

| | |
|---|------------|
| OA47904: ++HOLD for DOC DFSMSHsm Common Recover Queue and Message Simplification | |
| Original Creation Date: | 01/14/2016 |
| Last Modified Date: | 09/09/2016 |
| Version: | Version 1 |

Table of Contents

| | |
|---|--------------------------|
| 1 Introduction..... | 3 |
| 1.1 Overview..... | 3 |
| 2 Publication Updates..... | 5 |
| 2.1 DFSMSHsm Storage Administration Guide..... | 5 |
| 2.1.1 Availability management of SMS-managed storage..... | 5 |
| 2.1.1.1 Collecting fast replication messages in message data sets..... | 5 |
| 2.1.1.2 Specifying the maximum number of volume recovery from dump tasks..... | 5 |
| 2.1.2 Other availability management considerations..... | 6 |
| 2.1.2.1 Common Recover Queue..... | 6 |
| 2.1.3 Operating procedures for DFSMSHsm..... | 16 |
| 2.1.3.1 Altering the priority of DFSMSHsm queued requests..... | 16 |
| 2.2 DFSMSHsm Storage Administration Reference..... | 16 |
| 2.2.1 ALTERPRI command: Altering the priority of queued requests..... | 16 |
| 2.2.2 CANCEL command: Canceling a queued or active DFSMSHsm request..... | 16 |
| 2.2.3 FRRECOV command: Requesting a fast replication recovery..... | 17 |
| 2.2.3.1 Syntax of the FRRECOV command..... | 17 |
| 2.2.3.2 Optional parameters of the FRRECOV command..... | 17 |
| 2.2.4 HOLD command: Preventing DFSMSHsm functions from running..... | 18 |
| 2.2.4.1 COMMONQUEUE: Preventing common queue functions..... | 18 |
| 2.2.4.2 FRRECOV: Preventing fast replication recovery processing..... | 18 |
| 2.2.4.3 RECOVER: Preventing recovery and restore..... | 19 |
| 2.2.5 QUERY command: Displaying the status of parameters, statistics, and requests..... | 19 |
| 2.2.5.1 Syntax of the QUERY command..... | 19 |
| 2.2.5.2 COMMONQUEUE: Displaying common queue information..... | 20 |
| 2.2.6 RELEASE command: Releasing all or some DFSMSHsm functions for processing..... | 20 |
| 2.2.6.1 DUMP: Releasing dump..... | 20 |
| 2.2.6.2 FRRECOV: Releasing the fast replication recover function..... | 20 |
| 2.2.6.3 RECOVER: Releasing recovery and restore..... | 21 |
| 2.2.7 SETSYS command: Establishing or changing the values of DFSMSHsm control parameters..... | 21 |
| 2.2.7.1 Syntax of the SETSYS command..... | 21 |
| 2.2.7.2 SETSYS fast-path for DFSMSHsm functions..... | 21 |
| 2.2.7.3 Optional parameters of the SETSYS command..... | 22 |
| 2.2.7.4 MAXDUMPRECOVERTASKS: Specifying the maximum number of concurrent volume recovery from dump tasks..... | 24 |
| 2.2.7.5 Examples of how to code the SETSYS command..... | 24 |
| 2.2.8 Using the QUERY command..... | 25 |
| 2.3 z/OS DFSMS Using the new functions..... | 25 |
| 2.3.1 Using the DFSMSHsm Enhancements in z/OS V2R2..... | 25 |
| 2.3.2 Administrating..... | 26 |
| 2.4 z/OS DFSMSHsm implementation and customization guide..... | 28 |

