (reason ReasonCode) occurred while attempting to get a list of all attached aggregates.
-------------------	--

Explanation:

An error occurred while retrieving an aggregate list. See Return codes (errno) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, "Reason codes," on page 241](#) for a description of the reason code.

System action:

The program ends.

Administrator response:

Try to resolve the error and reissue the command.

Severity:

svc_c_sev_error

IOEZ00210E	Error <i>ErrorCode</i> (reason <i>ReasonCode</i>) occurred while attempting to get status for aggregate <i>Aggrname</i>.
-------------------	---

Explanation:

An error occurred while getting status for the specified aggregate. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, “Reason codes,”](#) on page 241 for a description of the reason code.

System action:

The program ends.

Administrator response:

Try to resolve the error and reissue the command.

Severity:

svc_c_sev_error

IOEZ00211E	Unable to allocate <i>ByteCount</i> bytes of storage for a file system list.
-------------------	---

Explanation:

zFS was unable to allocate working storage for a file system list.

System action:

The program ends.

Administrator response:

Enter the command again specifying a larger virtual storage size.

Severity:

svc_c_sev_error

IOEZ00212E	Error <i>ErrorCode</i> (reason <i>ReasonCode</i>) occurred while attempting to get a list of file systems for the aggregate <i>Aggrname</i>.
-------------------	---

Explanation:

An error occurred while retrieving a file system list. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, “Reason codes,”](#) on page 241 for a description of the reason code.

System action:

The program ends.

Administrator response:

Try to resolve the error and reissue the command.

Severity:

svc_c_sev_error

IOEZ00213E	Error <i>ErrorCode</i> (reason <i>ReasonCode</i>) occurred while attempting to get status for file system <i>Filesystem</i> in aggregate <i>Aggrname</i>.
-------------------	--

Explanation:

An error occurred while getting status for the specified file system. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, “Reason codes,”](#) on page 241 for a description of the reason code.

System action:

The program ends.

Administrator response:

Try to resolve the error and reissue the command.

Severity:

svc_c_sev_error

IOEZ00214E

The file system name *Filesystem* is too long. The maximum length for a file system name is *length*.

Explanation:

The file system name specified is too long. File system names cannot exceed the maximum length specified.

System action:

The program ends.

Administrator response:

Reissue the command specifying a valid zFS file system name.

Severity:

svc_c_sev_error

IOEZ00229I

***ProgramName* Commands are:.**

Explanation:

When the **help** subcommand or **-help** parameter is the only argument supplied for a command, a list of the valid subcommands is returned.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_warning

IOEZ00230I

***ProgramName* Unknown topic '*SubCommand*'.**

Explanation:

The **-help** parameter is not a valid subcommand for this command.

System action:

The program ends.

Administrator response:

Enter the command again specifying a valid subcommand name.

Severity:

svc_c_sev_warning

IOEZ00231I

***ProgramName* Ambiguous topic '*SubCommand*'; use 'apropos' to list.**

Explanation:

The **-help** parameter is ambiguous for *SubCommand* because more than one subcommand or topic can match the string. Use *apropos* to list them.

System action:

The program ends.

Administrator response:

Enter the command again using an appropriate topic that is not abbreviated.

Severity:

svc_c_sev_warning

IOEZ00232I

***ProgramName*: Type '*ProgramName* -help' for help.**

Explanation:

The *ProgramName* command was entered without an argument; at least one is required.

System action:

The program ends.

Administrator response:

Enter the command again with one or more arguments.

Severity:

svc_c_sev_warning

IOEZ00233I

ProgramName: Type '*ProgramName* help' or '*ProgramName* help -topic <command_name...>' for help.

Explanation:

The *ProgramName* command suite was entered without the required command name.

System action:

The program ends.

Administrator response:

Enter the command suite again specifying an appropriate command name.

Severity:

svc_c_sev_warning

IOEZ00234I

UnkownAmbig operation '*SubCommand*'; type '*ProgramName* help' for list.

Explanation:

The subcommand *SubCommand* is either ambiguous or not valid. Commands cannot be abbreviated because multiple subcommands might match the abbreviation. Type *ProgramName* help for a list of valid subcommands.

System action:

The program ends.

Administrator response:

Enter the *ProgramName* command again specifying a valid, unabbreviated subcommand.

Severity:

svc_c_sev_warning

IOEZ00235I

UnkownAmbig operation '*SubCommand*'; type '*ProgramName* help' for list.

Explanation:

The subcommand *SubCommand* is either ambiguous or not valid. Commands cannot be abbreviated because multiple subcommands might match the abbreviation. Type *ProgramName* help for a list of valid subcommands.

System action:

The program ends.

Administrator response:

Enter the *ProgramName* command again specifying a valid, unabbreviated subcommand.

Severity:

svc_c_sev_warning

IOEZ00236I

ProgramName -help'.

Explanation:

The command was entered with a parameter that was ambiguous or not valid.

System action:

The program ends.

Administrator response:

Enter the command again using a valid parameter.

Severity:

svc_c_sev_warning

IOEZ00237I **'ProgramName SubCommand -help'.****Explanation:**

The parameter entered for the subcommand was ambiguous or not valid.

System action:

The program ends.

Administrator response:

Enter the subcommand with a valid parameter.

Severity:

svc_c_sev_warning

IOEZ00238I **ProgramName Too many arguments.****Explanation:**

Too many arguments were entered for the *ProgramName* command.

System action:

The program ends.

Administrator response:

Enter the command again specifying fewer arguments.

Severity:

svc_c_sev_error

IOEZ00239E **ProgramName Too many values after parameter *Parameter*.****Explanation:**

Too many values were entered for parameter *Parameter*. *Parameter* is a single-valued parameter and more than one value was provided on the command line.

System action:

The program ends.

Administrator response:

Enter the command with only one value for the *Parameter* parameter.

Severity:

svc_c_sev_error

IOEZ00240E **ProgramName Missing required parameter '*Parameter*'.****Explanation:**

A required parameter for this command was not entered.

System action:

The program ends.

Administrator response:

Enter the command specifying a value for the required parameter.

Severity:

svc_c_sev_error

IOEZ00241I **ProgramNameAmbigUnk switch '*Switch*'; type the following for detailed help.****Explanation:**

The switch name entered was either ambiguous, abbreviated, or an incorrect switch.

System action:

The program ends.

Administrator response:

Enter the command with a valid, unabbreviated switch name.

Severity:

svc_c_sev_warning

IOEZ00242I

The parameters "*Parameter1*" and "*Parameter2*" cannot be used together.

Explanation:

The parameters *Parameter1* and *Parameter2* are mutually exclusive, both cannot be specified on same command line.

System action:

The request fails.

Administrator response:

Enter the command with either *Parameter1* or *Parameter2*, but not both.

Severity:

svc_c_sev_warning

IOEZ00243I

Usage: *ProgramNameSubCommand*.

Explanation:

This message displays the command syntax when the -help parameter is entered.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00244E

StringType String is longer than the maximum length of *MaxLength*.

Explanation:

The specified string is too long.

System action:

The request fails.

Administrator response:

Correct the string in error and try the operation again.

Severity:

svc_c_sev_error

IOEZ00245E

More than *MaxValues* values are specified for parameter *Parm*. At most *MaxValues* values are allowed.

Explanation:

There are too many values specified.

System action:

The request fails.

Administrator response:

Reduce the number of values and retry.

Severity:

svc_c_sev_error

IOEZ00246E

The value *Value* of the *Parm* parameter is not numeric or is greater than 4294967295.

Explanation:

The specified value is not numeric or is too large. A numeric value less than 4294967295 is required.

System action:

The request fails.

Administrator response:

Correct the value in error and retry.

Severity:

svc_c_sev_error

IOEZ00247E

The *ParmList* parameters are mutually exclusive. Both *Parm1* and *Parm2* have been specified. Only one may be specified.

Explanation:

More than one mutually exclusive parameter has been specified.

System action:

The request fails.

Administrator response:

Ensure only one of the parameters is specified and retry.

Severity:

svc_c_sev_error

IOEZ00248I

VSAM linear dataset *Dataset* successfully created.

Explanation:

The data set was successfully created.

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00249E

There were problems creating VSAM linear dataset *Dataset*. Code = *ErrorCode*, reason = *ReasonCode*. More messages may follow.

Explanation:

The named data set could not be created. See Return codes (errno) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, "Reason codes,"](#) on page 241 for a description of the reason code.

System action:

The request fails.

Administrator response:

Correct any problems and retry.

Severity:

svc_c_sev_error

IOEZ00250E

The primary space specification *Primary* is invalid. A number greater than 0 is required.

Explanation:

The primary space specification is not valid.

System action:

The program continues.

Administrator response:

Specify a valid primary space.

Severity:

svc_c_sev_error

IOEZ00251E

Insufficient memory to call IDCAMS.

Explanation:

The memory required to build the parameter list for IDCAMS could not be obtained. This memory must reside below the 16M line.

System action:

The request fails.

Administrator response:

Ensure that sufficient memory is available, then try the request again.

Severity:

svc_c_sev_error

IOEZ00252E	Aggregate name '<i>aggrname</i>' contains invalid characters. Operation terminated.
-------------------	--

Explanation:

The aggregate name contains characters which are not allowed. Operation terminated.

System action:

The request fails.

Administrator response:

See the **ioeagfmt** command in the *z/OS File System Administration* for the list of characters allowed in an aggregate name and try the operation again.

Severity:

svc_c_sev_error

IOEZ00253E	File system name '<i>Filesystem</i>' contains invalid characters. Operation terminated.
-------------------	--

Explanation:

The file system name contains characters which are not allowed. The operation is terminated.

System action:

The request fails.

Administrator response

See the **zfsadm create** command in the *z/OS File System Administration* for the list of characters allowed in a file system name and try the operation again.

Severity:

svc_c_sev_error

IOEZ00300I	Successfully set <i>Option</i> to <i>Value</i>.
-------------------	--

Explanation:

The zFS kernel successfully changed the indicated configuration option.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00301E	Could not set <i>Option</i> to <i>Value</i>, error code <i>ErrorCode</i>, reason code <i>ReasonCode</i>.
-------------------	---

Explanation:

The zFS kernel encountered an error while attempting to set the configuration option. See Return codes (errno) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See Appendix A, "Reason codes," on page 241 for a description of the reason code.

System action:

The program continues.

Administrator response:

Consult the provided reason code. One possible cause is a value was larger than the zFS defined maximum or smaller than the zFS defined minimum for that option.

Severity:

svc_c_sev_error

IOEZ00303E **Error *ErrorCode* reason *ReasonCode* while attempting to extend the *CacheName* cache dataspace.**

Explanation:

The zFS kernel could not extend the specified log file cache dataspace and thus could not attach an aggregate.

System action:

The program continues

Administrator response:

Consult *z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN* for a description of the DSPSERV macros and reason codes.

Severity:

svc_c_sev_error

IOEZ00304E **Error *ErrorCode* reason *ReasonCode* while attempting to create the *CacheName* cache dataspace.**

Explanation:

The zFS kernel could not create the indicated cache dataspace and thus could not initialize.

System action:

The program ends.

Administrator response:

Consult the *z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN* for a description of the DSPSERV macros and reason codes.

Severity:

svc_c_sev_error

IOEZ00308E **Aggregate *aggrname* failed dynamic grow, (by user *userid*).**

Explanation

Dynamic growth of the specified aggregate, caused by the actions of the specified user, failed. All subsequent attempts to dynamically grow the aggregate will also fail. This condition can be changed by either doing a **zfsadm grow** command, or by detaching and re-attaching the aggregate.

It is possible to see multiple copies of this message. For additional information, see [Dynamically growing a compatibility mode aggregate in z/OS File System Administration](#).

System action:

The program continues.

Administrator response:

Determine if the failure code represents a problem. If so, and the problem persists, contact the service representative.

Severity:

svc_c_sev_error

IOEZ00309I **Aggregate *aggrname* successfully dynamically grown (by user *username*).**

Explanation

Dynamic growth of the specified aggregate, caused by the actions of the specified user, was successful. It is possible to see multiple copies of this message. For additional information, see [Dynamically growing a compatibility mode aggregate in z/OS File System Administration](#).

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00312I	Dynamic growth of aggregate <i>aggrname</i> in progress, (by user <i>username</i>).
-------------------	--

Explanation

Dynamic growth of the specified aggregate, caused by the actions of the specified user, has been initiated. A subsequent message will be issued to indicate the success or failure of this operation.

It is possible to see multiple copies of this message. For additional information, see [Dynamically growing a compatibility mode aggregate in z/OS File System Administration](#).

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00314E	The file system name <i>filesystem</i> is not unique. Its aggregate name must be specified.
-------------------	--

Explanation:

An aggregate name must be specified to uniquely identify the file system.

System action:

The program ends.

Administrator response:

Ensure that the aggregate name is specified and retry.

Severity:

svc_c_sev_error

IOEZ00315I	There are now <i>Number</i> file systems with the name <i>Filesystem</i> in use.
-------------------	---

Explanation:

This message is issued for informational purposes only. There are zFS file systems in different aggregates that have the same name. As of z/OS V1R9, the `allow_duplicate_filesystems` configuration option is removed and always acts like it is on (that is, duplicate zFS file system names are allowed in different aggregates).

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00317I	The value for configuration option <i>ConfigOption</i> is <i>ConfigValue</i>.
-------------------	--

Explanation:

The named configuration option has the indicated value.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00318I

Attempting to add file system *filesystem* whose name is already in use.

Explanation:

The file system has the same name as one already in use.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00320I

No options have been specified.

Explanation:

A **zfsadm** command option has not been specified.

System action:

The program ends.

Administrator response:

Specify a valid option for the command and retry.

Severity:

svc_c_sev_notice

IOEZ00321E

Could not get value for *Option*, error code *ErrorCode* reason code *ReasonCode*.

Explanation:

The zfs kernel encountered an error while attempting to get the value of the configuration option. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, "Reason codes," on page 241](#) for a description of the reason code.

System action:

The program continues.

Administrator response:

Correct any problems and retry.

Severity:

svc_c_sev_error

IOEZ00322E

The file system name *filesystemname* contains invalid characters and cannot be used.

Explanation:

The named file system name cannot be attached because its name contains invalid characters. The name may not actually show invalid characters because only the first 44 characters are printed.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_warning

IOEZ00323I**Attempting to extend *aggregate_name* to *new_size* control_interval_size byte control intervals.****Explanation**

The aggregate is being extended to the indicated size.

In the message text:

aggregate_name

Name of the aggregate.

control_interval_size

Control interval size of the aggregate.

new_size

New size of the data set.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00324I**Formatting to *control_interval_size* byte control interval *new_size* for secondary extents of *aggregate_name*.****Explanation**

The aggregate has been extended into one or more secondary extents and one or more of those secondary extents is now being formatted. The data set will be the specified size when complete.

In the message text:

aggregate_name

Name of the aggregate.

control_interval_size

Control interval size for the data set.

new_size

New control interval size of the data set.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00324I**Formatting to 8K block number *Size* for secondary extents of *dataset_name*.****Explanation:**

The aggregate has been extended into one or more secondary extents and one or more of those secondary extents is now being formatted. The data set will be the specified size when complete.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00325E

Error Error formatting secondary extents for *dataset_name*.

Explanation:

The attempt to format one or more secondary extents for the aggregate has failed. The z/OS UNIX System Services return code is shown. See [Return codes \(errno\)](#) in [z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program continues.

Administrator response:

Verify that the data set exists and that the application can write to it.

Severity:

svc_c_sev_error

IOEZ00326E

Error Error extending *dataset_name*.

Explanation:

The attempt to extend the aggregate has failed. The z/OS UNIX System Services return code is shown. See [Return codes \(errno\)](#) in [z/OS UNIX System Services Messages and Codes](#) for a description of the return code.

System action:

The program continues.

Administrator response:

Verify that the data set exists and that the application can write to it. Ensure that there is enough DASD space on the data sets associated DASD volumes for the data set to be grown. Examine the syslog for any IEC070I or IOEZ00445E messages that might indicate the reason for the error.

Severity:

svc_c_sev_error

IOEZ00327I

Done. *dataset_name* is now a zFS aggregate.

Explanation:

zFS has finished formatting an aggregate.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00328E

Errors encountered making *dataset_name* a zFS aggregate.

Explanation:

Errors were encountered while attempting to format a zFS aggregate.

System action:

The program continues.

Administrator response:

This message should be preceded by other messages that describe the errors encountered.

Severity:

svc_c_sev_error

IOEZ00329I

Attempting to extend *dataset_name* by a secondary extent.

Explanation:

The aggregate is being extended by a secondary extent.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00331A

File system name *PFSName* is invalid, must be ZFS.

Explanation:

An incorrect name for the z/OS File System was specified in the BPXPRM file FILESYSTYPE statement. This name must be the characters ZFS.

System action:

The program ends.

Administrator response:

Update the FILESYSTYPE statement and ensure that the TYPE parameter specifies ZFS, example: TYPE (ZFS).

Severity:

svc_c_sev_error

IOEZ00334I

Return code and reason code for dump is *RSNcodeRC*.

Explanation

This message shows the return and reason code returned by SVC dump processing. The return code from SDUMPX processing is in bits 24 - 31 (the last byte). If the return code is 8, the reason code is contained in bits 16 - 23 (the third byte). See [z/OS MVS Programming: Authorized Assembler Services Reference LLA-SDU](#) for additional information.

In the message text:

RSNcodeRC

The return and reason code of SDUMPX processing.

Note: In some circumstances, a dump can be suppressed by the operating system.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00336I

***AggrName* could not be marked as a zFS aggregate in the catalog,
rc=*ReturnCode* rsn=*ReasonCode*.**

Explanation

After successfully attaching, or formatting a zFS aggregate, a call to the catalog service marks *AggrName* as a zFS aggregate. This operation failed. The return and reason codes are from the catalog service. This failure itself does not prevent the aggregate from being attached or formatted correctly. However, other system processing that depends on the catalog entry may not function properly. DFSMS uses this bit during backup processing (ADRDSSU dump/restore) and enables the backup utility to quiesce the aggregate automatically before the backup starts, and unquiesce the aggregate when the backup ends. If the catalog bit is not set, the automatic quiesce will not occur and backup processing will fail.

In the message text:

AggrName

The name of the aggregate.

ReturnCode

The catalog service return code.

ReasonCode

The catalog service reason code.

System action:

The program continues.

Administrator response:

Find the specified catalog return and reason codes documented in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#) (message IDC3009I), and determine the cause of the problem. This message can be the result of mounting a zFS file system using the VSAM PATH entry. Do not use a path entry as the file system name in the MOUNT command (see the topic about DEFINE PATH in z/OS DFSMS Access Method Services Commands). After correcting the problem, to get this aggregate marked as a zFS aggregate in the catalog, unmount the file system and then mount it again. If the catalog bit is not set during the format step, the zFS task attempts to set the catalog bit during the mount processing for the zFS file system.

Severity:

svc_c_sev_notice

IOEZ00337E

zFS kernel: non-terminating exception *AbendCode* occurred, reason *ReasonCode*abend psw *PSW1 PSW2*.

Explanation

A zFS kernel encountered an exception. A dump will be issued and the internal trace table will be printed, if possible. This exception is non-terminating; the zFS kernel will continue to run, though there may be errors encountered for a given file or file system. The abend code is *AbendCode*, and the psw is given by *PSW1* and *PSW2*. See [z/OS MVS System Codes](#) for additional information about abend codes.

In the message text:

AbendCode

The z/OS abend code received.

ReasonCode

The z/OS reason code received, or register 15 at time of abend.

PSW1

The first word of the abend PSW.

PSW2

The second word of the abend PSW.

System action:

The program continues.

Administrator response:

This message will be accompanied by additional messages indicating more about the problem. Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00338A

zFS kernel: restarting exception *AbendCode* occurred, reason *ReasonCode*abend psw *PSW1 PSW2*.

Explanation

The zFS kernel encountered an exception. A dump will be issued and the internal trace table will be printed. The zFS kernel will be internally stopped and restarted. File systems will remain mounted but some programs with open files might see I/O errors. If severe exceptions are encountered during the restart, zFS will be terminated and the system will display a BPXF032D message that must be replied to before zFS can be reinitialized. See [z/OS MVS System Codes](#) for more information about abend codes.

In the message text:

AbendCode

The z/OS abend code received.

ReasonCode

The zFS abend reason code received, or register 15 at time of abend.

PSW1

The first word of the abend PSW.

PSW2

The second word of the abend PSW.

System action

The program (IOEFSCM) continues

Administrator response:

This message will be followed by additional messages indicating more about the problem. Contact your service representative.

Severity:

svc_c_sev_fatal

IOEZ00340E**Potential zFS hang detected. Taking informational dump...**

Explanation:

The zFS Physical File System detected at least one user task that has a request that has not been satisfied for approximately three minutes. zFS considers this to be a potential hang and dumps the zFS address space for diagnosis purposes.

System action:

The program continues.

Administrator response:

Contact your service representative and supply the dump.

Severity:

svc_c_sev_error

IOEZ00341E**zfsadm query *OptionName* failed, rc = *ReturnCode*, reason = *ReasonCode*.**

Explanation

The indicated **zfsadm query** command failed. The return and reason codes may indicate the cause of the error. If possible, correct the error and retry. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, "Reason codes," on page 241](#) for a description of the reason code.

In the message text:

OptionName

The query option name.

ReturnCode

The z/OS UNIX System Services error code received.

ReasonCode

The zFS reason code.

System action:

The program continues.

Administrator response:

Correct any errors and retry the command.

Severity:

svc_c_sev_error

IOEZ00342I **Aggregate is attached. Cannot format.****Explanation:**

The aggregate is attached by zFS on another system. It cannot be formatted.

System action:

The program continues.

Administrator response:

Detach the aggregate before attempting to format.

Severity:

svc_c_sev_notice

IOEZ00350I **Successfully joined group *GroupName*.****Explanation**

zFS is initializing its sysplex support and has successfully joined the named XCF group. In the message text:

GroupName

The XCF group name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00351E **Invalid message, command code *code* received from system *SysName*.****Explanation**

An incorrect request packet was received from the remote zFS system. In the message text:

Code

The internal zFS command code.

SysName

The remote system name.

System action:

The program continues.

Administrator response:

If the problem persists, contact IBM service.

Severity:

svc_c_sev_error

IOEZ00353E **Error *ErrorCode* reason *ReasonCode* received from system *SysName* during *OpName* for *AggrName_FileSysName*.****Explanation**

An unexpected error was received while attempting to notify a remote system of the operation that was being performed. The specified operation was being performed against the specified aggregate or file system. See Return codes (errno) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See Appendix A, "Reason codes," on page 241 for a description of the reason code.

In the message text:

ErrorCode

The z/OS UNIX System Services error code received.

ReasonCode

The zFS reason code.

SysName

The remote system name.

OpName

The operation that was being performed.

AggrName **FileSysName**

The name of the aggregate or file system being operated on.

System action:

The program continues.

Administrator response:

If the problem persists, contact IBM service.

Severity:

svc_c_sev_error

IOEZ00354E

Sysplex request *ServiceRequested* failed with return code *ReturnCode* and reason code *ReasonCode*.

Explanation

zFS was attempting to use sysplex services and received a failure from the specified IXC macro. It failed with the specified return and reason codes.

In the message text:

ServiceRequested

The IXC macro name.

ReturnCode

The IXC macro return code.

ReasonCode

The IXC macro reason code.

System action

The request fails and zFS continues operating normally.

Severity

```
svc c sev error
```

Administrator Response

Determine the cause of the failure and re-start zFS, if necessary. See *z/OS MVS Programming: Sysplex Services Reference* for additional information about IXC return and reason codes.

IOEZ00356E

Failed joining group *GroupName* as member *MemberName*, return code *ReturnCode* reason code *ReasonCode*.

Explanation

zFS was issuing the sysplex service IXCJOIN and received the failure detailed in the message. It failed with the specified return and reason codes.

In the message text:

GroupName

The sysplex group name.

MemberName

The system member name.

ReturnCode

The IXC macro return code.

ReasonCode

The IXC macro reason code.

System action

The program ends.

Severity

svc_c_sev_error

Administrator Response

Determine the cause of the failure and restart zFS, if necessary. You might have an incorrect XCF group name in the IOEFSPRM group option. See *z/OS MVS Programming: Sysplex Services Reference* for additional information about IXC return and reason codes.

IOEZ00357I**Successfully left the sysplex group.****Explanation**

zFS was issuing the sysplex service IXCLEAVE to leave the sysplex group. The call was successful and this system will no longer process sysplex requests.

System action

The program continues.

Administrator response:

None.

Severity

svc_c_sev_notice

IOEZ00358E**The attempt to create the *DaemonName* daemon failed with return code *ReturnCode*.****Explanation**

The creation of a thread for the daemon failed. See *z/OS MVS Programming: Authorized Assembler Services Reference ALE-DYN* for additional information about the ATTACH macro and its return codes. In the message text:

DaemonName

The name of the daemon.

ReturnCode

The MVS attach return code.

System action:

The program ends.

Severity:

svc_c_sev_notice

Administrator Response:

Determine why the thread could not be created by using the return code from the ATTACH macro. Correct the problem and re-start zFS. If the problem persists, contact your service representative.

IOEZ00359E **Error *ErrorCode* (reason *ReasonCode*) occurred while attempting to get a list of all systems in the XCF group for zFS.**

Explanation

An error occurred while retrieving the system names list. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, “Reason codes,” on page 241](#) for a description of the reason code.

In the message text:

ErrorCode

The z/OS UNIX System Services error code.

ReasonCode

The zFS reason code.

Explanation:

Try to resolve the error and reissue the command.

System action:

The program ends.

Severity:

svc_c_sev_error

IOEZ00360I **There are no systems in the XCF group for zFS.**

Explanation:

There are no systems currently in the XCF group for zFS.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00361I **A total of *count* systems are in the XCF group for zFS.**

Explanation

The message shows the total number of systems in the XCF group for zFS. It is followed by a list of the system names.

In the message text:

count

The number of systems currently in the XCF group for zFS.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00362E **The system name *SysName* is too long.**

Explanation

A system name must be no longer than 8 in length.

Determine the cause of the failure and re-start zFS, if necessary. See [z/OS MVS Programming: Sysplex Services Reference](#) for additional information about IXC return and reason codes.

Explanation

An error occurred while attempting to allocate memory needed to process the command.

In the message text:

NumBytes

The number of bytes.

System action:

The request fails, and zFS continues operating normally.

Severity:

svc_c_sev_error

Administrator Response:

Ensure that sufficient memory is available, then try the request again.

IOEZ00373E	Error <i>ReturnCode</i>, reason <i>ReasonCode</i> received opening parmlib.
-------------------	--

Explanation

An error was received while attempting to open the parmlib. The return code from the IEFPRMLB macro is shown. See [z/OS MVS Programming: Authorized Assembler Services Reference EDT-IXG](#) for information about the IEFPRMLIB macro and its reason codes.

In the message text:

ReturnCode

The return code received from the IEFPRMLB macro.

ReasonCode

The reason code received from the IEFPRMLB macro.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Ensure that the parmlib is accessible to the program and try the operation again.

IOEZ00374I	No IOEZPRM DD specified in <i>name</i> proc. Parmlib search being used.
-------------------	--

Explanation:

This message notifies the user that parmlib search is being used. *name* is the name of the address space where zFS is running.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00380E	Specified vnode cache size <i>Config_vnode_size</i> is invalid. Using value <i>vnode_size</i>.
-------------------	---

Explanation

During initialization, zFS found that the value specified in the configuration data set for **vnode_cache_size** is not valid. The value must be a number in the range of 32 to 500000. zFS initialization continues.

In the message text:

Config_vnode_size

The configuration data set specification for the vnode size.

vnode_size

The value being used for the size.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Ensure that zFS is running with acceptable values for both the size of the vnode cache and the limit of the vnode cache. If not, update the configuration data set variables.

IOEZ00381E

Error *ErrorCode* reason *ReasonCode* received while attempting to move aggregate *AggrName* to system *System*.

Explanation

An unexpected error was received while attempting to move aggregate *AggrName* to system *System*.

In the message text:

ErrorCode

The error code received.

ReasonCode

The reason code received.

AggrName

The aggregate name.

System

The system name.

System action:

The program ends.

Severity:

svc_c_sev_error

Administrator Response:

Try to resolve the error and reissue the command.

IOEZ00382I

Aggregate *AggrName* is now owned by system *System*.

Explanation

Aggregate *AggrName* was successfully moved to *System*.

In the message text:

AggrName

The aggregate name.

System

The system name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00383E

Unrecoverable error encountered for inode *Inode* uniquifier *Uniquifier* in filesystem *FileSetName*.

Explanation

A fatal error was encountered for an object with inode *Inode* in filesystem *FileSetName*. Future accesses to the file system object may fail and updates may be lost.

In the message text:

Inode

The inode of the object in error.

Uniquifier

The uniquifier of the object in error.

FileSetName

File system that contains the object.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00385E

Sysplex sharing error encountered for aggregate *AggrName*.

Explanation

An internal error was encountered for aggregate *AggrName*. Recent updates might be lost. Future accesses to the aggregate might fail. Unmount and remount the aggregate. If that does not clear the problem, restart zFS. If the problem still exists, restart zFS on all the systems in the sysplex.

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00387I

System *SysName* has left group *GroupName*, aggregate recovery in progress.

Explanation

A zFS system has detected that another system in the sysplex has gone down. The remaining systems in the sysplex will begin aggregate recovery to attach aggregates owned by the down system.

In the message text:

SysName

The system name.

GroupName

The sysplex group name.

System action:

The program continues.

Administrator response:

None.

Severity:

IOEZ00388I**Aggregate takeover being attempted for aggregate *AggrName*.****Explanation**

A zFS system has either been asked to assume ownership for an aggregate or it has detected that another system in the sysplex went down and is attempting to takeover ownership of the aggregate.

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00389I**Error *ErrorCode* reason *ReasonCode* occurred while attempting takeover of *AggrName*.****Explanation**

A zFS system has encountered an error while attempting to assume ownership of an aggregate. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, “Reason codes,” on page 241](#) for a description of the reason code.

In the message text:

ErrorCode

The z/OS UNIX System Services error code received.

ReasonCode

The zFS reason code.

AggrName

The aggregate name.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_notice

IOEZ00390I**recovery statistics:.****Explanation**

This message notifies the user that an aggregate underwent recovery. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS.

System action

The program continues.

Administrator response

None. See [z/OS File System Administration](#) for additional information about log files.

Severity

svc_c_sev_notice

IOEZ00391I

Elapsed time was *time* ms.

Explanation

This message indicates the amount of time aggregate recovery took. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS.

In the message text:

time

The elapsed time in milliseconds.

System action

The program continues.

Administrator response

None. See [z/OS File System Administration](#) for additional information about log files.

Severity

svc_c_sev_notice

IOEZ00392I

***NumPages* log pages recovered consisting of *NumRecords* records.**

Explanation

This message indicates the number of log pages and records read while recovering a zFS aggregate. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS.

In the message text:

NumPages

The number of pages.

NumRecords

The number of records.

System action

The program continues.

Administrator response

None. See [z/OS File System Administration](#) for additional information about log files.

Severity

svc_c_sev_notice

IOEZ00393I

Modified *NumBlocks* data blocks.

Explanation

This message indicates the number of metadata blocks on disk updated by the recovery process. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS. In the message text:

NumBlocks

The number of data blocks.

System action

The program continues.

Administrator response

None. See [z/OS File System Administration](#) for additional information about log files.

Severity

svc_c_sev_notice

IOEZ00394I

***NumDataRecords* redo-data records, *NumFillRecords* redo-fill records.**

Explanation

This message indicates the number of transaction records committed to disk. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS.

In the message text:

NumDataRecords

The number of redo data records.

NumFillRecords

The number of redo fill records.

System action

The program continues.

Administrator response

None. See [z/OS File System Administration](#) for additional information about log files.

Severity

svc_c_sev_notice

IOEZ00395I

***NumDataRecords* undo-data records, *NumFillRecords* undo-fill records.**

Explanation

This message indicates the number of transaction records rolled back. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS.

In the message text:

NumDataRecords

The number of undo data records.

NumFillRecords

The number of undo fill records.

System action

The program continues.

Administrator response

None. See [z/OS File System Administration](#) for additional information about log files.

Severity

svc_c_sev_notice

IOEZ00396I***NumBlocks* not written blocks.**

Explanation

This message indicates the number of transaction records not committed because the block was reused to contain user-file data. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS.

In the message text:

NumBlocks

The number of blocks.

System action

The program continues.

Administrator response

None. See [z/OS File System Administration](#) for additional information about log files.

Severity

svc_c_sev_notice

IOEZ00397I**recovery statistics for *AggrName*:**

Explanation

This message notifies the user that an aggregate underwent recovery. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS.

In the message text:

AggrName

The aggregate name.

System action

The program continues.

Administrator response

None. See [z/OS File System Administration](#) for additional information about log files.

Severity

svc_c_sev_notice

IOEZ00398E**Specified vnode cache size *Config_vnode_size* is greater than the vnode cache limit. Using value *vnode_size*.**

Explanation

During initialization, zFS found that the value specified in the configuration data set for *vnnode_cache_size* is greater than the value for maximum size allowed for the vnode cache. The value must be a number in the range of 32 to 500000. zFS initialization continues.

In the message text:

Config_vnode_size

The configuration data set specification for the vnode size.

vnnode_size

The value being used for the size.

System action

The program continues.

Severity

svc_c_sev_error

Administrator Response

Ensure that zFS is running with an acceptable size for the vnode cache size. If not, update the configuration dataset variable *vnnode_cache_size*.

IOEZ00400I***NumBlocks* blocks zeroed.**

Explanation

This message indicates the number of disk blocks zeroed out because the block was used to contain user-file data, but the I/O had not yet completed, or was canceled. This was done because the new block security (NBS) configuration option was in effect for the aggregate. This is normal if the aggregate was not detached cleanly before a stop or an outage of zFS.

In the message text:

NumBlocks

The number of blocks.

System action

The program continues.

Administrator response:

None.

Severity

svc_c_sev_notice

IOEZ00401I**No report option specified to be reset.**

Explanation:

A **zfsadm query** command was entered, with the **-reset** specified. However, no other options were specified along with it. No statistics were reset.

System action

The program ends.

Severity

svc_c_sev_notice

Administrator Response

Specify a valid statistic option for the command and retry.

IOEZ00405I	<i>ProgramName</i> Option <i>Option</i> has been processed. This specification is ignored.
-------------------	---

Explanation

The specified option was already processed. The option was either explicitly specified previously or implicitly assumed according to the option omission rules. We are ignoring this specification.

In the message text:

ProgramName

The program name.

Option

The option.

System action:

The program ends.

Severity:

svc_c_sev_warning

Administrator Response:

Enter the command with only one value for the *option* option.

IOEZ00410I	Shareoptions for aggregate <i>AggrName</i> altered. New value is (3,3).
-------------------	--

Explanation

The shareoptions for the indicated aggregate were altered.

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00411I	Alter for shareoptions failed for aggregate <i>AggrName</i>.
-------------------	---

Explanation

The shareoptions for the indicated aggregate could not be altered to (3,3).

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Severity:

svc_c_sev_warning

Administrator Response:

The trace will contain IDCAMS messages. Correct the problem and retry.

IOEZ00412I

Catalog search failed for aggregate *AggrName*. Shareoptions are not altered.

Explanation

The shareoptions for the indicated aggregate could not be determined; they are not altered.

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_warning

IOEZ00413I

IGGCSI00 could not be loaded. Shareoptions for aggregate *AggrName* are not altered.

Explanation

The shareoptions for the indicated aggregate could not be determined because the catalog search routine could not be loaded. The shareoptions are not altered.

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Severity:

svc_c_sev_warning

Administrator Response:

Ensure IGGCSI00 is installed and retry.

IOEZ00416I

Aggregate *AggrName* moved to system *SystemName* at shutdown.

Explanation

The zFS kernel owning the specified aggregate moved the aggregate to another system to allow shutdown to continue.

In the message text:

AggrName

The aggregate name.

SystemName

The system name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

Explanation

The zFS kernel owning the specified aggregate tried to move the specified aggregate to the specified system but failed. Shutdown will continue by attempting to move the aggregate to a different system name. If all systems fail, then the aggregate will be detached, but some file updates may be lost. See Return codes (errno) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, “Reason codes,”](#) on page 241 for a description of the reason code.

In the message text:

ErrorCode

The z/OS UNIX System Services error code received.

ReasonCode

The zFS reason code received.

AggrName

The aggregate name.

SystemName

The system name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_warning

Explanation

The zFS kernel encountered a problem associating the specified mount name to the indicated file system. The only affect this will have is that mount name options of the various zFS administrative commands will not work correctly. This has no bearing on whether the file system mount succeeded. Check for other messages indicating the success or failure of the mount process.

In the message text:

mountname

The mount name.

sysname

The system name.

System action:

The program continues.

Administrator response:

Do not use the mount name options on the administrative commands that normally allow them to be used. Instead, full specifications in terms of file system name and aggregate name will need to be used.

Severity:

svc_c_sev_notice

Explanation

A syntax error was found while parsing the parmlib member containing the zFS configuration parameters. The line in error is shown as an aid to the administrator.

In the message text:

String

The text in error.

StartLine

The starting line number of the statement in error.

EndLine

The ending line number of the statement in error.

StartParmlibSuffix

The first parmlib member suffix.

EndParmlibSuffix

The last parmlib member suffix.

System action:

The program continues.

Administrator response:

Correct the line in error and try the operation again.

Severity:

svc_c_sev_error

IOEZ00421E

Insufficient memory to allocate *NumBytes* bytes.

Explanation

An error occurred while attempting to allocate memory needed to process the command.

In the message text:

NumBytes

The number of bytes.

System action

The request fails, and zFS continues operating normally.

Administrator response:

Allocate more storage and retry the request.

Severity:

svc_c_sev_error

IOEZ00422E

Aggregate *Aggrname* disabled.

Explanation

An error was found while processing requests for the specified aggregate. As a result, all reads and writes to the aggregate are prevented. In the message text:

AggrName

The aggregate name.

System action

zFS continues operating normally for all other aggregates; the specified aggregate is disabled.

Administrator response

If there were reported IO errors, perhaps due to lost connectivity to DASD or device failures, ensure that the hardware is working properly and then unmount the filesystem and run IOEAGSLV to repair any corruption that may have occurred. IOEAGSLV will attempt to repair the aggregate but it may not always be successful. Then, try reattaching or remounting the filesystem. If a remount does not work, zFS must be stopped and

restarted to allow for the aggregate to be attached in write mode. If this condition persists, contact your service representative.

If there were no reported IO errors, there is less probability of a corruption. You may allow zFS to attempt to automatically reenable the aggregate; however, if the disabled aggregate was corrupt, the reenabled aggregate will remain corrupt. To definitively determine if the aggregate is corrupt, follow the procedure previously stated for the IO error case and run IOEAGSLV. While IOEAGSLV is running, the aggregate will be offline. Larger aggregates, will take longer to inspect and repair and, hence, will remain offline longer.

For more information, see the topic about [Diagnosing disabled aggregates in z/OS File System Administration](#).

Severity:

svc_c_sev_error

IOEZ00424E	Internal lock error, aggregate <i>AggrName</i> must be detached and the format restarted.
-------------------	--

Explanation

An internal lock could not be obtained to complete the request to format the specified aggregate. The aggregate must be detached and the format restarted.

In the message text:

AggrName

The aggregate name.

System action:

zFS continues operating normally for all other aggregates. The specified aggregate cannot be formatted.

Administrator response:

Detach the aggregate and attempt to format again.

Severity:

svc_c_sev_error

IOEZ00425E	UNQUIESCE FAILURE: rc = <i>ReturnCode</i> rsn = <i>ReasonCode</i>.
-------------------	---

Explanation

An error occurred while attempting a console command aggregate unquiesce. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the error code. See [Appendix A, "Reason codes," on page 241](#) for a description of the reason code. In the message text:

ReturnCode

The return code.

ReasonCode

The reason code.

System action:

zFS continues operating normally for all other aggregates, the specified aggregate remains quiesced.

Administrator response:

Check the return code (rc) and the reason code (rsn) and respond accordingly. If the return code is 129 (ENOENT) and the reason code is EFxx6775 (the xx can be anything), then the aggregate is not owned on this system. Issue the MODIFY UNQUIESCE from the owning system. Use the z/OS UNIX **zfsadm lsaggr** command to determine which system owns the aggregate.

Severity:

svc_c_sev_error

IOEZ00426E	UNQUIESCE FAILURE: too many input parameters.
-------------------	--

Explanation:

The user input an extraneous token when they typed the aggregate unquiesce command.

System action:

zFS continues operating normally for all other aggregates, the specified aggregate remains quiesced.

Administrator response:

Check the syntax and enter the command again.

Severity:

svc_c_sev_error

IOEZ00433E

Internal error, aggregate *AggrName* cannot be detached from zFS.

Explanation

An internal error was found while processing a request to detach the specified aggregate. The detach problem could not be internally fixed. The aggregate can no longer be detached and hence is unusable until zFS is restarted.

In the message text:

AggrName

The aggregate name.

System action:

zFS continues operating normally for all other aggregates, the specified aggregate can no longer be detached or attached.

Administrator response:

zFS will need to be restarted to clear this error condition. A dump will likely be obtained by zFS, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00434E

Error creating zFS kernel trace, error code=*ErrorCode* reason code=*ReasonCode*.

Explanation

An unexpected error occurred creating zFS kernel trace. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, “Reason codes,” on page 241](#) for a description of the reason code.

In the message text:

ErrorCode

The error code that was returned from the zFS kernel.

ReasonCode

The reason code that was returned from the zFS kernel.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00435E

Error aborting the zFS kernel, error code=*ErrorCode* reason code=*ReasonCode*.

Explanation

An unexpected error occurred aborting the zFS kernel. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, “Reason codes,” on page 241](#) for a description of the reason code.

In the message text:

ErrorCode

The error code that was returned from the zFS kernel.

ReasonCode

The reason code that was returned from the zFS kernel.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00436E

Disabling IO to DASD volume *VolumeName*.

Explanation

An internal error has forced zFS to temporarily disable IO to a particular DASD volume. All queued IO is stopped with error and IO to the DASD volume remains suppressed until in-progress IO to the DASD volume completes. This message will normally be preceded by a dump to provide necessary service information. Once all in-progress IO completes, the DASD IO queue is re-initialized and future IOs will be allowed from that point in time.

In the message text:

VolumeName

The name of the DASD volume.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00437I

IO to DASD volume *VolumeName* has been enabled.

Explanation

An internal error had forced zFS to temporarily disable IO to a particular DASD volume; the internal error state has been cleaned up and IO to the DASD volume is now allowed.

In the message text:

VolumeName

The name of the DASD volume.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00438I

Starting Query Command *QueryName*.

Explanation

Starting to run a query command.

In the message text:

QueryName

The name of the query being run.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00439E

Read-only aggregate *aggregate_name* is also mounted read/write on system *system_name*.

Explanation

The aggregate is either mounted or attached read-only on a local system and is also mounted or attached in read/write mode on another system. The local system might encounter errors on subsequent reads.

In the message text:

aggregate_name

The aggregate name.

system_name

The read/write system name.

System action:

The program continues.

Administrator response:

Unmount or detach the aggregate from all systems and then remount or reattach it in the desired mode. If you unmount read/write aggregate first, you might get error messages on the read-only systems.

Severity:

svc_c_sev_warning

IOEZ00440E

Internal error found while processing read-only aggregate *AggrName*.

Explanation

An internal error was found while processing requests for the specified aggregate. The aggregate may be attached read-write on another system. A changed aggregate will most likely result in errors on the read-only system.

In the message text:

AggrName

The aggregate name.

System action:

zFS continues operating normally. Other errors may be encountered for this aggregate if it is not unmounted and remounted.

Administrator response:

Try quiescing the aggregate on the read-write system and then unmounting and remounting the file system.

Severity:

svc_c_sev_error

IOEZ00441E

Operation failed for aggregate *AggrName*. ZFS does not support striped data sets.

Explanation

The Aggregate name could not be formatted.

Explanation

The attempt to extend the aggregate has failed. The DFSMS return code and PDF code are shown. Refer to message IEC070I and IEC161I in [z/OS MVS System Messages, Vol 7 \(IEB-IEE\)](#) for additional information.

In the message text:

AggrName

The aggregate name.

DFSMScode

The DFSMS return code.

PDFcode

The Problem Determination Function (PDF) return code.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response

Verify that the data set exists and that the application can write to it and ensure that there is enough DASD space on the data sets associated with the DASD volumes.

If the aggregate was not defined with extended addressability, also ensure that the size of the aggregate does not exceed the maximum allowed size of 4GB. For other errors, check the system log for any IEC messages for this aggregate that contain the DFSMS and PDF codes specified in the message. If any are found, follow the diagnostic procedures listed for these codes in the IEC messages.

IOEZ00447E	System <i>SystemName</i> went down or did not reply to <i>OperationName</i> for <i>Name</i>.
-------------------	---

Explanation

An unexpected error was received while attempting to notify a remote system of operation *OperationName*. The operation was being done for the aggregate or file system *Name*.

In the message text:

SystemName

The remote system name.

OperationName

The operation that was being performed.

Name

The name of the aggregate or file system name being operated on.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00451E	Error in allow striped, error code=<i>ErrorCode</i> reason code=<i>ReasonCode</i>.
-------------------	---

Explanation

An unexpected error occurred allowing use of striped data sets. See [Return codes \(errnos\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, "Reason codes," on page 241](#) for a description of the reason code.

In the message text:

ErrorCode

The error code that was returned from the zFS kernel.

ReasonCode

The reason code returned from the zFS kernel.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00453I

Processing xcf trace message from system *SystemName*.

Explanation

Producing a trace at the request of the indicated system.

In the message text:

SystemName

The remote system name which sent the xcf message.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00500I

Converting *AggrName* for fast mount processing.

Explanation

The zFS kernel is converting an aggregate for fast mount processing. Once the conversion is done, all future mounts or attaches of this aggregate will occur faster.

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00501E

Internal error converting bitmap file in *AggrName* for fast mount.

Explanation

The zFS kernel was converting an aggregate for fast mount processing and encountered an internal error while attempting to convert the bitmap file.

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Obtain a dump of the zFS address space by using the F ZFS,DUMP command, then contact your service representative.

IOEZ00502E

Internal error salvaging bitmap, could not get status, code=ErrorCode.

Explanation

The **ioeagslv** program failed to obtain status information for the bitmap file. The free space numbers recorded on disk will be incorrect and the aggregate cannot be salvaged. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the return code.

In the message text:

ErrorCode

The error code from the getstatus operation.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

IOEZ00503E

Bitmap computed blocksFree=CompBlocks fragsFree= CompFrag, on disk blocksFree=NumBlocks, fragsFree=NumFrag.

Explanation

The **ioeagslv** program determined the recorded number of free blocks and fragments on disk does not match the number computed by the salvage program.

In the message text:

CompBlocks

The computed number of free blocks in the aggregate.

CompFrag

The computed number of free fragments in the aggregate.

NumBlocks

The number of free blocks recorded in the aggregate.

NumFrag

The number of free fragments recorded in the aggregate.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

If the **ioeagslv** program was run with the default options or was run with the **-salvageonly** option, then the **ioeagslv** program will correct the free space numbers recorded on disk; otherwise, the free space numbers recorded on disk are incorrect and the **ioeagslv** program should be run with the default options or the **-salvageonly** option to correct the information on disk.

IOEZ00504E

Internal error salvaging bitmap, could not set status, code=ErrorCode.

Explanation

The `ioeagslv` program failed to correct the status and free space information for the bitmap file. This means the free space information on disk cannot be corrected and the aggregate cannot be salvaged. See [Return codes \(errno\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the return code.

In the message text:

ErrorCode

Error code from the `setstatus` operation.

System action:

The program ends.

Administrator response:

Contact your service representative.

Severity:

`svc_c_sev_error`

IOEZ00505I

Corrected free space information, `blocksFree=NumBlocks`
`fragsFree=NumFrag`s.

Explanation

The `ioeagslv` program successfully corrected the free space information recorded for the bitmap file. This means the bitmap file has been successfully restored on disk.

In the message text:

NumBlocks

The corrected number of free blocks on disk.

NumFrag

The corrected number of free fragments on disk.

System action:

The program continues.

Administrator response:

None.

Severity:

`svc_c_sev_notice`

IOEZ00506E

FILESYSTYPE PARM too long, `length = length`.

Explanation

The FILESYSTYPE PARM is too long. The length of the specified string is provided in the message. The maximum length is 1024 characters.

In the message text:

length

The length of the string.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

`svc_c_sev_error`

IOEZ00507E

Error in FILESYSTYPE PARM specification *string*.

Explanation

There is an error in the FILESYSTYPE PARM, as specified in the message.

In the message text:

string

The FILESYSTYPE PARM string.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00508E

Suffix at position *position* too long.

Explanation

The suffix starting in the specified column position is too long.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00509E

Comma or right paren expected at position *position*.

Explanation

A comma or a right parenthesis should have been in the specified column position, but is not.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00510E

Suffix at position *position* not two alphanumeric chars.

Explanation

The specified suffix should be two characters long, but is not. The suffix specification should be in the specified column position.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00511E

No suffixes found at position *position*.

Explanation

No valid parmlib suffix specification was found at the specified position.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00512E

No closing paren found at position *position*.

Explanation

A closing parenthesis was not found where one was expected, in the specified column position.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00513E

No open paren found at position *position*.

Explanation

Parmlib suffix specifications require parentheses around them. There is no opening parenthesis found. It is expected in the specified column position.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00514E

No equals sign found at position *position*.

Explanation

Parmlib suffix specifications require an equals sign. It is expected in the specified column position.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00515E

Extra characters were found at position *position*.

Explanation

Extra characters are not allowed. They start in the specified column.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00516E

PRM not found at position *position*.

Explanation

No PRM was found for parmlib suffix specifications. It is expected in the specified column.

In the message text:

position

The position of the error.

System action:

The program ends.

Administrator response:

Correct the problem and restart zFS. If the problem persists, or there is no error, contact your service representative.

Severity:

svc_c_sev_error

Explanation

The zFS kernel is converting the file system to a format that will allow for fast mounting. Fast mounting will allow for the efficient handling of a large number of objects in the file system.

In the message text:

FileSystem

Name of the file system.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

Explanation

The zFS kernel could not convert the file system to a version that will allow for fast mounts. The file system will remain in the prior format. The file system should still mount successfully.

In the message text:

ErrorCode

The internal error code received.

FileSystem

Name of the file system.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Contact your service representative.

Explanation

The zFS salvage program could not correct the file system control information. The file system may be permanently damaged.

In the message text:

ErrorCode

The internal error code received.

FileSystem

Name of the file system.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

Explanation

The zFS salvage program is correcting the file system free space information on disk due to either a problem found with this information or there were corrupted files in the file system which required corrections to the file system control information.

In the message text:

FileSystem

Name of the file system.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_notice

IOEZ00522E	Filesystem <i>FileSystem</i> is not in a compat aggregate. It cannot be mounted.
-------------------	---

Explanation

Mounts for file systems which do not reside in a compat aggregate are not allowed.

In the message text:

FileSystem

Name of the file system.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Ensure that the file system resides in a compat aggregate before retrying the mount. Also ensure that the file system name specified on the MOUNT statement in the BPXPRMxx parmlib member is in uppercase, not mixed-case or lowercase.

IOEZ00523I	zFS no longer supports the stop command. Please issue <code>f omvs,stoppfs=zfs</code>.
-------------------	---

Explanation:

The stop command is not supported. Issue the **omvs stoppfs** command.

System action:

The program continues.