| | | |
|------------------------|--|--------------------|
| 2.4.1 | Specifying commands that define your environment..... | 28 |
| 2.4.2 | Customizing in a sysplex environment..... | 28 |
| 2.4.3 | Patch for FRRECOV COPYPOOL FROMDUMP performance to bypass the EXCLUSIVE NONSPEC ENQ..... | 30 |
| 2.5 | z/OS MVS System Messages, Vol 2 (ARC-ASA)..... | 30 |
| 2.5.1 | ARC1500I..... | 30 |
| 2.5.2 | ARC1542I..... | 31 |
| 2.5.3 | ARC1560I..... | 31 |
| 2.5.4 | ARC1562I..... | 33 |
| 2.5.5 | ARC1563I..... | 33 |
| 2.5.6 | ARC1564E..... | 33 |
| 2.5.7 | ARC1565E..... | 34 |
| 2.5.8 | ARC1566I..... | 34 |
| 2.5.9 | ARC1567I..... | 34 |
| 2.5.10 | ARC1568I..... | 34 |
| 2.5.11 | ARC1569I..... | 35 |
| 2.5.12 | ARC1127I..... | 35 |
| 2.5.13 | ARC1132I..... | 36 |
| 2.5.14 | ARC1829I..... | 36 |
| 2.5.15 | ARC1870I..... | 37 |
| 2.5.16 | ARC1873I..... | 39 |
| 2.6 | z/OS Summary of Messages and Interface Changes..... | 39 |
| 2.6.1 | DFSMSHsm summary of interface changes..... | 39 |
| 2.7 | z/OS DFSMSHsm Data Areas..... | 40 |
| 2.7.1 | MWE - Management Work Element..... | 40 |
| 2.7.2 | MWE Control Block Cross-Reference..... | 40 |
| 2.7.3 | FRGCB—Fast Replication Global Control Block..... | 41 |
| 2.7.4 | FRGCB Data Area Cross-Reference..... | 41 |
| 2.7.5 | JGCB – HSMplex Global Control Block..... | 42 |
| 2.7.6 | JGCB Data Area Cross-Reference..... | 43 |
| 2.7.7 | XCFP – XCF Implementation macro..... | 43 |
| 2.7.8 | XCFP Data Area Cross-Reference..... | 44 |

1 Introduction

1.1 Overview

The document describes the changes provided with APAR OA47904 that affected the IBM publications. It also includes the changes for SYSROUTE'd APARs OA47931 and OA47933.

Common Recover Queue (CVQ)

This is follow on support of the Common Dump Queue function provided in z/OS V2R2.

Today, a recover of a single copy pool must be processed on the same DFSMSHsm host that initiated the request. In a copy pool environment, a dump of a single copy pool can consist of 1000s of volumes. In order to have an expectation that this number of volumes can be processed within a reasonable amount of time, the workload needs to be distributed across multiple DFSMSHsm hosts. Distributing the workload across multiple DFSMSHsm hosts increases the number of available tasks and improves throughput by distributing the workload instead of concentrating it on a single DFSMSHsm address space. This enhancement will exploit the architecture provided with the Common Dump Queue (CDQ) function in z/OS V2R2. This support will be referred to as the Common Recover Queue (CVQ).

CVQ Scope

| Command/Function | Qualifications | Supported in CVQ |
|---|--|------------------|
| FRRECOV COPYPOOL | None | YES |
| | From dump determined by HSM for latest version | |
| | FROMDUMP specified | YES |
| | FROMDASD specified or latest version from DASD determined by HSM | N/A |
| FRRECOV TOVOLUME (Individual volume restore) | FROMCOPYPOOL | YES |
| | From dump determined by HSM for latest version | |
| | FROMCOPYPOOL | YES |
| | FROMDUMP specified | |
| | FROMCOPYPOOL FROMDUMP | YES |
| | APPLYINCREMENTAL specified | |
| | FROMCOPYPOOL | N/A |
| | FROMDASD specified or latest version from DASD determined by HSM | |
| FRRECOV DSNAME | none | NO |
| RECOVER DSNAME | FROMDUMP | NO |
| RECOVER DSNAME | From dump not specified | NO |
| RECOVER * | TOVOLUME FROMDUMP | YES |
| | TOVOLUME FROMDUMP APPLYINCREMENTAL | YES |

To summarize the CVQ scope table, FRRECOV and RECOVER volume restores and its optional APPLYINCREMENTAL are eligible to be processed in the CVQ. The DSNAME function of either command is not eligible.

Message Simplification support for Recover in a Common Recover Queue, plus TOVOLUME support

This is follow on support of the message simplification support for FRBACKUP provided in z/OS V2R2.

Today, error analysis of DFSMSHsm Fast Replication errors is difficult. For the most part, the fast replication function is invoked by DB2 Administrators. When errors occur in DFSMSHsm, they get a very simple RC=xx from DB2 to indicate that there was an error in DFSMSHsm processing. They must then explore the various Backup, Dump and Recover activity logs to try to find the error associated with their particular copy pool.

This enhancement will simplify the analysis process by dynamically allocating a fast replication message data set related to the processing of a particular copy pool (FRRECOV COPYPOOL FROMDUMP) invocation in a CVQ environment or an individual volume restore from a copy pool (FRRECOV TOVOLUME FROMDUMP) invocation in either environment. The FRRECOV COPYPOOL message simplification support in a non-CVQ environment was previously provided. No DSNAMES support is provided.

All DFSMSHsm and DFSMSdss messages related to the processing of a particular copy pool or individual volume restore from a copy pool will be written to that copy pool specific data set. Messages from other components will not be captured. Messages related to the QUERY command will not be captured either. The data set will be named in such a way that it can be uniquely identified. This will provide a single data set that will contain all of the messages required to analyze a problem. Customers will be responsible for the management of these data sets, including retention, expiration.

Changes to Common Dump Queue (CDQ) with the Common Recover Queue (CVQ) support

Any QUERY, CANCEL or ALTERPRI command issued from a non-MS host that is disconnecting from the group will have the command sent to the Master Scheduler (MS) host and responses received on the host issuing the QUERY, CANCEL or ALTERPRI command. With CDQ, these commands issued from a disconnecting host were not sent to the MS.

Patches

In OA49227, HSM was modified to increase throughput during recovery of parallel FRRECOV COPYPOOL FROMDUMP requests. Originally, the support in that APAR was made available only after the user enabled the function with a patch to the MCVT. Given Common Recovery Queue's primary goal is to increase throughput of parallel volume restores, the enhancement in OA49227 will be enabled for all CVQ-submitted FRR CP FROMDUMP requests without the need for users to set the enabling patch.

2 Publication Updates

2.1 DFSMSHsm Storage Administration Guide

Updates needed by section in the DFSMSHsm Storage Administration Guide and Reference sections: SC23-6871-04

2.1.1 Availability management of SMS-managed storage

2.1.1.1 Collecting fast replication messages in message data sets

Added the data set name that gets allocated for FRRECOV TOVOLUME requests, and added the *volser* data set name description after the 'ppppppp' description.

Message Data Set Allocation:

Message data sets will have the following naming convention:

- For FRRECOV TOVOLUME requests: *msgdsprefix.FR.Vvolser.Dyyddd.Thhmmss.z*

where:

volser

[is the volser of the volume that was recovered.](#)

2.1.1.2 Specifying the maximum number of volume recovery from dump tasks

The following paragraphs added at the end of this section before the related reading list.

When the SETSYS COMMONQUEUE(RECOVER(CONNECT(*basename*))) is specified, the preference by the system will be to fully utilize the available volume restore tasks as defined by MAXDUMPRECOVERTASKS on each of the Common Recover Queue HSMplex members. When this SETSYS is not specified only local tasks where the command was issued will be utilized. This reduced parallelism might lengthen the total time needed for the volume recovery to complete. In either case, the number of concurrent tasks is further limited by the stacking of the volumes on the dump tapes. For example, if 100 volumes are being recovered are all on a single tape, then only a single recover task will be used.

If you want a common recover host to be able to submit volume restore work, but not process the work, issue the SETSYS MAXDUMPRECOVERTASKS(0) command on that host. This can be either a Master Scheduler host or a non-Master Scheduler host. In the case of a Master Scheduler host, this does not impact the Master Scheduler duties for the Common Recover Queue.

2.1.2 Other availability management considerations

Added new Common Recover Queue section after the Common Dump Queue section:

All of the following has been added.

2.1.2.1 Common Recover Queue

The Common Recover Queue (CVQ) is a queue of DFSMSHsm volume restore requests from FRRECOV COPYPOOL, FRRECOV TOVOLUME, and RECOVER TOVOLUME with FROMDUMP commands that is shared by multiple DFSMSHsm hosts. See Table 1, CVQ Scope, for details on which command variations are supported. This queue is implemented through the use of the Cross-system Coupling Facility (XCF) to balance the volume restore workload across the resources available in each DFSMSHsm host in an HSMplex. Individual DFSMSHsm hosts are defined to XCF as a member of a group to be used for volume restore processing. By distributing volume restore work to all the members in the group, the number of available tasks to perform the work is increased which also improves throughput of the volume restore function as the workload is no longer concentrated on a single host's address space.