Severity:

svc_c_sev_notice

Administrator Response:

Issue **f omvs,stoppfs=zfs**.

IOEZ00524I	zFS has a potentially hanging thread caused by: <i>UserList</i>.
-------------------	---

Explanation

zFS Hang Detector found a thread that could be hanging.

In the message text:

UserList

This a list of address space IDs and TCB addresses causing the hang

System action:

The program continues.

Severity:

svc_c_sev_notice

Administrator Response

Monitor the potential hang. You can enter **f zfs,query,threads** command to see the thread states. If the problem persists, try cancelling the hung address space IDs. If the problem continues, contact your IBM service representative (if you have set up the trace data set, the **f zfs,trace,print** command can gather trace information). You can also enter **f zfs,hangbreak**, but prepare for additional, extraneous dumps.

For complete details for message and hang detection, see *z/OS File System Administration*.

IOEZ00536E

First block of *FirstBlock* (block size *BlockSize*) starts the file system at byte *FirstBlock* * *BlockSize*, which is after the canonical superblock address of *SuperAddr*.

Explanation

An invalid initial empty value has been specified. The number specified is too large. You cannot specify a number larger than 8.

In the message text:

FirstBlock

The size of the specified block.

BlockSize

The 8K block size.

SuperAddr

The address of the superblock (64K).

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Specify a valid initial empty value and retry.

IOEZ00537E

Log has bad checksum length.

Explanation:

Recovery failed. The checksum length recorded on the disk does not match the checksum version recorded on the disk.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Run salvage and retry.

IOEZ00538I

Log has unknown checksum type on disk. Ignoring checksum check.

Explanation:

Ignoring unknown checksum type on disk.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00539E**Unable to retrieve checksum in log page.****Explanation:**

Recovery failed.

System action:

The program continues.

Severity:

svc_c_sev_error

Administrator Response:

Run salvage and retry.

IOEZ00540I**Log is using unknown checksum version *CurrentVersion*, Changing to use version *NewVersion*.****Explanation**

Setting the checksum version.

In the message text:

CurrentVersion

The version in use.

NewVersion

The new version.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00545E**A conversion of aggregate *AggrName* to version 3 is not complete. The conversion must be complete before the aggregate is usable.****Explanation**

Use IOEAGSLV to complete the conversion to version 3 and retry the request.

In the message text:

AggrName

The aggregate name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_error

IOEZ00547I**zFS has a potentially hanging XCF request on systems:*Systemnames*.**

Explanation

The zFS hang detector identified that a thread sent a message to another member of the sysplex, as indicated in the list, and is possibly hanging because it is waiting for the reply, because XCF is stalled or delayed, or because of some other reason.

In the message text:

Systemnames

The list of system names.

System action

The program continues.

Administrator response

Monitor the potential hang. Check the other members of the sysplex specified in the list to see if there is an obvious reason for the seemingly long wait for the reply. For complete details for messages and hang detection, see the topic about [Understanding zFS hang detection in z/OS File System Administration](#).

Severity

svc_c_sev_notice

IOEZ00548I	Requesting that <i>MemberName</i> takeover aggregate <i>AggrName</i> (requests: local <i>LocalReqs</i>, new owner <i>NewOwnerReqs</i> total <i>TotalReqs</i>).
-------------------	---

Explanation

This message can be issued when either of the following situations occurs:

- The local system is requesting that the named member take over the aggregate to reduce XCF communication.
- The aggregate has become disabled and zFS is attempting to automatically recover the aggregate.

In the message text:

MemberName

The member name.

AggrName

The aggregate name.

LocalReqs

The number of requests processed from the current owning system during the measurement period.

NewOwnerReqs

The number of requests processed from *MemberName* during the measurement period.

TotalReqs

The total number of requests processed from all systems during the measurement period.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00549E	Media Manager return code <i>ErrorCode</i> for aggregate <i>Aggr_name</i>.
-------------------	---

Explanation

The Media Manager return code was non zero for the specified aggregate.

Administrator response:

None.

Severity:

svc_c_sev_error

IOEZ00555E

-system cannot be specified without -all.

Explanation:

A **zfsadm detach** command specified -system without -all. -all is required when using -system.

System action:

The program ends.

Administrator response:

Retry the command specifying valid options.

Severity:

svc_c_sev_error

IOEZ00557E

zfsadm command failed, rc = *ReturnCode*, reason = *ReasonCode*.

Explanation

The indicated **zfsadm** command failed. The return and reason codes might indicate the cause of the error. If possible, correct the error and retry.

In the message text:

ReturnCode

The return code.

ReasonCode

The reason code.

System action:

The program continues.

Administrator response:

Correct any errors and retry the command.

Severity:

svc_c_sev_error

IOEZ00558E

An internal error occurred waiting for retry of a user task dump for *asid*.

Explanation

zFS attempted to take a dump of a user task in the *asid*. The dump might have failed because the address space was taking another dump. In an attempt to get a dump, zFS retries the SDUMPX call on the zFS task. If an internal error occurs while attempting to wait for the zFS task to complete, then the dump might not finish before the user task completes the zFS request. The dump might not prove useful in resolving the reason for dumping; contact your IBM service representative.

In the message text:

asid

The user task address space ID.

System action:

The program continues.

Administrator response:

Contact your service representative.

Severity:

svc_c_sev_error

Explanation

This message is issued when the named zFS program starts. It identifies the product name, feature name, version, release, modlevel, service level, creation date and asid of the zFS program.

In the message text:

progrname

The program name.

prodname

The product name.

featurename

The feature name.

ver

The product version.

rel

The product release.

mod

The product modification level.

slv

The product service level.

date

The date and time the zFS program was created.

asid

The asid of the address space.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

Explanation

This message is issued on a sysplex member that is receiving messages from a member trying to come up. The sysplex member that is already up requires the minimum service level indicated. The system initializing is not at the minimum service level. An error will be returned to the system initializing.

In the message text:

SystemName

The name of the remote system initializing.

ServiceLevel

The service level required on the remote system.

System action

The program continues.

Administrator response

Apply service and then restart zFS. For more information, see the topic about applying required APARs in [z/OS File System Administration](#).

Severity

svc_c_sev_error

IOEZ00579I

Restarting *Program*, Restart count *Count*.

Explanation

This message is issued when the named zFS program restarts. It identifies the restart count of the zFS program.

In the message text:

Program

The program name.

Count

The number of restarts zFS has taken.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00581E

There are quiesced zFS aggregates.

Explanation

This message appears on a system that has at least one zFS aggregate that is quiesced. There is a time delay between when the aggregate is quiesced and when the message appears. When there are no quiesced zFS aggregates on the system, this message is removed. There is also a delay between when the last aggregate is unquiesced and when the message is removed. The message is handled by a thread that wakes up every 30 seconds and checks to see if there are any quiesced aggregates owned by this system. It is possible for an aggregate to be quiesced and unquiesced in the 30 second sleep window of the thread and no quiesce message to appear. The message remains if one aggregate is unquiesced and another is quiesced within the 30 second sleep window.

The purpose of the message is to aid in diagnosis when there is an apparent hang. The message indicates that there is at least one quiesced aggregate and the apparent hang might be due to a quiesced aggregate. File I/O to a quiesced aggregate is held up until the aggregate is unquiesced. An aggregate is normally quiesced during backup and unquiesced when the backup is done.

To determine which aggregates are quiesced, use the z/OS UNIX **zfsadm lsaggr** command. After you determine which aggregates are quiesced, you can use the **zfsadm aggrinfo -long** command to determine who quiesced the aggregate (the job name, system, and time stamp). You can also use the MVS command **F ZFS,QUERY,FILE,ALL** to determine which aggregates are quiesced, especially in automation scripts that issue commands to the console. It is possible that the output of the command might not exactly match the message (message displayed with no quiesced aggregates owned by this system or no message with quiesced aggregates owned by this system).

System action

The program continues.

Administrator response

Determine whether the aggregate should be unquiesced. See the topic about [Understanding zFS hang detection in z/OS File System Administration](#).

Severity

svc_c_sev_notice

IOEZ00587I	zFS Release level cannot change during a recycle, zFS must terminate.
-------------------	--

Explanation:

zFS cannot change its release level without fully terminating the address space. zFS will assert and go down after issuing this message. It can then be restarted at any desired release.

System action:

The program ends.

Administrator response:

Do not attempt to change the current zFS release by an internal restart.

Severity:

svc_c_sev_notice

IOEZ00589E	Aggregate <i>Aggrname</i> failed to reopen. <i>Membername</i> cannot become the zFS owner of this aggregate.
-------------------	---

Explanation

The aggregate failed to reopen locally, possibly due to a DASD outage, some other connectivity difficulty due to the DASD or a hardware failure. This member cannot become the owner of this aggregate. z/OS UNIX will continue to send file operations to this member, expecting zFS to forward them to the zFS owner.

In the message text:

Aggrname

The name of the aggregate.

Membername

The name of the member.

System action:

The program continues.

Administrator response:

If all members fail to reopen the aggregate, unmount and remount the aggregate.

Severity:

svc_c_sev_notice

IOEZ00590E	Aggregate <i>Aggrname</i> was successfully reopened. <i>Membername</i> is now eligible to become a zFS owner.
-------------------	--

Explanation

The aggregate successfully reopened locally. This member is now eligible to become the zFS owner.

In the message text:

Aggrname

The name of the aggregate.

Membername

The name of the member.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00591I

**zFS is encountering delays in XCF message transmission to systems:
Systemnames.**

Explanation

The zFS Hang Detector found that a thread has sent a message to another member of the sysplex, as indicated in the list, and is possibly hanging due to the message not arriving at the target system. This could be because XCF is stalled or delayed for some other reason, or due to some other system resource that is constrained.

In the message text:

Systemnames

The list of system names.

System action:

The program continues.

Administrator response:

Monitor the potential hang. You can enter **f zfs,query,threads** to see the thread states. This information will appear in the system log, and should include which systems from which a reply is expected. The **f zfs,trace,print** command will gather trace information for the support team if your trace data set is configured. Check the other members of the sysplex specified in the list to see if there is an obvious reason for the seemingly long wait for the reply. In the case of actual XCF problems, XCF documentation may yield some other diagnostic techniques. If the problem persists, contact your service representative. You can also enter **f zfs,hangbreak**, but be prepared for additional extraneous dumps.

Severity:

svc_c_sev_notice

IOEZ00592I

**zFS is encountering delays with XCF replies from systems:
Systemnames.**

Explanation

The zFS Hang Detector found that one or more threads have sent a message to another member of the sysplex, as indicated in the list, and is experiencing delays in either replying to those messages or in receiving those replies on the local system. This could be because XCF is stalled or delayed for some other reason, or due to some other system resource that is constrained.

In the message text:

Systemnames

The list of system names.

System action:

The program continues.

Administrator response:

Monitor the potential hang. You can enter **f zfs,query,threads** to see the thread states. This information will appear in the system log, and should include which systems from which a reply is expected. The **f zfs,trace,print** command will gather trace information for the support team if your trace data set is configured. Check the other members of the sysplex specified in the list to see if there is an obvious reason for the seemingly long wait for the reply. In the case of actual XCF problems, XCF documentation may yield some other diagnostic techniques. If the problem persists, contact your service representative. You can also enter **f zfs,hangbreak**, but be prepared for additional extraneous dumps.

Severity:

svc_c_sev_notice

IOEZ00598E	<i>ErrorCode reason ReasonCode received formatting aggregate AggrName.</i>
-------------------	---

Explanation

An aggregate could not be formatted due to an unexpected error.

In the message text:

ErrorCode

The error code.

ReasonCode

The reason code.

AggrName

The name of the aggregate.

System action:

The program ends.

Administrator response:

Try to correct the problem and retry.

Severity:

svc_c_sev_error

IOEZ00599E	<i>ErrorCode during TSR for aggregate Name.</i>
-------------------	--

Explanation

An error occurred during token state recovery during aggregate takeover operation.

In the message text:

ErrorCode

The error code.

Name

The name of the aggregate.

System action:

The program continues.

Administrator response:

Contact the service representative.

Severity:

svc_c_sev_error

IOEZ00604I	<i>Task asid=Asid tcb=TCB is delayed outside zFS while Action_object.</i>
-------------------	--

Explanation

The specified task running in zFS has called a system service to perform an operation. The operation is, for example, opening or closing the specified data set. This operation has not returned to zFS for at least 3 intervals of the hang detector.

In the message text:

Asid

The asid of the task.

TCB

The TCB address of the task.

Action_object

The action phrase and object of action. The `osi_ctl` interface calls for `dub`, `getmntstatus`, `quiesce`, `updatefilesys`, and `unquiesce`. For more information, see [z/OS UNIX System Services Planning](#).

The code limits a maximum of 15 of these messages per hang detector cycle. There can be more tasks than messages listed.

System action

The program continues.

Administrator response

Try to correct the problem by following documented trouble shooting procedures for the operation specified. If the *Action_object* is:

- Allocating dataset, see information about the following topics:
 - [Monitoring the space used by the control data set in z/OS DFSMSrmm Implementation and Customization Guide](#)
 - [Allocating system-managed data sets in z/OS DFSMS Using Data Sets](#).
- Altering LDS, see [z/OS DFSMSdftp Storage Administration](#) or [z/OS DFSMS Access Method Services Commands](#).
- Calling catalog search interface for *datasetname*, see [z/OS DFSMS Managing Catalogs](#).
- Closing dataset, see the following topics in [z/OS DFSMS Using Data Sets](#):
 - Using VERIFY to process improperly closed data sets
 - Recovering from errors due to an improperly closed VSAM data set
 - Checking for problems in catalogs and data sets.
- Defining LDS, see [z/OS DFSMSdftp Storage Administration](#) and [z/OS DFSMS Using Data Sets](#) for more information.
- Dequeuing, see [z/OS MVS Planning: Global Resource Serialization](#)
- Growing dataset, see [z/OS DFSMSshm Diagnosis](#) or [z/OS DFSMSshm Managing Your Own Data](#).
- Locating dataset, see the topic about [Checking for problems in catalogs and data sets in z/OS DFSMS Using Data Sets](#).
- Performing trace dataset close, see the topic about [Processing a partitioned data set extended \(PDSE\) in z/OS DFSMS Using Data Sets](#).
- Performing trace dataset unallocate, see the topic about [Processing a partitioned data set extended \(PDSE\) in z/OS DFSMS Using Data Sets](#).
- Opening dataset, see the topic about [Opening a data set in z/OS DFSMS Using Data Sets](#).
- Recalling dataset, see [z/OS DFSMSshm Storage Administration](#) and [z/OS DFSMSshm Managing Your Own Data](#).
- Remounting dataset, see the topic about problem diagnosis for a shared system in [z/OS MVS Diagnosis: Reference](#).
- stimer wait, see the topic about understanding namespace validation and correction in [z/OS File System Administration](#).
- Updating high used relative byte address (HURBA), see [Requesting that DFSMSdss double-check data set high used RBA values for LDS data sets in z/OS DFSMSdss Storage Administration](#).
- Waiting for I/O, see [z/OS DFSMS Using Data Sets](#).

SMF record type 92 reports the activity of mounted file systems and files. For more information about [SMF record type 92](#) see [z/OS UNIX System Services Planning](#).

Severity

svc_c_sev_notice

IOEZ00605I

Task asid=Asid tcb=TCB appears to be delayed.

Explanation

The specified user task running in zFS has been running for at least three intervals of the hang detector. There can be many reasons for this behavior that are expected. For example, the task has a low dispatch priority that is preventing its progress, or the task is swapped out, or the system has a resource constraint such as an auxiliary storage shortage. It could also be doing a long running task, or looping. Another possibility is that zFS has called another system service (not covered in message IOEZ00604I) and it has not returned.

In the message text:

Asid

The asid of the task.

TCB

The TCB address of the task.

The code limits a maximum of 15 of these messages per hang detector cycle. There can be more tasks than messages listed.

System action

The program continues.

Administrator response

Try to correct the problem using the troubleshooting procedures for the task identified. Also see the topic about [Understanding zFS hang detection in z/OS File System Administration](#).

Severity

svc_c_sev_error

IOEZ00606E	Error <i>ErrorCode</i> reason <i>ReasonCode</i> received while attempting to setauditfid for aggregate <i>AggrName</i> .
------------	--

Explanation

An unexpected error was received while attempting to setauditfid for aggregate *AggrName*.

In the message text:

ErrorCode

The error code.

ReasonCode

The reason code.

AggrName

The name of the aggregate.

System action

The program ends.

Administrator response

Try to resolve the error and reissue the command.

Severity

svc_c_sev_error

IOEZ00607I	Auditfid set for aggregate <i>AggrName</i> .
------------	--

Explanation

The auditfid for aggregate *AggrName* was successfully set.

In the message text:

AggrName

The name of the aggregate.

System action

The program continues.

Administrator response:

None.

Severity

svc_c_sev_notice

IOEZ00608I

Aggregate *AggrName* has new format auditfid. Specify -force or -old to override.

Explanation

Aggregate *AggrName* has a new format auditfid. A new format auditfid cannot be reset by default. Specify -force or -old.

In the message text:

AggrName

The name of the aggregate.

System action

The program ends.

Administrator response

Issue the command again with -force or -old.

Severity

svc_c_sev_notice

IOEZ00609I

Parameter *parm* is too long.

Explanation

Parameter *parm* has too many characters in it.

In the message text:

parm

The name of the parameter.

System action

The program ends.

Administrator response

Use a shorter name for the parameter.

Severity

svc_c_sev_notice

IOEZ00610I	zFS Name Space Validation could not obtain ENQ on <i>ResName</i>, waiting <i>Seconds</i> seconds.
-------------------	--

Explanation

zFS is unable to perform name space validation because it cannot obtain a GRS ENQ on the specified resource name. zFS will retry the ENQ in *Seconds* seconds. Name space validation occurs either when zFS detects an internal error or there are XCF communication problems in the sysplex.

In the message text:

ResName

The resource name.

Seconds

The number of seconds.

System action

The program continues.

Administrator response

Examine the state of the sysplex and take corrective action as appropriate. See [Performance and debugging in z/OS File System Administration](#).

Severity

svc_c_sev_notice

IOEZ00611I	zFS Name Space Validation could not ping system <i>SysName</i>, waiting <i>Seconds</i> seconds.
-------------------	--

Explanation

zFS is unable to perform name space validation because it cannot communicate with the specified system. zFS will retry communications in *Seconds* seconds. Name space validation occurs either when zFS detects an internal error or there are XCF communication problems in the sysplex.

In the message text:

SysName

The name of the system.

Seconds

The number of seconds.

System action:

The program continues.

Administrator response

Examine the state of the sysplex and take corrective action as appropriate. See [Performance and debugging in z/OS File System Administration](#).

Severity

svc_c_sev_warning

IOEZ00612I	zFS Name Space Validation is running due to a detected internal error.
-------------------	---

Explanation

zFS has detected an internal error and is performing name space validation and correction. zFS will correct the problem for any aggregate that has an inconsistent state across the sysplex by internally re-mounting the aggregate on all systems or restarting zFS on one or more sysplex members, or both.

System action

The program continues.

Administrator response

Examine the state of the sysplex and take corrective action as appropriate. See [Performance and debugging in z/OS File System Administration](#).

Severity

svc_c_sev_warning

IOEZ00613I	zFS Name Space Validation is running due to a detected XCF communication failure or message timeout.
-------------------	---

Explanation

zFS has detected either an XCF communication failure or had a name space related message time out. zFS will correct the problem for any aggregate that has an inconsistent state across the sysplex by internally remounting the aggregate.

System action

The program continues.

Examine the state of the sysplex and take corrective action as appropriate. See [Performance and debugging in z/OS File System Administration](#).

Severity

svc_c_sev_warning

IOEZ00614A	zFS has detected an incompatible interface level <i>IntLevel</i> for member <i>Sysname</i>.
-------------------	--

Explanation

zFS has detected that there are other members in the sysplex that have incompatible release or service levels with the current system that is attempting to initialize. Initialization ends.

In the message text:

Intlevel

The interface level of the other named member.

Sysname

The name of the other system.

System action:

The program ends.

Administrator response

Ensure that you have the correct release and service levels of zFS in your sysplex. To determine service levels, see the topic about [Determining service levels in z/OS File System Administration](#).

Also, see the topic about applying required APARs in [z/OS File System Administration](#).

Severity

svc_c_sev_fatal

IOEZ00615E	zFS encountered error <i>Rcode</i> reason=<i>RsnCode</i> obtaining GRS serialization for aggregate <i>Aggregate</i>.
-------------------	---

Explanation

zFS could not obtain the proper GRS serialization to detach an aggregate that was no longer mounted in the sysplex. You might experience failures while attempting to mount file systems for the aggregate.

In the message text:

Rcode

The GRS return code.

RsnCode

The GRS reason code.

Aggregate

The aggregate name.

System action:

The program continues.

Administrator response

Determine the source of the GRS problem. Issue `zfsadm detach` to detach the aggregate from zFS.

Severity

svc_c_sev_error

IOEZ00616E	A timeout has occurred attempting to establish a connection with system <i>Target_system</i>.
-------------------	--

Explanation

While attempting to establish a connection with the target system, a timeout has occurred. However, it is not known if the target system is truly down or not as it is still a member of the XCF group. There will be a maximum of six attempts at establishing the connection. If it is still unsuccessful, then this system will terminate.

In the message text:

Target_system

The target system.

System action

zFS will attempt to establish a connection six times, and terminate if the connection cannot be established within these attempts.

Administrator response

Determine the source of the XCF related problem, and then use XCF diagnostic techniques to correct it. If necessary, restart zFS.

Severity:

svc_c_sev_error

IOEZ00617I	zFS is running sysplex <i>SysLevel</i> with interface level <i>IntLevel</i>.
-------------------	---

Explanation

The message indicates which zFS sysplex functionality is enabled and the sysplex communications interface level. Beginning in z/OS V1R13, zFS uses interface level 4.

In the message text:

SysLevel

The sysplex functionality name.

IntLevel

The zFS XCF protocol interface level.

4

The z/OS V2R5 level (long-running command support)

4

The z/OS V2R4 level (long-running command support)

4

The z/OS V2R3 level (long-running command support)

4

The z/OS V2R2 level (enhanced log and enhanced status APIs)

4

The z/OS V2R1 level (extended directory)

4

The z/OS V1R13 level (enhanced connect)

System action

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00618E

SY1 owns the long term enq for aggregate Aggrname or is the owner as defined by the first entry found. System SY2 thinks SY3 owns the aggregate.

Explanation

While validating the zFS namespace, zFS found a cache entry for an aggregate that does not have the correct owner. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The system owning the long term enq.

Aggrname

The aggregate name.

SY2

The system with the bad cache entry.

SY3

The system the bad cache entry indicates is the owner.

System action

zFS will attempt to correct the bad cache entry.

Administrator response

None, if zFS can correct the bad cache entry. Otherwise, restart zFS on the system with the bad cache entry.

Severity

svc_c_sev_error

IOEZ00619E	Aggregate <i>Aggrname</i> is unowned. System <i>SY1</i> thinks <i>SY2</i> owns the aggregate.
-------------------	--

Explanation

While validating the zFS namespace, zFS found a cache entry for an unowned aggregate that indicates the aggregate is owned. Correction may cause zFS some systems to be restarted.

In the message text:

Aggrname

The aggregate name.

SY1

The system with the bad cache entry.

SY2

The system with the bad cache entry indicates is the owner.

System action

zFS will attempt to correct the bad cache entry.

Administrator response

None, if zFS can correct the bad cache entry. Otherwise, restart zFS on the system with the bad cache entry.

Severity

svc_c_sev_error

IOEZ00620E	<i>SY1</i> is the owner of RW aggregate <i>Aggrname</i> and is not sysplex aware for RW. There is another cache entry for this aggregate on system <i>SY2</i>
-------------------	--

Explanation

RW aggregates that are owned by systems that are not sysplex aware for RW aggregates can only have one cache entry in the entire sysplex. Correction may cause zFS on some systems to be restarted.

In the message text:

Aggrname

The aggregate name.

SY1

The system owning the long term enq.

SY2

Another system with a cache entry.

System action

zFS will attempt to correct the error.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity

svc_c_sev_error

IOEZ00621E

There is only one cache entry for aggregate *Aggrname* on system *System* but this entry indicates there are other systems connected.

Explanation

Connected systems must have cache entries for the aggregate, but there is only one entry for the aggregate in the sysplex. Correction may cause zFS on some systems to be restarted.

In the message text:

Aggrname

The aggregate name.

System

The system name.

System action

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity

svc_c_sev_error

IOEZ00622E

There is only one cache entry for aggregate *Aggrname* on system *SY1* but this entry indicates that this system is connected to another system.

Explanation

Connect flags indicate another system owns the aggregate, but there is only one entry for the aggregate in the sysplex. Correction may cause zFS on some systems to be restarted.

In the message text:

Aggrname

The aggregate name.

SY1

The system name.

System action

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity

svc_c_sev_error

IOEZ00623E

Aggregate *Aggrname* is unowned and growing on system *System*.

Explanation

A growing aggregate should always be owned. Correction may cause zFS on some systems to be restarted.

In the message text:

Aggrname
The aggregate name.

System
The system name.

System action

zFS will attempt to correct.

Administrator response:
None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity
svc_c_sev_error

IOEZ00624E	Aggregate <i>Aggrname</i> is on system <i>SY1</i> with flags <i>Flag1</i> and also on system <i>SY2</i> with flags <i>Flag2</i>.
-------------------	---

Explanation

The two specified systems have mismatching flags. Correction may cause zFS on some systems to be restarted.
In the message text:

Aggrname
The aggregate name.

SY1
The system 1 name.

Flag1
The system 1 flag name.

SY2
The system 2 name.

Flag2
The system 2 flag name.

System action

zFS will attempt to correct.

Administrator response:
None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity
svc_c_sev_error

IOEZ00625E	Aggregate <i>Aggrname</i> is on system <i>System1</i> with <i>Field1 Currstate1</i> and also on system <i>System2</i> with <i>Field2 Curststate2</i>.
-------------------	--

Explanation

This is an internal error. The two specified systems have mismatching flags. Correction may cause zFS on some systems to be restarted.
In the message text:

Aggrname
The aggregate name.

System1

The system1 name.

Field1

The field1 name.

Currstate1

The field1 value.

System2

The system2 name.

Field2

The field2 name.

Curstate2

The field2 value.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity

svc_c_sev_error

IOEZ00626E

Aggregate *Aggrname* is on system *System1* with *Field1val1*,,*val2* and also on system *System2* with *Field2 val3*,,*val4*.

Explanation

The two specified systems have mismatching flags. Correction may cause zFS on some systems to be restarted.

In the message text:

Aggrname

The aggregate name.

System1

The system1 name.

Field1

The field1 ID.

val1

The ID field1 value.

val2

The ID field1 value.

System2

The system2 name.

Field2

The field2 ID.

val3

The ID field2 value.

val4

The ID field2 value.

System action

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity

svc_c_sev_error

IOEZ00627E

Aggregate *Aggrname* on system *System* is unowned but its cache entry indicates other systems are connected.

Explanation

An unowned aggregate should not have other systems connected to it. Correction may cause zFS on some systems to be restarted.

In the message text:

Aggrname

The aggregate name.

System

The system name.

System action

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity

svc_c_sev_error

IOEZ00628E

Aggregate *Aggrname* is unowned, but system *System* has a cache entry indicating a connection to an owning system.

Explanation

An unowned aggregate should not have connections. Correction may cause zFS on some systems to be restarted.

In the message text:

Aggrname

The aggregate name.

System

The system name.

System action

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00629E

The owner *SY1* of aggregate *Aggrname* thinks client *SY2* is *status* to the *fstype* file system but the connect flag on system *SY3* is *state*.

Explanation

The connection status of an aggregate on the owner does not match its status on the client. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The system1 name.

Aggrname

The aggregate name.

SY2

The client system name.

status

The connect status.

fstype

The file system type.

SY3

The client system name.

state

The connect flag state.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00630E

The owner SY1 of aggregate Aggrname thinks client SY2 is not connected, but the client has a cache entry.

Explanation

The aggregate is connected to the client. Clients that are not connected must not have cache entries. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The owner system name.

Aggrname

The aggregate name.

SY2

The client system name.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00631E

The client SY1 of aggregate Aggrname thinks it is not connected, yet it has a cache entry.

Explanation

The aggregate is connected to the client. Clients that are not connected must not have cache entries. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The owner system name.

Aggrname

The aggregate name.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00632E

The owner SY1 of aggregate *Aggrname* thinks client SY2 is not connected to the *fstype* filesystem, but the client thinks otherwise.

Explanation

The connection status of an aggregate on the owner system does not match its status on the client. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The owner system name.

Aggrname

The aggregate name.

SY2

The client system name.

fstype

The file system type.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00633E

The client SY1 of aggregate *Aggrname* thinks it is not connected to the *fstype* filesystem, but the owner thinks otherwise.

Explanation

The connection status of an aggregate on the owner does not match its status on the client. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The client system name.

Aggrname

The aggregate name.

fstype

The file system type.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00634E	The owner <i>SY1</i> of aggregate <i>Aggrname</i> thinks system <i>SY2</i> is connected but it does not have a cache entry.
-------------------	--

Explanation

The owner connect mask indicates the client should have a cache entry but it does not. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The owner system name.

Aggrname

The aggregate name.

SY2

The client name.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00635E	System <i>SY1</i> owns aggregate <i>Aggrname</i> but does not have a cache entry for it.
-------------------	---

Explanation

The owner of an aggregate does not have a cache entry for it. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The owner system name.

Aggrname

The aggregate name.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00636E	System <i>SY1</i> owns aggregate <i>Aggrname</i> but its cache entry says the aggregate is owned by <i>SY2</i>.
-------------------	--

Explanation

The owner of an aggregate does not have a cache entry for it. Correction may cause zFS on some systems to be restarted.

In the message text:

SY1

The owner system name.

Aggrname

The aggregate name.

SY2

The other owner system name.

System action:

zFS will attempt to correct.

Administrator response:

None, if zFS can correct this error. Otherwise, remount the associated file system.

Severity:

svc_c_sev_error

IOEZ00639I	<i>progrname: prodname featurenameVersion ver.rel.mod</i>
	Service Level <i>slv</i>.
	Created on <i>date</i>.
	<i>syslevel</i>

Explanation

This message is issued when the level of zFS is queried using the **query** command. The system returns the product name, feature name, version, release, modification level, service level, and creation date of the zFS program.

In the message text:

progrname

Program name.

prodname

Product name.

featurename

Feature name.

ver

Product version.

rel

Product release.

mod

Product modification level.

slv

Product service level.

date

Date and time the daemon was created.

syslevel

Sysplex functionality name/interface level. The zFS XCF protocol interface levels are:

4

The z/OS V2R5 level (long-running command support)

4

The z/OS V2R4 level (long-running command support)

4

The z/OS V2R3 level (long-running command support)

4

The z/OS V2R2 level (enhanced log and enhanced status APIs)

4

The z/OS V2R1 level (extended directory)

4

The z/OS V1R13 level (enhanced connect)

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00640E

**Failed sending the *msgname* initialization message to system *sysname*.
Please restart.**

Explanation

zFS could not send a required initialization XCF message to the specified system. Initialization terminates. Restart zFS by responding R to console message BPXF032D.

In the message text:

msgname

The message name. This is an internal XCF message defined by zFS. This information can be used by IBM service.

sysname

The system name zFS attempted to send the internal XCF message to.

System action

The program ends. zFS file systems might be unmounted or moved in a sysplex.

Administrator response:

Examine the console log for system *sysname* to see if there are any outstanding operator responses required. Try to determine why system *sysname* is not able to receive XCF messages. After zFS is restarted, you might need to mount any zFS file systems that were unmounted.

Severity:

svc_c_sev_error

IOEZ00641I

Cannot attach aggregate *Aggrname* in R/O mode due to incomplete clone.

Explanation

zFS cannot attach the specified aggregate in R/O mode because there is an incomplete clone from a prior mount of the aggregate.

In the message text:

Aggrname

Aggregate name.

System action:

Mount of the aggregate fails.

Administrator response:

The aggregate needs to be mounted/attached in R/W mode to allow the incomplete clone to be deleted. The delete of the incomplete clone will automatically be attempted by zFS when the aggregate is mounted R/W.

Severity:

svc_c_sev_notice

IOEZ00642I

Mounting filesystem *filesystemname* is delayed, SMS is not active.

Explanation

The Storage Management Subsystem is not active. zFS will wait for up to one minute for SMS to become active, and then the mount will fail if SMS is still not active.

In the message text:

filesystemname

File system being mounted.

System action:

The mount fails if SMS is not active after one minute.

Administrator response:

If SMS is not active, consult SMS diagnostics to determine the cause. This can be a temporary situation due to a long running SMS user exit. This situation may clear itself. If the mount fails, you will need to mount the file system after SMS becomes active to make it available.

Severity:

svc_c_sev_warning

IOEZ00643I

**The value for configuration option *ConfigOption* is: *Prognam*e: *Prodname*
Featurename *Version* *Version* *Service Level* *SLV* *Created on* *Date*.**

Explanation

This message is issued when the level of zFS is queried using the CONFIGQUERY command. The system returns the product name, feature name, version, release, modification level, service level and creation date of the zFS program.

In the message text:

ConfigOption

Option name.

***Prognam*e**

Program name.

Featurename

Feature name.

Prodname

Product name.

Version

Product *version.release.prodnam*e.*modification* level.

SLV

Product service level.

Date

Date and time the Daemon was created.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00644I

**The value for configuration option *ConfigOption* is: *Prognam*e: *Prodname*
Featurename *Version* *Version* *Service Level* *SLV* *Created on* *Date*
Syslevel.**

Explanation

This message is issued when the level of zFS is queried using the CONFIGQUERY command. The system returns the product name, feature name, version, release, modification level, service level and creation date of the zFS program.

In the message text:

ConfigOption

Option name.

Progrname

Program name.

Prodnamd

Product name.

Featurename

Feature name.

Version

Product *version.release.prodname.modification* level.

SLV

Product service level.

Date

Date and time the Daemon was created.

Syslevel

Sysplex functionality name\interface level.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00645A

Error code initializing the zFS kernel control program.

Explanation

The zFS kernel control program IOEFSCM encountered an unexpected error during initialization.

In the message text:

code

Error code for the initialization failure.

System action:

The program ends.

Administrator response:

None; the internal error would be reported by prior messages and dumps.

Severity:

svc_c_sev_fatal

IOEZ00646I

zFS Kernel is restarting *numfilesys* file systems.

Explanation

zFS had an internal failure and is recovering from that failure. This message will remain on the operator's console until zFS has internally re-processed every file system that was mounted at the time of the failure.

In the message text:

numfilesys

Number of mounted file systems.

System action:

The program continues.

Administrator response:

None; the internal error would be reported by prior messages and dumps.

Severity:

svc_c_sev_warning

IOEZ00650I**Successfully changed the attribute of aggregate *aggregate_name*.****Explanation**

The aggregate attribute was updated successfully.

aggregate_name

Name of the aggregate.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00651E**The -aggregate option must be specified.****Explanation**

The -aggregate option is a required option for the **zfsadm chaggr** command.

System action:

The program ends.

Administrator response:

Retry the command specifying the -aggregate option.

Severity:

svc_c_sev_error

IOEZ00652E***aggregate_name* is not mounted.****Explanation**

zFS attributes can only be changed for aggregates which are mounted.

In the message text:

aggregate_name

Name of the aggregate.

System action:

The program ends.

Administrator response:

Ensure that the name of the aggregate is entered correctly, and that it is the name of a mounted HFS-compatibility mode aggregate. Try the command again.

Severity:

svc_c_sev_error

IOEZ00653E**The attribute for aggregate *aggregate_name* could not be updated.
The operation fails with error code=*ErrorCode* and reason
code=*ReasonCode*.****Explanation**

An error was encountered during the attempt to update the aggregate attribute.

In the message text:

aggregate_name

Name of the aggregate.

ErrorCode

The error code.

ReasonCode

The reason code.

System action:

The program ends.

Administrator response:

If the error cannot be resolved, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00655E

Syntax error in the argument of the -aggrgrow option - *argument*.

Explanation

There is a syntax error in the value that is specified for the argument on the -aggrgrow option of the **zfsadm chaggr** command. The incorrect value is shown in the message.

In the message text:

argument

The argument in error.

System action:

The program continues.

Administrator response:

Correct the argument and retry the command.

Severity:

svc_c_sev_error

IOEZ00657E

Syntax error with aggregate -aggrfull parameter - *parm_name*.

Explanation

An aggregate attribute could not be updated because the -aggrfull option of the **zfsadm chaggr** command is improperly specified. See [z/OS File System Administration](#) for a description of the valid values for -aggrfull.

In the message text:

parm_name

The -aggrfull parameter in error.

System action:

The program ends.

Administrator response:

Correct the value and try the command again.

Severity:

svc_c_sev_error

IOEZ00658E

Timeout exceeded for *Operation* operation. A reply was not received from system *System*.

Explanation

A reply was not received within the timeout period for an operation.

In the message text:

Operation

Name of the operation.

System

Name of the system.

System action

The program continues.

Administrator response

Use the information in the message that refers to the system that is not responding and determine the status of that system. Check the console of the system not responding for outstanding messages and act accordingly.

Severity

svc_c_sev_error

IOEZ00659E	Timeout exceeded for the <i>Operation</i> operation on <i>Name</i>. A reply was not received from system <i>System</i> operation.
-------------------	--

Explanation

A reply was not received within the timeout period for an operation.

In the message text:

Operation

Name of the operation.

Name

Name of the file system or aggregate.

System

Name of the system.

System action

The program continues.

Administrator response

Use the information in the message that refers to the system that is not responding and determine the status of that system. For problem determination, use [z/OS MVS Diagnosis: Reference](#).

Severity

svc_c_sev_error

IOEZ00660I	There is <i>Numtasks</i> task(s) delayed outside zFS while <i>Action_object</i>.
-------------------	---

Explanation

One or more tasks running in zFS have called a system service to perform an operation. The operation is, for example, enqueue waiting for the specified data set. However, this operation has not yet returned to zFS and the delay has caused the hang detector to suspect a problem.

In the message text:

Numtasks

Number of tasks.

Action_object

The action phrase and object of action. The `osi_ctl` interface calls for `dub`, `getmntstatus`, `quiesce`, `updatefilesystem`, and `unquiesce`. For more information, see [z/OS UNIX System Services Planning](#).

System action:

The program continues.

Administrator response

Try to correct the problem using the troubleshooting procedures for the operation specified. If the *Action_object* is:

- *calling ISGQUERY*, see *z/OS MVS Planning: Global Resource Serialization*
- *calling an MVS sysplex service macro*, zFS has called a sysplex service macro, such as IXCMSGI, IXCMSGO, IXCDELET, IXCLEAVE, IXCJOIN or IXCQUERY. See *z/OS MVS Programming: Sysplex Services Reference* .
- *deleting ACEE, getting ACEE, or setting ACEE*, see *z/OS Security Server RACF Callable Services*.
- *enqueuing dataset*, see *z/OS MVS Planning: Global Resource Serialization*.
- *freeing storage or obtaining storage*, see *z/OS MVS Diagnosis: Tools and Service Aids*. For information about the STORAGE macro, see *z/OS MVS Programming: Assembler Services Reference IAR-XCT*.
- *performing security check*, see *z/OS Security Server RACF Diagnosis Guide* or the corresponding documentation, if you are using a different security product.
- *writing to the console*, see *z/OS MVS Diagnosis: Tools and Service Aids*. For information about the WTO macro, see *z/OS MVS Programming: Assembler Services Reference IAR-XCT*.

Severity:

svc_c_sev_notice

IOEZ00661I

zFS has a slow running XCF request on systems: *Systemnames*.

Explanation

The zFS hang detector identified that a thread sent a message to another member of the sysplex, as indicated in the list, and the thread has been running slowly for a time period that causes the hang detector to suspect a problem.

In the message text:

Systemnames

List of system names.

System action

This message is written to the console and then deleted (DOMed) when the target system responds.

Administrator response

Monitor the situation by issuing **f zfs,query,threads** to view thread states. This information appears in the system log and displays system status. If your trace data set is set up, issue **f zfs,trace,print** to gather trace information. Check the other members of the sysplex for an obvious reason for the wait. If the problem persists, contact your service representative. For information about setting up the trace data set, see the topic about *Performance and debugging in z/OS File System Administration*.

Severity

svc_c_sev_notice

IOEZ00662I

ZFS is low on storage.

Explanation:

zFS has determined that less than 60 MB of storage is remaining in the address space. While in this condition, storage requests related to the creation of new vnodes (file structures) and new mounts of file systems will fail.

System action:

The program continues. New mounts of zFS file system will fail.

Administrator response

It is possible to see this message the first time you define zFS or as the result of changes to the IOEFSPRM cache size options. If this is the case, check your initial or new IOEFSPRM settings, lower the cache sizes, and then try restarting zFS. See the settings for [IOEFSPRM](#) in [z/OS File System Administration](#).

Issue **f zfs,query,storage** to get a storage report and determine what resources are using unnecessary storage. Take steps to limit the storage being used by that area. Consider restarting zFS to return to a more stable environment. Continue to monitor the situation. If the situation becomes unacceptable, take dumps of the z/OS UNIX and zFS address spaces and call IBM Service. For information about limiting storage and obtaining a dump, see the topic about [Performance and debugging](#) in [z/OS File System Administration](#).

Severity:

svc_c_sev_warning

IOEZ00663I**ZFS is critically low on storage.**

Explanation:

zFS has determined that less than 20 MB of storage is remaining in the address space. While in this condition, all zFS storage requests related to the creation of new vnodes (file structures), or new mounts of file systems will fail.

System action:

The program continues. New mounts of zFS file system will fail.

Administrator response

It is possible to see this message the first time you define zFS or as the result of changes to the IOEFSPRM cache size options. If this is the case, check your initial or new IOEFSPRM settings, lower the cache sizes, and then try restarting zFS. See the settings for [IOEFSPRM](#) in [z/OS File System Administration](#).

Issue **f zfs,query,storage** to get a storage report and determine what resources are using unnecessary storage. Take steps to limit the storage being used by that area. Consider restarting zFS to return to a more stable environment. Continue to monitor the situation. If the situation becomes unacceptable, take dumps of the z/OS UNIX and zFS address spaces and call IBM Service. For information about limiting storage and obtaining a dump, see the topic on [Performance and debugging](#) in [z/OS File System Administration](#).

Severity:

svc_c_sev_warning

IOEZ00664E**Space monitoring initialization for *Name* has failed.**

Explanation

zFS either has an internal error or a lack of storage that was detected when attempting to initialize the necessary environment for space monitoring.

In the message text:

Name

Aggregate name or file system name.

System action

The program continues. The mounted file system's space usage will not be monitored.

Administrator response

If you see messages IOEZ00662I or IOEZ00663I before this message is issued, try lowering storage usage and remounting. For more information about lowering storage, see the topic about [Performance and debugging](#) in [z/OS File System Administration](#).

If you do not see messages IOEZ00662I or IOEZ00663I before this message being issued, contact IBM Service.

Severity:

svc_c_sev_error

IOEZ00665E**Restarting local system *Sysname* for correction.****Explanation**

A local system is performing correction and has determined that it must restart itself. In the message text:

Sysname

Name of the local system performing validation.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_error

IOEZ00666E**zFS kernel will be restarted automatically for a correction request from *Sysname*.****Explanation**

A remote system performing validation has requested that the local system restart to correct a problem.

In the message text:

Sysname

Name of the remote system performing validation.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_error

IOEZ00667I**ZFS suppressing a dump for an error condition.****Explanation**

zFS has encountered an error condition that normally causes a system dump. However, zFS has determined that it is already taking a dump on another task. Therefore, the dump for this error is being suppressed.

System action

The program continues.

Administrator response

Ensure a zFS dump is produced automatically. If not, see the troubleshooting procedures in [z/OS File System Administration](#).

Severity

svc_c_sev_notice

IOEZ00668I**zFS Configuration option *configoption* is obsolete and is not used.****Explanation**

The specified configuration option is obsolete. It is not needed anymore and is ignored by zFS. Values can be set and queried, but they are not used by zFS processing.

In the messages text:

configoption

zFS configuration option that is obsolete.

System action:

The program continues.

Administrator response:

There is no need to specify or query this configuration option.

Severity:

svc_c_sev_notice

IOEZ00670I **Starting FSINFO command.**

Explanation:

The FSINFO command is now processing. There will be an IOEZ00849I message indicating that the display from this modify command is completed.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00671E **Bad file system name *filesys_name* passed to FSINFO.**

Explanation:

zFS could not open the specified file system for FSINFO processing. It was not a valid file system name.

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_error

IOEZ00674E **zFS failed to obtain critical ENQ on resource *ResName*. Rc and rsn are *rc* and *rsn*. zFS terminating.**

Explanation

zFS is unable to perform initialization because it cannot obtain a GRS ENQ on the specified resource name with qname SYSZIOEZ. To fix this problem you must determine the cause. For example, another zFS in the sysplex could be holding the ENQ or there could be XCF communication problems in the sysplex.

ResName

The resource name that could not be obtained.

rc

zFS return code.

rsn

zFS reason code.

System action:

The program ends.

Administrator response:

Examine the state of the sysplex and take corrective action as appropriate.

Severity:

svc_c_sev_warning

IOEZ00675E

zFS will terminate... system *Sysname1* does not support the following feature(s) used by this system (*Sysname2*) *feature_list*.

Explanation

The initializing system is using a feature that is not supported by the named system. zFS on the initializing system terminates. This message is issued on the initializing system.

Sysname1

The remote system that is already active, but does support the features required by the system trying to initialize.

Sysname2

The local system that is trying to initialize.

feature_list

The list of features used by the initializing system (*Sysname2*), which are not supported by the remote system. If the feature name is *sysplex=filesys*, see the topic about [Using zFS in a shared file system environment in *z/OS File System Administration*](#).

System action:

zFS initialization terminates.

Administrator response:

Either apply service to *Sysname1* and retry, or remove the use of the indicated features from *Sysname2* and retry. For more information, see the topic on applying required APARs in [z/OS File System Administration](#).

Severity:

svc_c_sev_warning

IOEZ00676E

zFS on *Sysname1* should be terminated... it does not support the following feature(s) used by this system (*Sysname2*) *feature_list*.

Explanation

The named system is initializing and does not support a feature used by this system. The zFS on the initializing system terminates. Otherwise, zFS on the initializing system should be terminated with the MODIFY OMVS,STOPPFS=ZFS operator command. This message is issued on the remote system.

Sysname1

The initializing system that does not have support.

Sysname2

The local system using features that are not supported by the initializing system.

feature_list

The list of features used by the local system (*Sysname2*), which are not supported by the remote (initializing) system. If the feature name is *sysplex=filesys*, see [Using zFS in a shared file system environment in *z/OS File System Administration*](#).

System action:

zFS initialization terminates.

Administrator response:

If zFS on the initializing system does not terminate on its own, then stop zFS with the MODIFY OMVS,STOPPFS=ZFS operator command. Apply the necessary service to the initializing system (*Sysname1*) and retry. If you need to run *Sysname1* without the necessary service, do not use the feature in *feature_list* on *Sysname2*.

Severity:

svc_c_sev_warning

IOEZ00677E

zFS will terminate... system *Sysname1* uses feature(s) not supported by the initializing system (*Sysname2*).

Explanation

The named system (*Sysname1*) is using a feature that is not supported by the initializing system (*Sysname2*). The zFS on the initializing system terminates. This message is issued on the initializing system (*Sysname2*).

Sysname1

The remote system using feature that is not supported.

Sysname2

The local (initializing) system that does not have support for the feature that *Sysname1* is using.

System action:

zFS initialization terminates.

Administrator response:

Inspect the *feature_list* field of message IOEZ00676E on the remote system to determine which support is missing. Apply service to the initializing system to provide the missing support and retry.

Severity:

svc_c_sev_warning

IOEZ00678E

Aggregate *Aggregate_name* has reached the maximum architected size.

Explanation

The specified aggregate has reached the maximum architected size and can no longer be grown. The maximum architected size for a version 1.4 aggregate is approximately 4TB. The maximum architected size for a version 1.5 aggregate is approximately 16 TB.

Aggregate_name

Is the name of the aggregate

System action

The program continues.

Administrator response

If more space is needed for the aggregate, you should consider moving some of its files and directories to another file system. If the aggregate is version 1.4, you can also consider converting the aggregate to version 1.5. For more information about maximum aggregate sizes, see the section on minimum and maximum file system sizes in [z/OS File System Administration](#).

Severity

svc_c_sev_warning

IOEZ00700E

Aggregate *aggregate_name* is disabled for writing.

Explanation

The specified aggregate is disabled for writing. The aggregate must be unmounted before it can be salvaged. The aggregate is disabled because there were I/O errors to the aggregate or because the zFS kernel had an internal error that affected the aggregate structures on disk. It is important to successfully unmount the aggregate so that no corrupted structures exist in the zFS kernel for the aggregate and to allow log file recovery to run. Doing this ensures that the aggregate starts without errors on the next mount.

aggregate_name

The name of the aggregate

System action

The salvage operation terminates.

Administrator response:

Unmount the aggregate and then run **ioeagslv** utility against the aggregate.

Severity:

svc_c_sev_error

IOEZ00701E

Could not read 8K page *block_number* from aggregate *aggregate_name*.

Explanation

An error occurred that prevented the salvage program from reading the specified block. This probably means that the aggregate cannot be repaired, or possibly that significant amounts of information will be lost in the aggregate.

block_number

block number

aggregate_name

The name of the aggregate

System action:

The salvage operation terminates.

Administrator response:

If the problem is due to an I/O error, zFS issues additional messages. Correct the I/O errors before proceeding.

Severity:

svc_c_sev_error

IOEZ00702E

Errors encountered verifying AFL on aggregate *aggregate_name*.

Explanation

An error occurred while the salvage program was verifying the AFL (aggregate file system list) on the aggregate. The aggregate cannot be salvaged. This message is intended for IBM service personnel and is preceded by messages that indicate what errors occurred while verifying the AFL.

aggregate_name

The name of the aggregate

System action

The salvage operation terminates.

Administrator response:

Restore the aggregate from a backup copy.

Severity:

svc_c_sev_error

IOEZ00703E

Corruption: *description*.

Explanation

The salvage program found that the aggregate is corrupted. The description provided in the message is intended for IBM service. Additional messages indicate the action taken by the salvage program; for example, whether the salvage program continues or not based on the severity of the error.

description

Text describing the corrupted data

System action:

The salvage operation continues.

Administrator response:

Examine subsequent messages provided by the salvage utility to determine whether action is required.

Severity:

svc_c_sev_error

IOEZ00704E**Expected: text.****Explanation**

The salvage program found that the aggregate is corrupted. A preceding message shows the corrupted values in a region of the aggregate. This message shows the expected values in the specified region of the aggregate. This message is intended for IBM service. Additional messages follow that indicate the action taken by the salvage utility.

In the message text:

text

Text describing the expected values

System action:

The salvage operation continues.

Administrator response:

Examine subsequent messages provided by the salvage utility to determine whether action is required.

Severity:

svc_c_sev_error

IOEZ00705I

Formatted *aggregate_version* aggregate size *formatted_size* 8K blocks, data set size *VSAM_size* 8K blocks, *HFS_migration_status* migrated from HFS.

Explanation

This message displays the aggregate version, formatted size, in 8 K blocks, of the aggregate (the amount that the zFS format utility has formatted), the current initialized (used) size of the data set that contains the zFS aggregate, and the HFS migration state of the aggregate.

aggregate_version

The aggregate version number.

formatted_size

The formatted size.

HFS_migration_status

Was or was not migrated from HFS.