Table 1: CVQ Scope

| Command/Function | Qualifications | Supported in CVQ |
|---|--|------------------|
| FRRECOV COPYPOOL | None From dump determined by HSM for latest version | YES |
| | FROMDUMP specified | YES |
| | FROMDASD specified or latest version from DASD determined by HSM | N/A |
| FRRECOV TOVOLUME (Individual volume restore) | FROMCOPYPOOL From dump determined by HSM for latest version | YES |
| | FROMCOPYPOOL FROMDUMP specified | YES |
| | FROMCOPYPOOL FROMDUMP APPLYINCREMENTAL specified | YES |
| | FROMCOPYPOOL FROMDASD specified or latest version from DASD determined by HSM | N/A |
| FRRECOV DSNAME | none | NO |
| RECOVER DSNAME | FROMDUMP | NO |
| RECOVER DSNAME | From dump not specified | NO |
| RECOVER * | TOVOLUME FROMDUMP | YES |
| | TOVOLUME FROMDUMP APPLYINCREMENTAL | YES |
| To summarize the CVQ scope table, FRRECOV and RECOVER volume restores and its optional APPLYINCREMENTAL are eligible to be processed in the CVQ. The DSNAME function of either command is not eligible. | | |

Figure 1 shows that the CVQ consists of one or more DFSMSHsm hosts that are connected to an XCF defined group, where submitted volume restore requests will be distributed and processed among the volume restore

resources available in the CVQ group, with results returned back to the host where the request originated.





Figure 2 shows that the CVQ group allows for flexible configurations. This provides the capability to define multiple queues in the same HSMplex, allow group members to both receive and process requests, just process requests and receive requests but not process them.



In CVQ environments

- A CVQ can not span an HSMplex
- Priority of requests is supported in a CVQ
- A host can only be connected to one CVQ
- Progress and status messages are recorded on the host processing the request
- Can coexist with a CDQ environment with the same or subset of hosts

Results are returned to the submitting host but progress and status messages are recorded on the host processing the request. By using the message simplification support along with CVQ processing, each CVQ group member will write all DFSMSshsm and DFSMSdss messages related to the processing of a particular copy pool (FRRECOV COPYPOOL FROMDUMP) or individual volume restore from a copy pool (FRRECOV TOVOLUME FROMDUMP) to a single shared message data set. This simplifies diagnostics as all messages are collected from each host into a single data set. The data set is managed by the Master Scheduler (MS) host for the CVQ. See the SETSYS FASTREPLICATION(MESSAGE DATASET(YES)) option.

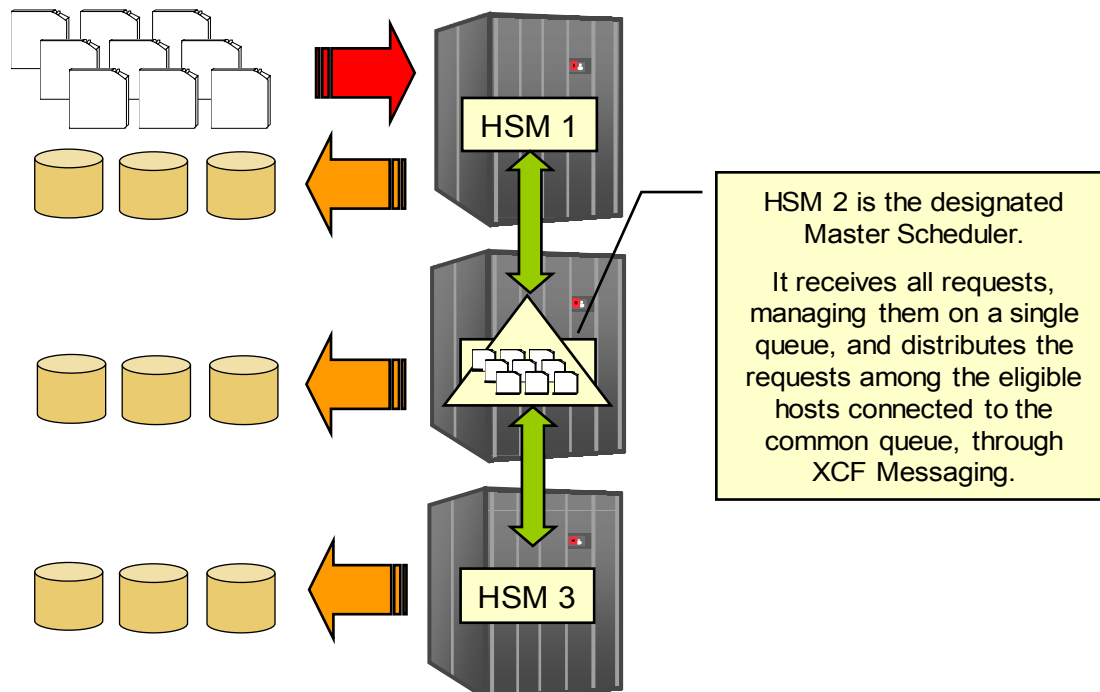
The Master Scheduler (MS) is the DFSMSshsm host that manages all of the volume restore requests in the Common Recover Queue (CVQ). All volume restore requests are sent to it and it assigns work to all of the other hosts eligible in the group to process requests from the CVQ. Each group needs a Master Scheduler to manage CVQ processing for the group, otherwise volume restore requests will be processed locally. The other main types of DFSMSshsm hosts include the Submitting Host and the Processing Host.