VSAM_size

The size of the VSAM data set,

System action:

The salvage operation continues.

Severity:

svc_c_sev_notice

IOEZ00707I

Log file size *number_blocks* 8K blocks, verified correct.

Explanation

The log file was found to be correct during an aggregate salvage operation.

number_blocks

The size of the log file in 8 KB blocks

System action:

The salvage operation continues.

Severity:

svc_c_sev_notice

IOEZ00708E

Log file on aggregate *aggregate_name* is corrupted.

Explanation

The salvage operation found that the log file on the aggregate is corrupted. The salvage operation cannot complete and the aggregate will need to be restored from a backup. This message is intended for IBM service personnel and is preceded by messages that indicate what errors occurred while verifying the log file.

In the message text:

aggregate name

The name of the aggregate

System action:

The salvage operation ends.

Administrator response:

Restore the aggregate from a backup copy.

Severity:

svc_c_sev_error

IOEZ00709I**Bitmap size *number_blocks* 8K blocks, verified correct.**

Explanation

An aggregate salvage operation found the bitmap to be correct.

In the message text:

number_blocks

The size of the bitmap in 8-KB blocks

System action:

The salvage operation continues.

Severity:

svc_c_sev_notice

IOEZ00710E**Bitmap on aggregate *aggregate_name* is corrupted.**

Explanation

The block allocation bitmap on the aggregate is corrupted and the salvage operation cannot be completed. The aggregate must be restored from a backup. This message is intended for IBM service personnel and is preceded by messages that indicate what errors occurred while the salvage operation was verifying the bitmap structure.

aggregate_name

The name of the aggregate

System action:

The salvage operation ends.

Administrator response:

Restore the aggregate from a backup copy.

Severity:

svc_c_sev_error

IOEZ00711I**Aggregate *aggregate_name* successfully *operation_name*.**

Explanation

An aggregate was successfully verified or repaired by the **zfsadm salvage** command.

aggregate_name

The name of the aggregate.

operation_name

The operation that was performed.

System action:

The program ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00712E	Could not salvage aggregate <i>aggregate_name</i> because it could not be found.
-------------------	---

Explanation

The specified aggregate either does not exist, cannot be accessed, or is not a zFS aggregate.

aggregate_name

The name of the aggregate.

System action:

The program ends.

Administrator response:

Provide the name of a zFS aggregate and try the program again.

Severity:

svc_c_sev_notice

IOEZ00713E	Error <i>ErrorCode</i> reason code <i>ReasonCode</i> received salvaging aggregate <i>aggregate_name</i>.
-------------------	---

Explanation

An aggregate could not be salvaged due to an unexpected error, or possibly the aggregate was busy with another command or operation.

aggregate_name

The name of the aggregate.

ErrorCode

The error code that was received.

ReasonCode

The reason code that was received.

System action:

The program ends.

Administrator response:

If there are no other commands or operations running against the aggregate, and the problem continues, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00715E	Open for aggregate <i>aggregate_name</i> failed with code=<i>ErrorCode</i> reason code=<i>ReasonCode</i>.
-------------------	--

Explanation

An attempt to open the specified aggregate failed with the specified return code and reason code.

aggregate_name

The name of the aggregate.

ErrorCode

The error code that was received.

ReasonCode

The reason code that was received.

System action:

The program ends.

Administrator response:

Correct the problem based on the reason code. If the problem cannot be resolved, contact your service representative.

Severity:

svc_c_sev_error

IOEZ00718I**ZFS is low on recovery tasks.****Explanation:**

zFS has determined that 80%, or more, of the recovery tasks in the address space are actively performing recovery actions.

System action:

The program continues.

Administrator response:

A large number of zFS operations have been canceled, causing recovery actions to take place. Either these actions are not being done in a timely manner, or there is a system problem that will soon drastically affect the ability of zFS to properly perform its operations. You should begin looking for the cause of all these cancel actions.

Severity:

svc_c_sev_warning

IOEZ00719I**ZFS is critically low on recovery tasks.****Explanation:**

zFS has determined that 95%, or more, of the recovery tasks in the address space are actively performing recovery actions. While in this condition, all zFS operations will be failed with return code EBUSY (114) and reason EF8D6A2A. This will continue until the situation is resolved.

System action:

The program continues.

Administrator response:

A large number of zFS operations have been canceled, causing recovery actions to take place. Either these actions are not being done in a timely manner, or there is a system problem that will soon drastically affect the ability of zFS to properly perform its operations. You should begin looking for the cause of all these cancel actions.

Severity:

svc_c_sev_warning

IOEZ00720I**Initializing system *SysName* will not be allowed to join the sysplex. It is not running sysplex=filesys.****Explanation**

There is at least one member in the sysplex running z/OS R13 or higher. zFS requires all other members to be either at z/OS R13 or higher or to be running sysplex=filesys.

In the message text:

SysName

Initializing system that is not running sysplex=filesys.

System action:

The program continues.

Administrator response:

zFS on the initializing system will terminate. To correct this problem, either remove all z/OS R13 or higher members from the sysplex or specify sysplex=filesys for the initializing system and retry the operation.

Severity:

svc_c_sev_warning

IOEZ00721I**Sysplex member *SysName* is not running sysplex=filesys. zFS on this initializing member will terminate.****Explanation**

The initializing member is running z/OS R13 or later. It has found that the named member is not running sysplex=filesys. zFS does not allow a z/OS R13 or higher member to join a sysplex where there are members that are not running sysplex=filesys. All z/OS R13 or later members run sysplex=filesys. z/OS R12 members must specify sysplex=filesys.

In the message text:

SysName

Active system that is not running sysplex=filesys.

System action

zFS initialization terminates

Administrator response:

zFS on the initializing z/OS R13 or later system will terminate. To correct this problem, either remove all z/OS R12 members that are not running sysplex=filesys or restart them with sysplex=filesys and retry the operation.

Severity:

svc_c_sev_warning

IOEZ00722I**Primary file system size size 8K blocks, verified correct.****Explanation**

The primary file system was found to be correct during an aggregate salvage operation.

In the message text:

size

The size of the primary file system in 8-KB blocks

System action:

The salvage operation continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00723E**Primary file system on aggregate *aggregate_name* is corrupted.****Explanation**

The primary file system is corrupted. See message IOEZ00734E or IOEZ00735I to determine if the aggregate has repairable (minor) or severe (major) corruptions. For repairable (minor) corruptions, run the salvage program without any options to repair the aggregate. For severe (major) corruptions, the salvage operation cannot be completed. The aggregate must be restored from a backup copy. This message is intended for IBM service personnel and is preceded by messages that indicate what errors occurred while the salvage operation was verifying the primary file system.

In the message text:

aggregate_name

The name of the aggregate

System action:

The salvage operation ends if the aggregate has severe (major) corruptions. The salvage operation continues if the aggregate has repairable (minor) corruptions.

Administrator response:

Restore the aggregate from a backup copy if the aggregate has severe (major) corruptions.

Severity:

svc_c_sev_error

IOEZ00724I **Could not automatically re-enable disabled aggregate *AggrName*.**

Explanation

An attempt was made to dynamically re-enable a disabled aggregate, but this failed. This could be due to the aggregate containing more than one file system, or necessary storage is not available to re-enable the file system, or the re-enablement failed for some reason.

In the message text:

AggrName

Name of the disabled aggregate that could not be remounted

System action:

The program continues.

Administrator response:

If message IOEZ00662I “IOEZ00662I” on page 112 or IOEZ00663I “IOEZ00663I” on page 113 is present, then refer to the Administrator Action section of the message for the required steps to take. If not, then the aggregate may have to be manually unmounted and mounted.

Severity:

svc_c_sev_error

IOEZ00725I **Automatic re-enablement of file system *FileSys* complete.**

Explanation

After the owning aggregate for the specified file system became disabled, the file system was successfully re-enabled.

In the message text:

FileSys

Name of the file system

System action:

The program continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00726I **Waiting for recovery of *number* in-progress operations for filesystem *FilesystemName*.**

Explanation

The zFS kernel program IOEFSKN encountered a terminating exception. The zFS kernel program will wait up to 5 minutes for *number* user file requests to recover before proceeding with kernel program termination. The controller program IOEFSCM will automatically restart the zFS kernel program after the user file requests have completed their recovery. If the user file operations recovery is not completed the controller program will terminate the kernel program and request that z/OS UNIX System Services restart the zFS PFS.

In the message text:

number

The number of in-progress operations

FilesystemName

The name of the file system

System action:

The program continues.

Severity:

svc_c_sev_notice

IOEZ00727I**Pre-processing *Numvnodes* vnodes for file system restart.****Explanation**

zFS has restarted after a terminating exception, and is preprocessing the list of vnodes present at the time of zFS termination. A vnode is the internal structure that represents a file or directory in zFS. Because z/OS UNIX has direct references to these vnodes, zFS has to maintain them across an internal restart. zFS has to inspect each vnode; if z/OS UNIX is no longer referencing it, then it is discarded. If it is still in use by z/OS UNIX, then it is re-initialized and will be re-cached into the zFS vnode cache for later processing after its corresponding file system is internally remounted.

In the message text:

Numvnodes

Number of vnodes that were present at zFS termination.

System action:

The program continues.

Administrator response:

None; the internal error would be reported by prior messages and dumps.

Severity:

svc_c_sev_notice

IOEZ00728I**Resuming user operations for file system *FileSys* Waking *NumTasks* waiting tasks Recached *Numvnodes* vnodes for this file system.****Explanation**

zFS has restarted after a terminating exception, and has re-readied the named file system for operation; user operations are resumed for this file system.

In the message text:

FileSys

File system being resumed.

NumTasks

Number of tasks that were waiting to access the file system.

Numvnodes

Number of vnodes that had to be recached into the zFS vnode cache for this file system.

System action:

The program continues.

Administrator response:

None; the internal error would be reported by prior messages and dumps.

Severity:

svc_c_sev_notice

IOEZ00729I**Verification of aggregate *aggregate_name* started.**

Explanation

A salvage operation has started to verify the aggregate with the name *aggregate_name*.

In the message text:

aggregate_name

The name of the aggregate.

System action:

The salvage operation continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00730I

Verification of aggregate *aggregate_name* completed, no errors found.

Explanation

A salvage operation has finished verifying the aggregate with the name *aggregate_name* and found no errors.

In the message text:

aggregate_name

The name of the aggregate.

System action:

The salvage operation ends.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00731I

zFS internal restart complete.

Explanation:

zFS has restarted after a terminating exception, and has re-readied all mounted file systems for operation, all user operations are resumed and zFS is fully operational.

System action:

The program (IOEFSCM) continues.

Administrator response:

None.

Severity:

svc_c_sev_notice

IOEZ00733I

Verification of aggregate *aggregate_name* completed and the aggregate is in error.

Explanation

A salvage operation has finished verifying the aggregate with the name *aggregate_name*. It found some errors with the aggregate. The errors can be repaired.

In the message text:

aggregate_name

The name of the aggregate

System action:

The salvage operation ends.

Administrator response:

This message should have been preceded by other messages identifying the errors that were found. Run the salvager program with default options to repair the aggregate.

Severity:

svc_c_sev_notice

IOEZ00734E

Verification of aggregate *aggregate-name* stopped with return code *ReturnCode* due to major error(s).

Explanation

The salvager program encountered one of the following errors during verification and could not continue. The return code is *ReturnCode*. If the return code is EIO(122), the salvager program could not read or write the DASD. If the return code is ENOMEM(132), the salvager program ran out of storage. If the return code is EMVSERR(157), the salvager program had an internal error. A return code of 12 indicates the salvager program found an unrepairable error in the disk structures.

In the message text:

aggregate-name

The name of the aggregate

ReturnCode

The catalog service return code.

System action:

The salvage operation terminates.

Administrator response:

Use the return code to determine why salvager failed. If the return code is ENOMEM(132), then try to increase MEMLIMIT of the job and rerun salvager. If the return code is EMVSERR(157), then contact IBM service personnel. If the return code is 12, then restore the aggregate from a backup copy.

Severity:

svc_c_sev_error

IOEZ00735I

Salvage found *number* minor errors in aggregate *aggregate_name*.

Explanation

A salvage operation found *number* repairable errors on the aggregate named *aggregate_name*.

In the message text:

number

The number of errors found

aggregate_name

The name of the aggregate

System action:

The salvage operation ends.

Administrator response:

Run the salvager program with default options to repair the aggregate.

Severity:

svc_c_sev_notice

IOEZ00736I

zFS kernel initiated at *StartTime* Current status: *Status* Internal restart count *RestartCount* *RestartTime*.

Explanation

This message is a response to the MODIFY ZFS,QUERY,STATUS command. It shows the time (in GMT) when the zFS address space was started, the current status of the zFS address space, how many times zFS has internally restarted due to a severe internal error or due to the MODIFY ZFS,ABORT command, the time (in GMT) when it was last restarted if RestartCount is greater than zero.

In the message text:

StartTime

GMT date/time when the zFS address space was created

Status

Current status of zFS: initializing, shutting down, aborting, active, or internal remount

RestartCount

Number of times zFS has internally restarted

RestartTime

GMT date and time of the last zFS internal restart, if the internal restart count is greater than zero.

System action

The program (IOEFSCM) continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00739I	Salvage processed <i>d_pages</i> directory pages, <i>anodes</i> anodes, <i>i_blocks</i> indirect blocks and <i>a_pages</i> anode table pages.
-------------------	--

Explanation

The salvage operation displays the numbers of directory pages, anodes, indirect blocks and anode table pages it found.

In the message text:

d_pages

The number of directory pages found

anodes

The number of anodes found

i_blocks

The number of indirect blocks found

a_pages

The number of anode table pages found

System action

Salvage operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00740E	Internal remount of file system <i>FileSys</i> failed.
-------------------	---

Explanation

zFS has restarted and an attempt to internally remount the named file system has failed.

In the message text:

FileSys

Name of the file system.

System action

The program continues.

Administrator response

Remount the file system without switching modes. If the problem persists, restart ZFS.

Severity

svc_c_sev_error

IOEZ00741I

zFS is currently running at sysplex=filesys. You have specified sysplex=on. You should change your zFS configuration options file specification to sysplex=filesys and also specify sysplex_filesys_sharemode=rwshare.

Explanation

zFS in z/OS V1R13 or later always runs sysplex=filesys, independently of your sysplex= specification in your IOEFSPRM. In addition, since you have specified sysplex=on and you have not specified sysplex_filesys_sharemode, zFS is running with sysplex_filesys_sharemode=rwshare in order to be similar to sysplex=on (all zFS read-write file systems are mounted sysplex-aware by default).

System action

The program continues.

Administrator response

Update the zFS configuration options file, as indicated in the message text.

Severity

svc_c_sev_notice

IOEZ00742I

zFS is currently running at sysplex=filesys. You have specified sysplex=off. You should change your zFS configuration options file specification to sysplex=filesys.

Explanation

zFS is running sysplex=filesys. Specify sysplex=filesys so that your IOEFSPRM is not misleading. In addition, since you specified sysplex=off and you did not specify sysplex_filesys_sharemode, you should verify that the default sysplex_filesys_sharemode=norwshare is what you want.

System action

The program continues.

Administrator response

Verify the settings in the zFS configuration options file, as indicated in the message text.

Severity

svc_c_sev_notice

IOEZ00743I

zFS is currently running at sysplex=filesys. You have not specified sysplex=filesys. You should change your zFS configuration options file specification to sysplex=filesys.

Explanation

zFS in z/OS V1R13 or later always runs sysplex=filesys, independently of your sysplex= specification in your IOEFSPRM. In addition, you have specified sysplex_filesys_sharemode, so zFS is running with that option.

System action

The program continues.

Administrator response

Change the settings in the zFS configuration options file, as indicated in the message text.

Severity

svc_c_sev_error

IOEZ00744E

Time limit exceeded for recovery of in-progress operations for filesystem *FileSys*.

Explanation

zFS has taken a terminating exception, and the zFS kernel program IOEFSKN has waited 5 minutes for in-progress user file operations to recover before proceeding with IOEFSKN termination. Since recovery was not performed within the 5 minute window, IOEFSKN termination will proceed but zFS will stop and be restarted by z/OS UNIX instead of internally restarting.

In the message text:

FileSys

Name of the file system.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_error

IOEZ00745E

Automatic re-enablement of file system *FileSys* failed, rc=*ReturnCode* rsn=*ReasonCode*.

Explanation

The aggregate containing the specified file system has become disabled. zFS attempted to internally remount the file system with samemode. This internal remount failed with the specified values. zFS will attempt this remount several times. If it continues to fail, then message IOEZ00746E will be issued.

In the message text:

FileSys

Name of the file system.

ReturnCode

Return code of the remount request.

ReasonCode

Reason code of the remount request.

System action

The salvage operation terminates.

Administrator response

None.

Severity

The program continues.

IOEZ00746E

Automatic re-enablement of file system *FileSys* halted after repeated failures.

Explanation

zFS has repeatedly attempted to internally remount the disabled aggregate with samemode or move the aggregate to another sysplex member, and each attempt has failed. zFS will no longer attempt to remount or move it.

In the message text:

FileSys

Name of the file system.

System action

The salvage operation terminates.

Administrator response

You should unmount the file system, and then run `ioeagslv` against this aggregate, as it could be corrupted. Then, mount the file system again.

Severity

The program continues.

IOEZ00747I

Automatically re-enabling file system *FileSys*.

Explanation

zFS is attempting to internally remount the disabled aggregate with samemode.

In the message text:

FileSys

Name of the file system.

System action

The program continues.

Administrator response

The aggregate is temporarily disabled due to an internal zFS error. File requests will fail while the aggregate is disabled. zFS will attempt to re-enable the aggregate automatically. When the aggregate is re-enabled, the file can be closed and reopened and the request can be attempted again. If zFS cannot re-enable the aggregate automatically, you will need to unmount/mount or remount the aggregate before the file request can be attempted again.

Severity

svc_c_sev_notice

IOEZ00750E	Support for multi-file system aggregates has been removed; attach or mount of <i>aggrname</i> is denied.
-------------------	---

Explanation

Multi-file system aggregates are no longer supported.

In the message text:

aggrname

The name of the aggregate.

System action

The salvage operation terminates.

Administrator response

Specify a compatibility mode aggregate and retry the operation.

Severity

The program continues.

IOEZ00751I	<i>mountparm</i> is no longer a supported mount parm. It is ignored.
-------------------	---

Explanation

The specified mount parameter is no longer supported.

In the message text:

mountparm

The mount parameter that is no longer supported.

System action

The mount parameter is ignored.

Administrator response

Either ignore this message or remove the mount parameter specified.

Severity

svc_c_sev_notice

IOEZ00752E	Verification of aggregate terminated because error limit <i>max_errors</i> exceeded.
-------------------	---

Explanation

A salvage operation encountered more errors than the salvage parameter MAX_ERRORS allows, and cannot continue.

In the message text:

max_errors

The value of the MAX_ERRORS salvage parameter

System action

The salvage operation ends.

Administrator response

Restore the aggregate from a backup or increase the value of the MAX_ERRORS parameter in the IOEFSPRM file.

Severity

svc_c_sev_error

IOEZ00753I Salvage is repairing anode table pages.**Explanation**

A salvage operation found corrupted anode table pages and is repairing them.

System action

The salvage repair operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00754I Salvage is repairing partially free list.**Explanation**

A salvage operation found that the anode table partially free list is corrupted, and is repairing it.

System action

The salvage repair operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00755I Salvage is repairing the ZLC list.

Explanation

A salvage operation found that the zero link count list (ZLC) is corrupted and is repairing it.

System action

The salvage repair operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00756I

Salvage is repairing file, directory, and ACL objects.

Explanation

A salvage operation found errors with one or more objects in the file system and is repairing those objects.

System action

The salvage repair operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00757I

Salvage is repairing bitmap pages.

Explanation

A salvage operation found errors in the contents of the bitmap, and is repairing the pages of the bitmap found to be in error.

System action

The salvage repair operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00758I

Salvage repaired *anode* anode table pages, *p_free* partially free pages, *zlc* *zlc* pages, *indirect* indirect blocks.

Explanation

A salvage operation repaired the numbers of anode table pages, partially free pages, zlc pages and indirect blocks indicated.

In the message text:

anode

The number of anode table pages repaired

p_free

The number of partially free pages repaired

zlc

The number of zlc pages repaired

indirect

The number of indirect blocks repaired

System action

The salvage operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00759I

Salvage found *number* corruptions of security attributes.

Explanation

A salvage operation found damaged security objects stored on disk; for example, ACLs and stored SAF security packets. After repair, some users' access to some file system objects could change: a user might have access to an object that was originally denied, or might not have access to an object that was originally allowed.

In the message text:

number

The number of security objects that were damaged

System action

The salvage operation continues.

Administrator response

Based on the file system being repaired and the level of security required for that file system, decide whether to use the repaired file system.

Severity

svc_c_sev_notice

IOEZ00760I

No IOEZPRM DD specified. Parmlib search being used.

Explanation

The JCL used to run the salvager program did not include a IOEZPRM DD statement. Therefore, the parmlib concatenation is being searched for the IOEPRMxx members that contain the zFS configuration options for your system. The configuration options can include options for the salvager program.

System action

The salvage operation continues.

Administrator response

Based on the file system being repaired and the level of security required for that file system, decide whether to use the repaired file system.

Severity

svc_c_sev_notice

IOEZ00761E Specified PARM too long, length = *length*.

Explanation

The PARM string is too long. The length of the specified string is given in the message. The maximum length is 1024 characters.

In the message text:

length

The length of the string.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00762E Error in PARM specification *string*.

Explanation

There is an error in the PARM, as specified in the message. In the message text:

string

The PARM string.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00763E Suffix at position *position* too long.

Explanation

The specified suffix starting in the specified column position is too long.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00764E**Comma or right paren expected at position *position*.****Explanation**

Either a comma or a right parenthesis should have been in the specified column position, but is not.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00765E**Suffix at position *position* not two alphanumeric chars.****Explanation**

The specified suffix should be two characters long, but is not. The suffix specification should be in the specified column position.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00766E

No suffixes found at position *position*.

Explanation

No valid parmlib suffix specification was found at the specified position.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00767E

No closing paren found at position *position*.

Explanation

A closing parenthesis was not found where one was expected, in the specified column position.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00768E

No open paren found at position *position*.

Explanation

Parmlib suffix specifications require parentheses around them. There is no opening parenthesis found. It is expected in the specified column position.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00769E No equals sign found at position *position*.

Explanation

Parmlib suffix specifications require an equals sign. It is expected in the specified column position.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00770E Extra characters were found at position *position*.

Explanation

Extra characters are not allowed. They start in the specified column.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00771E PRM not found at position *position*.

Explanation

No PRM was found for parmlib suffix specifications. It is expected in the specified column.

In the message text:

position

The position of the error.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00773E Failed to initialize storage, ending.

Explanation

Failed to set up for storage management. The program terminates.

System action

The program ends.

Administrator response

Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

svc_c_sev_error

IOEZ00774E Open for aggregate *aggregate_name* failed with code=*return_code*. Aggregate may be attached.

Explanation

An attempt to open the specified aggregate failed with the specified return code. The problem might be that the aggregate is already mounted or attached.

In the message text:

aggregate_name

The name of the aggregate

return_code

The return code

System action

The salvage operation ends.

Administrator response

Use the **zfsadm lsaggr** command to see if the aggregate is attached. If it is, use the UNMOUNT command to unmount the file system, and resubmit the job. If the problem persists, or if the aggregate is not attached, contact your service representative.

Severity

svc_c_sev_error

IOEZ00775E

The required SYSPRINT DD card has a problem code=*return_code*.

Explanation

Either the SYSPRINT DD statement is missing from the JCL used to run the salvager program, or the salvager program has found a problem with the SYSPRINT DD statement.

In the message text:

return_code

The return code

System action

The salvage operation ends.

Administrator response

Correct the JCL and resubmit the job. If the error persists, or you cannot find an error in the SYSPRINT DD statement, contact your service representative.

Severity

svc_c_sev_error

IOEZ00776I

Salvage is repairing the totally free page stack.

Explanation

A salvage operation is in the process of repairing the totally free page stack.

System action

The salvage repair operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00777A

**Restarting exception *a_code* occurred, reason *r_code* abend psw
<*PSW_1 PSW_2*> <*PSW_3 PSW_4*>.**

Explanation

A zFS program encountered an exception running in AMODE64. A dump is issued and the internal trace table is printed. If the zFS kernel abends, it stops and restarts. File systems remain mounted, but files that are currently open might have I/O errors. If severe exceptions are encountered during the restart, zFS will be terminated and

the system will display a BPXF032D message that must be replied to before zFS can be reinitialized. See [z/OS MVS System Codes](#) for more information about abend codes.

In the message text:

a_code

The z/OS abend code received.

r_code

The z/OS reason code received, or register 15 at time of abend.

PSW_1

The high half of the first doubleword of the abend PSW.

PSW_2

The low half of the first doubleword of the abend PSW.

PSW_3

The high half of the second doubleword of the abend PSW.

PSW_4

The low half of the second doubleword of the abend PSW.

System action

The program continues.

Administrator response

This message is followed by additional messages giving more information about the problem. Contact your service representative.

Severity

svc_c_sev_fatal

IOEZ00778I

General Registers Rreg_num: <r1 r2> <r3 r4> <r5 r6> <r7 r8>.

Explanation

A zFS program running in AMODE64 encountered an exception. This message shows the 64-bit registers at time of error. The message displays the contents of four registers, beginning with the register indicated by *reg_num*. The message is repeated, showing a different set of registers each time, until all 64-bit registers have been displayed.

In the message text:

reg_num

The number of the first register whose contents are displayed

r1

The high half of the 64-bit register *reg_num*

r2

The low half of the 64-bit register *reg_num*

r3

The high half of the 64-bit register *reg_num + 1*

r4

The low half of the 64-bit register *reg_num + 1*

r5

The high half of the 64-bit register *reg_num + 2*

r6

The low half of the 64-bit register *reg_num + 2*

r7

The high half of the 64-bit register *reg_num + 3*

r8

The low half of the 64-bit register *reg_num + 3*

System action

The program continues.

Administrator response

This message is followed by additional messages giving more information about the problem.

Severity

svc_c_sev_warning

IOEZ00780E

Aggregate *aggregate_name* contains zero or more than one file system and cannot be salvaged.

Explanation

The specified aggregate has either zero file systems or more than one file system. The salvager does not support multi-file system aggregates or HFS-compatibility mode aggregates with a clone (.bak file).

In the message text:

aggregate_name

The name of the aggregate

System action

The salvage operation ends.

Administrator response

If the aggregate has a .bak file, remove it and restart the salvager. If the aggregate is a multi-file system aggregate, you need to run a salvager program from a release prior to z/OS V2R1.

Severity

svc_c_sev_error

IOEZ00781I

Salvage repaired *directory* directory entries and *bitmap* bitmap pages.

Explanation

This message summarizes the number of directory entries and bitmap pages repaired by a salvage operation.

In the message text:

directory

The number of directory entries

bitmap

The number of bitmap pages

System action

The salvage operation ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00782I	Salvage has verified <i>verified_num</i> of <i>total_num</i> object in the <i>file_system_object_name</i> .
-------------------	---

Explanation

During a salvage operation, this message shows the progress of the operation by indicating how many items in the described file system object have been verified.

In the message text:

verified_num

The number of objects processed.

total_num

The number of objects in the file system structure.

object

A description of the object

file_system_object_name

A description of the file system structure that is being verified.

System action

The salvage operation continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00783E	Aggregate <i>aggregate_name</i> is damaged.
-------------------	---

Explanation

During a salvage operation, the specified aggregate was found to be damaged. A successful repair has not been completed.

In the message text:

aggregate_name

The name of the aggregate

System action

The salvage operation continues.

Administrator response

Try running the salvage operation again. If subsequent repair attempts are not successful, restore the aggregate from a backup copy.

Severity

svc_c_sev_error

IOEZ00784E

Non-terminating exception *AbendCode* occurred, reason *ReasonCode* abend psw *PSW1H PSW1L PSW2H PSW2L*.

Explanation

A zFS kernel sub-routine encountered an exception while running AMODE64. A dump is issued and the internal trace table is printed. This exception is non-terminating. The zFS kernel continues to run, though errors might occur for a file or file system.

In the message text:

AbendCode

The z/OS abend code received.

ReasonCode

The z/OS reason code received, or register 15 at time of abend.

PSW1H

The high half of the first doubleword of the abend PSW.

PSW1L

The low half of the first doubleword of the abend PSW.

PSW2H

The high half of the second doubleword of the abend PSW.

PSW2L

The low half of the second doubleword of the abend PSW.

System action

The program continues.

Administrator response

This message is followed by additional messages with more information about the problem. Contact your service representative.

Severity

svc_c_sev_error

IOEZ00785I

Log file contents corrupted, repairing log contents and continuing with salvage.

Explanation

The salvage program found that the content of the aggregate log file is wrong. Salvage resets the log file and continues its processing.

System action

The salvage operation continues.

Administrator response

None.

Severity

svc_c_sev_error

IOEZ00786I	Aggregate <i>aggregate_name</i> was not unmounted cleanly. Last update time stamp=<i>time_stamp</i> sysname=<i>system_name</i>
-------------------	---

Explanation

At the beginning of the salvager run, the aggregate was found to have not been unmounted cleanly after last use. The time stamp indicates the last periodic update time. The system name is where the last update occurred.

In the message text:

aggregate_name

The name of the aggregate that was not unmounted cleanly.

time_stamp

The time stamp indicating the approximate time that zFS was canceled or the system went down.

system_name

The name of the system where the last update occurred.

System action

The salvage operation continues.

Severity

svc_c_sev_notice

IOEZ00787I	Repair of aggregate <i>aggregate_name</i> completed. All minor errors have been repaired.
-------------------	--

Explanation

During the salvage run, the aggregate was found to be damaged and a repair of all minor errors has been completed. This message does not indicate whether the repair was successful.

In the message text:

aggregate_name

The name of the aggregate that was repaired

System action

The salvage operation ends.

Severity

svc_c_sev_notice

IOEZ00788E	Open for aggregate <i>aggregate_name</i> failed with code=<i>return_code</i>. The aggregate was not found.
-------------------	---

Explanation

An attempt to open the specified aggregate failed with the specified return code. The aggregate could not be found.

In the message text:

aggregate_name

The name of the aggregate.

return_code

The return code.

System action

The program ends.

Administrator response

Ensure that the aggregate exists and resubmit the job.

IOEZ00789E

Open for aggregate *aggregate_name* failed with code=*return_code*. The aggregate is busy and may be mounted or attached.

Explanation

An attempt to open the specified aggregate failed with the specified return code.

In the message text:

aggregate_name

The name of the aggregate.

return_code

The return code.

System action

The program ends.

Administrator response

Check to see if the aggregate is attached via the `zfsadm lsaggr` command. Correct the problem and resubmit the job. If the problem persists, or there is no error, contact your service representative.

Severity

`svc_c_sev_error`

IOEZ00790I

Aggregate *aggrname* is already at version 5.

Explanation

A request has been made to convert an aggregate to version 5 and the aggregate is already version 5.

In the message text:

aggrname

The name of the aggregate.

System action

The program ends.

Administrator response

None.

Severity

`svc_c_sev_notice`

IOEZ00791I

Successfully converted directory *name* to version 5 format.

Explanation

The directory was successfully converted to version 5.

In the message text:

name
The name of the directory.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00792E **Failed to prepend the current working directory, error=errorCode reason=reasonCode.**

Explanation

Either BPX1GCW failed to return the current working directory or the length of the full directory names exceeds 1024.

errorCode
The error code.

reasonCode
The reason code.

System action

The program continues.

Administrator response

Correct the problem and retry the command. If the problem persists, contact your service representative.

Severity

svc_c_sev_error

IOEZ00793E **Error=ReturnCode reason=ReasonCode received while attempting to convert directory name to version 5 format.**

Explanation

The directory failed to convert to version 5 format.

In the message text:

ReturnCode
The error code.

ReasonCode
The reason code.

name
The directory name.

System action

The program continues.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00794E	Failed to get information for directory <i>name</i>, error=<i>returnCode</i> reason=<i>reasonCode</i>.
-------------------	---

Explanation

Failed to get information for the specified object.

In the message text:

name

The directory name.

returnCode

The error code.

reasonCode

The reason code.

System action

The program ends.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00795E	Open for aggregate <i>aggregate_name</i> failed with code=<i>return_code</i>.
-------------------	--

Explanation

An attempt to open the specified aggregate failed with the specified return code. The problem may be storage shortage or internal errors.

In the message text:

aggregate_name

The name of the aggregate.

return_code

The return code.

System action

The program ends.

Administrator response

If the return code indicates a storage shortage, then increase storage and re-submit the job. Message IOEZ00823 may also have been previously issued. If the problem persists or there are internal errors, contact your service representative.

Severity

svc_c_sev_notice

IOEZ00797I	Skip page <i>logical_page</i> (physical page <i>physical_page</i>) of v5 directory (index=<i>index_number</i>).
-------------------	---

Explanation

The indicated v5 directory page was too corrupted to continue. The salvager program will skip processing the rest of the entries on this page and move on to the next directory page.

In the message text:

logical_page

The logical page number of the directory page.

physical_page

The physical page number of the directory page.

index_number

The index number of the directory.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00798I	Found more than <i>number_entries</i> entries in directory (index=<i>index_number</i>) and skipping duplicate name checking for any additional entries.
-------------------	--

Explanation

Salvager has found more than specified number of valid entries in a directory indicated by the index number. To avoid any potential storage shortage, the salvager program will suspend duplicate name checking for any additional entries found in the same directory.

In the message text:

number_entries

The maximum number of entries allowed in duplicate name checking.

index_number

The index number of the directory.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00799A

zFS stand-alone utility ran out of storage above bar, total of *<high_megabytes,,low_megabytes>* bytes used, size attempted was *<high_storage,,low_storage>*, MEMLIMIT was *<high_memlimit,,low_memlimit>*, class *class_number* type *storage_type*.

Explanation

zFS stand-alone utility ran out of storage above the 2G bar when attempting to obtain *high_storage,,low_storage* bytes. *high_megabytes, low_megabytes* is the number of bytes that has been obtained and assigned by the utility above the bar already. *high_memlimit,,low_memlimit* shows the system defined MEMLIMIT value that the stand-alone utility uses. *class_number* and *storage_type* is information for IBM service personnel.

In the message text:

high_megabytes

The high-half of the total number of megabytes used by the zFS utility

low_megabytes

The low-half of total number of megabytes used by zFS utility.

high_storage

The high-half of amount of storage that zFS utility attempted to obtain.

low_storage

The low-half of amount of storage that zFS utility attempted to obtain.

high_memlimit

The high-half of the MEMLIMIT value used by zFS utility.

low_memlimit

The low-half of the MEMLIMIT value used by zFS utility.

class_number

The zFS utility sub-component class number.

storage_type

The zFS utility storage type.

System action

The program ends.

Administrator response

Try restarting the stand-alone utility with a bigger MEMLIMIT value. For more information about setting MEMLIMIT and how the system determines what value to use, see [z/OS MVS Programming: Extended Addressability Guide](#). If problem continues, contact your service representative.

Severity

svc_c_sev_fatal

IOEZ00800I

Conversion has processed *number* of *total* objects in the anode table.

Explanation

During the conversion of a zFS aggregate, the stand-alone conversion program has processed the indicated number of objects in the file system. The conversion process has to examine each object, and if its a directory that needs to be re-formatted, then it needs to convert it to the new version format.

In the message text:

number

The number of objects processed.

total

The number of objects in the file system structure.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00801I

Directory conversion begun for inodenum*ber*, directory size is *total* pages.

Explanation

During the conversion of a zFS aggregate, the stand-alone conversion program found that a directory with the indicated inode number is in the wrong format and needs to be converted. The original size of the directory is shown in pages to provide an indicator of the amount of pages that need to be read while doing the conversion.

In the message text:

number

The inode number.

total

The size of the directory in 8K pages.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00802I

Directory conversion completed, directory size now *total* pages.

Explanation

During the conversion of a zFS aggregate, the stand-alone conversion program finished re-formatting the directory and shows the new size of the directory in 8K pages.

In the message text:

total

The size of the directory in 8K pages.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00803I

***dirs* directories were found, *dirs4* converted, *v4Pages* directory pages converted to *v5Pages* directory pages.**

Explanation

During the conversion of a zFS aggregate, the stand-alone conversion program has found the indicated number of directories thus far in the anode table scan. *dirs4* indicates the number of directories that required conversion. *v4Pages* is the total number of pages that these directories occupied and *v5Pages* is the number of pages the newly converted directories occupy.

In the message text:

dirs

The number of directories found so far during the scan.

dirs4

The number of directories requiring conversion.

v4Pages

The number of 8K pages occupied by the directories in their original format.

v5Pages

The number of 8K pages now occupied by the converted directories.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00804E

**Failed to get file system information for *aggregate_name*,
error=*error_code* reason=*reason_code*.**

Explanation

An error occurred while trying to get information for the specified aggregate. See [Return codes \(errno\) in z/OS UNIX System Services Messages and Codes](#) for a description of the error code. See [Appendix A, “Reason codes,” on page 241](#) for a description of the reason code.

System action

If **zfsadm fsinfo** is used, the program ends.
If MODIFY FSINFO is used, the program continues.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00805A Aggregate size *size* 4K pages is too large to convert to version 4.

Explanation

During the conversion of a version 5 zFS aggregate to a version 4 zFS aggregate, the stand-alone conversion program has found the aggregate size is too large. A version 4 aggregate cannot be larger than 4TB in size.

In the message text:

size

The size of the aggregate dataset in 4K pages.

System action

The program ends.

Administrator response

Create a new version 4 aggregate and copy the data from the version 5 aggregate manually, in sections.

Severity

svc_c_sev_fatal

IOEZ00806A Directory size *sizeK* bytes, sub-directory count *SubDirs* cannot be converted.

Explanation

During the conversion of a version 5 zFS aggregate to a version 4 zFS aggregate, the stand-alone conversion program has found that a version 5 directory has more than 65535 sub-directories or is larger than 4GB.

In the message text:

size

The size of the directory in K bytes.

SubDirs

The number of sub-directories in the directory.

System action

The program ends.

Administrator response

Create a new version aggregate and copy the data from the version 5 aggregate manually, in sections.

Severity

svc_c_sev_fatal

IOEZ00807I

In a wait to verify that aggregate *aggregate_name* has no other writers. Member *member_name* in sysplex *sysplex_name* last wrote to the aggregate on *time*.

Explanation

When an aggregate is cleanly unmounted, zFS can quickly determine that there are no other members writing to the aggregate. This aggregate was not cleanly unmounted. Another possible reason is that it is a copy that was made without first unmounting the aggregate, and now the copy is being mounted. zFS needs to wait 65 seconds to verify that there are no other members writing to the aggregate.

In the message text:

aggregate_name

The name of the aggregate.

member_name

The member name.

sysplex_name

The sysplex name. This field might not be provided.

time

The time of the last recorded write to the aggregate.

System action

The program continues.

Administrator response

Wait for zFS to finish the mount operation. See [z/OS File System Administration](#) for information about unmounting zFS file systems before copying or moving an aggregate.

Severity

svc_c_sev_notice

IOEZ00808I

Successfully converted all directories in aggregate *aggrname* to version *version*.

Explanation

All directories in the aggregate were successfully converted to the specified version.

In the message text:

aggrname

The aggregate name.

version

The version.

System action

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00809I

Error *error_code* reason *reason_code* received while attempting to convert all directories in aggregate *aggrname* to version *version*.

Explanation

An unexpected error was received while converting the directories of the aggregate to the specified version.

In the message text:

error_code

The error code that was received.

reason_code

The reason code that was received.

aggrname

The aggregate name.

version

The version.

System action

The program ends.

Administrator response

Try to resolve the error and re-issue the command.

Severity

svc_c_sev_notice

IOEZ00810I

Successfully changed aggregate *aggrname* to version *version*.

Explanation

The aggregate version was successfully changed.

In the message text:

aggrname

The aggregate name.

version

The version.

System action

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00811E

Error *return_code* reason *reason_code* received while attempting to change aggregate *aggrname* to version *version*.

Explanation

An error was encountered changing the aggregate version.

In the message text:

return_code

The return code.

reason_code

The reason code.

aggrname

The aggregate name.

version

The version.

System action

The program ends.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00812I	Successfully changed aggregate <i>aggrname</i> to version 1.<i>version</i>.
-------------------	--

Explanation

The aggregate was changed to version 1.5 successfully.

In the message text:

aggrname

The aggregate name.

version

The version.

System action

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00813I	Error <i>error_code</i> reason <i>reason_code</i> received while attempting to change aggregate <i>aggrname</i> to version 1.<i>version</i>.
-------------------	---

Explanation

An unexpected error was received while changing the aggregate to the specified version.

In the message text:

error_code

The error code that was received.

reason_code

The reason code that was received.

aggrname

The aggregate name.

version

The version.

System action

The program ends.

Administrator response

Try to resolve the error and reissue the command.

Severity

svc_c_sev_notice

IOEZ00814E

Must specify either -path or -aggrversion.

Explanation

A **zfsadm convert** command specified neither -path nor -aggrversion. Either -path or -aggrversion must be specified.

System action

The program ends.

Administrator response

Retry the command specifying valid options.

Severity

svc_c_sev_error

IOEZ00815E

Enqueue *queue_name*, held by *system_name*, could not be obtained for shutdown processing. zFS will terminate abnormally.

Explanation

The specified enqueue is needed for shutdown processing and is not available. It is held by the named system. zFS will terminate abnormally.

In the message text:

queue_name

The enqueue name.

system_name

The name of the system holding the enqueue.

System action

zFS will terminate abnormally.

Administrator response

Try to determine why the enqueue is being held by the named system and take appropriate actions to get it released. For problem determination, see [z/OS File System Administration](#). If the problem persists, contact your service representative.

Severity

svc_c_sev_error

IOEZ00822E	Logsize too large with the number of physical disk blocks available.
-------------------	---

Explanation

The logsize specified is too large relative to the size of the VSAM linear dataset.

System action

The program ends.

Administrator response

Specify a smaller logsize and issue the command again.

Severity

svc_c_sev_error

IOEZ00823E	Dynamic allocation error occurred for dataset <i>DatasetName</i>. rc=<i>Retcode</i> err=<i>Reterr</i> info=<i>Retinfo</i> SMSreason=<i>Retrsn</i> func=<i>func</i>
-------------------	---

Explanation

A dynamic allocation error occurred while processing the specified dataset. *Retcode*, *Reterr*, *Retinfo* and *Retrsn* are the error codes supplied by dynamic allocation. *Function_code* describes the dynamic allocation function attempted.

In the message text:

DatasetName

The dataset name.

Retcode

The return code.

Reterr

The return error.

Retinfo

The return info.

Retrsn

The return reason.

func

The function code.

System action

The program continues.

Administrator response

Use the dynamic error codes to determine the reason for the failure. Correct the error and retry the operation.

Severity

svc_c_sev_error

IOEZ00824E	Aggregate too small. Increase the size of the VSAM linear dataset.
-------------------	---

Explanation

The size of the VSAM linear dataset is too small.

System action

The program continues.

Administrator response

Increase the size of the VSAM linear dataset and issue the command again.

Severity

svc_c_sev_error

IOEZ00825E	Aggregate contains file system. Specify the -overwrite option to format.
-------------------	---

Explanation

The VSAM linear data set already contains a zFS aggregate. Specify the -overwrite option to format the new aggregate.

FileSys

Name of the file system.

System action

The program continues.

Administrator response

If the existing file system is indeed supposed to be reformatted, then specify the -overwrite option and reissue the command.

Severity

svc_c_sev_error

IOEZ00826A	zFS stand-alone utility out of storage, total of <i>num</i>M bytes used, size attempted was <i>amount_storage</i>, class <i>class_number</i> type <i>storage_type</i>.
-------------------	---

Explanation

The zFS stand-alone utility ran out of storage when attempting to obtain *amount_storage* bytes. *num* is the number of bytes that has been obtained by the zFS utility. *class_number* and *storage_type* is information for IBM service personnel.

In the message text:

num

The total number of megabytes used by zFS.

amount_storage

The amount of storage that zFS attempted to obtain.

class_number

The zFS kernel subcomponent class number.

storage_type

The kernel storage type.

System action

The program ends.

Administrator response

Try restarting the stand-alone utility with a bigger MEMLIMIT value. For more information about setting MEMLIMIT and how the system determines what value to use, see [z/OS MVS Programming: Extended Addressability Guide](#). If the problem continues, contact your service representative.

Severity

svc_c_sev_fatal

IOEZ00827I

zfsadm query *name* completed successfully. There is no information to display.

Explanation

The indicated **zfsadm query** command completed successfully. There was no data to display.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00828E

***ProgName*: MODIFY command - *Parm* rejected: RC= EBUSY.**

Explanation

All zFS modify command threads are busy processing a long-running modify command, or zFS is being shut down.

In the message text:

ProgName

The program name.

Parm

The program parameters on the MODIFY command that were passed to zFS.

System action

The program continues.

Administrator response

If zFS is shutting down, then reissue the command after it is restarted. Otherwise, reissue the command after one of the other modify commands is completed. You can also try adding another modify command thread with the **zfsadm config** command.

Severity

svc_c_sev_error

IOEZ00829E	<i>ProgName</i> : Modify command failed: <i>Parm</i> failed: rc = return_code.
------------	--

Explanation

An error occurred while attempting a console command.

In the message text:

Parm

The program parameters on the MODIFY command that were passed to zFS.

ProgName

The program name.

return_code

The return code.

System action

The program continues.

Administrator response

Check the return code (rc) and respond accordingly.

Severity

svc_c_sev_error

IOEZ00830E	File system <i>FileSystem</i> has been quiesced for at least <i>seconds</i> by <i>Operation</i> operation. It was quiesced by <i>asid</i> on <i>SystemName</i> at time (GMT). It is zFS owned by <i>OwningSystem</i>. There may be other quiesced file systems.
------------	--

Explanation

The named file system is quiesced and there is at least one task waiting to access it. Access is allowed when the file system is unquiesced. There might be other quiesced file systems.

In the message text:

FileSystem

The name of the file system.

seconds

The minimum number of seconds that a file system must be quiesced before it can appear in this message.

Operation

The name of the operation that is causing the quiesce.

asid

Either the ASID and job that is causing the quiesce or an indication that zFS is causing the quiesce.

SystemName

The name of the system that is causing the quiesce.

time

The time of the quiesce.

OwningSystem

The name of the owning system.

System action

The program continues.

Administrator response

To determine all the file systems that are quiesced, use either the z/OS UNIX **zfsadm aggrinfo -long**, the F ZFS,QUERY,FILE, QUIESCED or the D OMVS,F operator commands. To unquiesce the file systems, use ISPF SHELL (ISHELL) to Reset unmount or quiesce from the Work with Mounted File Systems panel. Note that for a shared file system configuration, the attempt to unquiesce a quiesced sysplex root file system will fail if the authorized user ID that you use was defined with an OMVS HOME directory, and the user ID is not already active. If the condition persists, contact your service representative.

Severity

svc_c_sev_error

IOEZ00831E	stat() failed for pathname <i>name</i>, rc = <i>error_code</i>.
-------------------	--

Explanation

An error occurred while attempting to get status information of the specified path name.

System action

The program ends.

Administrator response

Check the error code and respond accordingly.

Severity

svc_c_sev_error

IOEZ00832E	Aggregate name <i>aggregate_name</i> mounted on pathname <i>path_name</i> is not a zFS aggregate.
-------------------	--

Explanation

The aggregate is not a valid zFS aggregate.

In the message text:

aggregate_name

The name of the aggregate.

path_name

The path name of the aggregate.

System action

The program ends.

Administrator response

Specify a zFS aggregate name.

Severity

svc_c_sev_error

IOEZ00833E	Failed to get information on pathname <i>path_name</i>. error=<i>error_code</i> reason=<i>reason_code</i>.
-------------------	---

Explanation

BPX1GMN failed to return file system information.

In the message text:

error_code

The error code that was received.

path_name

The path name of the aggregate.

reason_code

The reason code that was received.

System action

The program ends.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00834E	Mount for file system <i>Filesystem</i> failed because the dataset does not refer to the same dataset as the owner <i>Owner</i>. <i>LocalSystem</i> has the file system on VOLSER <i>LocalVolser</i> whereas the owner has it on <i>OwnerVolser</i>.
-------------------	---

Explanation

The named file system does not refer to the same dataset as on the owning system. The local mount fails and z/OS UNIX will function ship operations to the z/OS UNIX owner.

In the message text:

Filesystem

The name of the file system.

Owner

The name of the system that owns the file system.

LocalSystem

The name of the local system.

LocalVolser

The local VOLSER.

OwnerVolser

The owner VOLSER.

System action

The program continues.

Administrator response

If z/OS UNIX function shipping is not desired, ensure that the file system refers to the same dataset. Then unmount and remount.

Severity

svc_c_sev_error

IOEZ00835E	Mount for file system <i>Filesystem</i> failed because the dataset does not refer to the same dataset as the owner <i>Owner</i> . The file system is on VOLSER <i>Volser</i> but <i>LocalSystem</i> has it at CCHH <i>LocalCCHH</i> whereas the owner has it at <i>OwnerCCHHr</i>.
-------------------	---

Explanation

The named file system does not refer to the same dataset as on the owning system. The local mount fails and z/OS UNIX will function ship operations to the z/OS UNIX owner.

In the message text:

Filesystem

The name of the file system.

Owner

The name of the system that owns the file system.

Volser

The VOLSER.

LocalSystem

The name of the local system.

LocalCCHH

The local CCHH.

OwnerCCHHr

The owner CCHH.

System action

The program continues.

Administrator response

If z/OS UNIX function shipping is not desired, ensure that the file system refers to the same dataset. Then unmount and remount.

Severity

svc_c_sev_error

IOEZ00836I	<i>program_name</i>: MODIFY command - <i>command_name</i> accepted.
-------------------	--

Explanation

The MODIFY command *command_name* has been accepted.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00837E Invalid FSINFO sorting method *sorting_method*.

Explanation

The specified sorting method is invalid. Supported sorting methods are NAME, REQUESTS and RESPONSE.

System action

If **zfsadm fsinfo** is used, the program ends.

If MODIFY FSINFO is used, the program continues.

Administrator response

Correct the sorting method and try again.

Severity

svc_c_sev_error

IOEZ00838E Bad wildcard character after aggregate option *aggregate_name*.

Explanation

FSINFO only supports wildcard with * at the beginning or the ending of the aggregate name.

In the message text:

aggregate_name

The aggregate name with wildcard

System action

If **zfsadm fsinfo** is used, the program ends.

If MODIFY FSINFO is used, the program continues.

Administrator response

Correct the aggregate name and try again.

Severity

svc_c_sev_error

IOEZ00839E Invalid selection criteria after **-select** option.

Explanation

FSINFO encountered an invalid selection criteria.

System action

The program ends.

Administrator response

Correct the selection criteria and try again.

Severity

svc_c_sev_error

IOEZ00840E	FSINFO -system option is not allowed with -basic, -full, -owner or -reset option.
-------------------	--

Explanation

The FSINFO -system option is not allowed if the -basic, -full, -owner, or -reset options are specified.

System action

The program ends.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00841E	FSINFO -select or -exceptions option is not allowed with -path or -reset option.
-------------------	---

Explanation

FSINFO -select or -exceptions option is not allowed if either -path or -reset option is specified.

System action

The program ends.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00842E	FSINFO -sort option is not allowed with -reset option.
-------------------	---

Explanation

FSINFO -sort option is not allowed if the -reset option is specified.

System action

The program ends.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00843E	Failed to get file system information error=<i>error_code</i> reason=<i>reason_code</i>.
-------------------	---

Explanation

Getting file system information has failed due to the displayed return code and reason code.

In the message text:

error_code

The error code that was received.

reason_code

The reason code that was received.

System action

If **zfsadm fsinfo** is used, the program ends.

If MODIFY FSINFO is used, the program continues.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00844E	Invalid option <i>option_name</i>.
-------------------	---

Explanation

An invalid option was specified for the FSINFO command.

System action

The program continues.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00845E	Reset option cannot be used with the select, exceptions, or sort options.
-------------------	--

Explanation

The FSINFO command does not support the reset option if one of the select, exceptions, or sort options was also specified.

System action

The program continues.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00846E	Invalid criteria string <i>string</i> for FSINFO select option.
-------------------	--

Explanation

The select option of the FSINFO command contains an invalid criteria string.

In the message text:

string

The criteria string.

System action

The program continues.

Administrator response

Correct the criteria string and try again.

Severity

svc_c_sev_error

IOEZ00847E	Data set <i>aggregate_name</i> is not formatted as a zFS aggregate.
-------------------	--

Explanation

The displayed VSAM linear data set is not formatted as a zFS aggregate or its size is smaller than the minimum size of a zFS aggregate.

In the message text:

aggrname

The name of the aggregate.

System action

The salvage operation ends.

Administrator response

Ensure that the data set is a VSAM linear data set with share options (3,3) that is correctly formatted as a zFS aggregate. The size of the data set must be larger than the minimum zFS aggregate size.

Severity

svc_c_sev_error

IOEZ00848I	Could not get completed file system information error=<i>error_code</i> reason=<i>reason_code</i>.
-------------------	---

Explanation

Getting file system information is not completed due to the displayed return code and reason code. FSINFO displays as much information as it can.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00849I **FSINFO command done.**

Explanation

The FSINFO command completed processing.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00850I **File System Status:.**

Explanation

Indicates that the information following the message is the status of the zFS file system and the statistics for each file system displayed. There will be an IOEZ00849I message indicating that the display from this modify command is completed.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00851I **Processing zlc for *file_system*.**

Explanation

The specified file system is being attached. There are files in it that were marked as deleted while the file system was last being used. Those files had not yet been removed from disk. These files are now being removed from the disk. This is normal zero link count (zlc) processing. Removing a large number of zlc files or a large zlc file may take a significant amount of time to complete. During this time the file system may be unavailable for use.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00852E	Data set is not a VSAM linear data set.
-------------------	--

Explanation

The data set is not defined as a VSAM linear data set.

System action

The program ends.

Administrator response

Ensure that the data set is correctly defined as a VSAM linear data set.

Severity

svc_c_sev_error

IOEZ00853I	There are no free 8K blocks remaining in aggregate <i>aggrname</i>. Applications may fail.
-------------------	---

Explanation

Applications may be failing because there are no free 8K blocks in the aggregate. The aggregate must be grown or applications may continue to fail.

In the message text:

aggrname

The aggregate name.

System action

The program continues.

Administrator response

Use the **zfsadm grow** command to grow the aggregate.

Severity

svc_c_sev_notice

IOEZ00854I Salvage for *aggrname* successfully canceled.

Explanation

An online salvage for the aggregate was successfully canceled.

In the message text:

aggrname

The aggregate name.

System action

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00856E **FSINFO: select or exceptions option is not allowed if no wildcard is specified in aggregate name.**

Explanation

If a wildcard is not specified in the aggregate name string, you cannot use the select or exceptions options.

System action

If **zfsadm fsinfo** is used, the program ends.

If MODIFY FSINFO is used, the program continues.

Administrator response

Correct the error and try again.

Severity

svc_c_sev_error

IOEZ00857I No zFS aggregate is found matching the requested criteria.

Explanation

There were no zFS aggregates found that have the attributes or aggregate names specified with the command keywords.

System action

If **zfsadm** is used, the program ends.

If MODIFY is used, the program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00858E Invalid SMF interval string '*string*'.

Explanation

The parsed string was not valid. It did not have any of the required options.

In the message text:

string

The interval string. It must consist of one of the following options:

- ON
- ON,xx where xx is a number between 1 and 60
- OFF

System action

The system continues.

Administrator response

None.

Severity

svc_c_sev_error

IOEZ00859I FSINFO: reset has completed.

Explanation

Indicates that reset processing is done. This message does not guarantee that the specified data set exists or is a valid zFS aggregate.

System action

- If **zfsadm fsinfo** is used, the program ends.
- If **MODIFY FSINFO** is used, the program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00860E The attributes and mount parms for aggregate *aggrname* may not match.

Explanation

Mismatches occurred because errors were encountered while aggregate attributes or mount parms were being updated.

In the message text:

aggrname
The aggregate name.

System action

The program continues.

Administrator response

Issue a samemode remount to restore the aggregate attributes from the mount parms and clear this message.

Severity

svc_c_sev_error

IOEZ00861E	A request to <i>action_needed</i> attribute for aggregate <i>aggrname</i> failed because system <i>downlevel_system</i> is at release <i>downlevel_release</i> and needs to be at release <i>minimum_required_release</i> or above.
-------------------	--

Explanation

The request failed because there is a downlevel system in the sysplex.

In the message text:

action_needed
Action to take with the attribute

aggrname
The aggregate name.

downlevel_system
The downlevel system name.

downlevel_release
The release of the downlevel system.

minimum_required_release
The minimum release that is required.

System action

The program continues.

Administrator response

Try the operation again after removing all downlevel systems in the sysplex.

Severity

svc_c_sev_error

IOEZ00862E	Syntax error with aggregate - <i>rwshare</i> parameter - <i>parm_name</i>.
-------------------	---

Explanation

An aggregate attribute could not be updated because the `-rwshare` option of the **zfsadm chaggr** command is improperly specified. See [z/OS File System Administration](#) for a description of the valid values for `-rwshare`.

parm_name

The `-rwshare` parameter in error.

System action

The program ends.

Administrator response

Correct the value and try the command again.

Severity

svc_c_sev_error

IOEZ00866E	Aggregate <i>aggrname</i> was disabled <i>aggr_disabled</i> times and re-enabled <i>aggr_reenabled</i> times. zFS will try to re-enable the aggregate one more time after salvaging it.
-------------------	--

Explanation

zFS has re-enabled the aggregate a number of times by either moving the zFS owner or by remounting. But the aggregate continues to be disabled. This might indicate that a salvage is needed. zFS will initiate a salvage operation and then try one last time to re-enable the aggregate.

In the message text:

aggrname

The aggregate name.

aggr_disabled

The number of times the aggregate was disabled.

aggr_reenabled

The number of times the aggregate was re-enabled.

System action:

The program continues.

Administrator response

If the aggregate does become disabled again and message IOEZ00746E is displayed, unmount the aggregate and contact your service representative.

Severity

svc_c_sev_error

IOEZ00867I	Salvage for aggregate <i>aggrname</i> was interrupted.
-------------------	---

Explanation

Salvage for the aggregate was interrupted by a **zfsadm salvage -cancel** command or by an unmount.

In the message text:

aggrname

The aggregate name.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00870A	The local initializing system cannot coexist with system <i>system_name</i>. The local system terminates.
-------------------	--

Explanation

The release span between the two system is not supported. Initialization ends.

In the message text:

system_name
The name of the system.

System action

The program ends.

Administrator response

Ensure that you have the correct release and service levels of zFS in your sysplex.

Severity

svc_c_sev_fatal

IOEZ00872E	Error shrinking aggregate <i>aggrname</i>, error code=<i>error_code</i> reason code=<i>reason_code</i>.
-------------------	--

Explanation

The shrink of the aggregate failed due to an unexpected error, or possibly the aggregate ran out of free space or was busy with another command or operation.

In the message text:

aggrname
The aggregate name.

error_code
The error code.

reason_code
The reason code.

System action

The program ends.

Administrator response

Use message IOEZ00904E along with the return and reason codes to determine the reason for the failure, and try the command again if the shrink is still needed. If there are no other commands or operations running against the aggregate, and the problem continues, contact your service representative.

Severity

svc_c_sev_error

IOEZ00873I Aggregate *aggrname* successfully shrunk.

Explanation

An aggregate was successfully reduced in size by the **zfsadm shrink** command.

In the message text:

aggrname

The aggregate name.

System action

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00877I Aggregate *aggregate_name* is successfully decrypted.

Explanation

The aggregate was successfully decrypted by the **zfsadm** command.

aggregate_name

The name of the aggregate.

System action

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00879E Error encrypting aggregate *aggregate_name*, error code=*ErrorCode*
reason code=*ReasonCode*.

Explanation

The encryption of the aggregate failed due to an unexpected error, or possibly the aggregate was busy with another command or operation.

aggregate_name

The name of the aggregate.

ErrorCode

The error code that was received.

ReasonCode

The reason code that was received.

System action

The program ends.

Administrator response

Use the return and reason codes to determine the reason for the failure and try the command again. If there are no other commands or operations running against the aggregate and the problem continues, contact your service representative.

Severity

svc_c_sev_error

IOEZ00880E	Error decrypting aggregate <i>aggregate_name</i>, error code=<i>ErrorCode</i> reason code=<i>ReasonCode</i>.
-------------------	---

Explanation

The decryption of the aggregate failed due to an unexpected error, or possibly the aggregate was busy with another command or operation.

aggregate_name

The name of the aggregate.

ErrorCode

The error code that was received.

ReasonCode

The reason code that was received.

System action

The program ends.

Administrator response

Use the return and reason codes to determine the reason for the failure and try the command again. If there are no other commands or operations running against the aggregate and the problem continues, contact your service representative.

Severity

svc_c_sev_error

IOEZ00881I	<i>aggregate_name</i> is being shrunk.
-------------------	---

Explanation

zFS is in the process of shrinking the specified aggregate, which can take a long time. File and directory access is allowed throughout most of the shrink process. If a shutdown occurs or if the aggregate is unmounted, the shrink process is canceled. You can also use the **zfsadm shrink -cancel** command to cancel the shrink process.

In the message text:

program_name

The name of the zFS program.

file_name

The name of the trace output file.

System action:

The program continues.

Administrator response

Ensure that the **trace_dsn** parameter specifies a valid file name that can be opened for output.

Severity

svc_c_sev_error

IOEZ00885I Shrink request for aggregate *aggregate_name* successfully canceled.

Explanation

The shrinking of the specified aggregate was successfully canceled.

In the message text:

aggregate_name

The name of the aggregate.

System action:

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00886E The user is not authorized to run this command.

Explanation

Only the user who mounted the file system or those who have either a UID of 0 or READ access to the SUPERUSER.FILESYS.PFSCCTL profile in the UNIXPRIV class can issue this command.

System action:

The program ends.

Administrator response

None.

Severity

svc_c_sev_error

IOEZ00887E Aggregate *aggregate_name* is not found.

Explanation

The specified aggregate name is not found. It is either not a valid aggregate name or the name is incorrect.

In the message text:

aggregate_name

The name of the aggregate.

System action:

The program ends.

Administrator response

Specify the correct aggregate name and issue the command again.

Severity

svc_c_sev_error

IOEZ00888I

***aggregate_name* is being encrypted or decrypted.**

Explanation

zFS is in the process of encrypting or decrypting the specified aggregate, which can take a long time. File and directory access are allowed throughout most of the encrypt or decrypt process. If a shutdown occurs or if the aggregate is unmounted, the process is canceled. You can also use the **zfsadm encrypt -cancel** or **zfs decrypt -cancel** command to cancel the process.

In the message text:

aggregate_name

The name of the aggregate.

System action:

The program continues.

Administrator response

Allow the encryption or decryption process to continue. To cancel it, use the **zfsadm encrypt** or **zfsadm decrypt** command accordingly. See [z/OS File System Administration](#) for more details about encrypt or decrypt processing and how to cancel it.

Severity

svc_c_sev_notice

IOEZ00889E

Message ID %1\$d not found.

Explanation

zFS attempted to issue a message and the message could not be found.

messageID

The ID of the message zFS attempted to issue.

System action:

The program continues.

Administrator response

Contact your service representative.

Severity

svc_c_sev_error

IOEZ00890E

The local system failed with rc *return_code* and reason *reason_code* while retrieving the automove option for aggregate *aggregate_name*. The local system will not become the zFS owner of the aggregate.

Explanation

The automove option for the aggregate could not be retrieved. Without this information, the local system cannot determine whether it can become the zFS owner of the aggregate.

In the message text:

return_code

The return code.

reason_code

The reason code that was received.

aggregate_name

The name of the aggregate.

System action:

The program continues.

Administrator response

Contact your service representative.

Severity

svc_c_sev_error

IOEZ00891E

The local system failed with *return_code* and reason *reason_code* while retrieving the automove option for aggregate *aggregate_name*. The local system will not attempt to move ownership of the aggregate.

Explanation

The automove option for the aggregate could not be retrieved. Without this information, the local system cannot determine whether any remote system is a suitable owner. The aggregate will not be moved.

In the message text:

return_code

The return code.

reason_code

The reason code that was received.

aggregate_name

The name of the aggregate.

System action:

The program continues.

Administrator response

Contact your service representative.

Severity

svc_c_sev_error

IOEZ00892I**Aggregate *aggregate_name* encrypt or decrypt successfully canceled.**

Explanation

The encrypting or decrypting of the specified aggregate was successfully canceled.

In the message text:

aggregate_name

The name of the aggregate.

System action:

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00895E**Option -encrypt or -compress is only allowed if -version5 is specified.**

Explanation

Format with encryption or compression only applies to version 1.5 aggregates.

System action:

The program ends.

Administrator response

Specify the -version5 option in order to format with encryption or compression.

Severity

svc_c_sev_error

IOEZ00896I**Aggregate *aggregate_name* released *tracks* tracks, moved *blocks* blocks, and moved *fragments* fragments.**

Explanation

The specified number of tracks was successfully released from the specified zFS aggregate. The shrink processing moved the specified number of 8K blocks and 1K fragments.

In the message text:

aggregate_name

The name of the aggregate.

tracks

The number of tracks released.

blocks

The number of 8K blocks moved.

fragments

The number of 1K fragments moved.

System action:

The program continues.

Administrator response:

None.

Severity

svc_c_sev_notice.

IOEZ00897E

Error code *errorCode* reason *reasonCode* received while releasing space from aggregate *aggregate_name*. *blocks* blocks and *fragments* fragments moved.

Explanation

During a shrink operation, there was a failure releasing space from the specified aggregate data set. The error and reason codes are from the PARTREL service. The amount of space zFS will use in the aggregate is as specified by the shrink command options. However, the space in the unused portion of the data set is not released. Also shown is the number of 8K blocks and 1K fragments that were moved during the shrink processing.

In the message text:

errorCode

The PARTREL error code received.

reasonCode

The PARTREL or catalog diagnosis reason.

aggregate_name

The name of the aggregate.

blocks

The number of 8K blocks moved.

fragments

The number of 1K fragments moved.

System action:

The program continues.

Administrator response

The return and reason codes are described in [z/OS DFSMSdfp Diagnosis](#). The return code and diagnostic reason code specified is from the PARTREL service. Additional information for the PARTREL service can be found in [z/OS DFSMSdfp Advanced Services](#). Use this to determine the cause of the failure, and try another shrink operation. If the problem persists, contact your service representative.

Severity

svc_c_sev_error.

IOEZ00898I

***aggregate_name* is being compressed or decompressed.**

Explanation

zFS is in the process of compressing or decompressing the specified aggregate, which can take a long time. File and directory access are allowed throughout most of the compression or decompression process. If a shutdown occurs or if the aggregate is unmounted, the process is canceled. You can also use the **zfsadm compress -cancel** or **zfs decompress -cancel** command to cancel the process.

In the message text:

aggregate_name

The name of the aggregate.

System action:
The program continues.

Administrator response

Allow the compression or decompression process to end. To cancel it, use the **zfsadm encrypt** or **zfsadm decrypt** command accordingly. See *z/OS File System Administration* for more details about compression and decompression processing and how to cancel it.

Severity

svc_c_sev_notice

IOEZ00899I Aggregate *aggregate_name* is successfully compressed.

Explanation

An aggregate was successfully compressed by the **zfsadm** command.

In the message text:

aggregate_name
The name of the aggregate.

System action:
The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00900I Aggregate *aggregate_name* is successfully decompressed.

Explanation

An aggregate was successfully decompressed by the **zfsadm** command.

In the message text:

aggregate_name
The name of the aggregate.

System action:
The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00901E Error compressing aggregate *aggregate_name*, error code=*error_code* reason code=*reason_code*.

Explanation

The compression of the aggregate failed due to an unexpected error, or possibly the aggregate was busy with another command or operation.

In the message text:

aggregate_name

The name of the aggregate.

error_code

The error code that was received.

reason_code

The reason code that was received.

System action:

The program continues.

Administrator response

Use the return and reason codes to determine the reason for the failure and try the command again. If there are no other commands or operations running against the aggregate and the problem continues, contact your service representative.

Severity

svc_c_sev_error

IOEZ00902E	Error decompressing aggregate <i>aggregate_name</i>, error code=<i>error_code</i> reason code=<i>reason_code</i>.
-------------------	--

Explanation

The decompression of the aggregate failed due to an unexpected error, or possibly the aggregate was busy with another command or operation.

In the message text:

aggregate_name

The name of the aggregate.

error_code

The error code that was received.

reason_code

The reason code that was received.

System action:

The program continues.

Administrator response

Use the return and reason codes to determine the reason for the failure and try the command again. If there are no other commands or operations running against the aggregate and the problem continues, contact your service representative.

Severity

svc_c_sev_error

IOEZ00903I	Aggregate <i>aggregate_name</i> compress or decompress successfully canceled.
-------------------	--

Explanation

The compressing or decompressing of the specified aggregate was successfully canceled.

In the message text:

aggregate_name

The name of the aggregate.

System action:

The program ends.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00904E

**Aggregate *aggregate_name* shrink halted with error code=*error_code*.
Moved *blocks* blocks and *fragments* fragments.**

Explanation

The shrinking of the aggregate failed due to an unexpected error. The size of the aggregate has not been reduced. If the error code is 119 (EFBIG) or 133 (ENOSPC), there was insufficient free space to shrink the aggregate to the requested size. If active increase was allowed, the new size of the aggregate was increased back to its original size. The aggregate may not be a good candidate to be shrunk. If active increase was not allowed, file system activity used all the free blocks before the shrink process was completed. If the error code is 120 (EINTR), the shrink was interrupted by a shutdown, unmount force, or a shrink cancel request. The specified number of 8K blocks and 1K fragments were moved before shrink processing ended.

In the message text:

aggregate_name

The name of the aggregate.

error_code

The error code.

blocks

The number of 8K blocks moved.

fragments

The number of 1K fragments moved.

System action:

The program ends.

Administrator response

If the error code is 119 (EFBIG) or 133 (ENOSPC) and active increase was not allowed, try the **shrink** command with a larger new size. If active increase was allowed and you still want to shrink the aggregate, remove unwanted files and directories from the aggregate and try the **shrink** command again. For other error codes, try to determine the reason for the failure and reissue the **shrink** command. If the problem persists, contact your service representative.

Severity

svc_c_sev_error.

IOEZ00905E

File system *file_name* failed to mount or open due to unsupported structures on disk.

Explanation

An attempt was made to mount or use a file system that includes features or disk structures that are not supported at this release.

In the message text:

file_name

The name of the file.

System action:

The program continues.

Administrator response

Mount the file system only on a release that supports the file system structures.

Severity

svc_c_sev_error

IOEZ00906E

File system *file_name* connect rejected from system *system_name* due to unsupported structures on disk.

Explanation

An attempt was made to mount or use a file system that includes features or disk structures that are not supported at this release level. z/OS UNIX will function-ship operations for this file system to the z/OS UNIX owner.

In the message text:

file_name

The name of the aggregate.

system_name

The system name of the connecting system.

System action:

The program continues.

Administrator response

If the release level on this system should support the file system features and disk structures, contact IBM service.

Severity

svc_c_sev_error

IOEZ00907E

The local system cannot coexist with system *system_name* at release level *release_level*. *system_name* will terminate.

Explanation

The release span between the two system is not supported. The initializing system ends.

In the message text:

system_name

The name of the system.

release_level

The release level that cannot coexist with the local system.

System action:

The program continues.

Administrator response

Ensure that you have the correct release and service levels of zFS in your sysplex.

Severity

svc_c_sev_error

IOEZ00908I	An MVS catalog service call for <i>aggregate_name</i> failed with <i>rc=return_code</i> <i>rsn=reason_code</i>.
-------------------	--

Explanation

A catalog operation was required to determine whether the aggregate is eligible for certain zFS operations and it failed. For example, if a format at mount is desired or if encryption is desired for an existing aggregate, a catalog operation is required. The return and reason codes are from the MVS catalog service.

In the message text:

aggregate_name

The name of the aggregate.

return_code

The return code from the catalog service.

reason_code

The reason code from the catalog service.

System action:

The program continues.

Administrator response

Find the specified catalog return and reason codes documented in *z/OS MVS System Messages, Vol 6 (GOS-IEA)* under message IDC3009I, and determine the cause of the problem. After correcting the problem, try the operation again.

Severity

svc_c_sev_notice

IOEZ00942E	<i>aggregate_name</i> failed to mount read-only because recovery is required.
-------------------	--

Explanation

The mount for the aggregate failed because it was not cleanly unmounted the last time it was used. Aggregate recovery needs to be performed before it can be successfully mounted.

In the message text:

aggregate_name

The name of the aggregate.

System action

The program continues.

Administrator response

Mount the aggregate read/write, unmount it, and mount it read-only. Or, set the IOEFSPRM configuration variable ROMOUNT_RECOVERY=ON and mount the aggregate read-only.

Severity

svc_c_sev_error

IOEZ00943E

***aggregate_name* failed to mount because ICSF (Integrated Cryptographic Service Facility) is not started or cannot be accessed. Resolve the problem with ICSF and try the mount again.**

Explanation

The aggregate mount failed because the aggregate uses DFSMS VSAM encryption services but communication with ICSF has failed. ICSF may have issued messages that indicate the problem.

In the message text:

aggregate_name

The name of the aggregate.

System action

The program continues.

Administrator response

Check the system log for ICSF messages that indicate what the problem is and how to resolve it. Follow their recommendations and try mounting the aggregate again.

Severity

svc_c_sev_error

IOEZ00947A

Terminating exception *AbendCode* occurred, reason *ReasonCode* abend psw *PSW1* *PSW2*.

Explanation

A zFS kernel subroutine encountered an exception. A dump will be issued and the internal trace table will be printed. Severe problems encountered during initialization or during an internal restart will result in the termination of zFS and the system will display a BPXF032D message that must be replied to before zFS can be reinitialized. See [*z/OS MVS System Codes*](#) for more information about abend codes.

In the message text:

AbendCode

The z/OS abend code received.

ReasonCode

The zFS abend reason code received, or register 15 at time of abend.

PSW1

The first word of the abend PSW.

PSW2

The second word of the abend PSW.

System action

The program continues.

Administrator response

This message might be preceded by messages with more information about the problem. This message will be followed by additional informational messages. Contact your service representative.

Severity

svc_c_sev_fatal

IOEZ00948A

**Terminating exception *a_code* occurred, reason *r_code* abend psw
<PSW_1 PSW_2> <PSW_3 PSW_4>.**

Explanation

A zFS kernel subroutine encountered an exception while running in AMODE64. A dump will be issued and the internal trace table will be printed. Severe problems encountered during initialization or during an internal restart will result in the termination of zFS and the system will display a BPXF032D message that must be replied to before zFS can be reinitialized. See [z/OS MVS System Codes](#) for more information about abend codes.

In the message text:

a_code

The z/OS abend code received.

r_code

The z/OS reason code received, or register 15 at time of abend.

PSW_1

The high half of the first doubleword of the abend PSW.

PSW_2

The low half of the first doubleword of the abend PSW.

PSW_3

The high half of the second doubleword of the abend PSW.

PSW_4

The low half of the second doubleword of the abend PSW.

System action

The program continues.

Administrator response

This message might be preceded by messages that indicate more about the problem. This message is followed by additional informational messages. Contact your service representative.

Severity

svc_c_sev_fatal

IOEZ00949E

Mount for file system *filesystem_name* failed because the dataset does not refer to the same dataset as the owner *owner_name*. The client extent *volser*, *low_CI*, *hi_CI* cannot be found within any extent on the owner.

Explanation

The named file system does not refer to the same data set as on the owning system. The local mount fails and z/OS UNIX will function ship operations to the z/OS UNIX owner.

In the message text:

filesystem_name

The name of the file system.

owner_name

The name of the owner.

volser

The VOLSER in the client.

low_CI

The low CI in the client.

hi_CI

The high CI in the client.

System action

The program continues.

Administrator response

If z/OS UNIX function shipping is not desired, ensure that the file system refers to the same data set. Then unmount and remount.

Severity

svc_c_sev_error

IOEZ00950E

ProgName: MODIFY command - Parm rejected.

Explanation

The MODIFY command *Parm* has been rejected. The *Parm* command cannot be processed until the zFS kernel has completed initialization. An abort or internal restart in progress can also result in this error.

In the message text:

ProgName

The program name.

Parm

The program parameters on the MODIFY command that were passed to zFS.

System action

The program continues.

Administrator response

Retry the command after initialization completes. If initialization does not complete, zFS may need to be canceled. If the problem continues, contact your service representative.

Severity

svc_c_sev_error

IOEZ00951I

***Aggregate aggregate_name anode table length anode_table_size
LPI=LPI_value encrypt_status compress_status.***

Explanation

The encryption and compression attributes of the aggregate are displayed before the anode table pages are verified. The information is intended for IBM service personnel.

In the message text:

- aggregate_name**
The name of the aggregate.
- anode_table_size**
The anode table size in 8K.
- LPI_value**
The LPI value.
- encrypt_status**
The encryption status.
- compress_status**
The compression status.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00952E	File inode <i>inode_name</i> received snap-shot close error for file system <i>filesystem_name</i>.
-------------------	--

Explanation

A file opened for incremental backup could not be closed. This may make some space on-disk unusable and may make future backup attempts of the indicated file fail.

In the message text:

- inode_name**
The inode number of the file.
- filesystem_name**
The file system name.

System action

The program continues.

Administrator response

First make sure that there are no communication problems in the sysplex. Once those are resolved, an unmount remount without mode-switch can be used to clear the condition.

Severity

svc_c_sev_error

IOEZ00953E	Media Manager return code <i>ReturnCode</i> disconnecting aggregate <i>aggregate_name</i>.
-------------------	---

Explanation

The return code for the Media Manager disconnect was nonzero for the specified aggregate. This might occur during a failed mount or during an unmount of the specified aggregate.

In the message text:

ReturnCode

The hexadecimal return code from the Media Manager request.

aggregate_name

The name of the aggregate.

System action

The program continues.

Administrator response

Look up the return code in *z/OS DFSMSdfp Diagnosis* in the return codes section for the Media Manager. System resources obtained by Media Manager might still be held, potentially preventing future mounts and other operations for the aggregate from working correctly. Contact your service representative.

Severity

svc_c_sev_error

IOEZ00957E	zFS failed to assign key label <i>keylabel_name</i> to <i>aggregate_name</i> because it is not known to ICSF.
-------------------	--

Explanation

ISCF does not recognize the key label.

In the message text:

aggregate_name

The name of the aggregate.

keylabel_name

The name of the key label.

System action

The program continues.

Administrator response

Ensure that the correct key label is specified and try the command again.

Severity

svc_c_sev_error

IOEZ00958E	ICSF could not be contacted to verify key label <i>keylabel_name</i> for aggregate <i>aggregate_name</i>.
-------------------	--

Explanation

The key label must be known to ICSF before it can be used when defining or encrypting a zFS aggregate.

In the message text:

aggregate_name

The name of the aggregate.

keylabel_name

The name of the key label.

System action

The program continues.

Administrator response

Ensure that ICSF is running and verify that the key label is defined.

Severity

svc_c_sev_error

IOEZ00959E

ICSF returned error code *error_code* reason code *reason_code* while verifying key label *keylabel_name* for aggregate *aggregate_name*.

Explanation

A problem was encountered while using ICSF to verify a key label.

In the message text:

aggregate_name

The name of the aggregate.

error_code

The error code.

keylabel_name

The name of the key label.

return_code

The return code.

System action

The program continues.

Administrator response

Correct the problem and try the command again. Reason codes are documented in [z/OS Cryptographic Services ICSF Application Programmer's Guide](#).

Severity

svc_c_sev_error

IOEZ00960E

IOEFSKN termination hang, detaching IOEFSKN.

Explanation

IOEFSKN has started the termination process for a shutdown command or a zFS severe software error (abort). However, program IOEFSCM has detected a delay in that processing and considers that process is hung and is taking action to force termination to continue by detaching the IOEFSKN program. A detach abend might occur but termination now proceeds. The zFS restart prompt will likely be issued.

System action

The program continues.

Administrator response

Contact your service representative.

Severity

svc_c_sev_error

IOEZ00961E	IO driver termination problem detected, continuing termination.
-------------------	--

Explanation

The IOEFSKN main task has detected that there were problems encountered while terminating the zFS IO driver during a shutdown operation or an abnormal termination. Termination processing continues but zFS will end and the zFS restart prompt will likely be issued.

System action

The program continues.

Administrator response

Contact your service representative.

Severity

svc_c_sev_error

IOEZ00962E	Memory registration service FPZ4RMR failed with return code <i>return_code</i> reason code <i>reason_code</i>.
-------------------	---

Explanation

A problem was encountered when calling compression service FPZ4RMR.

In the message text:

return_code

The return code.

reason_code

The reason code.

System action

The program continues.

Administrator response

Correct the problem and try the command again. Return code and reason code are documented in [z/OS MVS Programming: Callable Services for High-Level Languages](#).

Severity

svc_c_sev_error

IOEZ00963E	The encryption type associated with key label <i>keylabel</i> is not supported by DFSMS.
-------------------	---

Explanation

DFSMS only supports AES-256 XTS protected key encryption.

In the message text:

keylabel

The key label.

System action

The program continues.

Administrator response

Try the command again with a different key label.

Severity

svc_c_sev_error

IOEZ00964E	<i>component</i> returned return code <i>return_code</i> reason code <i>reason_code</i> while trying to open aggregate <i>aggregate_name</i>.
-------------------	--

Explanation

A problem was encountered while opening an aggregate.

In the message text:

aggregate_name

The aggregate name.

component

The name of the component.

reason_code

The reason code.

return_code

The return code.

System action

The program continues.

Administrator response

See accompanying message IEC161I, correct the problem, and try the command again.

Severity

svc_c_sev_error

IOEZ00965E	The -keylabel parameter is not allowed because aggregate <i>aggregate_name</i> already has a key label.
-------------------	--

Explanation

The -keylabel parameter should only be used for an aggregate which does not have a key label.

In the message text:

aggregate_name

The aggregate name.

System action

The program continues.

Administrator response

Remove the -keylabel parameter and try the command again.

Severity

svc_c_sev_error

IOEZ00966E

The shrinking of aggregate *aggregate_name* is not allowed by *reason_failure*.

Explanation

The shrink of an SMS-managed aggregate is not allowed due to the specified reason.

In the message text:

aggregate_name

The aggregate name.

reason_failure

The reason for the failure.

System action

The program continues.

Administrator response

If the reason the shrink failed was because of a management class attribute, you should ensure you are using the correct management class for the aggregate. See [z/OS DFSMSdss Storage Administration](#) for information related to using management classes. For other reasons, contact your service representative.

Severity

svc_c_sev_error

IOEZ00981E

Dynamic allocation abend occurred for data set *dataset_name*. Abend code=*Abend_Code* func=*func*.

Explanation

Dynamic allocation abended while processing the specified data set. The abend code returned by the system is as shown. *func* describes the dynamic allocation function that was attempted.

In the message text:

Abend_Code

The z/OS abend code received.

dataset_name

The name of the data set.

func

The dynamic allocation function that was attempted.

System action

The program continues.

Administrator response

Review the system logs for error messages related to dynamic allocation and the specified data set. Correct the error and retry the operation.

Severity

svc_c_sev_error

IOEZ00982E	Offset <i>offset</i> of inode <i>inode_number</i> in file system <i>file_system</i> cannot be decompressed.
-------------------	--

Explanation

The data at the indicated offset of the listed file could not be decompressed due to an error returned by zEDC routine FPZ4ABC or there is incorrect data in that region of the file. If FPZ4ABC failed, additional messages will have preceded this message describing that error. System 2C3 dumps may also have preceded this message.

In the message text:

file_system

The name of the file system with the error.

inode_number

The inode number of the file in the file system.

offset

The offset into the file.

System action

The program continues.

Administrator response

The indicated file has one or more regions that cannot be decompressed. Delete or restore that file by using a method applicable to the type of data the file contained. Restoring from a backup is the typical method used to repair the file. Report the problem to your service representative.

Severity

svc_c_sev_error

IOEZ00983I	File with inode=<i>inode_number</i> (size=<i>inode_size</i>) in aggregate <i>aggregate_name</i> is being encrypted or decrypted.
-------------------	---

Explanation

zFS is in the process of encrypting or decrypting the file with the specified inode in the specified aggregate. The file size (in K) is also specified. The encryption or decryption process is delayed, and so might take a long time to complete. A possible cause of the delay is an application that is accessing the file during this operation. Use the **ls -i** command to find the file name and then use the FILEINFO command to monitor the percentage completion of the operation.

In the message text:

aggregate_name

The name of the aggregate.

inode_number
The number of the inode.

inode_size
The size of the inode in K.

System action

The program continues.

Administrator response

If the delay is not acceptable, the operation can be canceled. If you are encrypting the file, you can cancel the encryption by using the **zfsadm encrypt -cancel** command. If you are decrypting the file, use the **zfsadm decrypt -cancel** command. Both actions will cancel the encryption or decryption processing of the entire aggregate. If you want to determine which file this is, you can use the **ls -i** command.

Severity

svc_c_sev_notice

IOEZ00984I **File with inode=inode_number (size=inode_size) in aggregate aggregate_name is being compressed or decompressed.**

Explanation

zFS is in the process of compressing or decompressing the file with the specified inode in the specified aggregate. The file size (in K) is also specified. The compression or decompression process is delayed, and so might take a long time to complete. A possible cause of the delay is an application that is accessing the file during this operation. Use the **ls -i** command to find the file name and then use the FILEINFO command to monitor the percentage completion of the operation.

In the message text:

aggregate_name
The name of the aggregate.

inode_number
The number of the inode.

inode_size
The size of the inode in K.

System action

The program continues.

Administrator response

If the delay is not acceptable, the operation can be canceled. If you are compressing the file, you can cancel the compression by using the **zfsadm compress -cancel** command. If you are decompressing the file, use the **zfsadm decompress -cancel** command. Both actions will cancel the compression or decompression processing of the entire aggregate. If you want to determine which file this is, you can use the **ls -i** command.

Severity

svc_c_sev_notice

IOEZ00985I **Waiting for service_name to initialize.**

Explanation

The named service is not available. zFS is waiting to allow the service to initialize.

In the message text:

service_name

The name of the service.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ00986I Wait ending and *service_name* is not initialized.

Explanation

After waiting, the named service is still not initialized. Operations requiring this service will fail.

In the message text:

service_name

The name of the service.

System action

The program continues.

Administrator response

Determine why the named service will not initialize and restart the service.

Severity

svc_c_sev_warning

IOEZ00987I *service_name* is now initialized.

Explanation

The named service is now available.

In the message text:

service_name

The name of the service.

System action

The program continues.

Administrator response

Operations that failed because the service was not initialized can now be retried.

Severity

svc_c_sev_notice

IOEZ00988I	The wait for <i>service_name</i> initialization ended because zFS is being stopped or restarted.
-------------------	---

Explanation

zFS has stopped waiting for the named service to initialize.

In the message text:

service_name

The name of the service.

System action

The program continues.

Administrator response

When zFS is restarted, ensure that the named service is initialized before retrying operations that require the service.

Severity

svc_c_sev_notice

IOEZ00989I	The wait for <i>service_name</i> initialization is not started because zFS is being stopped or restarted.
-------------------	--

Explanation

Since zFS is being stopped or restarted, it does not need to wait for the initialization of the named service.

In the message text:

service_name

The name of the service.

System action

The program continues.

Administrator response

When zFS is restarted, ensure that the named service is initialized before retrying operations that require the service.

Severity

svc_c_sev_notice

IOEZ00990E	Mount for aggregate <i>aggregate_name</i> failed because ICSF is not initialized.
-------------------	--

Explanation

Opens that are waiting for ICSF to come up are failed when zFS is in the process of being stopped or restarted.

In the message text:

aggregate_name
The aggregate name.

System action

The program continues.

Administrator response

Try the operation again after ICSF is initialized.

Severity

svc_c_sev_error

IOEZ00991E	Mount for aggregate <i>aggregate_name</i> failed because zFS is being stopped or restarted.
-------------------	--

Explanation

Mounts that are waiting for ICSF to initialize are failed when zFS is being stopped or restarted.

In the message text:

aggregate_name
The aggregate name.

System action

The program continues.

Administrator response

Try the mount again when ICSF is initialized and zFS is not being stopped or restarted.

Severity

svc_c_sev_error

IOEZ00992E	Open for aggregate <i>aggregate_name</i> failed because ICSF is not initialized.
-------------------	---

Explanation

ICSF is not initialized. zFS cannot open an aggregate that has a key label if ICSF is not initialized

In the message text:

aggregate_name
The aggregate name.

System action

The program continues.

Administrator response

Try the operation again after ICSF is initialized.

Severity

svc_c_sev_error

IOEZ00993E	Open for aggregate <i>aggregate_name</i> failed because zFS is being stopped or restarted.
-------------------	---

Explanation

Opens that are waiting for a ICSF to initialize are failed when zFS is being stopped or restarted.

In the message text:

aggregate_name

The aggregate name.

System action

The program continues.

Administrator response

Try the operation again when zFS is not being stopped or restarted.

Severity

svc_c_sev_error

IOEZ00994I	The EDCFIXED option of the user cache size IOEFSPRM configuration option is ignored.
-------------------	---

Explanation

Since zEDC is not initialized, the EDCFIXED option cannot be processed. The user cache pages are not fixed for zEDC performance.

System action

The program continues.

Administrator response

Once zEDC is initialized, issue the **zfsadm config** -user_cache_size command with the EDCFIXED option.

Severity

svc_c_sev_notice

IOEZ00995E	Mount for aggregate <i>aggregate_name</i> failed because zEDC is not initialized.
-------------------	--

Explanation

zEDC is not initialized. zFS cannot mount an aggregate that is compressed if zEDC is not initialized.

In the message text:

aggregate_name

The aggregate name.

System action

The program continues.

Administrator response

Try the operation again after zEDC is initialized.

Severity

svc_c_sev_error

IOEZ00996E	The compression of aggregate <i>aggregate_name</i> failed because zEDC is not initialized.
-------------------	---

Explanation

Data cannot be compressed when zEDC is not initialized.

In the message text:

aggregate_name

The aggregate name.

System action

The program ends.

Administrator response

Try the operation again after zEDC is initialized.

Severity

svc_c_sev_error

IOEZ00997E	Aggregate takeover for aggregate <i>aggregate_name</i> failed because zEDC is not initialized.
-------------------	---

Explanation

This system cannot assume ownership because the aggregate contains compressed data and zEDC is not initialized.

In the message text:

aggregate_name

The aggregate name.

System action

The program continues.

Administrator response

Try the operation again after zEDC is initialized.

Severity

svc_c_sev_error

IOEZ00999I	zFS is waiting for ICSF initialization.
-------------------	--

Explanation

zFS is waiting for ICSF to initialize because the IOEFSPRM configuration option `initialize_for_encryption` is ON. zFS will continue initialization once it receives notification that ICSF is available.

System action

The program continues.

Administrator response

No action is required if you plan to use encrypted file systems. If zFS is still waiting when ICSF is initialized, then contact the service representative. If you do not plan on using encryption, then remove the configuration parameter before the next IPL of the system.

Severity

`svc_c_sev_notice`

IOEZ01000I	zFS is waiting for SMS initialization.
-------------------	---

Explanation

zFS is waiting for SMS to initialize because the IOEFSPRM configuration option `user_cache_size` has the `EDCFIXED` option specified. zFS will continue initialization once it receives notification that zEDC is available. SMS and zEDC are required products for zFS compression.

System action

The program continues.

Administrator response

No action is required if you plan to use compressed file systems. If zFS is still waiting when SMS is initialized, then contact the service representative. If you do not plan on using compression, then remove the configuration parameter before the next IPL of the system.

Severity

`svc_c_sev_notice`

IOEZ01001I	zFS is waiting for zEDC initialization.
-------------------	--

Explanation

zFS is waiting for zEDC to initialize because the IOEFSPRM configuration option `user_cache_size` has the `EDCFIXED` option specified. zFS will continue initialization once it receives notification that zEDC is available. SMS and zEDC are required products for zFS compression.

System action

The program continues.

Administrator response

No action is required if you plan to use compressed file systems. If zFS is still waiting when zEDC is initialized, then contact the service representative. If you do not plan on using compression, then remove the configuration parameter before the next IPL of the system.

Severity

svc_c_sev_notice

IOEZ01002E	The encryption of aggregate <i>aggregate_name</i> failed because ICSF services are not available.
-------------------	--

Explanation

File systems cannot be encrypted when ICSF services are not available.

In the message text:

aggregate_name

The name of the file system.

System action

The program ends.

Administrator response

Try the command again after ICSF is available.

Severity

svc_c_sev_error

IOEZ01003I	File system <i>file_system</i> received a compression failure due to lack of zEDC devices.
-------------------	---

Explanation

A compression file write for one or more files in the indicated file system failed due to lack of zEDC devices. The file data will be stored in noncompressed format. No application errors will occur but the file system might benefit from a decompress and a compress command to save maximum space.

In the message text:

file_system

The name of the file system.

System action

The program continues.

Administrator response

Once zEDC devices are available, you can decompress the file system and compress it again to determine if space savings can be made. This is likely only useful if the file system is large or has many large files.

Severity

svc_c_sev_notice

IOEZ01004E	File system <i>file_system</i> CI=<i>control_interval_number</i> (<i>data_type</i> inode=<i>inode_number</i>) is not stored in encrypted format.
-------------------	---

Explanation

A decryption read request of an extended format data set found that zFS has stored data that is supposed to be encrypted, in the clear. The type of data that is affected is enclosed in parentheses. It is one of the following

values: meta (internal file system metadata), compressed file, or uncompressed file. A dump will be generated to provide diagnostics for IBM service. If the data is file data, then the inode of the owning file is shown. For metadata, this value will be zero.

In the message text:

control_interval_number

The control interval number.

data_type

The type of affected data.

file_system

The name of the file system.

inode_number

The inode number of the affected file.

System action

The program continues.

Administrator response

If the data type is meta or compressed file data, you can use the salvager program to correct the condition. If the data is uncompressed file data, then *inode* identifies the file. That file can be rewritten with new contents. For example, use the **cp** command to copy the file to a new name and then use the **mv** command to change the new file to the original name. Then delete the old file. Report the problem to the service representative.

Severity

svc_c_sev_error

IOEZ01005I

**Successfully repaired corrupted data for file (with index *file_index*).
The compression group information: logpage *logical_page* total pages
pages_before_compression compressed pages *pages_after_compression*.**

Explanation

The salvager found corrupted encrypted data in the file with the specified index and repaired the corrupted data. The corrupted data is on an aggregate that is in compressed format. The compression group information that is associated with the corrupted data is also displayed and is intended for IBM service.

In the message text:

file_index

The file index.

logical_page

The first logical page in the compression group.

pages_before_compression

The number of pages before compression.

pages_after_compression

The number of pages after compression.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01006I **Successfully repaired corrupted data in physical page *physical_page*.**

Explanation

The salvager found corrupted encryption data in the specified physical page and repaired the corrupted data. This message is intended for IBM service.

In the message text:

physical_page

The physical page in the zFS aggregate.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01007I **Successfully repaired corrupted data in ACL with index *index*.**

Explanation

The salvager found corrupted encryption data in the specified ACL object and repaired the corrupted data. This message is intended for IBM service.

In the message text:

index

The ACL index.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01008I **File with inode=*inode_number* (size=*inode_size*) in aggregate *aggregate_name* is encountering delays during shrink processing.**

Explanation

zFS is in the process of shrinking the specified aggregate and has encountered the file or directory with the specified inode and the specified size (in K). Because the shrink process is being delayed, the process might take a long time to complete. Applications that are accessing the file or directory during the shrink process is a possible cause of this delay.

In the message text:

aggregate_name

The name of the file system.

inode_number

The inode number of the affected file.

inode_size

The size of the inode in K.

System action

The program continues.

Administrator response

If the delay is unacceptable, the shrink process can be canceled by using the **zfsadm shrink -cancel** command. If you want to determine which file or directory the inode represents, you can use the **ls -i** command. Otherwise, wait for the shrink command to complete.

Severity

svc_c_sev_notice

IOEZ01009I	Call to catalog service for dataset <i>aggregate_name</i> failed, rc=<i>ReturnCode</i> rsn=<i>ReasonCode</i>.
-------------------	--

Explanation

A call to the MVS catalog service for the specified data set failed. The return and reason codes are from the MVS catalog service. Other system processing that depends on the catalog entry may not function properly.

In the message text:

aggregate_name

The name of the file system.

ReturnCode

The catalog service return code.

ReasonCode

The catalog service reason code.

System action

The program continues.

Administrator response

Catalog return and reason codes are listed in message IDC3009I in [z/OS MVS System Messages, Vol 6 \(GOS-IEA\)](#). Find the specified return and reason codes in order to determine the problem.

Severity

svc_c_sev_notice

IOEZ01027E	Conversion to version 1.4 file system is disabled.
-------------------	---

Explanation

zFS no longer supports the creation of version 1.4 file systems and disallows the conversion to a version 1.4 file system.

System action

IOEFSUTL program ends.

Administrator response

Submit the JCL on a V2R2 or V2R3 system to convert the aggregate to version 1.4.

Severity

svc_c_sev_error

IOEZ01028E	Unexpected errors during takeover processing for HA file system <i>filesystem_name</i> with new owner <i>system_name</i>.
-------------------	--

Explanation

zFS encountered unexpected errors while resolving pending operations that were in progress when the original owner of the file system went down. This means the applications that requested those operations will receive EIO (122) errors from zFS.

In the message text:

filesystem_name

The name of the file system.

system_name

The name of the system.

System action

The program continues.

Administrator response

Contact your service representative.

Severity

svc_c_sev_error

IOEZ01029E	The HA attribute is not assigned to file system <i>file_system_name</i> because system <i>system_name</i> does not support high availability file systems.
-------------------	---

Explanation

The HA attribute was to be assigned to the specified file system either because the HA mount parameter was specified or the IOEFSPRM configuration option HA=ON is specified. The HA attribute can only be used when all members of a sysplex support high availability file systems.

In the message text:

file_system

The name of the file system.

system_name

The name of the system that does not support high availability file systems.

System action

The program continues.

Administrator response

To use the HA file system attribute, ensure that all members in the sysplex have OA57508 applied. Then you can enable the HA attribute for this file system by using the **zfsadm chaggr** command with the -ha option. Alternatively, you could unmount the file system, and mount it again with the HA mount parameter. Or, you can issue the **zfsadm config** command with -ha on specified, then unmount and mount the file system again.

Severity

svc_c_sev_error

IOEZ01030E	Initializing system <i>system_name</i> will not be allowed to join the sysplex because it does not support high availability file systems.
-------------------	---

Explanation

There is at least one high availability (HA) mounted file system in the sysplex. HA support requires that all members of a sysplex support HA.

In the message text:

system_name

Name of the initializing system that does not support HA.

System action

The program continues.

Administrator response

zFS on the initializing system will terminate. To correct this problem, either unmount or change the HA status of all file systems that are currently mounted HA. The **zfsadm chaggr** command can be used to change the HA status of a file system.

Severity

svc_c_sev_warning

IOEZ01031E	Invalid wildcard pattern in aggregate name <i>aggregate_name</i>.
-------------------	--

Explanation

zfsadm chaggr only supports wildcards that have '*' at either the beginning or the end of the aggregate name, or both.

In the message text:

aggregate_name

The name of the aggregate with the wildcard.

System action

The program ends.

Administrator response

Correct the aggregate name and try again.

Severity

svc_c_sev_error

IOEZ01040A

zFS kernel has reached internal restart limit *restart_limit*, will terminate due to excessive errors.

Explanation

The zFS kernel has restarted due to internal error more than the indicated restart limit and has repeated restarts in the past 24 hours. The zFS kernel will terminate to prevent excessive dumps and restart looping.

In the message text:

restart_limit

The internal restart limit.

System action

The program ends.

Administrator response

Contact your service representative.

Severity

svc_c_sev_fatal

IOEZ01042E

Invalid control interval size found while formatting *filesystem_name*.

Explanation

zFS attempted to format a file system but found the control interval size of the data set was not 4096.

In the message text:

filesystem_name

The name of the file system.

System action

The program ends.

Administrator response

Delete the mentioned data set and create a valid zFS file system linear data set that has a control interval size of 4096 and specifies the ZFS keyword.

Severity

svc_c_sev_fatal

IOEZ01060E

File system *filesystem_name* received an unexpected zEDC compression failure with rc=*return_code* and reason=*reason_code*.

Explanation

A compression file write for one or more files in the indicated file system failed unexpectedly in zEDC, the file data will be stored in non-compressed format. No application errors will occur but the file system might benefit from a decompress and a compress command to save maximum space.

In the message text:

filesystem_name

The file system name.

return_code

The return code from zEDC.

reason_code

The reason code from zEDC.

System action

The program continues.

Administrator response

Once zEDC software patches are available, you can decompress the file system and compress it again to determine if space savings can be made. This is likely only useful if the file system is large or has many large files.

Severity

svc_c_sev_error

IOEZ01061I

Log file recovery skipped for file system *filesystem_name*

Explanation

Log file recovery was skipped for the indicated file system. The log was not clean but the recovery component determined that the metadata for the file system was consistent and that recovery could be skipped.

In the message text:

filesystem_name

The file system name.

System action

The program continues.

Administrator response

This is an informational message. If the file system identified in this message was later found to be corrupted by the zFS salvager, then report this message to IBM service.

Severity

svc_c_sev_notice

IOEZ01076E

Incorrect value *Release* provided in the argument for the -release option.

Explanation

An incorrect argument value was provided for the -release option of the **zfsadm** command.

In the message text:

Release

The value that the user entered for the -release option argument.

System action

The program ends.

Administrator response

Change the `-release` option argument to an integer in the range 1-100 and reissue the **zfsadm** command.

Severity

svc_c_sev_error

IOEZ01077I

Internally remounting *filesystem* after salvage repair.

Explanation

The specified file system was repaired during a salvage operation. It is being internally remounted with the same mode.

In the message text:

filesystem

The name of the file system.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01078E

Internal remount of *filesystem* failed *rc=return_code* and *rsn=reason_code*.

Explanation

The internal remount of the specified file system failed with the indicated return and reason codes.

In the message text:

filesystem

The name of the file system.

return_code

The return code for the failure.

reason_code

The reason code for the failure.

System action

The program continues.

Administrator response

Try and determine why the remount failed based on the return and reason codes, and try to remount the file system again. If the file system is no longer mounted, it will need to be manually remounted. If the problem persists, contact your service representative.

Severity

svc_c_sev_error

IOEZ01079E

Invalid wildcard pattern in aggregate name *aggregate_name*.

Explanation

The **zfsadm** command only supports wildcards that have '*' at either the beginning or the end of the aggregate name, or both.

In the message text:

aggregate_name

The aggregate name with the wildcard.