Types of Hosts used in a Common Recover Queue (subsection):

- Submitting Host
 - Host that receives requests from commands and using XCF messaging sends it to the Master Scheduler (MS) host
 - Performs command complete processing and posts the user complete
- Master Scheduler (MS)
 - Accepts requests from submitting host members and locally submitted requests and maintains them as a single common queue of all requests for the CVQ function
 - The MS can be the submitting host. In this case the command is not sent
 - Assigns/distributes the requests among the available tasks across all of the eligible hosts (processing hosts) connected to the CVQ while balancing the requests among all the hosts.
 - Sends command complete message back to submitting host
- Processing Host
 - Receives assigned work requests from the MS, places it on its local queue to process it
 - Sends work complete messages back to MS with specific protocols to support the optimization of currently mounted tapes
 - The processing host could be the submitting host and/or the MS host

Figure 3 shows a CVQ group where:

- HSM1 is the submitting and processing host
- HSM2 is the MS and processing host
- HSM3 is a processing host



Connecting to a Common Recover Queue (subsection)

The SETSYS COMMONQUEUE(RECOVER(CONNECT(*basename*))) is used to define the CVQ group name to XCF and to connect this DFSMSHsm host to it. The *basename* is a four character suffix to the HSM designated group name prefix. It has the following format: ARCV*basename*. The first DFSMSHsm host issuing this connect option will define the XCF group and connect itself to this group, while subsequent DFSMSHsm hosts will connect to the XCF group defined by the first host.

See the syntax of this command for specifics on other options and defaults.

For example, a SETSYS CQUEUE(RECOVER(CONNECT(QUE1))) defines an DFSMSHsm host as a member of the XCF group named , ARCVQUE1. The Master Scheduler Candidate (MSC) status was defaulted to YES. The first member to join the CVQ group that is a Master Scheduler candidate will become the Master Scheduler host for the CVQ group. The current SETSYS setting and the MS status can be displayed with QUERY SETSYS (message ARC1500I) and QUERY ACTIVE (message ARC1560I) respectively.

The MSC setting can be defined in conjunction with the CONNECT option (SETSYS COMMONQUEUE(RECOVER(CONNECT(*basename*) MSC(YES | NO))) or separately to preset the MSC state before the CONNECT option or change it after (SETSYS COMMONQUEUE(RECOVER(MSC(Y))).

For example, an MSC(N) option on each host prior to the CONNECT will ensure that no host becomes a master schedule candidate. In this case the CVQ group status will be DISABLED and no CVQ processing will occur. All work will be processed locally. After all the CONNECT options have been issued, one or more hosts MSC state could be changed with the MSC(Y) option. This will allow one host to take responsibilities for Master Scheduler processing. A DFSMSHsm host with a MSC(Y) state will be eligible to take over when the current Master Scheduler host fails or disconnects from the group.

It is recommended that all your DFSMSHsm hosts are defined to be Master Scheduler candidates. This will ensure one host will take on the Master Scheduler responsibilities when a CVQ group is defined and ensure another host will take on the Master Scheduler responsibilities if the current one is removed from the group due to a disconnection or system error. Hosts that may not have the processing capacity to perform this work may be defined to not be a Master Scheduler candidate.

To display the current SETSYS setting, use the QUERY SETSYS command and look for message ARC1500I. To display the master scheduler status, use the QUERY ACTIVE command and look for message ARC1560I.

Using the DISPLAY XCF,GROUP command, the output for this command will identify the group as ARCVQUE1 and identify the member name as HOST#hsmid for this host.

Placing requests on the Common Recover Queue (subsection)

The Common Recover Queue (CVQ) is a queue of DFSMSHsm volume restore requests. These include all requests from RECOVER *, FRRECOV COPYPOOL, FRRECOV TOVOLUME when FROMDUMP is specified or determined by DFSMSHsm that the most current version is from a dump (ie volume restores).