System action

The program ends.

Administrator response

Correct the aggregate name and try again.

Severity

svc_c_sev_error

IOEZ01080I

Shrinking aggregate name *aggregate_name* to size *Block_size* 8K blocks.

Explanation

zFS is starting to shrink the file system for the specified aggregate name using the new size provided with the -size option or internally calculated by using the argument specified on the -release option.

In the message text:

aggregate_name

The aggregate name.

Block_size

The targeted size of the file system in 8K blocks.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01081I

Shrink line for aggregate *aggregate_name* adjusted to block *Block_number* after moves: *Shrink_line_moves*.

Explanation

zFS has adjusted the shrink line for the aggregate to the block number specified using the indicated number of shrink line moves.

In the message text:

aggregate_name

The aggregate name.

Block_number

The new shrink line of the file system specified as a block number.

Shrink_line_moves

The number of shrink line moves that occurred while performing the shrink operation.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01082I

Activity detected while shrinking aggregate *aggregate_name*. Blocks allocated *Blocks_allocated*, blocks freed *Blocks_freed*, blocks reserved by clients *Client_blocks*.

Explanation

zFS detected file system activity during shrink processing which could have caused the shrink line to be moved, and therefore could have affected the amount of space released from the zFS aggregate.

In the message text:

aggregate_name

The aggregate name.

Blocks_allocated

The number of 8K blocks allocated.

Blocks_freed

The number of 8K blocks freed.

Client_blocks

The number of 8K client blocks reserved.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01083I

Shrink processing for aggregate *aggregate_name* converted *Fragments_converted* fragments to *Blocks_created* 8K blocks.

Explanation

zFS detected that the file system contained legacy fragmented objects that are no longer kept as fragmented objects. These objects were converted to blocked objects. These conversions could have caused the shrink line to move, and therefore, could have affected the amount of space released from the zFS aggregate.

In the message text:

aggregate_name

The aggregate name.

Fragments_converted

The number of 1K fragments converted to blocks.

Blocks_created

The number of 8K blocks created from 1K fragments.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01084I

The -remountro option was provided for unmounted aggregate *aggregate_name*. Attach will be attempted.

Explanation

The aggregate name specified is not currently mounted. The -remountro option will be ignored and an attach will be attempted to complete the **zfsadm shrink** operation.

In the message text:

aggregate_name

The aggregate name.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01085E

The remount of aggregate *aggregate_name* as R/W failed prior to the shrink operation.

Explanation

The -remountro option of the **zfsadm shrink** command was specified. The file system is mounted R/O. The remount as R/W has failed. The shrink operation was not performed.

In the message text:

aggregate_name

The aggregate name.

System action

The program ends.

Administrator response

Determine why the remount as R/W failed. Reissue the **zfsadm shrink** command after attempting to resolve any conflicts. If the problem persists, contact your service representative.

Severity

svc_c_sev_error

IOEZ01086E

The remount of aggregate *aggregate_name* as R/O failed after completing the shrink operation.

Explanation

The **-remountro** option of the **zfsadm shrink** command was specified. The file system was mounted R/O at that time and was internally remounted as R/W in order to shrink the aggregate. The internal remount back to R/O either failed or was not able to be performed.

In the message text:

aggregate_name

The aggregate name.

System action

The program ends.

Administrator response

Manually remount the aggregate as R/O if R/W access is not desired.

Severity

svc_c_sev_error

IOEZ01087I

Aggregate *aggregate_name* was successfully remounted as R/O after completing the shrink operation.

Explanation

The aggregate name specified was returned to R/O access after completing the shrink operation.

In the message text:

aggregate_name

The aggregate name.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01088I

Aggregate *aggregate_name* was temporarily remounted as R/W prior to completing the shrink operation

Explanation

The `-remountro` option of the **zfsadm shrink** command was specified. The file system was mounted R/O at that time, and was internally remounted as R/W in order to shrink the aggregate. It will automatically be remounted as R/O once the shrink operation completes.

In the message text:

aggregate_name

The aggregate name.

System action

The program continues.

Administrator response

None.

Severity

svc_c_sev_notice

IOEZ01089E

Aggregate *aggregate_name* is no longer mounted. No shrink operation performed.

Explanation

A **zfsadm shrink** command was entered using a wildcard. An attempt to shrink the specified aggregate failed. The aggregate was mounted at the time of the request, but was unmounted prior to shrink operation completion. The shrink operation was not performed.

In the message text:

aggregate_name

The aggregate name.

System action

The program ends.

Administrator response

Verify that the provided aggregate is mounted for the duration of the wildcard shrink command and retry the request. Alternatively, if a shrink of an unmounted aggregate is required, retry the request without a wildcard and allow the shrink operation to complete using the auto-attach feature.

Severity

svc_c_sev_notice

Chapter 3. IOEZHnnnnnt: zFS Health Checker messages

This section contains the messages that result from zFS checks from IBM Health Checker for z/OS. See [IBM Health Checker for z/OS User's Guide](#) for more information about the IBM health checks.

IOEZH0001I	A zFS PFCTL failed. Command <i>command</i>, subcommand <i>subcommand</i> failed with errno=<i>errno</i>, errnojr=<i>errnojr</i>
-------------------	--

Explanation:

The check could not execute.

System action:

The system continues processing.

Operator response:

Report this problem to the system programmer.

System programmer response:

Search problem reporting data bases for a fix for the problem. If a fix does not exist, call the IBM Support Center.

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

[z/OS File System Administration](#)

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0002I	Health check received an unknown entry code of <i>entry code</i> from IBM Health Checker for z/OS.
-------------------	---

Explanation:

This is an internal error.

System action:

The system continues processing.

Operator response:

Report this problem to the system programmer.

System programmer response:

Search problem reporting data bases for a fix for the problem. If a fix does not exist, call the IBM Support Center.

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

[*IBM Health Checker for z/OS User's Guide*](#)

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0040I zFS is running with a default meta_cache_size of cur_size.

Explanation:

zFS is running with a default metadata cache size. No action is required.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [*IOEFSPRM*](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0041I zFS is running with a default_or_specified meta_cache_size of meta_size and a default_or_specified metaback_cache_size of metaback_size. The calcul_default_or_overridden meta_cache_size is meta_sizeM. The sum of the current two cache sizes sumM is greater than or equal to the default meta_cache_size calcul_default_or_overriddenM.

Explanation

zFS is running with the indicated metadata cache size and metadata backing cache size. The sum of the current cache sizes is greater than (or equal to) the default (or user-overridden) metadata cache size. No action is required.

If the user-specified metadata cache size is less than 1 M (minimum) or more than 64 G (maximum), zFS replaces it with the minimum or maximum value and displays the current metadata cache size with the new value. If the user-specified metadata backing cache size is less than 1 M (minimum) or more than 2048 M (maximum), zFS replaces it with the minimum or maximum value and displays the current metadata backing

cache size with the new value. If either case is true, check the value that is defined in `meta_cache_size` or `metaback_cache_size` option in the IOEFSPRM configuration file and verify that it is within the valid range.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHC1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0042I **zFS is running with a default user_cache_size of size.**

Explanation:

zFS is running with a default user_cache_size. No action is required.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHC1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0043I

zFS is running with a *default_or_specified* user_cache_size of *cur_size* and the user cache size is greater than or equal to the *calcul_default_or_overridden* size of *new_size* M.

Explanation

zFS is running with the indicated user cache size that is greater than or equal to the default or user-overridden user cache size. No action is required.

If the user-specified user cache size is less than 10M (minimum) or more than 65536M (maximum), zFS replaces it with the minimum or maximum value and displays the current size with the new value. If that is the case, check the value that is defined in the user_cache_size option in the IOEFSPRM configuration file and verify that it is within the valid range.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0044E

zFS is running with a *default_or_specified* meta_cache_size of *cache_size* and a *default_or_specified* metaback_cache_size of *cache_size*. The *calcul_default_or_overridden* meta_cache_size is *cache_size* M. The sum of the current two cache sizes *cur_sum* M is less than the *calcul_default_or_overridden* meta_cache_size *new_sum* M.

Explanation

zFS is running with the indicated metadata cache size and metadata backing cache size. The sum of the current two cache sizes is less than the default (or user-overridden) metadata cache size. Running with a very small metadata cache might affect zFS performance. See [Performance tuning](#) in *z/OS File System Administration* to determine whether the current settings might impact zFS performance.

If the user-specified metadata cache size is less than 1 M (minimum) or more than 64 G (maximum), zFS replaces it with the minimum or maximum value and displays the current metadata cache size with the new value.

If the user-specified metadata backing cache size is less than 1 M (minimum) or more than 2048 M (maximum), zFS replaces it with the minimum or maximum value and displays the current metadata backing cache size with the new value.

System action:

The system continues processing.

Operator response:

Report this problem to the system programmer.

System programmer response

For user-specified cache size, check the value that is defined in `meta_cache_size` or `metaback_cache_size` option in the IOEFSPRM configuration file and verify that it is within the valid range. Depending on the performance analysis, if the current settings do not perform as well as the default size, specify `meta_cache_size` with the default size in your IOEFSPRM configuration file and restart zFS. The `meta_cache_size` configuration option can also be dynamically updated using the **zfsadm config** command.

Otherwise, specify the current meta cache size with the user override check parameter `meta_cache` on the PARM statement (for either HZSPRMxx or MODIFY hzsproc) in order to make the check succeed.

Problem determination:

N/A

Module:

IOEZCHK1

Source:

z/OS File System

Reference documentation:

See [Performance tuning](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0045E	zFS is running with a <i>default_or_specified user_cache_size</i> of <i>cur_size</i> that is less than the <i>calcul_default_or_overridden</i> size of <i>new_size</i> M.
-------------------	--

Explanation

zFS is running with the indicated user cache size and the size is less than the default or user-overridden user cache size. Running with a very small user cache size could affect zFS performance. See [Performance tuning](#) in *z/OS File System Administration* to determine whether the current setting might affect zFS performance.

If the user-specified user cache size is less than 10M (minimum) or more than 65536M (maximum), zFS replaces it with the minimum or maximum value and displays the current size with the new value.

System action:

The system continues processing.

Operator response:

Report this problem to the system programmer.

System programmer response:

For a user-specified cache size, check the value that is defined in the `user_cache_size` option in the IOEFSPRM configuration file and verify that it is within the valid range. Depending on the performance analysis, if the current setting does not perform as well as the default value, specify `user_cache_size` with the default size in your IOEFSPRM configuration file and restart zFS. The `user_cache_size` configuration option can also be dynamically updated using the **zfsadm config** command. Otherwise, specify the current user cache size with the user override check parameter `user_cache` on the PARM statement (for either HZSPRMxx or MODIFY hzsproc) in order to make the check succeed.

Problem determination:

N/A

Module:
IOEZHC1

Source:
z/OS File System

Reference documentation:
See "Performance tuning" or "IOEFSPRM" in *z/OS File System Administration*.

Automation:
N/A

Routing code:
N/A

Descriptor code:
N/A

IOEZH0062I **zFS configuration option metaback_cache_size is not specified.**

Explanation:
zFS is currently running with the metaback_cache_size option unspecified. zFS IOEFSPRM configuration option metaback_cache_size is no longer needed. The entire size of the metadata cache should be specified in IOEFSPRM configuration file option meta_cache_size. It is not required, but is recommended to keep your IOEFSPRM configuration file clean of outdated specifications to avoid any future confusion. No action is required.

System action:
The system continues processing.

Operator response:
N/A

System programmer response:
N/A

Problem determination:
N/A

Module:
IOEZHC1

Source:
z/OS File System

Reference documentation:
See IOEFSPRM in *z/OS File System Administration*.

Automation:
N/A

Routing code:
N/A

Descriptor code:
N/A

IOEZH0063I **zFS configuration option tran_cache_size is not specified.**

Explanation:
zFS IOEFSPRM configuration option tran_cache_size is ignored. zFS is currently running with the tran_cache_size unspecified. No action is required.

System action:
The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0064I **zFS configuration option `client_cache_size` is not specified.**

Explanation:

zFS IOEFSPRM configuration option `client_cache_size` is ignored. zFS is currently running with the `client_cache_size` unspecified. No action is required.

System action:

The system continues processing.

Operator response

N/A

System programmer response

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0065E **zFS configuration option `metaback_cache_size` is specified.**

Explanation:

zFS IOEFSPRM configuration option `metaback_cache_size` is currently specified, but is no longer needed. The entire size of the metadata cache should be specified in IOEFSPRM configuration file option `meta_cache_size`. It is not required, but is recommended to keep your IOEFSPRM configuration file clean of outdated specifications to avoid any future confusion.

System action:

The system continues processing.

Operator response

Report this problem to the system programmer.

System programmer response:

You should assign the combined values of the `metaback_cache_size` and `meta_cache_size` options into just the `meta_cache_size` option to ensure that zFS will continue to have similar performance. You can remove the specification of the `metaback_cache_size` option in the IOEFSPRM configuration file to avoid any future confusion. Otherwise, specify the user override check parameter `METABACK(EXISTENCE)` on the PARM string (via `HZSPRMxx` or `MODIFY hzsproc`) in order to make the check succeed.

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [Performance tuning in z/OS File System Administration](#).

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0066E

zFS configuration option `tran_cache_size` is specified.

Explanation:

zFS is currently running with a specified `tran_cache_size`. zFS IOEFSPRM configuration option `tran_cache_size` is ignored.

System action:

The system continues processing.

Operator response

Report this problem to the system programmer.

System programmer response:

Since zFS IOEFSPRM configuration option **`tran_cache_size`** is ignored, you should remove specification of `tran_cache_size` in IOEFSPRM to avoid any future confusion. Otherwise, specify the user override check parameter `TRANS(EXISTENCE)` on the PARM string (via `HZSPRMxx` or `MODIFY hzsproc`) in order to make the check succeed.

Problem determination:

N/A

Module:

IOEZHCK1

Source:

Reference documentation

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0067E **zFS configuration option `client_cache_size` is specified.**

Explanation:

zFS is currently running with a specified `client_cache_size`. zFS IOEFSPRM configuration option `client_cache_size` is ignored.

System action:

The system continues processing.

Operator response:

Report this problem to the system programmer.

System programmer response:

Since zFS IOEFSPRM configuration option `client_cache_size` is ignored, you should remove specification of `client_cache_size` in IOEFSPRM to avoid any future confusion. Otherwise, specify the user override check parameter `CLIENT(EXISTENCE)` on the PARMs string (via `HZSPRMxx` or `MODIFY hzsproc`) in order to make the check succeed.

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0068I **zFS configuration option `metaback_cache_size` is specified.**

Explanation:

The check override parameter `METABACK(EXISTENCE)` is specified in the PARMs string (via `HZSPRMxx` or `MODIFY hzsproc`) to verify that the configuration option `metaback_cache_size` has been specified. The check found zFS is currently running with a specified `metaback_cache_size`. zFS IOEFSPRM configuration option `metaback_cache_size` is no longer needed. The combined values of the `metaback_cache_size` and `meta_cache_size` options should be assigned into just the `meta_cache_size` option to ensure that zFS will continue to have similar performance. You can remove the specification of the `metaback_cache_size` option in the IOEFSPRM configuration file to avoid any future confusion.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0069I **zFS configuration option tran_cache_size is specified.****Explanation:**

The check override parameter TRANS(EXISTENCE) is specified in the PARMS string (via HZSPRMxx or MODIFY hzsproc) to verify that the configuration option tran_cache_size has been specified. The check found zFS is currently running with a specified tran_cache_size. zFS IOEFSPRM configuration option tran_cache_size is ignored.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0070I **zFS configuration option client_cache_size is specified.****Explanation:**

The check override parameter CLIENT(EXISTENCE) is specified in the PARMS string (via HZSPRMxx or MODIFY hzsproc) to verify that the configuration option `client_cache_size` has been specified. The check found zFS is currently running with a specified `client_cache_size`. zFS IOEFSPRM configuration option `client_cache_size` is ignored.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

Remove the user override check parameter METABACK(EXISTENCE) from the PARMS string (via HZSPRMxx or MODIFY hzsproc) or specify the user override check parameter METABACK(ABSENCE) in order to make the check succeed.

Problem determination:

N/A

Module:

IOEZCHK1

Source:

z/OS File System

Reference documentation:

See [IOEFSPRM](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0071E	zFS configuration option metaback_cache_size is not specified.
-------------------	---

Explanation:

The check override parameter METABACK(EXISTENCE) is specified in the PARMS string (via HZSPRMxx or MODIFY hzsproc) to verify the configuration option `metaback_cache_size` has been specified. The check found `metaback_cache_size` is not specified. zFS IOEFSPRM configuration option `metaback_cache_size` is no longer needed. No action is required for the `metaback_cache_size` option.

System action:

The system continues processing.

Operator response

Report this problem to the system programmer.

System programmer response

Remove the user override check parameter METABACK(EXISTENCE) from the PARMS string (via HZSPRMxx or MODIFY hzsproc) or specify the user override check parameter METABACK(ABSENCE) in order to make the check succeed.

Problem determination:

N/A

Module:

IOEZCHK1

Source:

z/OS File System

See [IOEFSPRM](#) in *z/OS File System Administration*.

N/A

N/A

N/A

zFS configuration option `tran_cache_size` is not specified.

The check override parameter TRANS(EXISTENCE) is specified in the PARMS string (via HZSPRMxx or MODIFY hzsproc) to verify the configuration option tran_cache_size has been specified. The check found tran_cache_size is not specified. zFS IOEFSPRM configuration option tran_cache_size is ignored. No action is required for the tran_cache_size option.

The system continues processing.

Report this problem to the system programmer.

Remove the user override check parameter TRANS(EXISTENCE) from the PARMS string (via HZSPRMxx or MODIFY hzsproc) or specify the user override check parameter TRANS(ABSENCE) in order to make the check succeed.

N/A

IOEZHK1

z/OS File System

See IOEFSPRM in *z/OS File System Administration*.

N/A

N/A

N/A

zFS configuration option `client_cache_size` is not specified.

The check override parameter CLIENT(EXISTENCE) is specified in the PARMS string (via HZSPRMxx or MODIFY hzsproc) to verify the configuration option client_cache_size has been specified. The check found client_cache_size is not specified. zFS IOEFSPRM configuration option client_cache_size is ignored. No action is required for the client_cache_size option.

The system continues processing.

Report this problem to the system programmer.

System programmer response:

Remove the user override check parameter CLIENT(EXISTENCE) from the PARMS string (via HZSPRMxx or MODIFY hzsproc) or specify the user override check parameter CLIENT(ABSENCE) in order to make the check succeed.

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See IOEFSPRM in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0074E	<p>The zFS USER_CACHE_SIZE specification of <i>nn</i>M results in a user cache configured as <i>nn</i> cache spaces, each with <i>nn</i> 8K spaces. <i>nn</i> cache spaces are registered with the zEDC Express Service <i>nn</i> are registered <i>nn</i> are not registered <i>nn</i> file systems out of <i>nn</i> are compressed. Compressed file systems are better supported when all cache spaces are registered with the zEDC Express service.</p>
-------------------	--

Explanation:

zFS is operating with compressed aggregates and the user cache is not optimally configured. All cache spaces should be registered with the zEDC Express service for better performance.

System action:

The system continues processing.

Operator response

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See the topics on "User file cache", "Fixed storage," and "Compress" in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0075I

The zFS USER_CACHE_SIZE specification of *nn*M results in a user cache configured as *nn* cache spaces, each with *nn* 8K spaces.
nn cache spaces are registered with the zEDC Express Service
nn are registered
nn are not registered
nn file systems out of *nn* are compressed.
An exception condition has not been detected. Relative to zEDC support, zFS is appropriately configured.

Explanation:

zFS is operating with compressed aggregates and the user cache is not optimally configured. All cache spaces should be registered with the zEDC Express service for better performance.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See the topics on "User file cache", "Fixed storage," and "Compress" in [*z/OS File System Administration*](#).

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0076I

zFS does not have enough performance data to do an analysis.

Explanation:

zFS needs more time to collect more performance data for Health Checker to analyze.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0077I

File cache performance is good. No further action is needed.

Explanation:

Based on the collected statistics, overall file cache hit ratios are good. You do not need to take any action.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0078I

Meta cache performance is good. No further action is needed.

Explanation:

Based on the collected statistics, overall meta cache hit ratios are good. You do not need to take any action.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0079I**File cache total significant intervals = *number*
number% of collected intervals showed sub-optimal performance data.****Explanation**

File cache total significant intervals is the number of observed intervals where file cache request threshold are met. % of collected intervals showing suboptimal performance data is the ratio of the total collected intervals with sub-optimal file cache performance to the file cache total significant intervals. This message indicates file cache performance is not optimal. A report of the time periods with sub-optimal performance will be displayed after this message. No action is required.

zFS considers file cache performance good when the file cache hit ratio, (the ratio of requests to the cache that find the desired data in the cache divided by the total requests to the cache), is above 90%. In many cases, active systems that read file data stored in zFS file systems will achieve this hit ratio since many UNIX applications read the same files multiple times and zFS aggressively performs read-ahead for files. However, there are some situations where the cache hit ratio might be less, such as system startup when applications might be reading large volumes of data for the first time, or during administration actions such as moving ownership of file systems and other zFS administration commands. If data that is not read frequently is accessed, a temporary dip in the hit ratio might occur. zFS will list the time periods (if infrequent) when the hit ratio was smaller than desired to assist the administrator in determining if the situation is transient or persistent. Another situation where the zFS file cache hit ratio is consistently lower than 90% is when the client applications are aggressively caching outside of zFS. In this case, the zFS is a second-level cache and will generally not see frequent reads of cached data. In these situations the zFS cache hit ratio might be much lower than 90%. Typically, situations where customer applications cache zFS file data aggressively will see file cache hit ratios closer to 50-60%.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [Performance tuning](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZH0080I**Meta cache total significant intervals = *meta_cache_significant*
% of collected intervals showing sub-optimal performance data =
*percentage*****Explanation:**

Meta cache total significant intervals is the number of observed intervals where meta cache request threshold are met. % of collected intervals showing sub-optimal performance data is the ratio of the total collected intervals with sub-optimal meta cache performance to the meta cache total significant intervals. This message indicates meta cache performance is not optimal. A report of the associated intervals with sub-optimal performance will be displayed following this message. No action is required.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Automation:

N/A

Reference documentation:

See [Performance tuning](#) in *z/OS File System Administration*.

Routing code:

N/A

Descriptor code:

N/A

IOEZH0083I Found no zFS file system with exception status.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZHCK1

Source:

z/OS File System

Reference documentation:

See [zfsadm fsinfo](#) in *z/OS File System Administration*.

Routing code:

N/A

Descriptor code:

N/A

IOEZH0084I Found zFS file system(s) in exception state.

Explanation

Found zFS file systems that meet one of more of the exceptional states, such as low on space, disk I/O error, software error, XCF communication error, out of space error, dynamic grow failure, dynamic grow disabled, mounted with non-use of high availability option and RWSHARE, having at least one V4 directory or aggregate disk size exceeded space monitoring threshold.

A report of the affected file systems will be displayed after this message. In the report, the file system name, owning system name and the file system status will be shown.

The report includes the file system name, owning system name and the file system states. The states contain abbreviated values. For quick reference, the values are defined in a Legend string at the end of the report.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZCHK1

Source:

z/OS File System

Reference documentation:

See [Converting an existing v4 directory to an extended \(v5\) directory](#) in *z/OS File System Administration*.

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

IOEZHO088I

**Could not get completed file system information error=*error_code*,
reason=*reason_code*.**

Explanation:

Getting file system information is not completed due to the displayed return code and reason code. Health check displays as much information as it can.

System action:

The system continues processing.

Operator response:

N/A

System programmer response:

N/A

Problem determination:

N/A

Module:

IOEZCHK1

Source:

z/OS File System

Automation:

N/A

Routing code:

N/A

Descriptor code:

N/A

Appendix A. Reason codes

Reason codes are listed by hexadecimal value and include a description of the action required to correct the error. Each reason code is made up of four bytes, `cccc rrrr`, where `cccc` is a halfword reason code qualifier and `rrrr` is the halfword reason code.

The two high-order bytes of the reason codes that are returned contain a value that is used to qualify the contents of the two low-order bytes. Reason code qualifiers for zFS are found within the range `X'EF01' - X'FFFF'`. For information about reason codes outside this range, see [Description and location information](#) in *z/OS UNIX System Services Messages and Codes*.

DF01rrrr reason codes

0001	The File Exporter Exit Routine determined that an unmount was attempted but the file system is busy.
-------------	---

Action

Try the operation again. If it continues to fail, contact the service representative.

0002	The File Exporter Exit Routine request for an open token failed.
-------------	---

Action

An Open request failed. Retry the open. If it continues to fail, contact the service representative.

0003	The File Exporter Exit Routine request for a token failed.
-------------	---

Action

Try the operation again. If it continues to fail, contact the service representative.

0005	The File Exporter Exit Routine was unable to find its Glue Anchor.
-------------	---

Action

Try the operation again. If it continues to fail, contact the service representative.

0006	The File Exporter Exit Routine was unable to obtain a recovery block.
-------------	--

Action

Try the operation again. If it continues to fail, contact the service representative.

0007	The File Exporter Exit Routine was called with an invalid operation.
-------------	---

Action

Try your request again. If it continues to fail, contact the service representative.

0008	A vnode operation has no File Exporter Exit Routine functional support.
-------------	--

Action

Contact the service representative.

0009	There was not enough memory for the Glue Anchor in the UNIX System Services address space.
-------------	---

Action

Try your request again. If it continues to fail, contact the service representative.

000A	There was not enough memory to create a Cell Pool for recovery blocks.
-------------	---

Action

Try your request again. If it continues to fail, contact the service representative.

000B	There was not enough memory to create a Cell Pool for wait structure blocks.
-------------	---

Action

Try your request again. If it continues to fail, contact the service representative.

000D	The File Exporter Exit Routine detected an invalid recovery block id.
-------------	--

Action

Contact the service representative.

000E	The File Exporter Exit Routine detected a missing FID.
-------------	---

Action

Contact the service representative.

000F	The File Exporter Exit Routine detected an invalid osictl header id.
-------------	---

Action

Contact the service representative.

0010	The File Exporter Exit Routine detected that an osictl has no optional parameter.
-------------	--

Action

Contact the service representative.

0011	The File Exporter Exit Routine detected that an invalid server initialization parameter.
-------------	---

Action

Contact the service representative.

0012	The File Exporter Exit Routine detected an invalid revoke token argument length.
-------------	---

Action

Contact the service representative.

0013	The File Exporter Exit Routine detected an invalid async grant token argument length.
-------------	--

Action

Contact the service representative.

0014 **The File Exporter Exit Routine detected an invalid osictl request.**

Action

Contact the service representative.

0017 **The File Exporter Exit Routine detected an invalid change debug argument length.**

Action

Contact the service representative.

0018 **The File Exporter Exit Routine detected an invalid lockctl option parameter.**

Action

Contact the service representative.

DF02rrrr reason codes

0001 **The File Exporter Exit Routine sleep function could not get a wait structure from the Cell Pool.**

Action

Contact the service representative.

DF03rrrr reason codes

0001 **The File Exporter Exit Routine made an invalid request.**

Action

Contact the service representative.

DF04rrrr reason codes

0001 **DFSKERN vnode cache is too small for the number of file systems exported and the specified virtual memory cache maximum files.**

Action

The maximum number of file systems attached is the vnode cache size (`_IOE_VNODE_CACHE_SIZE` environment variable) minus VM cache maximum files (`_IOE_VM_MAX_FILES` environment variable) minus 320. Increase the vnode cache environment variable or decrease the VM files maximum. For details, see the topic on environment variables in *z/OS Distributed File Service SMB Administration*.

0002 **The File Exporter Exit Routine made an invalid request. This is not supported.**

Action

To export this zFS file system, you must run zFS specifying `sysplex=filesys` in the IOEFSPRM file and ensure that the zFS file system is mounted non-sysplex aware. See [z/OS File System Administration](#) for further details.

Action

Contact the service representative.

600A	File system not mounted.
-------------	---------------------------------

Action

Ensure that the file system is mounted. Retry the operation. If it continues to fail, contact the service representative.

600B	Vget failure not yet reported.
-------------	---------------------------------------

Action

Contact the service representative.

600C	Vget attributes fetch failure.
-------------	---------------------------------------

Action

Contact the service representative.

600D	Access failure.
-------------	------------------------

Action

Contact the service representative.

600E	Failure getting attr specified in OSI struct.
-------------	--

Action

Contact the service representative.

600F	Failure getting index on volstat.
-------------	--

Action

Contact the service representative.

6010	Bad unique on vol_vget.
-------------	--------------------------------

Action

Contact the service representative.

6011	No vnode found for FID.
-------------	--------------------------------

Action

This reason code, when coupled with a return code of ENOENT, simply means that an operation was presented to zFS for a previously deleted file. If a different return code is presented, then contact the service representative.

6012	Could not open anode for vnode.
-------------	--

Action

This could occur if a task was trying to open a file that was removed by another task at the same time. It can also occur if there was a recent system or zFS outage. If however, it repeatedly occurs and a recent system outage did not occur, then contact the service representative.

6013	Error from <code>osi_getvnode</code>.
Action	
Contact the service representative.	
6014	Error from <code>osi_getvnode</code> token.
Action	
Contact the service representative.	
6015	SAF CKACC returned error.
Action	
The user did not have the correct permission on the object for the operation requested. Change the permissions on the object or have a user that is authorized execute the operation.	
6016	SAF AUDIT returned error.
Action	
Contact the service representative.	
6017	Vm_schedule error.
Action	
Contact the service representative.	
6018	Bad input name from z/OS UNIX.
Action	
Ensure that the file/directory name is not longer than 256 characters. If the name length is valid, contact the service representative.	
6019	Creating a dir on <code>vn_create</code>.
Action	
Contact the service representative.	
601A	Directory operation on a vnode that is not a directory.
Action	
Contact the service representative.	
601B	Invalid operation on a zero-link count vnode.
Action	
Ensure that the object still exists. If it does, contact your service representative.	
601C	Error creating a new file.
Action:	
Try to resolve the problem using the return code. If the problem persists, contact the service representative.	
601D	Error inserting a name in a directory.

Action

Contact the service representative.

601E	Create failure.
-------------	------------------------

Action

Contact the service representative.

601F	Vnode operation failed to get vnode token in a directory.
-------------	--

Action

Contact the service representative.

6020	Link failure.
-------------	----------------------

Action

Contact the service representative.

6021	Try to hard link a directory, not allowed.
-------------	---

Action

You cannot link to a directory. Correct the object being linked and try again.

6023	Error creating link, the link name already exists.
-------------	---

Action

Ensure that the link name does not already exist.

6024	A write operation to a read only filesystem.
-------------	---

Action

Ensure that the file system is mounted read-write.

6025	Error looking up a file.
-------------	---------------------------------

Action

Ensure the object exists. If it does, contact the service representative.

6026	Error fetching file attributes.
-------------	--

Action

Contact the service representative.

6027	A lookup attempted on an object that was not a directory.
-------------	--

Action

Ensure that the name containing the object being looked up is a directory.

6028	A required name entry in a directory was not found.
-------------	--

Action

If the return code is 129 (ENOENT), correct the directory name and try again. Otherwise, contact the service representative.

6029	Entry in directory found but inode was stale.
-------------	--

Action

Contact the service representative.

602A	Unexpected error searching a directory.
-------------	--

Action

Contact the service representative.

602B	Error on mkdir request.
-------------	--------------------------------

Action

Contact the service representative.

602C	Name given by mkdir already exists.
-------------	--

Action

Ensure that the directory name being created does not already exist.

602D	Error creating a sub-directory anode.
-------------	--

Action

Ensure that you are not trying to create more than 65533 subdirectories in a version 4 directory. Otherwise, contact the service representative.

602E	Error adding . and . . entries to the directory.
-------------	---

Action

Contact the service representative.

6030	Error truncating file during open.
-------------	---

Action

Contact the service representative.

6031	Try to remove link from a directory, not allowed.
-------------	--

Action

Ensure that the pathname is not a directory.

6032	Check 2 owners error.
-------------	------------------------------

Action

Ensure that the user has permission to do the operation.

6033	Attempt to delete a non-empty directory.
-------------	---

Action

Ensure that the directory is empty before attempting to remove it.

6034	Error attempting to remove a name from a directory.
-------------	--

Action

Ensure that you are not attempting to remove the . or .. entry from the directory.

If the return code is 133 (ENOSPC), you are trying to remove an entry from a directory that was created before zFS V1R13 and is less than 7 K; therefore it is stored in fragments. However, the file system is owned on a zFS V1R13 or later system and needs a free 8 K block to do the remove. You must grow the file system to make some 8 K blocks available.

Otherwise, contact the service representative.

6035 General vn_remove error.

Action

Contact the service representative.

6036 Rmdir error.

Action

Contact the service representative.

6037	General rename error.
------	-----------------------

Action

Contact the service representative.

6038	Target directory is a descendant of source directory.
-------------	--

Action

Ensure that the target of an mv is not the descendent of the source directory. For example, mv /a /a/b.

6039	Target of rename is a directory while the source of rename is not.
------	--

Action

Ensure that you are not attempting to rename a file to an existing directory.

603A	Target of rename is a directory but was not empty.
------	--

Action

Ensure that the directory that is the target of a rename is empty.

603B	Error replacing name with new name in a directory.
-------------	---

Action

Contact the service representative.

603C	General setattr error.
------	------------------------

Action

Contact the service representative.

603D	SAF chgmode error.
-------------	---------------------------

Action

Ensure that you are authorized to change the permissions.

603E	SAF chgowner error.
-------------	----------------------------

Action

Ensure that you are authorized to change the owner and that the new owner and group is valid.

603F	Attempt to truncate a directory, not allowed.
-------------	--

Action

Ensure that you are not attempting to truncate a directory.

6040	SAF clrsetid error.
-------------	----------------------------

Action

Contact the service representative.

6041	SAF chaudit error.
-------------	---------------------------

Action

Contact the service representative.

6042	Guardtime check on setattr failed.
-------------	---

Action

Contact the service representative.

6043	Operation required owner privileges and user did not have them.
-------------	--

Action

Ensure that you are the owner of the object.

6044	Setattr truncation error.
-------------	----------------------------------

Action

Contact the service representative.

6045	Symlink error.
-------------	-----------------------

Action

Contact the service representative.

6046	Symlink name already exists.
-------------	-------------------------------------

Action

Ensure that the symlink name does not already exist.

6047	Error writing symlink contents.
-------------	--

Action

Contact the service representative.

6048	Attempt to read past EOF.
-------------	----------------------------------

Action

Ensure that you are not attempting to read past the end of the file.

6049	General file read error.
-------------	---------------------------------

Action

Contact the service representative.

604A	General file write error.
-------------	----------------------------------

Action

Check the return code to determine why the write failed.

604B	Uiomove failed, probably due to user address space errors.
-------------	---

Action

Contact the service representative.

604C	Process limit was exceeded on write request.
-------------	---

Action

Ensure that the user is not attempting to exceed the process file size limit.

604D	Process limit was exceeded because it was not allowed to increase size of file.
-------------	--

Action

Ensure that you are not attempting to exceed the process file size limit.

604E	General VM cache failure.
-------------	----------------------------------

Action

If the return code is ENOSPC (133), the file system or aggregate is full. If this is the case, make more space available. Otherwise, contact the service representative.

604F	Readlink failure.
-------------	--------------------------

Action

Contact the service representative.

6050	Readlink for a vnode that was not a symlink.
-------------	---

Action

Ensure that the target of a readlink is actually a symbolic link.

6051 **Readdir error.**

Action

Contact the service representative.

6052 **File truncate error.**

Action

Contact the service representative.

6053 **File sync error.**

Action

Contact the service representative.

6054 **Syntax error in PARM for MOUNT request.**

Action

Ensure that the PARM parameter string on the MOUNT command for the zFS file system is syntactically correct. MOUNT PARMs are case sensitive.

6055 **Error with issuing a LOCATE call on an HFS-compatible aggregate.**

Action

Ensure that the zFS file system named on the MOUNT command has the same name as the VSAM Linear Data Set (and the HFS compatibility mode aggregate) that contains the file system.

6056 **HFS-compatible mount, - name not VSAM LDS.**

Action

Ensure that the zFS file system named on the MOUNT command has the same name as the VSAM Linear Data Set (and the HFS compatibility mode aggregate) that contains the file system.

6057 **HFS-compatible mount, - syntax in IOEFSPRM file.**

Action

Ensure that the syntax of the AGGRFULL option in the IOEFSPRM file is valid.

6058 **HFS-compatible mount, - error attaching aggregate or was found not to be HFS-compatible.**

Action

Ensure that the HFS compatibility mode aggregate that is being MOUNTed has a single read-write file system and the file system name is the same as the aggregate name (that is, the VSAM Linear Data Set name). This can also mean that you attempted to move ownership of a zFS sysplex-aware read-write file system to a system that does not support zFS sysplex-aware (that is, it is z/OS V1R12 but is running sysplex=off). This is not allowed.

A mount of a zFS file system that fails with return code 138 (ENXIO) can indicate that the file system is already mounted on another system that is not in the same sysplex but is sharing DASD accessibility. zFS prevents data corruption by not allowing a zFS file system to be mounted read-write on two systems that are not in a shared

file system environment. In this case, verify that the file system is not already mounted read-write on another system. zFS prevents a system from writing to a zFS aggregate that is mounted read-write on another system. Another possible cause is the VSAM linear data set was not formatted as a ZFS file system. Also, look in the system log for messages from other z/OS components related to this aggregate (data set) name. Otherwise, contact the service representative.

605A **Attempt to mount a file system R/W on a R/O aggregate.**

Action

Ensure that a file system that is contained in an aggregate that has been attached read-only is mounted read-only. Or, attach the aggregate read-write.

605B **Attempt to mount a file system R/W that is inherently R/O.**

Action

Contact the service representative.

605C **Fsfull config option syntax error in config file.**

Action

Ensure that the syntax of the FSFULL option in the IOEFSPRM file is valid.

605D **No root directory found for file system.**

Action

Contact the service representative.

605E **File system busy or unable to be mounted.**

Action

There is another operation in progress that conflicts with this operation. Try your request again. If it continues to fail, contact the service representative.

6060 **Invalid PFS control request command.**

Action

Ensure that the application is specifying a valid pfscctl request.

6061 **Failed obtaining request block for cross memory call.**

Action

Contact the service representative.

6062 **Failed obtaining data buffer for cross memory call.**

Action

Contact the service representative.

6063 **Failed obtaining data from the user buffer.**

Action

Contact the service representative.

6064	Failed writing data into the user buffer.
Action	
Contact the service representative.	
6065	Failed obtaining user credential information by <code>osi_getcred</code>.
Action	
Contact the service representative.	
6066	The caller is not UID 0.
Action	
Ensure that you are properly authorized for this operation. This subcommand requires UID 0 or READ access to the SUPERUSER.FILESYS.PFSCTL profile in the z/OS UNIXPRIV class.	
6067	Failure in <code>osi_wait</code> service.
Action	
Contact the service representative.	
6068	PFS control command failed.
Action	
Check the return code to determine why the pfsctl function failed.	
6069	Insufficient space on aggregate.
Action	
Ensure that there is sufficient space on the aggregate.	
606A	Container does not have an allocation handle.
Action	
Contact the service representative.	
606B	Incorrect aggregate handle.
Action	
Contact the service representative.	
606C	Incorrect anode handle.
Action	
Contact the service representative.	
606D	Index is out of range or anode is not valid.
Action	
Contact the service representative.	
606E	Incorrect, undefined or inconsistent arguments.

Action

Verify that the new size on grow is greater than the current size. The **zfsadm aggrinfo** command displays the current size in 1K blocks. Divide this amount by 8 to get the current size in 8K blocks. If this is not the case for the reason code, contact the service representative.

6070	The acl or plist file handle is not valid.
------	--

Action

Contact the service representative.

6071	Incorrect, undefined or inconsistent command value.
------	---

Action

Contact the service representative.

6072	A volume table index is not valid.
------	------------------------------------

Action

Contact the service representative.

6073	Fid unique identifier does not match.
------	---------------------------------------

Action

Contact the service representative.

6074	Fid volume ID does not match.
------	-------------------------------

Action

Contact the service representative.

6075	Incorrect volume handle.
------	--------------------------

Action

Contact the service representative.

6076	Incorrect file handle.
------	------------------------

Action

Contact the service representative.

6077	File cannot have negative link count.
------	---------------------------------------

Action

Contact the service representative.

6078	No additional volumes are in the aggregate.
-------------	--

Action

Contact the service representative.

6079	No additional containers are in the volume.
------	---

Action

Contact the service representative.

607A	Block is past the allocated end of the container.
-------------	--

Action

Contact the service representative.

607B	Block address is too large.
-------------	------------------------------------

Action

Contact the service representative.

607C	Specified device already has an aggregate.
-------------	---

Action

Contact the service representative.

6080	Specified anode index already exists.
-------------	--

Action

Contact the service representative.

6081	A transaction must be specified.
-------------	---

Action

Contact the service representative.

6082	Block is interior to the container but is not allocated.
-------------	---

Action

Contact the service representative.

6083	Specified anode index does not exist.
-------------	--

Action

Contact the service representative.

6084	The anode does not contain a file.
-------------	---

Action

Contact the service representative.

6085	The container does not contain a volume table.
-------------	---

Action

Contact the service representative.

6086	The bitmap is not consistent with the superblock.
-------------	--

Action

Contact the service representative.

6087	The superblock does not lead to an AVL.
-------------	--

Action

Contact the service representative.

6088	The file does not point to an acl container.
-------------	---

Action

Contact the service representative.

6089	Container is stored inline or fragmented.
-------------	--

Action

Contact the service representative.

608A	Container is stored inline or blocked.
-------------	---

Action

Contact the service representative.

608B	Container is stored fragmented or blocked.
-------------	---

Action

Contact the service representative.

608C	An anode being deleted is still open.
-------------	--

Action

Contact the service representative.

608D	An anode is not empty.
-------------	-------------------------------

Action

Contact the service representative.

608E	Other containers using these blocks for copy-on-write purposes.
-------------	--

Action

Contact the service representative.

608F	Read operation extends past container length.
-------------	--

Action

Contact the service representative.

6090	Specified size of the status area is too large.
-------------	--

Action

Contact the service representative.

6091	Supplied status data extends too far.
-------------	--

Action

Contact the service representative.

6093	Insufficient quota on volume.
-------------	--------------------------------------

Action

Contact the service representative.

6094	Copy-on-write is not available for this container.
-------------	---

Action

Contact the service representative.

6095	Block allocation of log is in error.
-------------	---

Action

Contact the service representative.

6096	Block allocation for bitmap ended in error.
-------------	--

Action

Contact the service representative.

6097	Error in copy-on-write reference.
-------------	--

Action

Contact the service representative.

6098	Management of multiple quota procedures not yet implemented.
-------------	---

Action

Contact the service representative.

6099	Initialization entry point called more than once.
-------------	--

Action

Contact the service representative.

609A	Object or module not properly initialized.
-------------	---

Action

Contact the service representative.

609B	Tried to insert a block that is already allocated.
-------------	---

Action

Contact the service representative.

609D	Incorrect block insertion parameters.
-------------	--

Action

Contact the service representative.

609E	A rock was specified whose value was zero.
-------------	---

Action

Contact the service representative.

609F	A volume was referenced while going offline.
-------------	---

Action

Contact the service representative.

60A0	An aggregate was referenced while going offline.
-------------	---

Action

Contact the service representative.

60A2	The container does not contain a volume header.
-------------	--

Action

Contact the service representative.

60A3	Volume does not have a root directory.
-------------	---

Action

Contact the service representative.

60A5	Current operation is not finished.
-------------	---

Action

Contact the service representative.

60A6	Anode must be copy-on-write.
-------------	-------------------------------------

Action

Contact the service representative.

60A7	Write operation encountered an inconsistent state.
-------------	---

Action

Contact the service representative.

60A8	Existing fragment group cannot be extended.
-------------	--

Action

Contact the service representative.

60A9	Aggregate has some open volumes.
-------------	---

Action

Contact the service representative.

60AA	Volume has some open anodes.
-------------	-------------------------------------

Action

Contact the service representative.

60AB	A volume being deleted is still open.
-------------	--

Action

Contact the service representative.

60AC	New length is inconsistent with storage method.
-------------	--

Action

Contact the service representative.

60AF	Iterator value indicates previous is last.
-------------	---

Action

Contact the service representative.

60B0	Aux type inappropriate for non-directory files.
-------------	--

Action

Contact the service representative.

60B1	Anode block and offset pair is not valid.
-------------	--

Action

Contact the service representative.

60B9	Block reference to indirect block has inconsistent header.
-------------	---

Action

Contact the service representative.

60BC	Event already in set.
-------------	------------------------------

Action

Contact the service representative.

60BD	Event still part of some set.
-------------	--------------------------------------

Action

Contact the service representative.

60BE	Arguments are either incorrect, inconsistent, or not allowed.
-------------	--

Action

Contact the service representative.

60C0	Device does not exist.
-------------	-------------------------------

Action

Contact the service representative.

60C2	The log/buffer package was re initialized.
-------------	---

Action

Contact the service representative.

60C3	The specified memory region is too small.
-------------	--

Action

Contact the service representative.

60C4	The buffer is not valid.
-------------	---------------------------------

Action

Contact the service representative.

60C5	The log has not been recovered.
-------------	--

Action

Contact the service representative.

60C6	The specified log was not valid.
-------------	---

Action

Contact the service representative.

60C7	The specified transaction is not valid.
-------------	--

Action

Contact the service representative.

60C8	The specified transaction is no longer valid.
-------------	--

Action

Contact the service representative.

60C9	The specified buffer could not be deleted.
-------------	---

Action

Contact the service representative.

60CA	The specified transaction is active.
-------------	---

Action

Contact the service representative.

60CB	The specified transaction has ended.
-------------	---

Action

Contact the service representative.

60CC	The specified transaction has completed.
-------------	---

Action

Contact the service representative.

60CD	The specified transaction has committed.
-------------	---

Action

Contact the service representative.

60D1	No buffers were available to satisfy the request.
-------------	--

Action

Contact the service representative.

60D4	Cannot log additional devices.
-------------	---------------------------------------

Action

Contact the service representative.

60D5	The requested log already exists.
-------------	--

Action

Contact the service representative.

60D8	There is no available space on the log device.
-------------	---

Action

Contact the service representative.

60D9	One or more arguments are not valid.
-------------	---

Action

Contact the service representative.

60DA	Could not read log in for recovery.
-------------	--

Action

Contact the service representative.

60DD	An error was encountered.
-------------	----------------------------------

Action

Contact the service representative.

60DE	An incorrect reference was found.
-------------	--

Action

Contact the service representative.

60DF	An incorrect log record was found.
-------------	---

Action

Contact the service representative.

60E0	VAn internal system check ended in error.
-------------	--

Action

Contact the service representative.

60E1	An incorrect log page was encountered.
-------------	---

Action

Contact the service representative.

60E2	A null pointer reference was detected.
-------------	---

Action

Contact the service representative.

60E3	An incorrect pointer was detected.
-------------	---

Action

Contact the service representative.

60E4	An incorrect record type was found during redo.
-------------	--

Action

Contact the service representative.

60E5	An incorrect record type was found during undo.
-------------	--

Action

Contact the service representative.

60E6	An I/O error was detected during recovery.
-------------	---

Action

Contact the service representative.

60E7	Recovery is complete.
-------------	------------------------------

Action

Contact the service representative.

60E8	Recovery must be run on this aggregate.
-------------	--

Action

Contact the service representative.

60E9	The requested operation would block.
-------------	---

Action

Contact the service representative.

60EA	A log page with an incorrect pass number was found.
-------------	--

Action

Contact the service representative.

60EB	Aggregate is already attached.
-------------	---------------------------------------

Action

The aggregate specified is already attached. Ensure you have specified the correct aggregate name.

60ED	Aggregate is not attached.
-------------	-----------------------------------

Action

Contact the service representative.

60EE	Aggregate may need to be recovered.
-------------	--

Action

Contact the service representative.

60EF	File System is locally mounted.
-------------	--

Action

Contact the service representative.

60F2	Name is too long.
-------------	--------------------------

Action

Contact the service representative.

60FE	Out of memory.
-------------	-----------------------

Action

Contact the service representative.

6108	Defined parameters exceed the maximum.
-------------	---

Action

Contact the service representative.

6109	Internal parsing error.
-------------	--------------------------------

Action

Contact the service representative.

610A	Too many values specified after a CMD_SINGLE switch.
-------------	---

Action

Contact the service representative.

610B	Too many parameters specified.
-------------	---------------------------------------

Action

Contact the service representative.

610C	Two or more mutually exclusive parameters used.
-------------	--

Action

Contact the service representative.

610D	Specify the minimum number of arguments.
-------------	---

Action

Contact the service representative.

610E	Unrecognized or ambiguous command name.
-------------	--

Action

Contact the service representative.

610F	Unrecognized or ambiguous switch name.
-------------	---

Action

Contact the service representative.

6111	Specify the required number of parameters.
-------------	---

Action

Contact the service representative.

6112	Specify a valid argument value.
-------------	--

Action

Contact the service representative.

6113	Unable to find appropriate commands.
-------------	---

Action

Contact the service representative.

6114	Token too large.
-------------	-------------------------

Action

Contact the service representative.

6115	File system in process of deletion.
-------------	--

Action

Contact the service representative.

6117	File System busy with setstatus operation.
-------------	---

Action

Contact the service representative.

611B	File System is being moved.
-------------	------------------------------------

Action

Contact the service representative.

611C	File System operation in progress.
-------------	---

Action

Contact the service representative.

611D	File System has been deleted.
-------------	--------------------------------------

Action

Contact the service representative.

611E	File System busy with clone.
-------------	-------------------------------------

Action

Contact the service representative.

611F	File System is damaged.
-------------	--------------------------------

Action

Contact the service representative.

6120	File System is out of service.
-------------	---------------------------------------

Action

Contact the service representative.

6121 Aggregate is being detached.

Action

Contact the service representative.

6123

Fsync attempted on an object that is not a directory or a regular file.

Action

Ensure that the fsync operation is issued against a directory or a regular file.

6124	Directory sync attempted on an object that is not a directory.
------	--

Action

Ensure that the directory sync operation is issued against a directory.

6125	Fsync attempted on non-regular file.
------	--------------------------------------

Action

Ensure that the fsync operation is issued against a regular file.

6126	Buffer size is 0 on a readdir attempt.
------	--

Action

Ensure that the readdir buffer size is larger than 0.

6128 **The AGGR_FORMAT structure has invalid values.**

Action

Correct the values and issue the request again.

6130	First usable block past super block.
-------------	---

Action

Contact your service representative.

6132	Initialempty value too big.
------	-----------------------------

Action

Initial empty value should be ignored. Contact your service representative.

6133	Specified logsize too many blocks for aggregate size.
------	---

Action

Do not specify a logsize greater than the size of the aggregate.

6136	Error detaching the aggregate after the operation completed.
------	--

Action

Contact the service representative.

6137	Cannot open VSAM Linear Data Set for format.
-------------	---

Action

Ensure the VSAM Linear Data Set exists and is accessible before retrying the operation.

6138	Aggregate is already formatted.
-------------	--

Action

The aggregate specified cannot be formatted because it already contains a ZFS file system. Either specify overwrite (using caution) or if the incorrect aggregate was specified, correct the aggregate name.

6139	File system partition too small.
-------------	---

Action

Contact the service representative.

613A	Too few disk blocks for file system.
-------------	---

Action

Try decreasing the logsize value or allocating a larger VSAM Linear Data Set. If the problem persists, contact the service representative.

613B	Disk sizes too small.
-------------	------------------------------

Action

Contact the service representative.

613C	Failed to initialize the device.
-------------	---

Action

Contact the service representative.

613D	Error initializing the device.
-------------	---------------------------------------

Action

Contact the service representative.

613E	Aggregate attach failed.
-------------	---------------------------------

Action

Contact the service representative.

613F	Internal error creating the file system.
-------------	---

Action

Contact the service representative.

6140	Failed to obtain storage while formatting.
-------------	---

Action

Contact the service representative.

6141	Internal error locating super block.
-------------	---

Action

Contact the service representative.

6142	Error reading super block.
-------------	-----------------------------------

Action

Contact the service representative.

6143	Error creating or writing new aggregate structures.
-------------	--

Action

Contact the service representative.

6144	Error syncing device to disk.
-------------	--------------------------------------

Action

Contact the service representative.

6145	Internal error creating basic aggregate disk structures.
-------------	---

Action

Contact the service representative.

6146	Internal error locating aggregate structure.
-------------	---

Action

Contact the service representative.

6147	Error obtaining aggregate size information.
-------------	--

Action

Contact the service representative.

6148	Formatting an attached aggregate.
-------------	--

Action

Verify you are using the correct aggregate name. If the problem persists, contact the service representative.

614A	GetfACL request was for an undefined ACL type value.
-------------	---

Action

If you specified an invalid ACL type, correct it. Otherwise, contact the service representative.

614B	Internal error reading ACL from disk.
-------------	--

Action

Retry the operation and if it continues to fail, contact the service representative.

614C	Object is not a directory so a File Default or Directory Default ACL cannot be changed.
-------------	--

Action

Either use the correct directory or modify the command to operate on a File ACL.

614D	SetfACL request was for an undefined ACL type value.
-------------	---

Action

If you specified an invalid ACL type, correct it. Otherwise, contact the service representative.

614E	Internal error writing ACL to disk.
-------------	--

Action

Retry the operation and if it continues to fail, contact the service representative.

614F	An invalid command code was requested.
-------------	---

Action

Correct the invalid command code and retry.

6150	Lookup failure for volser for aggregate attach.
-------------	--

Action

Contact the service representative.

6151	Write past largest supported zFS file size or past the file allocation limit, which includes space for file location maps.
-------------	---

Action

Remove data from the file to make room for the new data, or write the new data to another file.

6161	Admin thread abended.
-------------	------------------------------

Action

An application has issued a pfsctl call (or a user has issued a zfsadm command that resulted in a pfsctl call) and the thread running the request in zFS has abended. Contact the service representative.

6162	Recovery routine given control.
-------------	--

Action

The user has been cancelled or has issued CTRL-escape. The operation has completed. No action is required.

6163	An error occurred in the recovery routine.
-------------	---

Action

Contact the service representative.

6164	An error occurred in the efs_getvntok call.
-------------	--

Action

Contact the service representative.

6165	An attempt to format an aggregate with a log size greater than 128M (16384 8K blocks).
-------------	---

Action

Specify a logsize of less than or equal to 16384 blocks or do not specify a logsize.

6166	The vnode_cache_limit is not between vnode_cache_size and 7340032.
-------------	---

Action

Specify a vnode_cache_limit between the current vnode_cache_size and 7340032.

6167	Could not obtain a vnode due to vnode_cache_limit being reached.
-------------	---

Action

Increase the vnode_cache_limit.

6168	The vnode cache size is not between 32 and the vnode cache limit.
-------------	--

Action

Specify a vnode_cache_size between 32 and the vnode_cache_limit.

6200	Filesys delete pfscctl did not find file system.
-------------	---

Action

File system was not found. Correct the file system name and try again.

6201	Unexpected file system open failure.
-------------	---

Action

Contact the service representative.

6202	File system is mounted.
-------------	--------------------------------

Action

Unmount the file system before deleting the file system.

6203	File system getstatus call failed.
-------------	---

Action

Contact the service representative.

6204	File system setstatus call failed.
-------------	---

Action

Contact the service representative.

6205	File system sync status call failed.
-------------	---

Action

Contact the service representative.

6206	File system unclone of forward file system failed.
-------------	---

Action

Contact the service representative.

6207	File system sync failed.
-------------	---------------------------------

Action

Contact the service representative.

6208	File system setstatus on 2 file systems failed.
-------------	--

Action

Contact the service representative.

6209	File system destroy of file systems failed.
-------------	--

Action

Contact the service representative.

620A	File system sync of aggregate failed.
-------------	--

Action

Contact the service representative.

620B	File system close failed.
-------------	----------------------------------

Action

Contact the service representative.

620C	File system was found busy with another operation.
-------------	---

Action

Another file system operation is in progress. Reissue this operation when the other file system operation has completed.

620D	Aggregate is attached R/O.
-------------	-----------------------------------

Action

The file system cannot be deleted because the aggregate is attached read-only.

620E	Aggregate specified does not match aggregate found.
-------------	--

Action

The aggregate specified for the file system to be deleted is incorrect.

6223	Could not open file system.
-------------	------------------------------------

Action

Contact the service representative.

6224 **File system busy with other operation.**

Action

Contact your service representative.

6225 **R/W or backup file system is mounted.**

Action

Contact your service representative.

6239 **Aggregate not found.**

Action

The aggregate specified cannot be found. Correct the aggregate name and try again.

623A **Bad unquiesce handle passed in.**

Action

The handle passed to unquiesce is incorrect. Use the quiesce handle received from the quiesce. This may also mean that the aggregate was not in a quiesced state.

623B **Error closing file system.**

Action

Contact the service representative.

623E **Aggregate not found.**

Action

The aggregate specified cannot be found. Correct the aggregate name and try again.

623F **Aggregate already quiesced or quiescing, or unquiescing.**

Action

The aggregate is already quiesced or a quiesce or unquiesce operation is already in progress for this aggregate. If you are doing a grow operation or a backup operation, the aggregate should not be already quiesced.

6238 **Cannot forward this PFCTL to a downlevel system.**

Action

The zFS owner of the aggregate is downlevel and cannot process this pfscctl.

6240 **Aggregate sync operation failed.**

Action

Contact the service representative.

6241 **One or more file systems busy with other operations.**

Action

Use the error code to determine why the file system could not be quiesced during the aggregate quiesce.

6246	File system is busy.
-------------	-----------------------------

Action

The file system is busy with another file system operation. Try again later.

6248	Failure retrieving status from opened file system.
-------------	---

Action

Contact the service representative.

624D	Invalid parms in request.
-------------	----------------------------------

Action

An invalid parameter was passed. Correct the parameter and try again.

624E	Aggregate not found.
-------------	-----------------------------

Action

The aggregate specified cannot be found. Correct the aggregate name and try again.

624F	Failure retrieving status from aggregate.
-------------	--

Action

Contact the service representative.

6252	Aggregate not found.
-------------	-----------------------------

Action

The aggregate specified cannot be found. Correct the aggregate name and try again.

6253	Detach cannot proceed because aggregate is quiesced.
-------------	---

Action

Unquiesce the aggregate before detaching it.

6254	Detach cannot proceed because file systems are mounted.
-------------	--

Action

Unmount the file systems before detaching the aggregate.

6255	Detach failed because some file systems could not be closed.
-------------	---

Action

Contact the service representative.

6256	Detach failed because a long running admin command is in process.
-------------	--

Action

Attempt to use the appropriate zfsadm command with the -cancel option to interrupt the log running command. Reissue the detach after the long running operation completes. If the problem persists, contact the service representative.

6258 **Aggregate not found.**

Action

The aggregate specified cannot be found. Correct the aggregate name and try again.

6259 **Aggregate could not be grown. It is attached in R/O mode.**

Action

The aggregate must be attached read-write to do a grow.

625C **Could not grow quota in compat aggregate R/W file system.**

Action

Contact the service representative.

625D **Unexpected error growing aggregate.**

Action

If the error code is 8, you may have specified a new size smaller than the current size, a new size of zero with a zero secondary allocation on the data set, or there may not be sufficient space on the volume. Otherwise, contact the service representative.

6260 **Aggregate not attached.**

Action

The aggregate must be attached before a file system can be created in it.

6261 **Aggregate is quiesced.**

Action

Unquiesce the aggregate before attempting to create a file system in it.

6262 **File system name to be created ends in .bak.**

Action

Do not use a file system name that ends in .bak. This is reserved for backup file systems created via clone. Use a file system name that does not end in .bak.

6263 **File system with name already exists.**

Action

Use a file system name that does not already exist.

6264 **Unexpected error attaching newly created file system.**

Action

Contact the service representative.

6265	Error adding file system to aggregate file system table on disk.
Action	
Contact the service representative.	
6266	Error creating root directory in new file system.
Action	
Contact the service representative.	
6267	Error opening root directory vnode in new file system.
Action	
Contact the service representative.	
6268	Error adding . and .. to root directory.
Action	
Contact the service representative.	
6269	Aggregate attached R/O, cannot create file systems.
Action	
The aggregate must be attached read-write to create a file system in it.	
626A	Aggregate is COMPAT aggregate, cannot create new file systems on it.
Action	
You cannot create another read-write file system in a compatibility mode aggregate.	
626B	Quota smaller than minimum allowed.
Action	
Specify a quota of at least 128 when creating a file system. (128 means 128K).	
626D	AGGR_ATTACH passed in is invalid.
Action	
Ensure that you did not specify both AGGR_NBS and AGGR_NONBS and try again.	
626E	Aggregate already attached.
Action	
None.	
626F	Aggregate attach failed.
Action	
Ensure that the aggregate has been formatted. Otherwise, ensure that you are not mounting or attaching an aggregate read-write when it is already mounted or attached read-write on another (non-sysplex coordinated) system. Otherwise, contact the service representative.	

629F **AGGR_FORMAT length invalid.**

Action

The length of the AGGR_FORMAT structure is incorrect. It must be 116 bytes.

62A0 **AGGR_FORMAT version invalid.**

Action

The version of the AGGR_FORMAT structure is incorrect. It must be the value 1.

62A1 **AGGR_FORMAT reserved field invalid.**

Action

The reserved field of the AGGR_FORMAT structure must contain zeros.

62A2 **AGGR_ID eyecatcher invalid.**

Action

The eyecatcher in the AGGR_ID structure is incorrect. It must be AGID.

62A3 **AGGR_ID length invalid.**

Action

The length of the AGGR_ID structure is incorrect. It must be 84 bytes.

62A4 **AGGR_ID version invalid.**

Action

See [z/OS File System Administration](#) for the versions of AGGR_ID that are supported.

62A5 **AGGR_ID reserved invalid.**

Action

AGGR_ID reserved fields must be set to binary zeros.

62A6 **AGGR_ID name too long.**

Action

The aggregate name in the AGGR_ID must be no longer than 44 characters.

62A7 **AGGR_ID name not specified.**

Action

The aggregate name in the AGGR_ID must be specified.

62AA **AGGR_ATTACH eyecatcher invalid.**

Action

The eyecatcher in the AGGR_ATTACH structure is incorrect. It must be AGAT.

62AB **AGGR_ATTACH length invalid.**

Action

The length of the AGGR_ATTACH structure is incorrect. It must be 268 bytes.

62AC **AGGR_ATTACH version invalid.**

Action

The version of the AGGR_ATTACH must be 1.

62AD **AGGR_ATTACH reserved invalid.**

Action

AGGR_ATTACH reserved fields must be set to binary zeros.

62AE **AGGR_ATTACH both NBS and NONBS specified.**

Action

You can specify NBS or NONBS or neither in the AGGR_ATTACH. You cannot specify both.

62AF **AGGR_ATTACH increment invalid.**

Action

The increment percent must be greater than or equal to 1 and less than or equal to 99.

62B0 **AGGR_ATTACH threshold invalid.**

Action

The threshold percent must be greater than or equal to 1 and less than or equal to 99.

62B3 **AGGR_STATUS eyecatcher invalid.**

Action

The AGGR_STATUS eyecatcher must be AGST.

62B4 **AGGR_STATUS length invalid.**

Action

The AGGR_STATUS length must be 172, 172 or 260 bytes for AGGR_STATUS, AGGR_STATUS2 or AGGR_STATUS3 respectively.

62B5 **AGGR_STATUS version invalid.**

Action

The AGGR_STATUS version must be 1.

62B8 **FS_ID eyecatcher invalid.**

Action

The FS_ID eyecatcher must be FSID.

62B9 **FS_ID length invalid.**

Action

The FS_ID length must be 140 bytes for an FS_ID or 200 bytes for an FS_ID2.

62BA **FS_ID version invalid.**

Action

The FS_ID version must be 1 or 2.

62BB **FS_ID reserved invalid.**

Action

FS_ID reserved fields must be binary zeros.

62BC **FS_ID aggregate name too long.**

Action

The FS_ID aggregate name must be less than or equal to 44 characters.

62BD **FS_ID aggregate name not specified.**

Action

The FS_ID aggregate name must be specified.

62BE **FS_ID fs name not specified.**

Action

The FS_ID file system name must be specified.

62BF **FS_ID fs name too long.**

Action

The FS_ID file system name must be less than or equal to 44 characters.

62C0 **FS_ID fs name contains invalid characters.**

Action

The FS_ID file system name can only contain a valid zFS file system name. See [*z/OS File System Administration*](#) for information on valid zFS file system names.

62C3 **FS_STATUS eyecatcher invalid.**

Action

The FS_STATUS eyecatcher must be FSST.

62C4 **FS_STATUS length invalid.**

Action

The FS_STATUS length must be 396 bytes.

62C5 **FS_STATUS version invalid.**

Action

The FS_STATUS version must be 1.

62C8 **FILESYS_DATA eyecatcher invalid.**

Action

The FILESYS_DATA eyecatcher must be FSDT.

62C9 **FILESYS_DATA length invalid.**

Action

The length of the FILESYS_DATA must be 92 bytes.

62CA **FILESYS_DATA version invalid.**

Action

The FILESYS_DATA version must be 1.

62CB **FILESET_DATA reserved invalid.**

Action

FILESYS_DATA reserved fields must be binary zeros.

62CC **FILESYS_DATA quota invalid.**

Action

The quota must be less than or equal to 4294967295. This is 1 K blocks.

62CF **Syscall_parmlist bad offset at parm0.**

Action

The offset at parms0 must be 32.

62D0 **Syscall_parmlist parms0 truncated.**

Action

The pfsctl argument must be increased to contain the entire required argument.

62D1 **Syscall_parmlist parms0 missing.**

Action

A non-zero parms0 must be specified.

62D2 **Syscall_parmlist bad offset at parms1.**

Action

The offset at parms1 must contain the correct offset for the required parameters.

62D3 **Syscall_parmlist parms1 truncated.**

Action

The pfscctl argument must be increased to contain the entire required argument.

62D4 **Syscall_parmlist parms1 missing.**

Action

A non-zero parms1 must be specified.

62D5 **Syscall_parmlist bad offset at parms2.**

Action

The offset at parms2 must contain the correct offset for the required parameters.

62D6 **Syscall_parmlist parms2 truncated.**

Action

The pfscctl argument must be increased to contain the entire required argument.

62D7 **Syscall_parmlist parms2 missing.**

Action

A non-zero parms2 must be specified.

62D8 **Syscall_parmlist bad offset at parms3.**

Action

The offset at parms3 must contain the correct offset for the required parameters.

62D9 **Syscall_parmlist parms3 truncated.**

Action

The pfscctl argument must be increased to contain the entire required argument.

62DA **Syscall_parmlist parms3 missing.**

Action

A non-zero parms3 must be specified.

62DB **Syscall_parmlist unexpected opcode.**

Action

Provide a valid subcommand opcode.

62DC **Syscall_parmlist unexpected command.**

Action

Provide a valid zFS pfscctl command.

62E1 **Anode open failed.**

Action

Contact the service representative.

62E2	Object at index is not volume.
-------------	---------------------------------------

Action

Contact the service representative.

62E3	Volume open failed.
-------------	----------------------------

Action

Contact the service representative.

62E4	Cannot get enough memory for aggregate list.
-------------	---

Action

Contact the service representative.

62E5	Cannot get non-growing aggregate list. Try later.
-------------	--

Action

The number of attached aggregates grew several times during a request to list attached aggregate names. Try again later.

62E6	A Syscall parmlist is invalid.
-------------	---------------------------------------

Action

Provide a valid subcommand opcode or ensure the argument is large enough to hold the syscall parmlist.

6300	Getacl request for default object acls, but not on a directory.
-------------	--

Action

Issue the request against a directory.

6301	zfsadm format error when -size > default size and grow error.
-------------	---

Action

Ensure that space is available on the volume(s). If space is available and it still fails, contact the service representative.

6310	no storage for idcams parms.
-------------	-------------------------------------

Action

If storage is not available, run IDCAMS as a batch job.

6311	idcams failed.
-------------	-----------------------

Action

Examine the messages in the message buffer for the reason for the failure.

If your application does not display the message buffer (that contains messages from the IDCAMS program), try running the zfsadm define command since it does display the message buffer. (The error may depend on the userid running the application/command or other conditions (for example, space on the volume) may change.)

6312 no aggregate in AGGR_DEFINE.**Action**

Specify the aggregate name in the AGGR_DEFINE.

6313 aggregate too long.**Action**

Specify an aggregate name of 44 characters or less.

6314 catalog too long.**Action**

Specify a catalog name of 44 characters or less.

6315 data class too long.**Action**

Specify a data class name of 8 characters or less.

6316 management class too long.**Action**

Specify a management class name of 8 characters or less.

6317 model too long.**Action**

Specify a model name of 44 characters or less.

6318 model catalog too long.**Action**

Specify a model catalog name of 44 characters or less.

6319 storage class too long.**Action**

Specify a storage class name of 8 characters or less.

631A invalid number of volumes.**Action**

Specify a number of volumes as a number between (and including) 0 to 59.

631B volume too long.**Action**

Specify volume names of 6 characters or less.

631C invalid primary.

Action

Specify a primary allocation size of greater than 0.

631D	invalid secondary.
------	--------------------

Action

Specify a non-negative secondary allocation size (0 is allowed).

631E	invalid space unit.
------	---------------------

Action

If a space unit is specified, it must be cylinders (1), Kilobytes (2), Megabytes (3), Records (4), or Tracks (5).

631F **name is blank.**

Action

Specify a non-blank volume for each of the count of volumes specified in the numVolumes field in AGGR_DEFINE.

6320 AGGR_DEFINE eyecatcher invalid.

Action

The eyecatcher for the AGGR_DEFINE must be AGDF.

6321 AGGR_DEFINE length invalid.

Action

The AGGR_DEFINE length must be 772.

6322 AGGR DEFINE version invalid.

Action

The AGGR_DEFINE version must be 1.

6323 AGGR_DEFINE reserved invalid.

Action

AGGR_DEFINE reserved fields must be binary zeros.

6324	z/OS UNIX FILESYSTEM name invalid.
------	------------------------------------

Action

Specify allowable zFS file system name characters. These are documented in [*z/OS File System Administration*](#) under ioeagfmt.

6325	Aggregate name invalid.
------	-------------------------

Action

Specify allowable zFS aggregate name characters. These are documented in [*z/OS File System Administration*](#) under ioagfmt.

6400 **USER** cache size too big.

Action

Ensure that the -user_cache_size value specified is less than 65536M.

6401	Could not create a data space.
-------------	---------------------------------------

Action

Contact the service representative.

6402	Could not extend a data space.
-------------	---------------------------------------

Action

Contact the service representative.

6403	Could not create a new admin thread.
-------------	---

Action

Contact the service representative.

6404	Tried to exceed log buffer cache size limit, which is 1024M bytes.
-------------	---

Action

Specify a log cache size of 1024M or less.

6405	Could not extend the log buffer dataspace.
-------------	---

Action

Contact the service representative.

6406	Attempt to set metacache too small, min is 1M.
-------------	---

Action

Specify a meta cache size of 1M or more.

6407	Attempt to set metacache too big, max is 64G.
-------------	--

Action

Specify a meta cache size of 64G or less.

6408	Attempt to change metacache size failed.
-------------	---

Action

The meta cache size could not be made smaller than the current size. Continue to run with the current size or try again later.

6409	Attempt to set tran cache too small, min is 200.
-------------	---

Action

Specify a tran cache size of 200 or more.

640A	Attempt to set tran cache too big, max is 10000000.
-------------	--

Action

Specify a tran cache size of 10000000 or less.

640B	Attempt to change tran cache size failed.
-------------	--

Action

The tran cache size could not be made smaller than the size specified. Continue to run with the current size or try again later.

640C	Attempt to set vnode cache too small, min is 1000.
-------------	---

Action

Specify a vnode cache size of 1000 or more.

640D	Attempt to set vnode cache too big, max is 10000000.
-------------	---

Action

Specify a vnode cache size that is less than or equal to the maximum of 10000000.

640E	Attempt to change vnode cache size failed.
-------------	---

Action

The vnode cache size could not be made smaller than the size specified. Continue to run with the current size or try again later.

640F	Attempt to set log cache size too small.
-------------	---

Action

Specify a log cache size of 2M or more.

6410	Attempt to set metacache backing too small, min is 1M.
-------------	---

Action

Specify a metaback cache size of 1M or larger.

6411	Attempt to set metacache backing too big, max is 2G.
-------------	---

Action

Specify a metaback cache size of 2048M or smaller.

6412	Attempt to change metacache backing size failed.
-------------	---

Action

Contact the service representative.

6413	CFG_OPTION structure invalid.
-------------	--------------------------------------

Action

Correct the CFG_OPTION structure.

6414	option value string (co_string) was not syntactically correct in CFG_OPTION struct.
-------------	--

Action

Correct the value for the co_string.

6415	Sync interval smaller than 11 second minimum.
-------------	--

Action

Specify a sync interval of 11 or greater.

6416	Attempt to make the USER cache too small.
-------------	--

Action

Specify a configuration option USER_CACHE_SIZE value of 10M or greater.

6417	Thread pool set too small.
-------------	-----------------------------------

Action

Specify a configuration option value of 1 or greater.

6418	Too many threads attempted.
-------------	------------------------------------

Action

Specify a configuration option value of 256 or less.

6419	Both file system and mount specified.
-------------	--

Action

Specify either z/OS UNIX FILE SYSTEM_NAME or zFS file system name, but not both.

641A	The value specified for the -long_cmd_threads option is invalid. The value must be two numbers that are separated by a comma. The first number must be in the range 1-3. The second number must be in the range 1-64. Refer to the description of the IOEFSPRM long_cmd_threads option in z/OS Distributed File Service zFS Administration for more information about what the two numbers represent.
-------------	--

Action

Correct the value for the -long_cmd_threads option and try the command again.

6420	Both aggregate and mount specified.
-------------	--

Action

Specify either z/OS UNIX FILE SYSTEM_NAME or zFS aggregate name, but not both.

6421	mount name not found.
-------------	------------------------------

Action

Specify the correct z/OS UNIX file system (mounted) name.

6422	filesystem not found.
-------------	------------------------------

Action

Specify the correct zFS file system name or the correct aggregate that contains the zFS file system.

6423	The specified aggregate could not be found.
-------------	--

Action

Specify the correct aggregate name, ensure it is mounted and try the operation again.

6424	File system name not unique.
-------------	-------------------------------------

Action

The zFS file system name is not unique. Specify the aggregate that the file system resides in.

6425	File system name not found.
-------------	------------------------------------

Action

The file system is not mounted or attached. Ensure that the file system name is correct. Then mount or attach the file system in accordance with the requirements of the command in use.

6426	syscall_parmlist bad offset at parm4.
-------------	--

Action

Specify the correct offset in parm4.

6427	syscall_parmlist parm4 truncated.
-------------	--

Action

Specify a size to include all the parms.

6428	syscall_parmlist parm4 missing.
-------------	--

Action

Specify parm4.

6429	Cannot honor request to enforce unique filesystem because there are already non-unique file systems.
-------------	---

Action

Either do not attempt to disallow duplicate file systems or ensure that file system names are unique.

642A	Cannot specify aggrgrow and noaggrgrow on the zfsadm attach command.
-------------	---

Action

You can specify AGGRGROW or NOAGGRGROW or neither in the AGGR_ATTACH. You cannot specify both.

642B	Mount name cannot be used.
-------------	-----------------------------------

Action

The fsid_mtname field was used in the FS_ID2 for Delete File System. This is not allowed. Use fsid_name. (The file system cannot be mounted.)

6440	Internal failure.
Action	
An internal failure occurred. Contact the service representative.	
6441	Internal failure.
Action	
An internal failure occurred. Contact the service representative.	
6443	Internal failure.
Action	
An internal failure occurred. Contact the service representative.	
6444	CLIENT cache size too big.
Action	
Ensure that the -client_cache_size value specified is not greater than 65536M.	
6445	CLIENT cache size too small.
Action	
Ensure that the -client_cache_size value specified is 10M or greater.	
6446	The threshold,increment pair is invalid. Each value in the pair must be in the range 1 to 99.
Action	
Correct the value and try the operation again.	
6500	Received XCF request from system we do not know about.
Action	
Wait a few minutes to see if the situation corrects itself. If the problem persists, contact the service representative.	
6501	Received XCF request about aggregate we do not know about.
Action	
Wait a few minutes to see if the situation corrects itself. If the problem persists, contact the service representative.	
6503	Received a bad AGNOTIFY packet from another system.
Action	
Retry the operation. If the problem persists, contact the service representative.	
6504	Could not make an AGNOTIFY communication to another system.

Action

Retry the operation. If the problem persists, contact the service representative. Look in the system log for messages indicating XCF communications problems.

6505	Received a bad FSNOTIFY packet from another system.
-------------	--

Action

Retry the operation. If the problem persists, contact the service representative.

6506	Received XCF request about file system we do not know about.
-------------	---

Action

Wait a few minutes to see if the situation corrects itself. If the problem persists, contact the service representative.

6508	Received an FSNOTIFY_REGISTER request for a file system that this system already knows about.
-------------	--

Action

Restart this system or the system that this system thinks owns the aggregate.

6509	Received an FSNOTIFY request for an aggregate that is marked pending.
-------------	--

Action

Wait a few minutes to see if the situation corrects itself. If the problem persists, restart the system that this system thinks owns the aggregate.

650A	Received an FSNOTIFY request for a file system that we do not have any information about.
-------------	--

Action

Unmount the file system or restart the system that owns the aggregate.

650B	Received an FSNOTIFY request to rename a file system and the new name already exists.
-------------	--

Action

Contact your service representative.

650C	Could not make an FSNOTIFY communication to another system.
-------------	--

Action

Retry the operation. If the problem persists, contact the service representative. Look in the system log for messages indicating XCF communications problems.

650D	Directory or file seclabel is not zero.
-------------	--

Action

When linking to a target file that is a special character file, neither the directory nor the target file can have a seclabel. Either remove the seclabel or do not attempt to create this link.

650F	Directory and file seclabels not equivalent during link.
-------------	---

Action

When linking a file, the target file's seclabel and the directory seclabel must be equivalent. Make the seclabels equivalent or do not attempt the link.

6512	List Systems buffer and buffer length must both be zero or both non-zero.
-------------	--

Action

The offsets at parms0 and parms1 must be both either zero or both non-zero. It is not valid for one to be zero and the other to be non-zero.

6513	List Systems buffer is too small.
-------------	--

Action

The buffer and length passed in parms0 and parms1 is too small to hold the result to be returned. Use the size field returned (offset to size is in parms2) to determine the size of the buffer needed.

6514	List Systems has no system names to return.
-------------	--

Action

There are no system names to be returned. Zero is returned in the size field (offset to size is in parms2).

6515	List Systems cannot get memory to hold list of systems.
-------------	--

Action

zFS could not get the storage to hold the list of systems. The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

6516	XCF communication to another system failed.
-------------	--

Action

Retry the operation. Look in the system log for messages indicating XCF communications problems. If the problem persists, contact the service representative.

6517	Received a bad pfscctl packet from another system.
-------------	---

Action

An invalid command or subcommand was issued to a remote system. Retry the operation. If the return code is 157 (EMVSERR), contact the service representative.

6518	System name too long.
-------------	------------------------------

Action

Specify a system name that is eight characters or less.

6519	Unknown system name.
-------------	-----------------------------

Action

Specify a valid system name.

651A	Buffer is too small for result.
-------------	--

Action

Retry the operation. If the problem persists, contact the service representative.

651B	Buffer and buffer length must both be zero or both non-zero.
-------------	---

Action

Retry the operation. If the problem persists, contact the service representative.

651C	Parms0 is not zero.
-------------	----------------------------

Action

The offset at parms0 must be set to zero.

651D	Parms1 is not set to zero.
-------------	-----------------------------------

Action

The offset at parms1 must be set to zero.

651E	Parms2 is not set to zero.
-------------	-----------------------------------

Action

The offset at parms2 must be set to zero.

651F	Parms3 is not set to zero.
-------------	-----------------------------------

Action

The offset at parms3 must be set to zero.

6520	Parms4 is not set to zero.
-------------	-----------------------------------

Action

The offset at parms4 must be set to zero.

6521	Parms5 is not set to zero.
-------------	-----------------------------------

Action

The offset at parms5 must be set to zero.

6522	Parms6 is not set to zero.
-------------	-----------------------------------

Action

The offset at parms6 must be set to zero.

652B	Could not load IRRSSB00.
-------------	---------------------------------

Action

The load module IRRSSB00 (normally located in SYS1.CSSLIB in z/OS 1.5 or later) cannot be loaded. Ensure that this load module exists.

652C	Unsupported function (change seclabel).
-------------	--

Action

The change seclabel function is not supported prior to z/OS 1.5.

652D	XCF communications failure forwarding the command to the zFS owning system.
-------------	--

Action

The current owning system has likely gone down while processing the command. Try the command again after ownership of the file system is moved. If the problem persists, contact the service representative.

6609	An attempt was made to set the token cache size smaller than the number of allocated tokens.
-------------	---

Action

Determine the current number of allocated tokens by issuing modify zfs,query,stk. Reissue the command specifying a token_cache_size larger than the current number of allocated tokens and less than the maximum of 2621440.

660A	An attempt to set the token cache size larger than the currently supported maximum size was attempted.
-------------	---

Action

Retry the operation with a value in the range 20480 to 20000000.

6611	A conflicting administration command is already running.
-------------	---

Action

You are attempting to run a zFS cache resize command and another operation is already running that interferes with it such as another resize or a zFS shutdown. Retry the resize command. If the problem persists, contact the service representative.

6617	Sysplex communication error, XCF send failed.
-------------	--

Action

Determine the source of the error such as system down, or system hang. If error cannot be determined, contact the service representative.

6619	A server system does not understand or support the packet sent by a client system. It may be that there is an incompatibility in the required service levels among members of a sysplex.
-------------	---

Action

If the service levels of all sysplex members are at the same level, contact the service representative.

6638	Aggregate is no longer owned by system.
-------------	--

Action

The request was sent to the zFS owning system but when the request arrived at that system, the aggregate was no longer owned by that system. Reissue the request again.

6639	Aggregate is busy with another administration operation.
-------------	---

Action

Wait until the aggregate operation is complete and the aggregate is unquiesced.

663B	Could not sync aggregate at original owner system.
-------------	---

Action

Contact the service representative.

663D	An unmount is pending.
-------------	-------------------------------

Action

Wait for unmount to complete before issuing request.

663E	Aggregate is damaged.
-------------	------------------------------

Action

Detach the aggregate and attempt to attach it again.

663F	Vnode is unusable, possibly due to an outage/restart of zFS on a system.
-------------	---

Action

Try restarting the application, if that fails to correct the problem then try re-mounting the file system.

6640	Could not attach aggregate at the new owner and original owner failed.
-------------	---

Action

Detach the aggregate and attempt to attach it again.

6641	New owner not ready for takeover.
-------------	--

Action

Wait until the new owner is ready or try another owner.

6649	An open file lost some updates due to a server crash.
-------------	--

Action

Contact the service representative.

664A	An aggregate/file system administration function failed because a system went down.
-------------	--

Action

Try the command again after the system completes termination or completes its restart.

664B	A sysplex config command was issued but this is not a sysplex environment.
-------------	---

Action

Do not issue a sysplex config command when you are not in a sysplex environment.

664C	STAT_API eye catcher is wrong.
-------------	---------------------------------------

Action

The eyecatcher for the STAT_API structure must be STAP.

664D **STAT_API version number is wrong.**

Action

The version number for the STAT_API structure must be 1 or 2.

664E **STAT_API size is wrong.**

Action

The size of the STAT_API structure must be at least 48 bytes.

664F **Buffer size is too small.**

Action

The size of the buffers for a statistics query were not large enough.

6650 **Report type is wrong.**

Action

The statistics subcommand is invalid.

6651 **Flag value invalid.**

Action

The STAT_API flags value is invalid. It must be either 0x80 (for reset) or 0x00 (for no reset).

6652 **The aggregate specified for the pfsctl operation is owned by a system which may be down.**

Action

A compatibility mode aggregate should be taken over by another system. Reissue the operation after it has been taken over. If it is not taken over by another system in a reasonable amount of time, you should unmount it and then mount it again.

6653 **Statistics Locking Query Parm0 value is invalid.**

Action

The Parm0 value for the Statistics Locking Query subcommand must be 32.

6654 **Statistics Storage Query Parm0 value is invalid.**

Action

The Parm0 value for the Statistics Storage Query subcommand must be 32.

6655 **Statistics Locking Query Parm1 value is invalid.**

Action

The Parm1 value for the Statistics Locking Query subcommand must be 80.

6656 **Statistics Storage Query Parm1 value is invalid.**

Action

The Parm1 value for the Statistics Storage Query subcommand must be 80.

6657 **Statistics Iocounts Query Parm0 value is invalid.**

Action

The Parm0 value for the Statistics Iocounts Query subcommand must be 32.

6658 **Statistics Iocounts Query Parm1 value is invalid.**

Action

The Parm1 value for the Statistics Iocounts Query subcommand must be 80.

6659 **Statistics Iobyaggr Query Parm0 value is invalid.**

Action

The Parm0 value for the Statistics Iobyaggr Query subcommand must be 32.

665A **Statistics Iobyaggr Query Parm1 value is invalid.**

Action

The Parm1 value for the Statistics Iobyaggr Query subcommand must be 80.

665B **Statistics Iobydasd Query Parm0 value is invalid.**

Action

The Parm0 value for the Statistics Iobydasd Query subcommand must be 32.

665C **Statistics Iobydasd Query Parm1 value is invalid.**

Action

The Parm1 value for the Statistics Iobydasd Query subcommand must be 80.

665D **Detach of an unowned aggregate failed.**

Action

Contact the service representative.

665E **Remount requested but aggregate is busy with another operation.**

Action

Wait for the aggregate operation to complete before issuing remount.

665F **client_reply_storage specified is too small.**

Action

client_reply_storage must be at least 2M.

6660 **client_reply_storage specified is too big.**

Action

client_reply_storage must not exceed 128M.

6661	Client reply pool not changed since this is not a sysplex aware environment.
-------------	---

Action

Only issue this command in a sysplex aware environment. If this is a sysplex aware environment, then contact your service representative.

6662	Insufficient storage available to resize the client reply storage pools.
-------------	---

Action

Use the zfsadm config command to decrease the size of other caches or retry the command specifying a smaller client reply storage size.

6663	Insufficient storage for the XCF storage pool.
-------------	---

Action

The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

666B	AGGR_STATUS reserved field not zero.
-------------	---

Action

The reserved fields in AGGR_STATUS must be 0.

666C	Internal failure during MOUNT.
-------------	---------------------------------------

Action

Contact the service representative.

666D	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

666F	Internal failure during a write lock operation.
-------------	--

Action

Contact the service representative.

6670	Internal failure during a read lock operation.
-------------	---

Action

Contact the service representative.

6671	Internal failure during a truncate operation.
-------------	--

Action

Contact the service representative.

6672	Internal failure during a lock operation.
-------------	--

Action

Contact the service representative.

6673	Internal failure during a rename operation.
-------------	--

Action

Contact the service representative.

6674	Internal failure during a quiesce operation.
-------------	---

Action

Contact the service representative.

6675	Reserved bytes must be zero.
-------------	-------------------------------------

Action

The sa_fill bytes in the STAT_API must be zeros.

6676	Reserved bytes must be zero.
-------------	-------------------------------------

Action

The sa_reserve bytes in the STAT_API must be zeros.

6677	Aggregate is quiesced but caller does not want to wait.
-------------	--

Action

The aggregate is quiesced but the caller has made a request that cannot wait. The call returns with a failure.

6680	Failure during attempt to add items to the buffer cache.
-------------	---

Action

A failure occurred while trying to get storage to add items to the buffer cache. The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

6681	Failure during attempt to add items back into the buffer cache.
-------------	--

Action

A failure occurred while trying to get storage to add items back into the buffer cache. The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

6682	Failure during attempt to add items to the metaback buffer cache.
-------------	--

Action

A failure occurred while trying to get storage to add items to the metaback buffer cache. The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

6683	Failure during attempt to add items to the log buffer cache.
-------------	---

Action

A failure occurred while trying to get storage to add items to the log buffer cache. The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

6684 Failure during attempt to add items to the transaction cache.
Action

A failure occurred while trying to get storage to add items to the transaction cache. The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

6685 Failure during attempt to add items back into the transaction cache.
Action

A failure occurred while trying to get storage to add items back into the transaction cache. The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

6686 Failure during attempt to add items to the vnode cache.
Action

A failure occurred while trying to get storage to add items to the vnode cache. The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes.

6689 Failure to obtain write lock.
Action

A failure occurred during an attempt to obtain a write lock during a mount, unmount or file operation. Contact the service representative.

668A Failure to obtain write lock.
Action

A failure occurred during an attempt to obtain a write lock during a mount operation. Contact the service representative.

6690 Out of storage during a mount operation.
Action

A failure to obtain storage occurred during a mount operation. Try to take actions to release storage and retry the mount.

6696 Internal failure during mount operation.
Action

An internal failure occurred during a mount operation. Contact the service representative.

6697 Internal failure during mount operation.
Action

An internal failure occurred during a mount operation. Contact the service representative.

6698 Administration operation failed because zFS is shutting down.
Action

An administration operation failed because zFS is shutting down. Retry the operation when zFS is restarted.

6699 Internal failure.

Action

An internal failure occurred. Contact the service representative.

66A0	Failure to obtain lock.
-------------	--------------------------------

Action

An internal failure occurred. Contact the service representative.

66A1	Failure to obtain lock.
-------------	--------------------------------

Action

An internal failure occurred. Contact the service representative.

66A4	Failure to obtain lock.
-------------	--------------------------------

Action

An internal failure occurred. Contact the service representative.

66A6	Internal failure during list aggregates operation.
-------------	---

Action

An internal failure occurred during a file system delete operation. Contact the service representative.

66A7	Internal failure during list file systems operation.
-------------	---

Action

An internal error occurred during a list file system operation. Contact the service representative.

66A8	Internal failure during quiesce operation.
-------------	---

Action

An internal failure occurred during an aggregate quiesce operation. Contact the service representative.

66AA	Internal failure during aggregate operation.
-------------	---

Action

An internal failure occurred during an aggregate operation. Contact the service representative.

66AC	Internal failure during attach operation.
-------------	--

Action

An internal failure occurred during an attach operation. Contact the service representative.

66AD	Internal failure during file system administration operation.
-------------	--

Action

An internal failure occurred during a file system administration operation. Contact the service representative.

66B1	Internal failure during mount operation.
-------------	---

Action

An internal failure occurred during a mount operation. Contact the service representative.

66B2	Internal error trying to place a hold on a vnode.
-------------	--

Action

Contact the service representative.

66B3	Internal error locking a vnode.
-------------	--

Action

Contact the service representative.

66B4	Could not initialize a vnode.
-------------	--------------------------------------

Action

If the aggregate is disabled, unmount and remount the file system. Retry the operation. If the problem persists, contact the service representative.

66B5	Could not clear isuid or isgid bits.
-------------	---

Action

If the aggregate is disabled, unmount and remount the file system. Retry the operation. If the problem persists, contact the service representative.

66B8	Could not change the size of the vnode cache.
-------------	--

Action

If there is not enough storage available for the new vnode cache size, then try decreasing the size of other caches. Otherwise, contact the service representative.

66B9	Internal error. Could not set link count of root directory.
-------------	--

Action

Contact the service representative.

66BA	Internal error setting vnode key.
-------------	--

Action

Contact the service representative.

66BB	Internal error creating the file system.
-------------	---

Action

Contact the service representative.

66BC	Internal error getting the status for a vnode.
-------------	---

Action

Contact the service representative.

66BD	Internal error updating directory link count.
-------------	--

Action

If the aggregate is disabled, unmount and remount the file system. Retry the operation. If the problem persists, contact the service representative.

66BE	Internal error updating mtime or ctime.
-------------	--

Action

If the aggregate is disabled, unmount and remount the file system. Retry the operation. If the problem persists, contact the service representative.

66C0	Internal error, user file cache, bad page pointers.
-------------	--

Action

Contact the service representative.

66C1	Internal error, user file cache, bad page counts.
-------------	--

Action

Contact the service representative.

66C2	Internal error, user file cache, schedule dirty data failed.
-------------	---

Action

Contact the service representative.

66C3	Internal error, user file cache, sync of pending IO failed.
-------------	--

Action

Contact the service representative.

66C4	Internal error, user file cache, could not truncate file.
-------------	--

Action

Contact the service representative.

66C5	Internal error, user file cache, could not resize cache.
-------------	---

Action

Contact the service representative.

66C6	Internal error, user file cache, could not read a file.
-------------	--

Action

Contact the service representative.

66C7	Internal error, user file cache, could not write a file.
-------------	---

Action

Contact the service representative.

66C8	Internal error, user file cache, could not get file attributes.
-------------	--

Action

Contact the service representative.

66C9	Internal error, user file cache, could not schedule file data at close.
-------------	--

Action

Contact the service representative.

66CA	Internal error, user file cache, could not fsync file data.
-------------	--

Action

Contact the service representative.

66CB	Internal error, user file cache, could not set file attributes.
-------------	--

Action

Contact the service representative.

66CC	Internal error, user file cache, error handling page fault.
-------------	--

Action

Contact the service representative.

66CD	Internal error, user file cache, could not obtain resize lock.
-------------	---

Action

Contact the service representative.

66CE	Internal error, user file cache, could not free cache data for file.
-------------	---

Action

Contact the service representative.

66CF	Internal error, user file cache, could not obtain file block information.
-------------	--

Action

Contact the service representative.

66D0	Internal error, aggregate component, could not detach aggregate.
-------------	---

Action

Contact the service representative.

66D1	Aggregate name longer than 44 characters.
-------------	--

Action

An aggregate name is longer than 44 characters. Change the aggregate name to be 44 characters or less.

66D2	Aggregate name contains invalid characters.
-------------	--

Action

An aggregate name contains invalid characters. Ensure that the aggregate name consists only of valid characters. See [z/OS File System Administration](#) for information on valid characters in an aggregate name.

66D3 **zFS File System name longer than 44 characters.**

Action

A zFS file system name is longer than 44 characters. Change the zFS file system name to be 44 characters or less.

66D4 **zFS File System name contains invalid characters.**

Action

A zFS file system name contains invalid characters. Ensure that the zFS file system name consists only of valid characters. See [z/OS File System Administration](#) for information on valid characters in a zFS file system name.

66D5 **z/OS UNIX file system name longer than 44 characters.**

Action

A z/OS UNIX file system name is longer than 44 characters. Change the z/OS UNIX file system name to be 44 characters or less.

66D6 **z/OS Unix File System name contains invalid characters.**

Action

A z/OS UNIX file system name contains invalid characters. Ensure that the z/OS UNIX file system name consists only of valid characters. See [z/OS File System Administration](#) for information on valid characters in a z/OS UNIX file system name.

66D7 **System name too long.**

Action

Specify a system name that is 8 characters or less.

66DF **Syscall_parmlist bad offset at parms1.**

Action

The STAT_API structure must be 48 bytes. The offset at parms1 must be 80 (32+48).

66E0 **Timestamp too large for 5 bytes.**

Action

Cannot store a timestamp that is larger than 5 bytes.

66E1 **Aggregate could not be opened on this system.**

Action

This can mean that you attempted to move ownership of a zFS sysplex-aware read-write file system to a system that does not support zFS sysplex-aware (that is, it is z/OS V1R12 but is running with zFS sysplex=off). This is not allowed.

66E4 **Aggregate possibly attached on another system.**

Action

Determine if the aggregate is mounted or attached on another system. If the problem persists contact your service representative.

66E5 Internal error during ACL processing.**Action**

An internal failure occurred during ACL processing. Contact the service representative.

66E6 Length for parm0 is bad.**Action**

Fix the length and retry.

66E7 Length for parm1 is bad.**Action**

Fix the length and retry.

66E8 Length for parm2 is bad.**Action**

Fix the length and retry.

66E9 Length for parm3 is bad.**Action**

Fix the length and retry.

66EA Length for parm4 is bad.**Action**

Fix the length and retry.

66EB Length for parm5 is bad.**Action**

Fix the length and retry.

66EC Length for parm6 is bad.**Action**

Fix the length and retry.

66ED Invalid parm to afscall_health.**Action**

An internal failure occurred. Contact the service representative.

66EE CQ_QUERY_DEFAULT structure invalid.

Action

Correct the CQ_QUERY_DEFAULT structure.

66EF	QUERY_COMPRESSION_HEALTH_STATS structure has a bad eyecatcher.
-------------	---

Action

Correct the QUERY_COMPRESSION_HEALTH_STATS structure.

66F0	QUERY_COMPRESSION_HEALTH_STATS structure has a bad length field.
-------------	---

Action

Correct the QUERY_COMPRESSION_HEALTH_STATS structure.

66F1	QUERY_COMPRESSION_HEALTH_STATS structure has a bad version.
-------------	--

Action

Correct the QUERY_COMPRESSION_HEALTH_STATS structure.

66F2	PERFORMANCE_TUNE_PARMS structure has a bad eyecatcher.
-------------	---

Action

Correct the PERFORMANCE_TUNE_PARMS structure .

66F3	PERFORMANCE_TUNE_PARMS structure has a bad length field.
-------------	---

Action

Correct the PERFORMANCE_TUNE_PARMS structure.

66F4	PERFORMANCE_TUNE_PARMS structure has a bad version.
-------------	--

Action

Correct the PERFORMANCE_TUNE_PARMS structure.

66F5	Failed to obtain the lock for ZFS_CACHE_PERFORMANCE health check.
-------------	--

Action

An internal failure occurred. Contact the service representative.

66F6	Invalid number of DATA_COLLECT_STATS slots for ZFS_CACHE_PERFORMANCE health check.
-------------	---

Action

An internal failure occurred. Contact the service representative.

6700	Could not set the parm value as requested.
-------------	---

Action

An internal failure occurred. Contact the service representative.

6701	Invalid parm to afscall_config.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6702	Could not obtain lock during vnode_cache_limit processing.
-------------	---

Action

An internal failure occurred. Contact the service representative.

6703	Could not obtain lock during vnode_cache_size processing.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6704	Internal failure during user cache statistics processing.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6705	Internal failure during aggregate format.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6706	Internal failure during aggregate format.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6707	Internal failure during transaction lock operation.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6708	Internal failure during start transaction operation.
-------------	---

Action

An internal failure occurred. Contact the service representative.

6709	Internal failure during fsync operation.
-------------	---

Action

An internal failure occurred. Contact the service representative.

670A	Internal failure obtaining memory for operation.
-------------	---

Action

An internal failure occurred. Contact the service representative.

670B	Internal failure during cache lock operation.
-------------	--

Action

An internal failure occurred. Contact the service representative.

670C	Internal failure during update link count operation.
-------------	---

Action

An internal failure occurred. Contact the service representative.

670D	Internal failure during remove operation.
-------------	--

Action

An internal failure occurred. Contact the service representative.

670E	Internal failure during rename operation.
-------------	--

Action

An internal failure occurred. Contact the service representative.

670F	Internal failure during vnode lookup operation.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6711	Internal failure during acl update.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6714	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

6715	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

6716	Internal failure during attach operation.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6717	Internal failure during filesystem deplete.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6718	Internal failure setting size of the metadata buffer cache.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6719	Internal failure setting size of the metadata backing cache.
-------------	---

Action

An internal failure occurred. Contact the service representative.

671A	Failure to shrink size of transaction cache.
-------------	---

Action

An internal failure occurred. Contact the service representative.

671B	Failure setting size of logfile cache.
-------------	---

Action

An internal failure occurred. Contact the service representative.

671D	Internal failure during setup operation.
-------------	---

Action

An internal failure occurred. Contact the service representative.

671E	Internal failure during ifs_close.
-------------	---

Action

An internal failure occurred. Contact the service representative.

671F	Internal failure during thread pool configuration.
-------------	---

Action

An internal failure occurred. Contact the service representative.

6720	Internal failure during thread pool resize.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6721	Internal failure in wait service.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6722	Striped VSAM Linear Data Set not supported for zFS aggregate.
-------------	--

Action

Striped VSAM Linear Data Sets are not supported. If you are formatting, define a VSAM LDS that is not striped and then format that VSAM LDS. If you are attaching (or mounting) a striped VSAM LDS that has been formatted as a zFS aggregate, you can only attach it (or mount it) in read-only mode. After it is attached (or mounted), copy the data into another zFS aggregate that is not a striped VSAM LDS. If you cannot attach (or mount) for read-only the zFS aggregate that is a striped VSAM LDS, call the service representative.

6723	Internal failure while performing sysplex filesystem operation.
Action	
An internal failure occurred. Contact the service representative.	
6724	Internal failure while performing sysplex mount/quiesce operations.
Action	
An internal failure occurred. Contact the service representative.	
6726	An attempt to change ownership of R/W aggregate is not allowed unless both the source and target system support sysplex file sharing.
Action	
You cannot change the ownership of a R/W aggregate unless both the source and target system support sysplex file sharing.	
6727	Internal failure while handling new owner request at target system.
Action	
An internal failure occurred. Contact the service representative.	
6728	Internal failure while attempting aggregate takeover.
Action	
An internal failure occurred. Contact the service representative.	
6729	An error occurred during extension of the zFS aggregate.
Action	
DFSMS failed the extension attempt. The DFSMS return code is returned in the second byte of the return code. The DFSMS Problem Determination Function (PDF) code is returned in the third byte of the return code. The meaning of these codes can be determined by looking at message IEC161I.	
672A	Internal failure. Aggregate system lock lost during define processing, define ok.
Action	
An internal failure occurred. Contact the service representative.	
672B	Internal failure. Aggregate system lock lost during define processing, define failed.
Action	
An internal failure occurred. Contact the service representative.	
6730	Internal failure during sysplex admin operation.
Action	
An internal failure occurred. Contact the service representative.	
6738	Internal failure. xcf stgpool setsize received a bad pool number.

Action

An internal failure occurred. Contact the service representative.

673D	Internal failure. agsys_send_quiesce, format_send_vectors failed.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6743	Error attempting to receive or reply to a sysplex message.
-------------	---

Action

Use the associated return code to diagnose the problem. See [Return codes \(errnos\)](#) in *z/OS UNIX System Services Messages and Codes* for a description of the return code.

6745	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

6746	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

6747	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

6748	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

6749	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

674A	Internal failure. Trying to send after shutdown or abort started.
-------------	--

Action

An internal failure occurred. Contact the service representative.

674B	Internal failure. User cache not initiated.
-------------	--

Action

An internal failure occurred. Contact the service representative.

674C	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

674D	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

674E	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

6750	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

6751	Internal failure. Invalid state for vnd_Create.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6752	Internal failure. Invalid last called for vnd_Create.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6753	Internal failure. Invalid last called for vnd_Delete.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6754	Internal failure. Error in vnd_CheckPath.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6755	Internal failure. Error in vnd_Checkpath.
-------------	--

Action

An internal failure occurred. Contact the service representative.