APPLYINCREMENTAL can be applied. However, data set recovers are not eligible for the CVQ. A host that is participating in CVQ processing will place its volume restore requests onto the CVQ for processing. The following items apply to placement status:

- To verify that a host is connected to the CVQ, issue the QUERY ACTIVE command, then look for CONNECTION STATUS of CONNECTED in message ARC1560I for the Common Recover Queue
- A HOLD RECOVER or HOLD FRRECOV(TAPE) issued on a host will prevent it tasks for either function from being selected by the Master Scheduler as a processing host.
- A HOLD RECOVER or HOLD FRRECOV(TAPE) issued on the host that issued the volume restore request will place the request on the CVQ if the Master Scheduler is not disabled (Master Scheduler status from QUERY ACTIVE not DISABLED(HELD)). Both WAIT and NOWAIT type requests will be placed as the Master Scheduler is available to assign and process requests in the CVQ on other hosts. This host tasks will subsequently not be assigned by the MS for processing.
- HOLD RECOVER and FRRECOV(TAPE) issued on the Master Scheduler host prevents it from assigning and processing requests. Note that NOWAIT request types will be placed on the CVQ even though the function was held on the Master Scheduler, while WAIT type requests will be failed. A HOLD of a specific type (RECOVER or FRRECOV) will affect only the assignment of that type by the Master Scheduler. To keep the Master Scheduler's responsibilities intact, don't issue the HOLD but rather issue an SETSYS MAXDUMPRECOVERTASKS(0) if you do not want tasks on the MS host from being utilized by the CVQ. Another option is to transition the MS responsibilities to another host with SETSYS CQUEUE(RECOVER(MSC(N))) prior to issuing the HOLD.
- A HOLD RECOVER and FRRECOV(TAPE), issued on all the hosts in the CVQ will place the Master Scheduler in the DISABLED state. Use the messages issued in response to the QUERY ACTIVE to verify this state. Note that the MASTER SCHEDULER STATUS is DISABLED(HELD). When DISABLED:
 - New WAIT request types will be failed locally
 - New NOWAIT request types will be queued locally
 - Queued WAIT request types in the CVQ Master Scheduler host will be failed
- A SETSYS MAXDUMPRECOVERTASKS(0) issued on the host that issued the volume restore request will place the request on the CVQ, but it tasks will subsequently not be assigned by the MS for processing.
- Any request that is not submitted to the Master Scheduler host, like due to errors, is placed on the local queue of the host that issued the request.
- Requests submitted on the Master Scheduler host are placed on its local queue
- Requests that are sent to the MS host from other members, are placed on the local MS queue and then selected to run locally on the MS or assigned to be processed on another member in the CVQ group.
- The default priority of WAIT volume restore requests over NOWAIT volume restore requests does not change with the CVQ. However, the ALTERPRI command does support change requests currently queued to the CVQ.

Selecting requests from the Common Recover Queue (subsection)

The Master Scheduler is the DFSMSHsm host that manages all of the volume restore requests in the CVQ. All volume restore requests are sent to it and it assigns work to all of the other hosts eligible to process requests from the Common Recover Queue. The following items apply to selecting requests to be processed by the Common Recover Queue:

- The Master Scheduler (MS) selects hosts to process requests based on the current utilization of available tape resources in each host relative to each other. The first host with the lowest utilization will be selected to process the request.
- A host's resources are not available if the HOLD RECOVER, FRRECOV or SETSYS EMERGENCY commands has been issued or if the SETSYS MAXDUMPRECOVERTASKS(0) has been issued. In each case the system will indicate back to the MS with XCF messaging that this host has no resources available.
- Selected hosts that are not the Master Scheduler are assigned the work with XCF messaging. The assigned processing host selects the task to use.
- When the Master Scheduler selects its own volume restore task to process a CVQ request, the task is attached and assigned the work.

Processing requests off of the Common Recover Queue (subsection)

The Master Scheduler selects a host member of the CVQ to process a request placed on the CVQ. This host could be the Master Scheduler itself, another member in the CVQ and even the member that placed the request on the CVQ that the request originated from. This selected member is the processing host and is responsible for selecting an available task to use for the copy to tape processing that it will perform. The following items apply to a processing host in the Common Recover Queue:

- The assigned processing host selects an available task to process the request.
- Status and progress messages along with diagnostic trace data that is generated are recorded on the processing host. A single copy pool recover request may generate many different volume restores with each one processed by individual CVQ group members. Messages and diagnostic data are placed on the host processing the request, making it difficult to put together the message flow of an individual volume restore from all the processing hosts.
- By using the message simplification support along with CVQ processing, each CVQ group member will write all DFSMSshm and DFSMSdss messages related to the processing of a particular copy pool to a single shared message data set that is managed by the MS host. This simplifies the understanding of the message flow for the original fast replication copy pool command as all messages are collected from each host into this single data set. XCF messages are used to perform this function. See the SETSYS FASTREPLICATION(MESSAGE DATASET(YES)) option.
- As volumes complete processing, the selected processing host working in conjunction with the MS host determines if another volume restore is needed by the currently mounted tape. If this is the case, that task is assigned to process the additional volume restore. This is accomplished with XCF messaging.
- The processing host could be the submitting host and/or the MS host or one that is neither the submitting host nor the MS host.
- When the SETSYS COMMONQUEUE(RECOVER(CONNECT(*basename*))) is specified, the preference by the system will be to utilize the available recover tasks as defined by MAXDUMPRECOVERTASKS on all the recover group members in the HSMplex. However, the number of concurrent tasks is limited by the stacking of the volumes on the dump tapes. For example, if 100 volumes are being recovered are all on a single tape, then only a single recover task will be used.

Disconnecting from the Common Recover Queue (subsection)

SETSYS COMMONQUEUE(RECOVER(DISCONNECT)) is used to request that a member of the CVQ group that this host is currently connected too is to leave the group. See the syntax of this command for specifics on other options and defaults. The following items apply to this disconnect process:

- All CVQ requests being processed on this host and requests that originated from this host and selected to be processed in the CVQ are waited upon for completion before disconnecting the host from the CVQ. This host will no longer be selected to process new volume restore work from the CVQ.
- A HOLD RECOVER or FRRECOV issued prior to the disconnect command will prevent queued requests from being selected to be processed. After the disconnect has completed, a RELEASE RECOVER or FRRECOV will allow the queued requests to be processed locally.
- A disconnection of the host that is currently the Master Scheduler will wait for all CVQ active and queued work to complete prior to disconnecting. Once disconnected another member in the CVQ group that has been previously defined as a Master Scheduler candidate will automatically take over the responsibilities for it. A HOLD RECOVER or FRRECOV prior to a disconnection will prevent queued requests from being selected for processing. A subsequent RELEASE RECOVER or FRRECOV after the disconnecting will allow these requests to be processed locally.
- The QUERY ACTIVE command response in message ARC1560I will verify the disconnection.
- Volume restore resources available on this host after the disconnection has completed are not known or available to the remaining CVQ members. All volume restore requests received will be placed on this host's local queue and processed locally.
- For shutdown procedures, a DISCONNECT should first be issued on each host being removed from the HSMplex to allow CVQ requests to complete first.
- The DISCONNECT FORCE option should be reserved for emergency situations. With FORCE specified, requests that originated from this host and selected to be processed in the CVQ are not waited upon for completion before disconnecting. Messages will be issued to indicate that a request has been failed but still may complete in the CVQ.

Using QUERY, CANCEL and ALTERPRI in the Common Recover Queue (subsection)

The QUERY, CANCEL and ALTERPRI commands support the Common Recover Queue (CVQ).

The following items apply to QUERY processing:

- QUERY SETSYS displays the current SETSYS settings for the CVQ on the host that issued the QUERY command. Message ARC1500I displays the SETSYS settings on this host.
- QUERY ACTIVE displays the current state of the CVQ and active requests running it.
 - The state and active requests reported on are those that originated from the host issuing the QUERY command
 - Active requests reported could be active in any member in the CVQ. XCF messaging is used to retrieve active requests by sending the QUERY command to the Master Scheduler. These requests could be running in any member of the CVQ group and must match the filter criteria of the query. In this case, the active requests must have originated from the host issuing the QUERY command. The generated messages on the MS are returned to the query host using XCF and in turn the query host issues them. The query host then goes on to report its own active requests.
 - When the query host is also the MS, no XCF messaging is used to obtain the active requests as the CVQ queues are located on the MS host.
- QUERY REQUEST or USER display requests that originated from the host issuing the QUERY command

that match the REQUEST or USER filter criteria.