675D	Internal failure. Error in vnd_Read of vnd_Read_mlsnvalid.
-------------	---

Action

An internal failure occurred. Contact the service representative.

675E	Internal failure. Invalid state in vnd_Read.
-------------	---

Action

An internal failure occurred. Contact the service representative.

675F	Internal error in vnd_Read.
-------------	------------------------------------

Action

An internal failure occurred. Contact the service representative.

6760	Internal failure. Error in vnd_dirCheck.
-------------	---

Action

An internal failure occurred. Contact the service representative.

6763	Error formatting vector for owner.
-------------	---

Action

Insufficient memory. Retry the operation. If the problem continues, contact the service representative.

6764	Original owning system is no longer active.
-------------	--

Action

zFS ownership of the aggregate will not be changed.

6769	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

676A	Internal failure.
-------------	--------------------------

Action

An internal failure occurred. Contact the service representative.

676B	Task was abended during operation.
-------------	---

Action

None.

676D	Request failed due to grow in process.
-------------	---

Action

The request failed because zFS is busy doing a grow operation. Retry the request when the grow completes.

676E	An unknown RACF® error occurred.
-------------	---

Action

An unknown RACF error occurred. Contact the service representative.

676F	RACF was passed invalid parameters.
-------------	--

Action

Verify whether RACF was passed invalid parameters. Contact the service representative.

6770	An internal error occurred during RACF processing.
-------------	---

Action

Internal RACF error. Contact the service representative.

6771	RACF not authorized.
-------------	-----------------------------

Action

RACF not authorized. Set up authorizations or contact the service representative.

6772	RACF not installed.
-------------	----------------------------

Action

RACF not installed. Install RACF or contact the service representative.

6775	Aggregate not owned on local system.
-------------	---

Action

Ensure that the aggregate exists and is owned on the system where the command is issued from. If the aggregate is owned on a different system, issue the Unquiesce from that system. Use the OMVS **zfsadm lsaggr** command to determine which system owns the aggregate.

6776	Aggregate is not locally owned.
-------------	--

Action

Since it is possible the aggregate moved from the previous owning system to another system in the sysplex while the command was being executed, reissue the command. If the problem persists, contact the service representative.

6800	Mount for multi-file system aggregate is not allowed. This error could also result if a compatibility mode aggregate was not entered in uppercase or if there were system problems unmounting a file system.
-------------	---

Action

Using a release of z/OS prior to z/OS V2R1, attach the aggregate, mount the file system and copy the file system data to a compatibility mode aggregate. If this is a compatibility mode aggregate, ensure the name is entered in uppercase.

6806	File system not locally mounted on owning system.
-------------	--

Action

Retry operation after file system is mounted.

6809	Aggregate is not locally mounted on the remote system.
-------------	---

Action

Take action only if the aggregate is supposed to be locally mounted on the remote system.

680A	The create of the dircache association failed.
-------------	---

Action

Internal error. Contact the service representative.

680B	The security server could not return a UTOKEN for the user.
-------------	--

Action

Internal error. Contact the service representative.

680C	A bad vnode/anode pointer state exists.
-------------	--

Action

Internal error. Contact your service representative.

6900	Internal failure during write operation.
-------------	---

Action

Contact the service representative.

6901	Internal failure during write operation.
-------------	---

Action

Contact the service representative.

6903	Internal failure during read operation.
-------------	--

Action

Contact the service representative.

6904	Internal failure during wait for I/O.
-------------	--

Action

Contact the service representative.

6905	Internal failure during a read.
-------------	--

Action

Contact the service representative.

6906	Internal failure during clear setid.
-------------	---

Action

Contact the service representative.

6907	Internal failure during write of dirty pages.
-------------	--

Action

Contact the service representative.

6908	Internal failure during write.
-------------	---------------------------------------

Action

Contact the service representative.

6909	Internal failure during write.
-------------	---------------------------------------

Action

Contact the service representative.

690A	Internal failure during readlink.
-------------	--

Action

Contact the service representative.

690B	Internal failure during setacl.
-------------	--

Action

Contact the service representative.

690C	Internal failure during read.
-------------	--------------------------------------

Action

Contact the service representative.

690D	Internal failure during trunc.
-------------	---------------------------------------

Action

Contact the service representative.

690E	Internal failure during trunc.
-------------	---------------------------------------

Action

Contact the service representative.

690F	Internal failure during fsync.
-------------	---------------------------------------

Action

Contact the service representative.

6910	Internal failure during close.
-------------	---------------------------------------

Action

Contact the service representative.

6911	Internal failure during inactivate.
-------------	--

Action

Contact the service representative.

6912	Internal failure during setattr.
-------------	---

Action

Contact the service representative.

6913	Internal failure during getattr.
-------------	---

Action

Contact the service representative.

6914	Internal failure during access check.
-------------	--

Action

Contact the service representative.

6915	Internal failure during audit.
-------------	---------------------------------------

Action

Contact the service representative.

6916	Internal failure during getacl.
-------------	--

Action

Contact the service representative.

6917	Internal failure during XCF communication.
-------------	---

Action

Contact the service representative.

6918	Internal failure during lock for write.
-------------	--

Action

Contact the service representative.

6919	Internal failure during lock for read.
-------------	---

Action

Contact the service representative.

691A	Internal failure during access check.
-------------	--

Action

Contact the service representative.

691B	Internal failure during audit.
-------------	---------------------------------------

Action

Contact the service representative.

691C	Internal failure during check two owners.
-------------	--

Action

Contact the service representative.

691D	Internal failure during getsecinfo.
-------------	--

Action

Contact the service representative.

691E	Internal failure during vget.
-------------	--------------------------------------

Action

Contact the service representative.

691F	Internal failure during readdir.
-------------	---

Action

Contact the service representative.

6920	Internal failure during lookup.
-------------	--

Action

Contact the service representative.

6921	Internal failure during create.
-------------	--

Action

Contact the service representative.

6922	Internal failure during link.
-------------	--------------------------------------

Action

Contact the service representative.

6923	Internal failure during remove.
-------------	--

Action

Contact the service representative.

6924	Internal failure during rename.
-------------	--

Action

Contact the service representative.

6925	Internal failure during write lock.
-------------	--

Action

Contact the service representative.

6926	Internal failure during write lock.
-------------	--

Action

Contact the service representative.

6927	Internal failure during token get.
-------------	---

Action

Contact the service representative.

6928	Internal failure during write lock.
-------------	--

Action

Contact the service representative.

6929	Internal failure during remove token.
-------------	--

Action

Contact the service representative.

692A	Internal failure during add token.
-------------	---

Action

Contact the service representative.

692B	Internal failure during create host.
-------------	---

Action

Contact the service representative.

692C	Internal failure during delete host.
-------------	---

Action

Contact the service representative.

692D	Internal failure during revoke reply.
-------------	--

Action

Contact the service representative.

692E	Internal failure during token removal.
-------------	---

Action

Contact the service representative.

692F	Internal failure during revoke tokens.
-------------	---

Action

Contact the service representative.

6930	Internal failure during revoke reply.
-------------	--

Action

Contact the service representative.

6931	Internal failure during remove message.
-------------	--

Action

Contact the service representative.

6932	Internal failure during XCF async send.
-------------	--

Action

Contact the service representative.

6933	Internal failure during send revoke.
-------------	---

Action

Contact the service representative.

6934	Internal failure during stkset clean.
-------------	--

Action

Contact the service representative.

6935	Internal failure during TSR.
-------------	-------------------------------------

Action

Contact the service representative.

6936	Internal failure during setattr.
-------------	---

Action

Contact the service representative.

6937	Internal failure during return tokens reply.
-------------	---

Action

Contact the service representative.

6938	Internal failure during store data reply.
-------------	--

Action

Contact the service representative.

6939	Internal failure during get tokens.
-------------	--

Action

Contact the service representative.

693A	Internal failure during readlink.
-------------	--

Action

Contact the service representative.

693B	Internal failure during setacl.
-------------	--

Action

Contact the service representative.

693C	Internal failure during wait for return tokens.
-------------	--

Action

Contact the service representative.

693D	Internal failure during fetch data reply.
-------------	--

Action

Contact the service representative.

693E	Internal failure during statfs.
-------------	--

Action

Contact the service representative.

693F	Internal failure during return tokens reply.
-------------	---

Action

Contact the service representative.

6940	Internal failure during get write token.
-------------	---

Action

Contact the service representative.

6941	Internal failure during write lock.
-------------	--

Action

Contact the service representative.

6942	Internal failure during async grant.
-------------	---

Action

Contact the service representative.

6943	Internal failure during zlc.
-------------	-------------------------------------

Action

Contact the service representative.

6944	Internal failure during assoc open.
-------------	--

Action

Contact the service representative.

6945	Internal failure during assoc iterate.
-------------	---

Action

Contact the service representative.

6946	Internal failure during item create.
-------------	---

Action

Contact the service representative.

6947	Internal failure during item add.
-------------	--

Action

Contact the service representative.

6948	Internal failure during lock.
-------------	--------------------------------------

Action

Contact the service representative.

6949	Internal failure during get tokens and acl.
-------------	--

Action

Contact the service representative.

694A	Internal failure during flush tokens.
-------------	--

Action

Contact the service representative.

694B	Internal failure during token granting call.
-------------	---

Action

Contact the service representative.

694C	Internal failure during grant race.
-------------	--

Action

Contact the service representative.

694D	Internal failure during end token grant call.
-------------	--

Action

Contact the service representative.

694E	Internal failure during get two tokens.
-------------	--

Action

Contact the service representative.

694F	Internal failure during remove two tokens.
-------------	---

Action

Contact the service representative.

6950	Internal failure during queue tokens.
-------------	--

Action

Contact the service representative.

6951	Internal failure during zlc list create.
-------------	---

Action

Contact the service representative.

6952	Internal failure no sender.
-------------	------------------------------------

Action

Contact the service representative.

6953	The chgowner command could not be processed because either the original or new owner does not support this function or the aggregate is not sysplex-aware.
-------------	---

Action

Contact the service representative.

6954	Internal failure - catchup on owner.
-------------	---

Action

Contact the service representative.

6955	Internal failure - could not obtain tokens for a file.
-------------	---

Action

Contact the service representative.

6956	Out of storage during a client file operation.
-------------	---

Action

A failure to obtain storage occurred during a client file operation. Try to take actions to release storage and retry the operation.

6957	Internal failure - cannot get agsyslock
-------------	--

Action

Contact the service representative.

6958	Error reading file.
-------------	----------------------------

Action

Contact the service representative.

6959	Could not truncate file.
-------------	---------------------------------

Action

Contact the service representative.

695A	Could not write file.
-------------	------------------------------

Action

Contact the service representative.

695B	Could not get anode after file read.
-------------	---

Action

Contact the service representative.

695C	Could not get anode after file write.
-------------	--

Action

Contact the service representative.

695D	Could not fsync file.
-------------	------------------------------

Action

Contact the service representative.

695E	NFS share reservations are not supported on this File System level.
-------------	--

Action

Contact the service representative.

695F	A file may not be opened with Deny flags if O_NONBLOCK is off.
-------------	---

Action

Change the program to set O_NONBLOCK on and to handle the open being rejected with EBUSY.

6960	This open or operation conflicts with a share reservation that has denied the access intended.
-------------	---

Action

Wait awhile and try again or end the program that has the file open with a reservation.

6961	The file is already open in a way that this open is trying to deny.
-------------	--

Action

Wait awhile and try again or end the program that has the file open.

6962	The open was waiting due to a conflict, and then was interrupted causing the open to fail.
-------------	---

Action

Wait awhile and try again or end the program that caused the interrupt.

6963	An attempt to refresh a sysplex client cache has failed unexpectedly.
-------------	--

Action

If the file system object was not recently deleted and there were no reported ZFS errors then contact the service representative.

6964	An internal error was found in sysplex server local user token management.
-------------	---

Action

Contact the service representative.

6965	An internal error was found in sysplex server remote user token management.
-------------	--

Action

Contact the service representative.

6967	An internal error was found trying to start a zfsadm cache resize operation.
-------------	---

Action

Contact the service representative.

6968	A cache resize operation could not be started because zFS is shutting down.
-------------	--

Action

Set the cache to the appropriate size after you restart zFS.

6969	A cache resize operation could not be started because zFS is just starting.
-------------	--

Action

Retry the operation after zFS has completed initialization.

696A	A cache resize operation could not be started due to other in-progress admin commands.
-------------	---

Action

Retry the operation after zFS has completed the in-progress admin commands.

696B	The program calling zFS terminated causing zFS to reach End Of Memory.
-------------	---

Action

zFS continues.

696C	Aggregate and filesystem cannot be used when mounting a compat aggregate.
-------------	--

Action

Remove the aggregate and/or filesystem parms and retry.

696D	Bulk Times attempted while quiesced.
-------------	---

Action

Contact the service representative.

697C	internal error - format_send_vectors returned error.
-------------	---

Action

Contact the service representative.

698A	An operation was attempted against an aggregate that is not owned by any system.
-------------	---

Action

Determine why aggregate could not be attached to one of the systems in the sysplex, possible reasons are DASD connectivity issues, DASD failures, or DFSMS problems.

698B	internal error - could not send grow-notify message to clients.
-------------	--

Action

Contact the service representative.

698C	A sysplex client could not re-open an aggregate.
-------------	---

Action

Takeover attempts from this system could fail. Determine why aggregate could not be opened.

698E	Remount is not allowed when there are NFS V4 share reservations on any file in the File System.
-------------	--

Action

After the files are closed from the NFS clients the remount can be tried again.

6994	Received a bad HANG Detect packet from another system.
-------------	---

Action

Continue to allow the operation. If the problem persists, contact the service representative.

6995	Received a bad HANG Detect operation field from another system.
-------------	--

Action

Continue to allow the operation. If the problem persists, contact the service representative.

6996	Initializing system is downlevel.
-------------	--

Action

This reason is used internally to indicate an initializing system is downlevel and must terminate.

699A	Clone or unclone in process.
-------------	-------------------------------------

Action

Retry the operation.

699B	File system not locally mounted on owning system.
-------------	--

Action

Retry operation after file system is mounted

699C	Aggregate not found.
-------------	-----------------------------

Action

Retry the operation with the correct aggregate name.

699D	Aggregate is dead or unowned.
-------------	--------------------------------------

Action

Contact the service representative.

699E	Aggregate owner cannot process setauditfid.
-------------	--

Action

Move the aggregate to a zFS which supports setauditfid.

699F	Aggregate is not in the attached state.
-------------	--

Action

Retry the operation later. The aggregate may be quiesced. The **lsaggr** command can be used to confirm this.

69A0	Aggregate is not RW.
-------------	-----------------------------

Action

Remount the aggregate RW and retry the operation.

69A1	Could not get storage for block 0.
-------------	---

Action

Release some storage and retry.

69A2	Read of block zero failed.
-------------	-----------------------------------

Action

Ensure there is connectivity to the aggregate and retry.

69A3	Aggregate has a non-zero auditfid.
-------------	---

Action

Retry with the force or old option.

69A4	Invalid force or old value.
-------------	------------------------------------

Retry with a valid force or old option.

69A5 asevent create, queue io or wait failed.

Action

Contact the service representative.

69A7	Not allowed to use this directory on this system.
------	---

Action

Retry operation from a system with the proper PTF level applied.

69A8	Aggregate could not be opened on this system.
------	---

Action

The aggregate is owned by a member running a release of zFS prior to z/OS V1R12 and/or this particular file system is running non-sysplex aware. This system is running zFS sysplex-aware. zFS fails the catchup mount because we cannot allocate the file system shared. The prior release zFS or non-sysplex aware zFS system has allocated the file system exclusive. This is normal. z/OS UNIX will function ship zFS requests from the z/OS V1R12 sysplex-aware systems to the owning system. Or, the aggregate is owned by a member running zFS sysplex-aware and you are trying to move the aggregate to a system where zFS is running non-sysplex aware.

69A9	Multi-file system operations not allowed in a sysplex.
-------------	---

Action

An operation was attempted that is targeted for a multi-filesystem aggregate or is otherwise not compatible with an HFS-compatibility mode aggregate. This is normal as multi-filesystem aggregates are not allowed in a parallel sysplex.

69AA	A long-running operation was interrupted by an unmount, a file system shutdown command, or a cancel command. This is normal for long-running operations.
------	--

Action

To continue the long-running operation, you will need to restart it.

69AB	A system on the ready list could not be found on the target of an NSQUERY.
------	--

Action

Internal error

69AC A system on the ready list is not ready on the target of an NSQUERY.

Action

Internal error

69AC **A system on the ready list is not ready on the target of an NSQUERY.**

Action

Internal error

69AD	A system missing from the ready list is ready on the target of an NSQUERY.
Action	
Internal error	
69AE	The GRS concurrent ENQ limit has been reached.
Action	
Increase your systems concurrent limit to ensure proper functionality.	
69AF	An internal error was encountered during a GRS ENQ call.
Action	
Contact IBM service.	
69B0	An internal error was encountered during an IGSQRY call.
Action	
Contact IBM service.	
69B1	Aggregate cannot be opened in R/W mode on client system.
Action	
The aggregate is attached in R/W mode and a sysplex client has received a catch-up mount for it and that client does not have the sysplex-aware support enabled. Therefore the mount is failed. This is normal. z/OS UNIX will function ship zFS requests from this system to the owning system.	
69B2	LSAGGR cannot obtain buffer for GRS QUERY call.
Action	
Contact IBM service.	
69B3	LSAGGR cannot obtain HASH table buffer.
Action	
Contact IBM service.	
69B4	zFS could not obtain an ENQ required for the quiesce operation.
Action	
Contact IBM service.	
69B5	An unmount was requested and the file system was busy with an administration command.
Action	
Either wait for the command to complete or issue unmount with the FORCE option to interrupt the administration command.	
69B6	An internal error was encountered during sysplex client processing of aggregate quiesce.

Action

Contact IBM service.

69B7	An internal error was encountered during sysplex client locking of name in directory.
-------------	--

Action

Contact IBM service.

69BA	The owner of a file system went down during the forwarding of a PFSCtl call or the file system is unowned.
-------------	---

Action

Retry the operation.

69BB	A sysplex level admin system cannot takeover R/W aggregate for READ purposes at this time.
-------------	---

Action

Retry the operation.

69BD	An internal error was encountered while finding aggregate owner.
-------------	---

Action

Contact IBM service.

69BE	SVC 26 Listcat call returned an error.
-------------	---

Action

Verify that the data set exists. Try issuing Listcat against the data set.

69BF	Multi aggregate cannot be detached because it is cloning.
-------------	--

Action

Reissue the detach after the clone completes.

69C0	A client thinks an aggr is unowned and the owner is still up.
-------------	--

Action

The owner fails the start takeover that the client sends.

69C1	zFS address space is terminating, all operations will fail.
-------------	--

Action

Contact IBM Service.

69C2	An aggregate cannot be formatted to a size greater than 4TB.
-------------	---

Action

Specify a smaller size or recreate the linear dataset.

69C3	New owner has requested to fail takeover.
-------------	--

Action

Verify that the dataset is properly cataloged and that the catalog is accessible.

6A00 **CHAGGR_REQ eyecatcher invalid.**

Action

The eyecatcher in the CHAGGR_REQ structure is incorrect. It must be CARQ.

6A01 **CHAGGR_REQ length invalid.**

Action

See the zFS application programming interface information in [z/OS File System Administration](#).

6A02 **CHAGGR_REQ version invalid.**

Action

See [z/OS File System Administration](#) for assistance.

6A03 **CHAGGR_REQ reserved invalid.**

Action

CHAGGR_REQ reserved fields must be set to binary zeros.

6A04 **CHAGGR_REQ increment invalid.**

Action

The increment percent must be greater than or equal to 1 and less than or equal to 99.

6A05 **CHAGGR_REQ threshold invalid.**

Action

The threshold percent must be greater than or equal to 1 and less than or equal to 99.

6A06 **An invalid value was specified for the aggrfull aggregate attribute.**

Action

Try again after changing the value to OFF or xx,yy where xx is a threshold percentage and yy is an increment percentage.

6A07 **Aggregate not found.**

Action

The aggregate specified cannot be found. Correct the aggregate name and try again.

6A08 **An invalid value was specified for the aggrgrow aggregate attribute.**

Action

For the **zfsadm chaggr** command you can specify ON or OFF. For the API you can specify 1 or 2. Try again after correcting the value.

6A09 **The aggregate attribute cannot be changed to both -aggrgrow on and -aggrgrow off.**

Action

Limit your request to one attribute value and try again.

6A0A	The message notifying clients about attribute changes for an aggregate could not be sent.
-------------	--

Action

Use the return code to resolve the problem and try the command again. If the problem persists, contact the service representative.

6A0C	Unexpected error attempting to sync the metadata for a file system object.
-------------	---

Action

Contact the service representative.

6A0D	Unexpected error attempting to reserve bitmap blocks for sysplex client.
-------------	---

Action

Contact the service representative.

6A0E	Unexpected error attempting to update indirect blocks for a file at a sysplex server.
-------------	--

Action

Contact the service representative.

6A0F	Unexpected error attempting to create indirect blocks for a file at a sysplex server.
-------------	--

Action

Contact the service representative.

6A10	Internal error during forward of a getattr operation.
-------------	--

Action

Contact the service representative.

6A11	Internal error during forward of an fsync_file operation.
-------------	--

Action

Contact the service representative.

6A12	Internal error during forward of a close operation.
-------------	--

Action

Contact the service representative.

6A13	Internal error during forward of a trunc operation.
-------------	--

Action

Contact the service representative.

6A17	Internal error during forward of a read operation.
-------------	---

Action

Contact the service representative.

6A18	Internal error during forward of a write operation.
-------------	--

Action

Contact the service representative.

6A19	Internal error during forward of a readlink operation.
-------------	---

Action

Contact the service representative.

6A1A	Internal error during forward of a link operation.
-------------	---

Action

Contact the service representative.

6A1B	Internal error during forward of a lookup operation.
-------------	---

Action

Contact the service representative.

6A1C	Internal error during forward of a remove operation.
-------------	---

Action

Contact the service representative.

6A1D	Internal error during forward of a create operation.
-------------	---

Action

Contact the service representative.

6A1E	Internal error during forward of a readdir operation.
-------------	--

Action

Contact the service representative.

6A1F	Internal error during forward of a rename operation.
-------------	---

Action

Contact the service representative.

6A24	Could not create an entry in a version 5 directory.
-------------	--

Action

Contact the service representative if the return code is EMVSERR (157).

6A25	Internal error - could not access a version 5 directory block.
-------------	---

Action

Contact the service representative.

6A27	Internal error: Indirect block is already allocated to the file.
-------------	---

Action

Contact the service representative.

6A28	Aggregate sync operation failed.
-------------	---

Action

Contact the service representative.

6A29	Aggregate is busy. It cannot perform the operation requested.
-------------	--

Action

The aggregate is busy with another operation that conflicts with this operation. Try your request again. If it continues to fail, contact the service representative.

6A2A	Too many ZFS recoveries are in progress. Operation not performed.
-------------	--

Action

Try the request again.

6A2B	The metadata backing cache does not exist so cannot be resized.
-------------	--

Action

Ensure there is a metaback_cache_size statement in your parm file and restart zFS.

6A2C	Directory limit exceeded, directory too big to add entries.
-------------	--

Action

Put names in another directory or prune entries in the directory.

6A2D	zFS was internally restarted and the file system failed to remount.
-------------	--

Action

Unmount/remount the file system to clear the condition.

6A2E	The operation failed due to a disabled aggregate. An aggregate can be temporarily disabled due to an I/O error, an internal zFS error, or a dataset reopen operation fails. If the aggregate is encrypted, this can be caused by ICSF problems. File requests will fail while the aggregate is disabled. zFS will attempt to re-enable the aggregate automatically.
-------------	--

Action

If the aggregate is encrypted, correct any ICSF problem so the aggregate can be re-enabled. When the aggregate is re-enabled, the file can be closed and reopened and the request can be attempted again. If zFS cannot re-enable the aggregate automatically, you will need to unmount/mount or remount the aggregate before the file request can be attempted again.

6A2F	Operation failed due to disabled aggregate being remounted.
-------------	--

Action

The aggregate is temporarily disabled due to an internal zFS error. Quiesce requests will fail while the aggregate is being internally remounted. Retry the quiesce after the remount completes.

6A30	Too many tasks in the ZFS address space are in progress. Operation not performed.
-------------	--

Action

Try the request again as other tasks may have ended.

6A31	ZFS storage is becoming limited, so a storage obtain was failed. Operation not performed.
-------------	--

Action

Try the request again as storage may have become available.

6A32	New version aggregates cannot be processed on this system.
-------------	---

Action

If this is a primary mount, re-issue the mount on a system that supports higher version aggregates. If this is a catchup mount, no action is needed. z/OS UNIX will function ship operations to the z/OS UNIX owner.

6A33	An unexpected error was detected while opening the aggregate VSAM linear dataset.
-------------	--

Action

One possible cause is that the client system is not using the same physical dataset as the owning system. Ensure that each system is referencing the same physical dataset. For more information about sharing file systems, see the sections on using the zFS read/write sysplex-aware file systems in [z/OS File System Administration](#). Otherwise contact the service representative.

6B00	Quiesce for clone not allowed.
-------------	---------------------------------------

Action

Clones are no longer supported.

6B01	Mount of .bak or aggregate with .bak not allowed.
-------------	--

Action

Clones are no longer supported. Attach the aggregate and delete the .bak file system. Then, try the mount again.

6B02	Clone not allowed.
-------------	---------------------------

Action

Clones are no longer supported.

6B03	Create not allowed.
-------------	----------------------------

Action

Multi-file system aggregates are no longer supported.

6B04	Rename not allowed.
-------------	----------------------------

Action

Multi-file system aggregates are no longer supported.

6B05	Mount or attach not allowed for multi-file systems.
-------------	--

Action

Multi-file system aggregates are not supported.

6B06	Setquota not allowed.
-------------	------------------------------

Action

Multi-file system aggregates are no longer supported.

6B0C	Error occurred during a directory read at a sysplex client for an extended (v5) directory.
-------------	---

Action

This reason code will be accompanied by a return code and preceding error messages that provide more information about the problem. Possible causes are transmission failures, running out of memory, or a zFS internal error.

6B0D	The specified aggregate cannot be changed to version 1.5 or the specified directory cannot be changed to version 5 because the owner is downlevel, there are downlevel members in the sysplex or the aggregate is mounted read-only. Because the aggregate version cannot be changed, any specified path will not be converted.
-------------	--

Action

Try the command again after removing all downlevel members or remounting the aggregate read-write.

6B0E	The input directory is already version 4.
-------------	--

Action

Specify a version 5 directory and retry the operation.

6B0F	The input directory is already version 5.
-------------	--

Action

Specify a version 4 directory and retry the operation.

6B10	The input FOBJ_INFO buffer is invalid.
-------------	---

Action

The application must provide a buffer large enough to hold an FOBJ_INFO structure, have the proper eye catcher set in the structure, a proper version field, and a length field that is at least as large as the FOBJ_INFO structure.

6B11 **Grow size invalid.**

Action

See *z/OS File System Administration* for the maximum allowable size when growing an aggregate. When using the grow API, how you specify the grow size depends on the AGGR_ID version that is used to specify the aggregate name.

6B12 **Conversion of aggregate to version 1.5 during a mount failed.**

Action

Retry the command after specifying an option. For detailed help issue **zfsadm chaggr -help**.

6B13 **An option was not specified for the zfsadm chaggr command or pfscctl.**

Action

Run the salvager program against the aggregate and retry the mount.

6B14 **The path specified is not a directory or is already a version 5 directory.**

Action

Retry the command specifying a version 4 directory.

6B15 **The aggregate could not be found.**

Action

Ensure the aggregate is not being formatted. Retry the command specifying a valid aggregate name.

6B16 **Could not reset the backup bit in the DSCB.**

Action

Retry the command after correcting any error conditions.

6B17 **AGGR_FORMAT has an invalid af_newauditfid value.**

Action

Valid af_newauditfid values are 0 and 1. Correct the value and retry.

6B18 **AGGR_FORMAT has an invalid af_aggrversion value.**

Action

Valid af_aggrversion values are 0, 4 or 5. Version 1.5 aggregates cannot be formatted in a sysplex containing members at z/OS V1R13 or prior. Correct the value and retry.

6B19 **A write operation was attempted against a directory with a corrupted tree.**

Action

Attempt to repair the corrupted directory by using zfsadm convert or the salvager program and try again.

6B1A	A directory conversion operation was attempted in a sysplex containing downlevel systems.
-------------	--

Action

Conversion not allowed in a sysplex containing downlevel systems. Retry when all sysplex members are at z/OS V2R1 or later.

6B1B	Internal failure during conversion.
-------------	--

Action

Contact the service representative.

6B1C	Either an error was encountered performing automatic conversion or a previous error was reported. Once an error is reported, all future automatic conversions for this directory will fail.
-------------	--

Action

Use zfsadm convert to perform the conversion or remount the file system to restore automatic conversion.

6B1E	Aggregate is already version 1.5.
-------------	--

Action

No action is needed.

6B1F	Failed changing aggregate to version 1.5.
-------------	--

Action

Ensure the correct directory is specified and retry.

6B20	Failed changing directories to version 5.
-------------	--

Action

Ensure the correct directory is specified and retry.

6B21	The input aggregate is not version 1.5.
-------------	--

Action

Specify a directory in a version 1.5 aggregate and retry the operation.

6B22	Failed changing aggregate to version 1.4 or directory to version 4.
-------------	--

Action

Verify the aggregate is a version 1.5 aggregate. The aggregate cannot be greater than 4T in size. A directory cannot be greater than 4G in size and cannot have a link count greater than 65535. Correct the error and try the operation again, if still needed.

6B23	Format failed because the zFS kernel is not available and an aggregate version was not specified.
-------------	--

Action

Start the zFS kernel or supply an aggregate version and try the operation again.

6B24	The value specified for the format size or grow option is invalid.
-------------	---

Action

Specify a value in the range [1..536870912] for version 4 aggregates and [1..2147483648] for version 5 aggregates and try again.

6B25	The value specified for the format logsize option is invalid.
-------------	--

Action

Specify a logsize value in the range [13..16384] and try again.

6B26	Sync interval bigger than 21474836 second maximum.
-------------	---

Action

Specify a sync interval smaller than 21474836.

6B27	Statistics Query Parm0 value is invalid.
-------------	---

Action

The Parm0 value for all Statistics Query subcommands must be 32.

6B28	Statistics Query Parm1 value is invalid.
-------------	---

Action

The Parm1 value for all Statistics Query subcommands must be 80.

6B48	Unable to get storage to perform the requested operation.
-------------	--

Action

Make more space available and reissue the request.

6B49	Internal failure during stkm statistics processing.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6B4A	Internal failure during DASD extent processing for the aggregate.
-------------	--

Action

An internal failure occurred. Contact the service representative.

6B4B	Could not increase log cache to requested size.
-------------	--

Action

Determine whether zFS is low on memory. If not, then contact the service representative.

6B4C	Could not decrease log cache to requested size.
-------------	--

Action

Determine whether zFS is low on memory. If not, then contact the service representative.

6B4D	FSINFO attach FAILED.
-------------	------------------------------

Action

Contact the service representative.

6B4E	Internal error while processing FSINFO request.
-------------	--

Action

Contact the service representative.

6B4F	Failed obtaining memory for an ACL dataspace operation.
-------------	--

Action

Determine whether zFS is low on memory. If not, then contact the service representative.

6B50	Bad parameter value passed in FSINFO_REQUEST structure.
-------------	--

Action

Correct the bad parameter and retry the operation.

6B51	There are too many in-progress PFCTL commands.
-------------	---

Action

zFS has reached its limit of concurrent PFCTL calls. This is generally not typical. Contact the service representative.

6B52	Not enough storage to perform read or write request.
-------------	---

Action

The zFS address space is most likely out of storage. Try to decrease storage usage by decreasing cache sizes. Or, you can try to break up the operation into smaller requests. Then retry the operation.

6B53	zFS is shutting down.
-------------	------------------------------

Action

zFS will not process the requested operation. If this situation persists, contact the service representative.

6B54	The FSINFO_REQUEST structure contains an incorrect resumeName. Only characters valid in a zFS aggregate name are allowed.
-------------	--

Action

Correct the name and retry the operation.

6B55	The FSINFO_REQUEST structure contains an incorrect patternName. Only characters valid in a zFS aggregate name or a wildcard character are allowed.
-------------	---

Action

Correct the name and retry the operation.

6B56	The data set is not a VSAM linear data set.
-------------	--

Action

Define the VSAM data set as linear and then retry the operation.

6B5D	An error was encountered processing the zero link count list.
-------------	--

Action

If the aggregate is mounted, it will be disabled. Salvage the aggregate to correct the error.

6B5E	This system cannot become the aggregate's zFS owner because it is precluded by the aggregate's automove option.
-------------	--

Action:

If this system is allowed to become the owner, modify the aggregate's automove option.

6B5F	Unable to get the automove option for the aggregate.
-------------	---

Action

Contact the service representative.

6B60	z/OS UNIX System Services routine osi_ctl encountered an error.
-------------	--

Action

Contact the service representative.

6B61	Hi-used-RBA is more than hi-allocated-RBA. This should be verified with the TSO LISTCAT command.
-------------	---

Action

Contact the service representative.

6C00	UNIX System Services has supplied an invalid ACL type to zFS.
-------------	--

Action

Contact the service representative.

6C04	There is no salvage operation to cancel. The aggregate is not currently being salvaged.
-------------	--

Action

If the specified aggregate is correct, there is no action to take. Otherwise specify the correct aggregate and try again.

6C0B	The mount parm for the aggregate attribute could not be updated. zFS will attempt a samemode remount to restore the aggregate attribute to the value of the mount parm.
-------------	--

Action

Determine whether the aggregate attribute and mount parm match. Issue **zfsadm fsinfo** to get the aggregate attributes and **D OMVS,F,N=** to get the mount parms. If they do not match, issue another samemode remount. If the problem persists, contact your service representative.

6C0C	Aggregate attributes cannot be changed in a sysplex where some systems are at a release prior to z/OS V2R3. The server console may contain more information.
-------------	---

Action

Try the command again when no systems of the sysplex are at a release prior to z/OS V2R3.

6C0D Only one aggregate attribute can be changed at a time.

Action

More than one aggregate attribute was specified. Try the command again with only one aggregate attribute.

6C0E The aggregate attributes cannot be changed at this time. The aggregate may not be in a consistent state across all members in the sysplex due to problems sending sysplex messages.

Action

Wait until the problem is fixed and then try the command again. If the problem persists, contact your service representative.

6C0F The request failed because the `-rwshare on` and `-rwshare off` attributes were both specified for the aggregate.

Action

Try again, specifying `-rwshare on` or `-rwshare off`.

6C10 An invalid value was specified for the `-rwshare` aggregate attribute.

Action

The `-rwshare` aggregate attribute takes a value of either ON or OFF. When using the `chaggr` API, use the corresponding value, which is either 1 (for ON) or 2 (for OFF).

6C11 The aggregate attribute cannot be changed to both `-aggrfull xx,yy` and `-aggrfull off`.

Action

Limit your request to one attribute value and try again.

6C12 Errors were encountered changing the aggregate space monitoring attribute. Aggregate space monitoring is turned off.

Action

The zFS address space is most likely out of storage. Try the operation again after making more storage available. If the problem persists, contact your service representative.

6C13 Errors were encountered changing the aggregate attribute.

Action

If the zFS address space is out of storage, try the operation again after making more storage available. If the problem persists, contact your service representative.

6C14 EDC buffer pool too big.

Action

Ensure that the `-edc_buffer_pool` value specified is not more than 1024M.

6C15 EDC buffer pool too small.

Action

Ensure that the `-edc_buffer_pool` value specified is at least 10M.

6C18	The sharemode mount parm could not be updated. The corresponding aggregate sharemode remains unchanged.
-------------	--

Action

zFS encountered an error changing the sharemode mount parm. Try again and if the problem persists, contact your service representative.

6C19	The aggregate sharemode could not be updated to match the updated mount parm. zFS will attempt to restore the original sharemode mount parm.
-------------	---

Action

Determine whether the sharemode attribute and mount parm match. Issue **zfsadm fsinfo** to get the sharemode attribute and issue **D OMVS,F,N=** to get the mount parm. If they do not match, issue another samemode remount. If the problem persists, contact your service representative.

6C1A	Invalid salvage type. The salvage type must be 1, 2, or 3, indicating verify only, verify and repair, or cancel.
-------------	---

Action

Specify the correct salvage type and try the operation again.

6C1B	AGGR_FORMAT has an invalid af_encrypt value. The value must be 0, 1, or 2.
-------------	---

Action

Correct the value and retry.

6C1C	AGGR_FORMAT has an invalid af_compress value. The value must be 0, 1, or 2.
-------------	--

Action

Correct the value and retry.

6C1D	A format operation with encryption or compression setting was attempted in a sysplex that contains earlier release systems.
-------------	--

Action

Encryption or compression is not allowed in a sysplex that contains earlier release systems. Retry when all sysplex members are at least at z/OS V2R3.

6C21	No thread available for the requested long-running command.
-------------	--

Action

Consider increasing the size of the `long_cmd_threads` pool. Retry the request when there is an available thread for the long-running command.

6C22	Aggregate could not be shrunk. It is mounted in R/O mode.
-------------	--

Action

The aggregate must be mounted read-write to be shrunk.

6C23	Unexpected error shrinking aggregate.
-------------	--

Action

Contact the service representative.

6C24	The new size is not smaller than the current aggregate size.
-------------	---

Action

Use the **zfsadm fsinfo** command to determine the size of the aggregate. Try the command again with a new size that is smaller than the current aggregate size.

6C25	Error finding file system information.
-------------	---

Action

Ensure that the file system is mounted and try the command again. If the problem persists, contact the service representative.

6C26	The eye catcher of the SH_REQ structure is not SHRQ.
-------------	---

Action

Correct the eye catcher and try the command again.

6C27	The sh_len field of the SH_REQ structure is incorrect.
-------------	---

Action

Correct the sh_len value and try the command again.

6C28	The sh_ver field of the SH_REQ structure is incorrect.
-------------	---

Action

Correct the sh_ver value and try the command again.

6C29	The sh_flags field of the SH_REQ structure is incorrect.
-------------	---

Action

Correct the sh_flags value and try the command again.

6C2A	The sh_command field of the SH_REQ structure is incorrect.
-------------	---

Action

Correct the sh_command value and try the command again.

6C2B	The sh_reserved field of the SH_REQ structure is not all zeros.
-------------	--

Action

Correct the sh_reserved value and try the command again.

6C2C	The value in the sh_length field of the SH_REQ structure is not allowed. This value is larger than the maximum aggregate size.
-------------	---

Action

Ensure that the specification of new size for the aggregate is correct and try again.

6C2D	The option to prevent active increase (-noai) is not allowed with the stop (-cancel) command.
-------------	--

Action

Correct the values and try the command again.

6C2F	The specified zFS operation was not canceled because it is not in progress.
-------------	--

Action

If the specified aggregate is correct, there is no action to take. Otherwise, specify the correct aggregate and try again.

6C30	Request failed due to salvage in process.
-------------	--

Action:

The request failed because zFS is busy doing a salvage operation. Retry the request when the salvage completes.

6C31	Request failed due to shrink in process.
-------------	---

Action:

The request failed because zFS is busy doing a shrink operation. Retry the request when the shrink completes.

6C32	Cannot change a sharemode attribute when not in a sysplex.
-------------	---

Action:

Do not specify -rwshare or -norwshare on the **zfsadm chaggr** command when not in a sysplex.

6C33	The aggregate log could not be recovered.
-------------	--

Action:

Contact your service representative.

6C34	Salvage found minor errors in the aggregate.
-------------	---

Action:

Issue **zfsadm salvage** without the -verifyonly option to correct the errors.

6C35	Salvage found major errors in the aggregate.
-------------	---

Action:

Contact your service representative.

6C36	The corruption could not be repaired.
-------------	--

Action:

Contact your service representative.

6C37	The corruption in version 1.5 directories could not be repaired.
-------------	---

Action

Contact your service representative.

6C3E	Request failed due to encryption in progress.
-------------	--

Action

The request failed because zFS is busy with an encryption operation. Retry the request when the encryption completes or cancel the encryption and reissue the command.

6C3F	Request failed due to decryption in progress.
-------------	--

Action

The request failed because zFS is busy with a decryption operation. Retry the request when the decryption completes or cancel the decryption and reissue the command.

6C40	Invalid encrypt or decrypt type. The type must be 1, 2, or 3, indicating encrypt, decrypt, or cancel.
-------------	--

Action

Specify the correct type and try the operation again.

6C41	Salvage must perform log recovery but cannot because the aggregate was mounted read-only. Salvage terminates.
-------------	--

Action

Mount the aggregate read-write and try the operation again.

6C42	Salvage found inconsistencies in the aggregate that cannot be repaired because the aggregate is mounted read-only.
-------------	---

Action

Mount the aggregate read-write and try the operation again.

6C43	The fr_version field in FSINFO_REQUEST is invalid.
-------------	---

Action

Provide an fr_version value acceptable for the zFS release and try the operation again.

6C44	The fr_selection field in FSINFO_REQUEST is invalid with the fr_version specified.
-------------	---

Action

Provide an fr_selection value consistent with fr_version and try the operation again.

6C45	Internal failure sending sync request.
-------------	---

Action

Contact the service representative.

6C46	Error during sync operation.
-------------	-------------------------------------

Action

Correct the problem and try the operation again. If the problem persists, contact the service representative.

6C47	Unknown error during shrink operation.
-------------	---

Action

Check the system log for message IOEZ00904E to get additional error information. If the problem cannot be resolved, contact the service representative.

6C48	There is no shrink operation to cancel. The aggregate is not currently shrinking.
-------------	--

Action

If the aggregate specified is correct, there is no action to take. Otherwise specify the correct aggregate and try again.

6C49	The space used in the aggregate is larger than the specified new size.
-------------	---

Action

Use the **zfsadm fsinfo** command to determine the size of the aggregate, the amount of free space, and the amount of space in use. Try the command again with a new size that is larger than the amount of space currently being used.

6C4A	The new size is smaller than the minimum aggregate size.
-------------	---

Action

Use the **zfsadm fsinfo** command to determine the size of the aggregate, the amount of free space, and the amount of space in use. Try the command again with a new size that is larger than the amount of space currently being used.

6C4F	Unexpected error resizing the pool used for encryption and compression.
-------------	--

Action

Contact the service representative.

6C50	Error finding file system information.
-------------	---

Action

Ensure that the file system is mounted and try the command again. If the problem persists, contact the service representative.

6C51	Encrypt, decrypt, compress, or decompress is not allowed for aggregate with version earlier than 1.5.
-------------	--

Action

Ensure that the aggregate version is at least version 1.5. To find the aggregate version, use either the **zfsadm fsinfo** or **MODIFY FSINFO** command. If the problem persists, contact the service representative.

6C52	Could not convert the aggregate to version 1.4 because it is partially encrypted or compressed.
-------------	--

Action

Ensure that the aggregate is fully decompressed and decrypted before converting it to version 1.4.

6C53	Error during aggregate space release operation.
-------------	--

Action

Check the system log for message IOEZ00897E to get additional error information. If the problem cannot be resolved, contact the service representative.

6C54	The VSAM LDS is not formatted, is not marked as intended for zFS, and cannot be formatted at mount.
-------------	--

Action

Format the VSAM LDS using the **zfsadm format** command and try the mount again.

6C55	The VSAM LDS is not formatted and zFS could not determine whether a format is allowed during mount. The MVS catalog service could not be loaded.
-------------	---

Action:

Format the VSAM LDS using the **zfsadm format** command and try the mount again.

6C56	An MVS catalog operation could not be performed on the VSAM LDS.
-------------	---

Action:

Take the action documented in message IOEZ00908I.

6C57	The VSAM LDS could not be marked as a zFS aggregate. The MVS catalog service failed.
-------------	---

Action:

Take the action documented in message IOEZ00336I.

6C58	The VSAM LDS could not be formatted at mount.
-------------	--

Action:

Correct the problem and format the VSAM LDS using the **zfsadm format** command, then try the mount again.

6C5B	Request failed due to compression in progress.
-------------	---

Action

The request failed because zFS is busy doing a compression operation. Retry the request when the operation completes.

6C5C	Request failed due to decompression in progress.
-------------	---

Action

The request failed because zFS is busy doing a decompression operation. Retry the request when the operation completes.

6C5D	Could not use the edcfixed option to fix the user file cache. The edcfixed option is limited to 14GB user file caches or less.
-------------	---

Action:

Either make the user file cache smaller or do not specify the edcfixed option.

6C5E	Could not allow the encryption request.
-------------	--

Action:

An encryption request is only allowed if all sysplex members are release z/OS V2R3 or later. Retry the command when all systems are migrated to z/OS V2R3 or later.

6C5F	Could not allow the compression request.
-------------	---

Action:

A compression request is only allowed if all sysplex members are release z/OS V2R3 or later. Retry the command when all systems are migrated to z/OS V2R3 or later.

6C60	The key label specified is too long.
-------------	---

Action

Specify a key label of length 64 or less and try again.

6C61	The offset specified on a backup-read request is not the next-in-line offset for the file or EOF was returned on a previous read of the file.
-------------	--

Action

Correct the program and try again.

6C62	The argument length for a backup ioctl request is less than the size of a BK_REQ structure or the BK_REQ argument itself contains a bad length, eye catcher, or version.
-------------	---

Action

Correct the argument length or the BK_REQ fields and try again.

6C63	Reserved fields of the BK_REQ structure are not 0 for a backup ioctl request.
-------------	--

Action

Correct the program and try again.

6C64	You cannot register a backup of a file that is already registered for backup.
-------------	--

Action

Correct the program and try again.

6C65	You cannot register a backup of a file that is not opened for read.
-------------	--

Action

Correct the program and try again.

6C66	You cannot register a backup for an object that is not a file.
-------------	---

Action

Correct the program and try again. A backup is only allowed for regular files.

6C67	A process is attempting to perform a backup-read for a file that it did not register with.
-------------	---

Action

Correct the program and try again. A backup-register must be performed by the process before it reads the file.

6C68	The bk_request field in the BK_REQ structure contains an invalid value for a backup ioctl request.
-------------	---

Action

Correct the program and try again.

6C69	When migrating a file system from HFS to zFS, an inode being created already exists.
-------------	---

Action

Contact the service representative.

6C6A	The bk_bufferSize specified is too small.
-------------	--

Action

Change the program to obtain a larger buffer. Then retry the operation with the larger buffer size.

6C6B	An abend occurred in the security product while processing this operation.
-------------	---

Action

Try to determine the cause of the security product failure, correct it, and then retry the operation. If the problem persists, contact the service representative.

6C6C	An attempt to format a file system for encryption failed because the file system data set is not eligible for encryption.
-------------	--

Action

Refer to the section on defining the encryption management mechanism in *z/OS DFSMSdfp Storage Administration* for information about how to define the data set in an SMS class that is eligible for encryption and that has a valid assigned key label.

6C6E	The length of the key label exceeds 64 and zFS does not support key labels longer than 64.
-------------	---

Action

The key label can be specified when defining the aggregate or with the **zfsadm encrypt** command via the -keylabel option. Correct the error and retry the operation.

6C6F	The encryption operation cannot continue without a key label.
-------------	--

Action

The key label can be specified when defining the aggregate or with the **zfsadm encrypt** command -keylabel option. Correct the error and retry the operation.

6C71	-keylabel is not allowed when -system specifies a system at a release prior to z/OS V2R3.
-------------	--

Action

Either remove -keylabel, remove -system, or specify a system at release z/OS V2R3 or later.

6C72	Converting a migrated file system to a version 1.4 aggregate is not allowed.
-------------	---

Action

If this is not a migrated file system, contact the service representative.

6C73	The edcfixd option was specified with user_cache_size, but the user cache pages cannot be registered with the zEDC Express service.
-------------	--

Action

If the zEDC Express service is not available, ensure that the edcfixd option is not used with the user_cache_size option in the IOEFSPRM member. Also, ensure that it is not used with the **zfsadm config -user_cache_size** command.

6C74	Aggregate could not be encrypted or decrypted. It is mounted in R/O mode.
-------------	--

Action:

The aggregate must be mounted read-write to be encrypted or decrypted.

6C75	Aggregate could not be compressed or decompressed. It is mounted in R/O mode.
-------------	--

Action:

The aggregate must be mounted read-write to be compressed or decompressed.

6C76	ICSF could not find the specified key label.
-------------	---

Action

A key label must be known to ICSF before it can be used when defining or encrypting a zFS aggregate.

6C77	An attempt to format a file system failed because the file system cannot be encrypted or compressed with aggregate version 1.4.
-------------	--

Action

Either request formatting with the -version5 option or do not use the -encrypt or -compress option.

6C78	ICSF returned an error while validating a key label.
-------------	---

Action

Take the action documented in message IOEZ00959E.

6C79	A key label was specified but ICSF is not running.
-------------	---

Action

Ensure that ICSF is running and try the command again.

6C7A	Failed to set or resize the EDC_BUFFER_POOL due to zEDC error.
-------------	---

Action

Check error message IOEZ00962E to see the zEDC return code and reason code. Fix the error and try again.

6C7B	Failed to set or resize the user file cache due to page fix error.
-------------	---

Action

Check for error message IOEZ00962E to see the zEDC return code and reason code. Fix the error and try again.

6C7C	The key label specified is not associated with an encryption type supported by DFSMS.
-------------	--

Action

Specify another key label and try the command again.

6C7D	DFSMS returned an error while opening an aggregate.
-------------	--

Action

See messages IOEZ00964E and IEC161I at the console for more information.

6C7E	-keylabel1 is not allowed when an aggregate already has a key label.
-------------	---

Action

Remove `-keylabel` and try the command again.

6C7F	Shrinking of the aggregate was not allowed by a management class.
-------------	--

Action

See IOEZ00966E for further information about the aggregate shrink being disallowed. If a management class did not prevent the shrink, contact the service representative.

6C80	Encrypt or decrypt is not allowed if the aggregate is already in a partially compressed state.
-------------	---

Action

The **encrypt** and **decrypt** commands can be issued only when the aggregate is either not compressed or is fully compressed. Retry the command after the aggregate has been fully decompressed or compressed.

6C81	Compress or decompress is not allowed if the aggregate is already in a partially encrypted state.
-------------	--

Action

The **compress** and **decompress** commands can be issued only when the aggregate is either not encrypted or is fully encrypted. Retry the command after the aggregate has been fully decrypted or encrypted.

6C82	The target file system of a migration is not newly formatted with an allowed version.
-------------	--

Action:

Reformat the file system with an allowed version, using the `-overwrite` keyword if necessary. Retry the migration. If the problem persists, contact the service representative.

6C83	Invalid cancel type. The type must be 1 to cancel an encryption command or 2 to cancel a decryption command.
-------------	---

Action:

Specify the correct cancel type and try the operation again.

6C84	Invalid cancel type. The type must be 3 to cancel a compression command or 4 to cancel a decompression command.
-------------	--

Action:

Specify the correct cancel type and try the operation again.

6C85	Can only delete idle threads from the <code>long_cmd_threads</code> thread pool.
-------------	---

Action:

Try canceling a long-running command to make `long_cmd_threads` thread pool foreground threads idle, or wait for a long-running command to complete. After the necessary long-running commands are completed, try resizing the thread pool again.

6C86	<i>Parms4</i> must be zero if <i>parms1</i> indicates a non-cancel command, such as encrypt, decrypt, compress or decompress. To cancel an encrypt, decrypt, compress or decompress command, <i>parms1</i> must be 3.
-------------	--

Action:

Correct *parms4* or *parms1* and try the operation again.

6C87	Either you are not properly authorized for this operation or zFS could not verify your authorization. The subcommand requires UPDATE access to the aggregate.
-------------	---

Action:

Ensure that the proper access is granted and try the operation again.

6C88	An aggregate with a key label cannot be mounted because ICSF is not initialized.
-------------	---

Action:

Ensure ICSF is initialized and try the command again.

6C89	zEDC is not initialized. The compression or decompression operation is failed.
-------------	---

Action:

Ensure that zEDC is initialized and try the operation again.

6C8A	The mount or ownership change fails because zEDC is not initialized. The aggregate is compressed and requires zEDC to be ready before it can be mounted or owned on this system.
-------------	---

Action:

Try the operation again after zEDC is initialized.

6C8B	ICSF is not available. The encrypt or decrypt operation is failed.
-------------	---

Action:

Ensure that ICSF is initialized and try the operation again.

6C8C	A <code>zfsadm compress</code> command was attempted but no zEDC devices are available at the time the command was issued.
-------------	---

Action:

Try the operation again after zEDC devices are available.

6C8D	This function is currently not supported.
-------------	--

Action:

Contact your service representative to find out when the support will be available, and how to install it.

6C8E	The command failed because the aggregate is disabled. An aggregate can be temporarily disabled due to an I/O error, an internal zFS error, or a dataset reopen operation fails. If the aggregate is encrypted, this can be caused by ICSF problems. zFS will attempt to re-enable the aggregate automatically.
-------------	---

Action:

If the aggregate is encrypted, correct any ICSF problem so the aggregate can be re-enabled. When the aggregate is re-enabled, try the command again. If zFS cannot re-enable the aggregate automatically, you will need to unmount/mount or remount the aggregate before the command can be attempted again.

6C8F	The <code>bk_attrBufferLen</code> specified is too small.
-------------	--

Action:

Change the program to obtain a larger buffer. Then retry the operation.

6C90	The <code>bk_aclBufferLen</code> specified is too small.
-------------	---

Action:

Change the program to obtain a larger buffer. Then retry the operation.

6C92	Backup of a file is not allowed when it is contained in an aggregate that is shrinking.
-------------	--

Action:

Try the operation again after the shrink completes.

6C93	Shrink of an aggregate is not allowed if files in it are being backed up.
-------------	--

Action:

Try the operation again after the file backups complete.

6C94	Internal error committing a transaction. However, the operation may have succeeded.
Action: If needed, try the operation again. If the problem persists, contact the service representative.	
6C95	Internal error while backing up a file.
Action: Contact the service representative.	
6C96	Backup of a file is not allowed when it is contained in an aggregate that is owned by a system where zFS does not support file backups.
Action: Retry the command after ensuring that file backups are supported on every system that can become the zFS owner of the aggregate. If this problem persists, contact your service representative.	
6C97	Backup of a file in a NORWSHARE aggregate is not allowed when it is initiated from a system that is not the zFS owner of the aggregate containing the file.
Action: Use the zfsadm fsinfo or MODIFY FSINFO command to determine which system is the zFS owning system of the aggregate and run the application which uses the zFS backup API from that system. Alternatively, make the local system the zFS owner of the aggregate by using chmount .	
6C98	Performance moves for an aggregate are suspended when files in the aggregate are being backed up.
Action: No user action. zFS may attempt the performance move again after all file backups for the aggregate have concluded.	
6C99	Backup snapshot is no longer valid. This may be due to the original zFS owner going down.
Action: Start the backup again after the zFS owner is established.	
6C9A	Backup of a file is not allowed when it is contained in an aggregate that is encrypting or decrypting.
Action: Try the operation again after the encryption or decryption completes.	
6C9B	Backup of a file is not allowed when it is contained in an aggregate that is compressing or decompressing.
Action: Try the operation again after the compression or decompression completes.	
6C9C	Encrypting or decrypting of an aggregate is not allowed if files in it are being backed up.
Action: Try the operation again after the file backups complete.	
6C9D	Compressing or decompressing of an aggregate is not allowed if files in it are being backed up.
Action: Try the operation again after the file backups complete.	
6C9E	You cannot register a backup of a file that is on a version 1.4 aggregate.

Action

Use the **zfsadm convert -aggrversion** command to convert the aggregate version to 1.5 and then try again. See [z/OS File System Administration](#) for details on converting an existing aggregate to version 1.5.

6C9F **zFS has reached its internal limit for concurrent file backups.**

Action:

Try again after in-progress file backups are complete.

6CA0 **A sysplex client without high availability mount support (HA) is trying to connect to an HA mounted file system.**

Action:

Retry the operation when the client system is at a level that supports high availability mount.

6CA2 **An invalid value was specified for the -ha aggregate attribute.**

Action:

The -ha aggregate attribute takes a value of either ON or OFF. When using the Change Aggregate Attributes API, use the corresponding value, which is either 1 (for ON) or 2 (for OFF).

6CA3 **The aggregate attribute cannot be changed to both -ha and -noha.**

Action:

Limit your request to one attribute value and try again.

6CA5 **The HA aggregate attribute cannot be used in a sysplex where some systems do not support high availability file systems.**

Action:

Ensure that APAR OA57508 is applied to all systems in the sysplex and try the command again.

6CA6 **For a single aggregate query, the buffer passed to FSINFO is too small to contain all of the requested file system statistics.**

Action:

Use a larger buffer size and try the FSINFO request again. Use a buffer size of at least 145K for a single aggregate FSINFO query.

6CA7 **The shrink release percentage value is incorrect. It should be a integer value between 1 and 100.**

Action:

Try the shrink release command again with the percentage of free space to be released in the range of 1 to 100.

6D0A **An internal error occurred during the migration of a zFS file system to another zFS file system.**

Action:

Contact the service representative.

6D0B **An aggregate cannot be migrated if it is already being salvaged, shrunk, encrypted, decrypted, compressed, decompressed, moved, converted to v5, or being migrated.**

Action:

Try the migration again after the mutually exclusive operation is done.

6D0C **A salvage, shrink, encrypt, decrypt, compress, decompress, ownership change, aggregate version conversion, file backup, or another migration operation is not allowed while an aggregate is the source of a migration.**

Action:

Try the operation again after the migration is done.

6D0D	A migrating file system is in an unexpected state during a migration stop request.
-------------	---

Action:

If this is unexpected, contact the service representative.

6D0E	An aggregate cannot be migrated if it is not at version 1.4 or later.
-------------	--

Action:

Try the migration again after the aggregate has been converted to version 1.4 or later.

6D11	Backup of a file is not allowed when it is contained in an aggregate that is being migrated.
-------------	---

Action:

Try the backup operation again when the migration is complete.

6D12	Sharemode attributes cannot be changed when the aggregate is the source of a migration.
-------------	--

Action:

Try the command again when the migration is done.

Appendix B. 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).

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

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Index

Numerics

0001 [241](#), [243](#)
0002 [241](#), [243](#)
0003 [241](#)
0005 [241](#)
0006 [241](#)
0007 [241](#)
0008 [241](#)
0009 [241](#)
000A [242](#)
000B [242](#)
000D [242](#)
000E [242](#)
000F [242](#)
0010 [242](#)
0011 [242](#)
0012 [242](#)
0013 [242](#)
0014 [243](#)
0017 [243](#)
0018 [243](#)
6001 [244](#)
6002 [244](#)
6003 [244](#)
6004 [244](#)
6005 [244](#)
6006 [244](#)
6007 [244](#)
6008 [244](#)
6009 [244](#)
600A [245](#)
600B [245](#)
600C [245](#)
600D [245](#)
600E [245](#)
600F [245](#)
6010 [245](#)
6011 [245](#)
6012 [245](#)
6013 [246](#)
6014 [246](#)
6015 [246](#)
6016 [246](#)
6017 [246](#)
6018 [246](#)
6019 [246](#)
601A [246](#)
601B [246](#)
601C [246](#)
601D [246](#)
601E [247](#)
601F [247](#)
6020 [247](#)
6021 [247](#)
6023 [247](#)
6024 [247](#)

6025 [247](#)
6026 [247](#)
6027 [247](#)
6028 [247](#)
6029 [248](#)
602A [248](#)
602B [248](#)
602C [248](#)
602D [248](#)
602E [248](#)
6030 [248](#)
6031 [248](#)
6032 [248](#)
6033 [248](#)
6034 [249](#)
6035 [249](#)
6036 [249](#)
6037 [249](#)
6038 [249](#)
6039 [249](#)
603A [249](#)
603B [249](#)
603C [249](#)
603D [250](#)
603E [250](#)
603F [250](#)
6040 [250](#)
6041 [250](#)
6042 [250](#)
6043 [250](#)
6044 [250](#)
6045 [250](#)
6046 [250](#)
6047 [251](#)
6048 [251](#)
6049 [251](#)
604A [251](#)
604B [251](#)
604C [251](#)
604D [251](#)
604E [251](#)
604F [251](#)
6050 [251](#)
6051 [252](#)
6052 [252](#)
6053 [252](#)
6054 [252](#)
6055 [252](#)
6056 [252](#)
6057 [252](#)
6058 [252](#)
605A [253](#)
605B [253](#)
605C [253](#)
605D [253](#)
605E [253](#)
6060 [253](#)

[6061 253](#)
[6062 253](#)
[6063 253](#)
[6064 254](#)
[6065 254](#)
[6066 254](#)
[6067 254](#)
[6068 254](#)
[6069 254](#)
[606A 254](#)
[606B 254](#)
[606C 254](#)
[606D 254](#)
[606E 254](#)
[6070 255](#)
[6071 255](#)
[6072 255](#)
[6073 255](#)
[6074 255](#)
[6075 255](#)
[6076 255](#)
[6077 255](#)
[6078 255](#)
[6079 255](#)
[607A 256](#)
[607B 256](#)
[607C 256](#)
[6080 256](#)
[6081 256](#)
[6082 256](#)
[6083 256](#)
[6084 256](#)
[6085 256](#)
[6086 256](#)
[6087 257](#)
[6088 257](#)
[6089 257](#)
[608A 257](#)
[608B 257](#)
[608C 257](#)
[608D 257](#)
[608E 257](#)
[608F 257](#)
[6090 257](#)
[6091 258](#)
[6093 258](#)
[6094 258](#)
[6095 258](#)
[6096 258](#)
[6097 258](#)
[6098 258](#)
[6099 258](#)
[609A 258](#)
[609B 258](#)
[609D 259](#)
[609E 259](#)
[609F 259](#)
[60A0 259](#)
[60A2 259](#)
[60A3 259](#)
[60A5 259](#)
[60A6 259](#)
[60A7 259](#)
[60A8 259](#)

[60A9 260](#)
[60AA 260](#)
[60AB 260](#)
[60AC 260](#)
[60AF 260](#)
[60B0 260](#)
[60B1 260](#)
[60B9 260](#)
[60BC 260](#)
[60BD 260](#)
[60BE 261](#)
[60C0 261](#)
[60C2 261](#)
[60C3 261](#)
[60C4 261](#)
[60C5 261](#)
[60C6 261](#)
[60C7 261](#)
[60C8 261](#)
[60C9 261](#)
[60CA 262](#)
[60CB 262](#)
[60CC 262](#)
[60CD 262](#)
[60D1 262](#)
[60D4 262](#)
[60D5 262](#)
[60D8 262](#)
[60D9 262](#)
[60DA 262](#)
[60DD 263](#)
[60DE 263](#)
[60DF 263](#)
[60E0 263](#)
[60E1 263](#)
[60E2 263](#)
[60E3 263](#)
[60E4 263](#)
[60E5 263](#)
[60E6 263](#)
[60E7 264](#)
[60E8 264](#)
[60E9 264](#)
[60EA 264](#)
[60EB 264](#)
[60ED 264](#)
[60EE 264](#)
[60EF 264](#)
[60F2 264](#)
[60FE 264](#)
[6108 265](#)
[6109 265](#)
[610A 265](#)
[610B 265](#)
[610C 265](#)
[610D 265](#)
[610E 265](#)
[610F 265](#)
[6111 265](#)
[6112 265](#)
[6113 266](#)
[6114 266](#)
[6115 266](#)
[6117 266](#)

[611B 266](#)
[611C 266](#)
[611D 266](#)
[611E 266](#)
[611F 266](#)
[6120 266](#)
[6121 267](#)
[6123 267](#)
[6124 267](#)
[6125 267](#)
[6126 267](#)
[6128 267](#)
[6130 267](#)
[6132 267](#)
[6133 267](#)
[6136 267](#)
[6137 268](#)
[6138 268](#)
[6139 268](#)
[613A 268](#)
[613B 268](#)
[613C 268](#)
[613D 268](#)
[613E 268](#)
[613F 268](#)
[6140 268](#)
[6141 269](#)
[6142 269](#)
[6143 269](#)
[6144 269](#)
[6145 269](#)
[6146 269](#)
[6147 269](#)
[6148 269](#)
[614A 269](#)
[614B 269](#)
[614C 270](#)
[614D 270](#)
[614E 270](#)
[614F 270](#)
[6150 270](#)
[6151 270](#)
[6161 270](#)
[6162 270](#)
[6163 270](#)
[6164 270](#)
[6165 271](#)
[6166 271](#)
[6167 271](#)
[6168 271](#)
[6200 271](#)
[6201 271](#)
[6202 271](#)
[6203 271](#)
[6204 271](#)
[6205 271](#)
[6206 272](#)
[6207 272](#)
[6208 272](#)
[6209 272](#)
[620A 272](#)
[620B 272](#)
[620C 272](#)
[620D 272](#)

[620E 272](#)
[6223 272](#)
[6224 273](#)
[6225 273](#)
[6238 273](#)
[6239 273](#)
[623A 273](#)
[623B 273](#)
[623E 273](#)
[623F 273](#)
[6240 273](#)
[6241 273](#)
[6246 274](#)
[6248 274](#)
[624D 274](#)
[624E 274](#)
[624F 274](#)
[6252 274](#)
[6253 274](#)
[6254 274](#)
[6255 274](#)
[6256 274](#)
[6258 275](#)
[6259 275](#)
[625C 275](#)
[625D 275](#)
[6260 275](#)
[6261 275](#)
[6262 275](#)
[6263 275](#)
[6264 275](#)
[6265 276](#)
[6266 276](#)
[6267 276](#)
[6268 276](#)
[6269 276](#)
[626A 276](#)
[626B 276](#)
[626D 276](#)
[626E 276](#)
[626F 276](#)
[6270 277](#)
[6271 277](#)
[6273 277](#)
[6274 277](#)
[6277 277](#)
[6278 277](#)
[6279 277](#)
[627A 277](#)
[629E 277](#)
[629F 278](#)
[62A0 278](#)
[62A1 278](#)
[62A2 278](#)
[62A3 278](#)
[62A4 278](#)
[62A5 278](#)
[62A6 278](#)
[62A7 278](#)
[62AA 278](#)
[62AB 278](#)
[62AC 279](#)
[62AD 279](#)
[62AE 279](#)

[62AF 279](#)
[62B0 279](#)
[62B01 337](#)
[62B3 279](#)
[62B4 279](#)
[62B5 279](#)
[62B8 279](#)
[62B9 279](#)
[62BA 280](#)
[62BB 280](#)
[62BC 280](#)
[62BD 280](#)
[62BE 280](#)
[62BF 280](#)
[62C0 280](#)
[62C3 280](#)
[62C4 280](#)
[62C5 280](#)
[62C8 281](#)
[62C9 281](#)
[62CA 281](#)
[62CB 281](#)
[62CC 281](#)
[62CF 281](#)
[62D0 281](#)
[62D1 281](#)
[62D2 281](#)
[62D3 281](#)
[62D4 282](#)
[62D5 282](#)
[62D6 282](#)
[62D7 282](#)
[62D8 282](#)
[62D9 282](#)
[62DA 282](#)
[62DB 282](#)
[62DC 282](#)
[62E1 282](#)
[62E2 283](#)
[62E3 283](#)
[62E4 283](#)
[62E5 283](#)
[62E6 283](#)
[6300 283](#)
[6301 283](#)
[6310 283](#)
[6311 283](#)
[6312 284](#)
[6313 284](#)
[6314 284](#)
[6315 284](#)
[6316 284](#)
[6317 284](#)
[6318 284](#)
[6319 284](#)
[631A 284](#)
[631B 284](#)
[631C 284](#)
[631D 285](#)
[631E 285](#)
[631F 285](#)
[6320 285](#)
[6321 285](#)
[6322 285](#)

[6323 285](#)
[6324 285](#)
[6325 285](#)
[6400 285](#)
[6401 286](#)
[6402 286](#)
[6403 286](#)
[6404 286](#)
[6405 286](#)
[6406 286](#)
[6407 286](#)
[6408 286](#)
[6409 286](#)
[640A 286](#)
[640B 287](#)
[640C 287](#)
[640D 287](#)
[640E 287](#)
[640F 287](#)
[6410 287](#)
[6411 287](#)
[6412 287](#)
[6413 287](#)
[6414 287](#)
[6415 288](#)
[6416 288](#)
[6417 288](#)
[6418 288](#)
[6419 288](#)
[641A 288](#)
[6420 288](#)
[6421 288](#)
[6422 288](#)
[6423 289](#)
[6424 289](#)
[6425 289](#)
[6426 289](#)
[6427 289](#)
[6428 289](#)
[6429 289](#)
[642A 289](#)
[642B 289](#)
[6440 290](#)
[6441 290](#)
[6443 290](#)
[6444 290](#)
[6445 290](#)
[6446 290](#)
[6500 290](#)
[6501 290](#)
[6503 290](#)
[6504 290](#)
[6505 291](#)
[6506 291](#)
[6508 291](#)
[6509 291](#)
[650A 291](#)
[650B 291](#)
[650C 291](#)
[650D 291](#)
[650F 291](#)
[6512 292](#)
[6513 292](#)
[6514 292](#)

[6515 292](#)
[6516 292](#)
[6517 292](#)
[6518 292](#)
[6519 292](#)
[651A 292](#)
[651B 293](#)
[651C 293](#)
[651D 293](#)
[651E 293](#)
[651F 293](#)
[6520 293](#)
[6521 293](#)
[6522 293](#)
[652B 293](#)
[652C 293](#)
[652D 294](#)
[6609 294](#)
[660A 294](#)
[6611 294](#)
[6617 294](#)
[6619 294](#)
[6638 294](#)
[6639 294](#)
[663B 295](#)
[663D 295](#)
[663E 295](#)
[663F 295](#)
[6640 295](#)
[6641 295](#)
[6649 295](#)
[664A 295](#)
[664B 295](#)
[664C 295](#)
[664D 296](#)
[664E 296](#)
[664F 296](#)
[6650 296](#)
[6651 296](#)
[6652 296](#)
[6653 296](#)
[6654 296](#)
[6655 296](#)
[6656 296](#)
[6657 297](#)
[6658 297](#)
[6659 297](#)
[665A 297](#)
[665B 297](#)
[665C 297](#)
[665D 297](#)
[665E 297](#)
[665F 297](#)
[6660 297](#)
[6661 298](#)
[6662 298](#)
[6663 298](#)
[666B 298](#)
[666C 298](#)
[666D 298](#)
[666F 298](#)
[6670 298](#)
[6671 298](#)
[6672 298](#)

[6673 299](#)
[6674 299](#)
[6675 299](#)
[6676 299](#)
[6677 299](#)
[6680 299](#)
[6681 299](#)
[6682 299](#)
[6683 299](#)
[6684 300](#)
[6685 300](#)
[6686 300](#)
[6689 300](#)
[668A 300](#)
[6690 300](#)
[6696 300](#)
[6697 300](#)
[6698 300](#)
[6699 300](#)
[66A0 301](#)
[66A1 301](#)
[66A4 301](#)
[66A6 301](#)
[66A7 301](#)
[66A8 301](#)
[66AA 301](#)
[66AC 301](#)
[66AD 301](#)
[66B1 301](#)
[66B2 302](#)
[66B4 302](#)
[66B5 302](#)
[66B8 302](#)
[66B9 302](#)
[66BA 302](#)
[66BB 302](#)
[66BC 302](#)
[66BD 302](#)
[66BE 303](#)
[66C0 303](#)
[66C1 303](#)
[66C2 303](#)
[66C3 303](#)
[66C4 303](#)
[66C5 303](#)
[66C6 303](#)
[66C7 303](#)
[66C8 303](#)
[66C9 304](#)
[66CA 304](#)
[66CB 304](#)
[66CC 304](#)
[66CD 304](#)
[66CE 304](#)
[66CF 304](#)
[66D0 304](#)
[66D1 304](#)
[66D2 304](#)
[66D3 305](#)
[66D4 305](#)
[66D5 305](#)
[66D6 305](#)
[66D7 305](#)
[66DF 305](#)

66E0 305	6745 312
66E1 305	6746 312
66E4 305	6747 312
66E5 306	6748 312
66E6 306	6749 312
66E7 306	674A 312
66E8 306	674B 312
66E9 306	674C 312
66EA 306	674D 313
66EB 306	674E 313
66EC 306	6750 313
66ED 306	6751 313
66EE 306	6752 313
66EF 307	6753 313
66F0 307	6754 313
66F1 307	6755 313
66F2 307	675D 313
66F3 307	675E 313
66F4 307	675F 314
66F5 307	6760 314
66F6 307	6763 314
6700 307	6764 314
6701 307	6769 314
6702 308	676A 314
6703 308	676B 314
6704 308	676D 314
6705 308	676E 314
6706 308	676F 314
6707 308	6770 315
6708 308	6771 315
6709 308	6772 315
670A 308	6775 315
670B 308	6776 315
670C 309	6800 315
670D 309	6806 315
670E 309	6809 315
670F 309	680A 315
6711 309	680B 316
6714 309	680C 316
6715 309	6900 316
6716 309	6901 316
6717 309	6903 316
6718 309	6904 316
6719 310	6905 316
671A 310	6906 316
671B 310	6907 316
671D 310	6908 316
671E 310	6909 317
671F 310	690A 317
6720 310	690B 317
6721 310	690C 317
6722 310	690D 317
6723 311	690E 317
6724 311	690F 317
6726 311	6910 317
6727 311	6911 317
6728 311	6912 317
6729 311	6913 318
672A 311	6914 318
672B 311	6915 318
6730 311	6916 318
6738 311	6917 318
673D 312	6918 318
6743 312	6919 318

[691A 318](#)
[691B 318](#)
[691C 318](#)
[691D 319](#)
[691E 319](#)
[691F 319](#)
[6920 319](#)
[6921 319](#)
[6922 319](#)
[6923 319](#)
[6924 319](#)
[6925 319](#)
[6926 319](#)
[6927 320](#)
[6928 320](#)
[6929 320](#)
[692A 320](#)
[692B 320](#)
[692C 320](#)
[692D 320](#)
[692E 320](#)
[692F 320](#)
[6930 320](#)
[6931 321](#)
[6932 321](#)
[6933 321](#)
[6934 321](#)
[6935 321](#)
[6936 321](#)
[6937 321](#)
[6938 321](#)
[6939 321](#)
[693A 321](#)
[693B 322](#)
[693C 322](#)
[693D 322](#)
[693E 322](#)
[693F 322](#)
[6940 322](#)
[6941 322](#)
[6942 322](#)
[6943 322](#)
[6944 322](#)
[6945 323](#)
[6946 323](#)
[6947 323](#)
[6948 323](#)
[6949 323](#)
[694A 323](#)
[694B 323](#)
[694C 323](#)
[694D 323](#)
[694E 323](#)
[694F 324](#)
[6950 324](#)
[6951 324](#)
[6952 324](#)
[6953 324](#)
[6954 324](#)
[6955 324](#)
[6956 324](#)
[6957 324](#)
[6958 324](#)
[6959 325](#)

[695A 325](#)
[695B 325](#)
[695C 325](#)
[695D 325](#)
[695E 325](#)
[695F 325](#)
[6960 325](#)
[6961 325](#)
[6962 325](#)
[6963 326](#)
[6964 326](#)
[6965 326](#)
[6967 326](#)
[6968 326](#)
[6969 326](#)
[696A 326](#)
[696B 326](#)
[696C 326](#)
[696D 327](#)
[697C 327](#)
[698A 327](#)
[698B 327](#)
[698C 327](#)
[698E 327](#)
[6994 327](#)
[6995 327](#)
[6996 327](#)
[699A 327](#)
[699B 328](#)
[699C 328](#)
[699D 328](#)
[699E 328](#)
[699F 328](#)
[69A0 328](#)
[69A1 328](#)
[69A2 328](#)
[69A3 328](#)
[69A4 328](#)
[69A5 329](#)
[69A7 329](#)
[69A8 329](#)
[69A9 329](#)
[69AA 329](#)
[69AB 329](#)
[69AC 329](#)
[69AD 330](#)
[69AE 330](#)
[69AF 330](#)
[69B0 330](#)
[69B1 330](#)
[69B2 330](#)
[69B3 330](#)
[69B4 330](#)
[69B5 330](#)
[69B6 330](#)
[69B7 331](#)
[69BA 331](#)
[69BB 331](#)
[69BD 331](#)
[69BE 331](#)
[69BF 331](#)
[69C0 331](#)
[69C1 331](#)
[69C2 331](#)

[69C3 331](#)
[69C5 332](#)
[69C6 332](#)
[69C7 332](#)
[69C8 332](#)
[69C9 332](#)
[69CA 332](#)
[69CB 332](#)
[69CC 332](#)
[69CD 332](#)
[6A00 333](#)
[6A01 333](#)
[6A02 333](#)
[6A03 333](#)
[6A04 333](#)
[6A05 333](#)
[6A06 333](#)
[6A07 333](#)
[6A08 333](#)
[6A09 333](#)
[6A0A 334](#)
[6A0C 334](#)
[6A0D 334](#)
[6A0E 334](#)
[6A0F 334](#)
[6A10 334](#)
[6A11 334](#)
[6A12 334](#)
[6A13 334](#)
[6A17 335](#)
[6A18 335](#)
[6A19 335](#)
[6A1A 335](#)
[6A1B 335](#)
[6A1C 335](#)
[6A1D 335](#)
[6A1E 335](#)
[6A1F 335](#)
[6A24 335](#)
[6A25 336](#)
[6A27 336](#)
[6A28 336](#)
[6A29 336](#)
[6A2A 336](#)
[6A2B 336](#)
[6A2C 336](#)
[6A2D 336](#)
[6A2E 336](#)
[6A2F 337](#)
[6A30 337](#)
[6A31 337](#)
[6A32 337](#)
[6A33 337](#)
[6B00 337](#)
[6B02 337](#)
[6B03 338](#)
[6B04 338](#)
[6B05 338](#)
[6B06 338](#)
[6B0C 338](#)
[6B0D 338](#)
[6B0E 338](#)
[6B0F 338](#)
[6B10 338](#)

[6B11 339](#)
[6B12 339](#)
[6B13 339](#)
[6B14 339](#)
[6B15 339](#)
[6B16 339](#)
[6B17 339](#)
[6B18 339](#)
[6B19 339](#)
[6B1A 340](#)
[6B1B 340](#)
[6B1C 340](#)
[6B1E 340](#)
[6B1F 340](#)
[6B20 340](#)
[6B21 340](#)
[6B22 340](#)
[6B23 340](#)
[6B24 340](#)
[6B25 341](#)
[6B26 341](#)
[6B27 341](#)
[6B28 341](#)
[6B48 341](#)
[6B49 341](#)
[6B4A 341](#)
[6B4B 341](#)
[6B4C 341](#)
[6B4D 341](#)
[6B4E 342](#)
[6B50 342](#)
[6B51 342](#)
[6B52 342](#)
[6B53 342](#)
[6B54 342](#)
[6B55 342, 343](#)
[6B5E 343](#)
[6B5F 343](#)
[6B60 343](#)
[6B61 343](#)
[6C00 343](#)
[6C04 343](#)
[6C0B 343](#)
[6C0C 343](#)
[6C0D 344](#)
[6C0E 344](#)
[6C0F 344](#)
[6C10 344](#)
[6C11 344](#)
[6C12 344](#)
[6C13 344](#)
[6C14 344](#)
[6C15 344](#)
[6C18 345](#)
[6C19 345](#)
[6C1A 345](#)
[6C1B 345](#)
[6C1C 345](#)
[6C1D 345](#)
[6C21 345](#)
[6C22 345](#)
[6C23 346](#)
[6C24 346](#)
[6C25 346](#)

[6C26](#) [346](#)
[6C27](#) [346](#)
[6C28](#) [346](#)
[6C29](#) [346](#)
[6C2A](#) [346](#)
[6C2B](#) [346](#)
[6C2C](#) [346](#)
[6C2D](#) [347](#)
[6C2F](#) [347](#)
[6C30](#) [347](#)
[6C31](#) [347](#)
[6C32](#) [347](#)
[6C33](#) [347](#)
[6C34](#) [347](#)
[6C35](#) [347](#)
[6C36](#) [347](#)
[6C37](#) [347](#)
[6C3E](#) [347](#)
[6C3F](#) [348](#)
[6C40](#) [348](#)
[6C41](#) [348](#)
[6C42](#) [348](#)
[6C43](#) [348](#)
[6C45](#) [348](#)
[6C46](#) [348](#)
[6C47](#) [348](#)
[6C48](#) [348](#)
[6C49](#) [349](#)
[6C4A](#) [349](#)
[6C4F](#) [349](#)
[6C50](#) [349](#)
[6C51](#) [349](#)
[6C52](#) [349](#)
[6C53](#) [349](#)
[6C54](#) [349](#)
[6C55](#) [350](#)
[6C56](#) [350](#)
[6C57](#) [350](#)
[6C58](#) [350](#)
[6C5B](#) [350](#)
[6C5C](#) [350](#)
[6C5D](#) [350](#)
[6C5E](#) [350](#)
[6C5F](#) [350](#)
[6C60](#) [350](#)
[6C61](#) [350](#)
[6C62](#) [351](#)
[6C63](#) [351](#)
[6C64](#) [351](#)
[6C65](#) [351](#)
[6C66](#) [351](#)
[6C67](#) [351](#)
[6C68](#) [351](#)
[6C69](#) [351](#)
[6C6A](#) [351](#)
[6C6B](#) [352](#)
[6C6C](#) [352](#)
[6C6E](#) [352](#)
[6C6F](#) [352](#)
[6C71](#) [352](#)
[6C72](#) [352](#)
[6C73](#) [352](#)
[6C74](#) [352](#)
[6C75](#) [353](#)

[6C76](#) [353](#)
[6C77](#) [353](#)
[6C78](#) [353](#)
[6C79](#) [353](#)
[6C7A](#) [353](#)
[6C7B](#) [353](#)
[6C7C](#) [353](#)
[6C7D](#) [353](#)
[6C7E](#) [353](#)
[6C7F](#) [354](#)
[6C80](#) [354](#)
[6C81](#) [354](#)
[6C82](#) [354](#)
[6C83](#) [354](#)
[6C84](#) [354](#)
[6C85](#) [354](#)
[6C86](#) [354](#)
[6C87](#) [354](#)
[6C88](#) [355](#)
[6C89](#) [355](#)
[6C8A](#) [355](#)
[6C8B](#) [355](#)
[6C8C](#) [355](#)
[6C8D](#) [355](#)
[6C8E](#) [355](#)
[6C8F](#) [355](#)
[6C90](#) [355](#)
[6C92](#) [355](#)
[6C93](#) [355](#)
[6C94](#) [356](#)
[6C95](#) [356](#)
[6C96](#) [356](#)
[6C97](#) [356](#)
[6C98](#) [356](#)
[6C99](#) [356](#)
[6C9A](#) [356](#)
[6C9B](#) [356](#)
[6C9C](#) [356](#)
[6C9D](#) [356](#)
[6C9E](#) [356](#)
[6C9F](#) [357](#)
[6CA0](#) [357](#)
[6CA2](#) [357](#)
[6CA3](#) [357](#)
[6CA5](#) [357](#)
[6CA6](#) [357](#)
[6CA7](#) [357](#)
[6D0A](#) [357](#)
[6D0B](#) [357](#)
[6D0C](#) [357](#)
[6D0D](#) [358](#)
[6D0E](#) [358](#)
[6D11](#) [358](#)
[6D12](#) [358](#)

[A](#)

[accessibility](#)
 [contact IBM](#) [359](#)
[assistive technologies](#) [359](#)

C

contact
z/OS [359](#)

I

[IOEZ00001E 5](#)
[IOEZ00002E 5](#)
[IOEZ00003E 5](#)
[IOEZ00004I 5](#)
[IOEZ00005I 6](#)
[IOEZ00006E 6](#)
[IOEZ00007A 6](#)
[IOEZ00008E 6](#)
[IOEZ00009I 6](#)
[IOEZ00010A 7](#)
[IOEZ00011A 7](#)
[IOEZ00012A 7](#)
[IOEZ00013A 7](#)
[IOEZ00014A 7](#)
[IOEZ00015A 8](#)
[IOEZ00016I 8](#)
[IOEZ00017A 8](#)
[IOEZ00018A 8](#)
[IOEZ00019E 9](#)
[IOEZ00020I 9](#)
[IOEZ00023E 9](#)
[IOEZ00024E 9](#)
[IOEZ00025I 10](#)
[IOEZ00032I 10](#)
[IOEZ00033E 10](#)
[IOEZ00034I 10](#)
[IOEZ00035I 10](#)
[IOEZ00036I 10](#)
[IOEZ00037I 11](#)
[IOEZ00038E 11](#)
[IOEZ00039E 11](#)
[IOEZ00040E 11](#)
[IOEZ00041I 11](#)
[IOEZ00042I 12](#)
[IOEZ00043I 12](#)
[IOEZ00044I 12](#)
[IOEZ00045I 12](#)
[IOEZ00046E 12](#)
[IOEZ00048I 13](#)
[IOEZ00050I 13](#)
[IOEZ00051I 13](#)
[IOEZ00052I 13](#)
[IOEZ00053E 13](#)
[IOEZ00054E 14](#)
[IOEZ00055I 14](#)
[IOEZ00057I 14](#)
[IOEZ00062A 14](#)
[IOEZ00064I 14](#)
[IOEZ00068E 15](#)
[IOEZ00069I 15](#)
[IOEZ00070E 15](#)
[IOEZ00077I 15](#)
[IOEZ00078E 15](#)
[IOEZ00079I 16](#)
[IOEZ00080A 16](#)
[IOEZ00081A 16](#)
[IOEZ00082A 16](#)

[IOEZ00083A 17](#)
[IOEZ00084E 17](#)
[IOEZ00085E 17](#)
[IOEZ00087I 17](#)
[IOEZ00088I 18](#)
[IOEZ00092E 18](#)
[IOEZ00093E 18](#)
[IOEZ00094E 18](#)
[IOEZ00095E 18](#)
[IOEZ00096E 19](#)
[IOEZ00100E 19](#)
[IOEZ00105I 19](#)
[IOEZ00106I 19](#)
[IOEZ00109E 20](#)
[IOEZ00110E 20](#)
[IOEZ00112E 20](#)
[IOEZ00117I 20](#)
[IOEZ00118I 20](#)
[IOEZ00119E 21](#)
[IOEZ00120E 21](#)
[IOEZ00122I 21](#)
[IOEZ00123I 21](#)
[IOEZ00124E 21](#)
[IOEZ00127I 22](#)
[IOEZ00129I 22](#)
[IOEZ00131E 22](#)
[IOEZ00132E 22](#)
[IOEZ00133E 23](#)
[IOEZ00134E 23](#)
[IOEZ00135E 23](#)
[IOEZ00136E 23](#)
[IOEZ00138E 23](#)
[IOEZ00139E 24](#)
[IOEZ00140I 24](#)
[IOEZ00141E 24](#)
[IOEZ00142E 24](#)
[IOEZ00143E 24](#)
[IOEZ00144E 25](#)
[IOEZ00157E 25](#)
[IOEZ00158E 25](#)
[IOEZ00159E 25](#)
[IOEZ00163I 26](#)
[IOEZ00164I 26](#)
[IOEZ00165E 26](#)
[IOEZ00166I 26](#)
[IOEZ00167I 26](#)
[IOEZ00168E 27](#)
[IOEZ00170E 27](#)
[IOEZ00171I 27](#)
[IOEZ00173I 28](#)
[IOEZ00175E 28](#)
[IOEZ00178I 28](#)
[IOEZ00179I 28](#)
[IOEZ00181E 29](#)
[IOEZ00182E 29](#)
[IOEZ00183E 29](#)
[IOEZ00184E 29](#)
[IOEZ00185E 30](#)
[IOEZ00186E 30](#)
[IOEZ00187I 30](#)
[IOEZ00188A 30](#)
[IOEZ00190E 31](#)
[IOEZ00191E 31](#)
[IOEZ00199E 31](#)

IOEZ00200E	31	IOEZ00342I	48
IOEZ00201E	32	IOEZ00350I	48
IOEZ00202E	32	IOEZ00351E	48
IOEZ00207E	32	IOEZ00353E	48
IOEZ00208E	32	IOEZ00354E	49
IOEZ00209E	32	IOEZ00356E	49
IOEZ00210E	33	IOEZ00357I	50
IOEZ00211E	33	IOEZ00358E	50
IOEZ00212E	33	IOEZ00359E	51
IOEZ00213E	33	IOEZ00360I	51
IOEZ00214E	34	IOEZ00361I	51
IOEZ00229I	34	IOEZ00362E	51
IOEZ00230I	34	IOEZ00363E	52
IOEZ00231I	34	IOEZ00366E	52
IOEZ00232I	34	IOEZ00368I	53
IOEZ00233I	35	IOEZ00369I	53
IOEZ00234I	35	IOEZ00370I	53
IOEZ00235I	35	IOEZ00371E	53
IOEZ00236I	35	IOEZ00373E	54
IOEZ00237I	36	IOEZ00374I	54
IOEZ00238I	36	IOEZ00380E	54
IOEZ00239E	36	IOEZ00381E	55
IOEZ00240E	36	IOEZ00382I	55
IOEZ00241I	36	IOEZ00383E	55
IOEZ00242I	37	IOEZ00385E	56
IOEZ00243I	37	IOEZ00387I	56
IOEZ00244E	37	IOEZ00388I	57
IOEZ00245E	37	IOEZ00389I	57
IOEZ00246E	37	IOEZ00390I	57
IOEZ00247E	38	IOEZ00391I	58
IOEZ00248I	38	IOEZ00392I	58
IOEZ00249E	38	IOEZ00393I	58
IOEZ00250E	38	IOEZ00394I	59
IOEZ00251E	38	IOEZ00395I	59
IOEZ00252E	39	IOEZ00396I	60
IOEZ00253E	39	IOEZ00397I	60
IOEZ00300I	39	IOEZ00398E	60
IOEZ00301E	39	IOEZ00400I	61
IOEZ00303E	40	IOEZ00401I	61
IOEZ00304E	40	IOEZ00405I	62
IOEZ00308E	40	IOEZ00410I	62
IOEZ00309I	40	IOEZ00411I	62
IOEZ00312I	41	IOEZ00412I	63
IOEZ00314E	41	IOEZ00413I	63
IOEZ00315I	41	IOEZ00416I	63
IOEZ00317I	41	IOEZ00417E	64
IOEZ00318I	42	IOEZ00418I	64
IOEZ00320I	42	IOEZ00420E	64
IOEZ00321E	42	IOEZ00421E	65
IOEZ00322E	42	IOEZ00422E	65
IOEZ00323I	43	IOEZ00424E	66
IOEZ00324I	43	IOEZ00425E	66
IOEZ00325E	44	IOEZ00426E	66
IOEZ00326E	44	IOEZ00433E	67
IOEZ00327I	44	IOEZ00434E	67
IOEZ00328E	44	IOEZ00435E	67
IOEZ00329I	44	IOEZ00436E	68
IOEZ00331A	45	IOEZ00437I	68
IOEZ00334I	45	IOEZ00438I	68
IOEZ00336I	45	IOEZ00439I	69
IOEZ00337E	46	IOEZ00440E	69
IOEZ00338A	46	IOEZ00441E	69
IOEZ00340E	47	IOEZ00442E	70
IOEZ00341E	47	IOEZ00443E	70

IOEZ00444E	70	IOEZ00613I	94
IOEZ00445E	70	IOEZ00614A	94
IOEZ00447E	71	IOEZ00615E	95
IOEZ00451E	71	IOEZ00616E	95
IOEZ00453I	72	IOEZ00617I	95
IOEZ00500I	72	IOEZ00618E	96
IOEZ00501E	72	IOEZ00619E	97
IOEZ00502E	73	IOEZ00620E	97
IOEZ00503E	73	IOEZ00621E	98
IOEZ00504E	73	IOEZ00622E	98
IOEZ00505I	74	IOEZ00623E	98
IOEZ00506E	74	IOEZ00624E	99
IOEZ00507E	74	IOEZ00625E	99
IOEZ00508E	75	IOEZ00626E	100
IOEZ00509E	75	IOEZ00627E	101
IOEZ00510E	75	IOEZ00628E	101
IOEZ00511E	76	IOEZ00629E	101
IOEZ00512E	76	IOEZ00630E	102
IOEZ00513E	76	IOEZ00631E	102
IOEZ00514E	77	IOEZ00632E	103
IOEZ00515E	77	IOEZ00633E	103
IOEZ00516E	77	IOEZ00634E	104
IOEZ00518I	78	IOEZ00635E	104
IOEZ00519E	78	IOEZ00636E	104
IOEZ00520E	78	IOEZ00639I	105
IOEZ00521I	78	IOEZ00640E	106
IOEZ00522E	79	IOEZ00641I	106
IOEZ00523I	79	IOEZ00642E	106
IOEZ00524I	79	IOEZ00643I	107
IOEZ00536E	80	IOEZ00644I	107
IOEZ00537E	80	IOEZ00645A	108
IOEZ00538I	80	IOEZ00646I	108
IOEZ00539E	81	IOEZ00650I	109
IOEZ00540I	81	IOEZ00651E	109
IOEZ00545E	81	IOEZ00652E	109
IOEZ00547I	81	IOEZ00653E	109
IOEZ00548I	82	IOEZ00655E	110
IOEZ00549E	82	IOEZ00657E	110
IOEZ00550E	83	IOEZ00658E	110
IOEZ00551E	83	IOEZ00659E	111
IOEZ00553E	83	IOEZ00660I	111
IOEZ00555E	84	IOEZ00661I	112
IOEZ00557E	84	IOEZ00662I	112
IOEZ00558E	84	IOEZ00663I	113
IOEZ00559I	85	IOEZ00664E	113
IOEZ00576E	85	IOEZ00665E	114
IOEZ00579I	86	IOEZ00666E	114
IOEZ00581E	86	IOEZ00667I	114
IOEZ00587I	87	IOEZ00668I	114
IOEZ00589E	87	IOEZ00670I	115
IOEZ00590E	87	IOEZ00671E	115
IOEZ00591I	88	IOEZ00674E	115
IOEZ00592I	88	IOEZ00675E	116
IOEZ00598E	89	IOEZ00676E	116
IOEZ00599E	89	IOEZ00677E	116
IOEZ00604I	89	IOEZ00678E	117
IOEZ00605I	90	IOEZ00700E	117
IOEZ00606E	91	IOEZ00701E	118
IOEZ00607I	91	IOEZ00702E	118
IOEZ00608I	92	IOEZ00703E	118
IOEZ00609I	92	IOEZ00704E	119
IOEZ00610I	93	IOEZ00705I	119
IOEZ00611I	93	IOEZ00707I	119
IOEZ00612I	93	IOEZ00708E	119

IOEZ00709I	120	IOEZ00783E	144
IOEZ00710E	120	IOEZ00784E	145
IOEZ00711I	120	IOEZ00785I	145
IOEZ00712E	121	IOEZ00786I	146
IOEZ00713E	121	IOEZ00787I	146
IOEZ00715E	121	IOEZ00788E	146
IOEZ00718I	122	IOEZ00789E	147
IOEZ00719I	122	IOEZ00790I	147
IOEZ00720I	122	IOEZ00791I	147
IOEZ00721I	123	IOEZ00792E	148
IOEZ00722I	123	IOEZ00793E	148
IOEZ00723E	123	IOEZ00794E	149
IOEZ00724I	124	IOEZ00795E	149
IOEZ00725I	124	IOEZ00797I	150
IOEZ00726I	124	IOEZ00798I	150
IOEZ00727I	125	IOEZ00799A	151
IOEZ00728I	125	IOEZ00800I	151
IOEZ00729I	125	IOEZ00801I	152
IOEZ00730I	126	IOEZ00802I	152
IOEZ00731I	126	IOEZ00803I	153
IOEZ00733I	126	IOEZ00804E	153
IOEZ00734E	127	IOEZ00805A	154
IOEZ00735I	127	IOEZ00806A	154
IOEZ00736I	127	IOEZ00807I	155
IOEZ00739I	128	IOEZ00808I	155
IOEZ00740E	128	IOEZ00809I	156
IOEZ00741I	129	IOEZ00810I	156
IOEZ00742I	129	IOEZ00811E	156
IOEZ00743I	130	IOEZ00812I	157
IOEZ00744E	130	IOEZ00813I	157
IOEZ00745E	130	IOEZ00814E	158
IOEZ00746E	131	IOEZ00815E	158
IOEZ00747I	131	IOEZ00822E	159
IOEZ00750E	132	IOEZ00823E	159
IOEZ00751I	132	IOEZ00824E	160
IOEZ00752E	132	IOEZ00825E	160
IOEZ00753I	133	IOEZ00826A	160
IOEZ00754I	133	IOEZ00827I	161
IOEZ00755I	133	IOEZ00828E	161
IOEZ00756I	134	IOEZ00829E	162
IOEZ00757I	134	IOEZ00830E	162
IOEZ00758I	134	IOEZ00831E	163
IOEZ00759I	135	IOEZ00832E	163
IOEZ00760I	135	IOEZ00833E	164
IOEZ00761E	136	IOEZ00834E	164
IOEZ00762E	136	IOEZ00835E	165
IOEZ00763E	136	IOEZ00836I	165
IOEZ00764E	137	IOEZ00837E	166
IOEZ00765E	137	IOEZ00838E	166
IOEZ00766E	138	IOEZ00839E	166
IOEZ00767E	138	IOEZ00840E	167
IOEZ00768E	138	IOEZ00841E	167
IOEZ00769E	139	IOEZ00842E	167
IOEZ00770E	139	IOEZ00843E	168
IOEZ00771E	139	IOEZ00844E	168
IOEZ00773E	140	IOEZ00845E	168
IOEZ00774E	140	IOEZ00846E	169
IOEZ00775E	141	IOEZ00847E	169
IOEZ00776I	141	IOEZ00848I	169
IOEZ00777A	141	IOEZ00849I	170
IOEZ00778I	142	IOEZ00850I	170
IOEZ00780E	143	IOEZ00851I	170
IOEZ00781I	143	IOEZ00852E	171
IOEZ00782I	144	IOEZ00853I	171

IOEZ00854I 172	IOEZ00984I 200
IOEZ00856E 172	IOEZ00985I 200
IOEZ00857I 172	IOEZ00986I 201
IOEZ00858E 173	IOEZ00987I 201
IOEZ00859I 173	IOEZ00988I 202
IOEZ00860E 173	IOEZ00989I 202
IOEZ00861E 174	IOEZ00990E 202
IOEZ00862E 174	IOEZ00991E 203
IOEZ00866E 175	IOEZ00992E 203
IOEZ00867I 175	IOEZ00993E 204
IOEZ00870A 176	IOEZ00994I 204
IOEZ00872E 176	IOEZ00995E 204
IOEZ00873I 177	IOEZ00996E 205
IOEZ00877I 177	IOEZ00997E 205
IOEZ00879E 177	IOEZ00999I 205
IOEZ00880E 178	IOEZ01000I 206
IOEZ00881I 178	IOEZ01001I 206
IOEZ00882E 179	IOEZ01002E 207
IOEZ00883E 179	IOEZ01003I 207
IOEZ00884E 179	IOEZ01004E 207
IOEZ00885I 180	IOEZ01005I 208
IOEZ00886E 180	IOEZ01006I 209
IOEZ00887E 180	IOEZ01007I 209
IOEZ00888I 181	IOEZ01008I 209
IOEZ00889E 181	IOEZ01009I 210
IOEZ00890E 182	IOEZ01027E 210
IOEZ00891E 182	IOEZ01028E 211
IOEZ00892I 183	IOEZ01029E 211
IOEZ00895E 183	IOEZ01030E 212
IOEZ00896I 183	IOEZ01031E 212
IOEZ00897E 184	IOEZ01040A 213
IOEZ00898I 184	IOEZ01042E 213
IOEZ00899I 185	IOEZ01060E 213
IOEZ00900I 185	IOEZ01061I 214
IOEZ00901E 185	IOEZ01076E 214
IOEZ00902E 186	IOEZ01077I 215
IOEZ00903I 186	IOEZ01078E 215
IOEZ00904E 187	IOEZ01079E 216
IOEZ00905E 187	IOEZ01080I 216
IOEZ00906E 188	IOEZ01081I 216
IOEZ00907E 188	IOEZ01082I 217
IOEZ00908I 189	IOEZ01083I 217
IOEZ00942E 189	IOEZ01084I 218
IOEZ00943E 190	IOEZ01085E 218
IOEZ00947A 190	IOEZ01086E 219
IOEZ00948A 191	IOEZ01087I 219
IOEZ00949E 191	IOEZ01088I 220
IOEZ00950E 192	IOEZ01089E 220
IOEZ00951I 192	IOEZH0001I 221
IOEZ00952E 193	IOEZH0002I 221
IOEZ00953E 193	IOEZH0040I 222
IOEZ00957E 194	IOEZH0041I 222
IOEZ00958E 194	IOEZH0042I 223
IOEZ00959E 195	IOEZH0043I 224
IOEZ00960E 195	IOEZH0044E 224
IOEZ00961E 196	IOEZH0045E 225
IOEZ00962E 196	IOEZH0062I 226
IOEZ00963E 197	IOEZH0063I 226
IOEZ00964E 197	IOEZH0064I 227
IOEZ00965E 197	IOEZH0065E 227
IOEZ00966E 198	IOEZH0066E 228
IOEZ00981E 198	IOEZH0067E 229
IOEZ00982E 199	IOEZH0068I 229
IOEZ00983I 199	IOEZH0069I 230

IOEZH0070I [230](#)
IOEZH0071E [231](#)
IOEZH0072E [232](#)
IOEZH0073E [232](#)
IOEZH0074E [233](#)
IOEZH0075I [234](#)
IOEZH0076I [234](#)
IOEZH0077I [235](#)
IOEZH0078I [235](#)
IOEZH0079I [236](#)
IOEZH0080I [236](#)
IOEZH0083I [237](#)
IOEZH0084I [237](#)
IOEZH0088I [238](#)

z/OS File System Messages and Codes (*continued*)
 messages, new [xi](#), [xiii](#)
 messages, no longer issued [xi](#), [xiii](#)
zFS messages [5](#)

K

keyboard
 navigation [359](#)
 PF keys [359](#)
 shortcut keys [359](#)

M

messages
 z/OS File System Messages and Codes [xi](#),
 [xiii](#)

N

navigation
 keyboard [359](#)

R

reason codes [241](#)

S

severity levels [2](#)
shortcut keys [359](#)
slip trap
 diagnosis data
 setting [2](#)
 setting [2](#)

T

trademarks [365](#)

U

user interface
 ISPF [359](#)
 TSO/E [359](#)

Z

z/OS File System Messages and
 Codes
 messages [xi](#), [xiii](#)
 messages, changed [xi](#), [xiii](#)



Product Number: 5655-ZOS

SC23-6885-70