- Active requests reported could be active in any member in the CVQ and queued in the MS. XCF messaging is used to retrieve these requests by sending the QUERY command to the Master Scheduler. These active requests could be running in any member of the CVQ group or queued in the MS and must match the filter criteria of the query. In this case, the requests must have originated from the host issuing the QUERY command and match the REQUEST or USER criteria on the QUERY command. The generated messages on the MS are returned to the query host using XCF messaging and in turn the query host reports them. The query host then goes on to report its own requests.
- When the query host is also the MS, no XCF messaging is used to obtain the queued and active requests as the CVQ queues are located on the MS host.
- QUERY WAITING displays the number of requests that are queued in the CVQ that originated from the host issuing the QUERY command.
 - Queued requests in the CVQ will be queued in the MS. XCF messaging is used to retrieve these queued request numbers by sending the QUERY command to the Master Scheduler. These queued requests in the MS must match the filter criteria of the query. In this case, the requests must have originated from the host issuing the QUERY command. The generated messages on the MS are returned to the query host using XCF messaging and in turn the query host reports them. The query host then goes on to determine its own queued request numbers, combines them with what was returned from the CVQ, then reports the totals.
 - When the query host is also the MS, no XCF messaging is used to obtain the queued request numbers as the CVQ queues are located on the MS host.
- QUERY COMMONQUEUE(RECOVER) displays all queued and active requests being managed by the CVQ. Use this query to get a snapshot of all CVQ activity from any host connected to the CVQ.
 - Queued and active requests in the CVQ are managed by the MS for the CVQ group. XCF messaging is used to retrieve these requests by sending the QUERY command to the Master Scheduler. These queued and active requests in the MS have no matching criteria to satisfy. All requests will be reported. The generated messages on the MS are returned to the query host using XCF messaging and in turn the query host reports them.
 - When the query host is also the MS, no XCF messaging is used to obtain the queued and active requests as the CVQ queues are located on the MS host. The filter criteria is all queued requests running in the CVQ regardless of where the volume restore request command originated.
 - Messages ARC1562I and ARC1563I are reported for the active and queued requests in the CVQ.
- The CANCEL REQUEST or USERID command cancels queued requests in the CVQ for requests originated in the host that issued the CANCEL command that also matches the REQUEST or USERID filter criteria.
 - Queued requests in the CVQ are managed by the MS for the CVQ group. XCF messaging is used to cancel these requests by sending the cancel command to the Master Scheduler. These queued requests in the MS must have originated from the host issuing the CANCEL command and match the filter criteria of the REQUEST and USERID parameters. When matching, the requests are cancelled. The generated messages on the MS are returned to the cancel host using XCF messaging and in turn the cancel host reports them.

- The QUERY COMMONQUEUE(RECOVER) command can be used to determine the requests running in the CVQ which will also identify the HOST id of where each request originated from.
 - When the cancel host is also the MS, no XCF messaging is used to cancel the queued requests as the CVQ queues are located on the MS host.
 - To cancel an active request in the CVQ, issue the CANCEL TCBADDRESS on the host that this request is active in.
- The ALTERPRI REQUEST or USERID command changes the priority of queued requests in the CVQ for requests originated in the host that issued the ALTERPRI command that also matches the REQUEST or USERID filter criteria.
 - Queued requests in the CVQ are managed by the MS for the CVQ group. XCF messaging is used to alter the priorities of these requests by sending the ALTERPRI command to the Master Scheduler. These queued requests in the MS must have originated from the host issuing the ALTERPRI command and match the filter criteria of the REQUEST and USERID parameters. When matching, the requests priority is altered. The generated messages on the MS are returned to the ALTERPRI host using XCF messaging and in turn the ALTERPRI host reports them.
 - The QUERY COMMONQUEUE(RECOVER) command can be used to determine the requests running in the CVQ which will also identify the HOST id of where each request originated from.
 - When the ALTERPRI host is also the MS, no XCF messaging is used to alter the priority of the queued requests as the CVQ queues are located on the MS host.

Processing of failures in the Common Recover Queue (subsection)

The Common Recover Queue (CVQ) group support uses XCF facilities to detect when members lose their connection to the CVQ. This could be due to some of the following reasons:

- When a host requested a disconnection with the SETSYS COMMONQUEUE(RECOVER(DISCONNECT))
- When an XCF related failure causes it to be removed from the group
- When an z/OS® image fails

In all these cases the host is considered as no longer defined to the group.

When the host that is now not defined was the Master Scheduler (MS), the remaining hosts in the CVQ group that are Master Scheduler candidates will start a race to take over the Master Scheduler responsibilities. All hosts would have been notified that the MS was the one that went not defined. For each of these, any work that was placed on the CVQ prior to this event will be placed on the CVQ, managed by the new MS host. This may result in some volumes being processed again.

When the host that is now not defined is not the Master Scheduler (MS), the MS host will be notified that this host is no longer defined to the group. The CVQ in this case will search for requests previously selected to run on this now not defined host and fail them with an ARC1127I or ARC1829I with a return code of 105. This indicates that a host in the common queue group became not defined with the return code indicating that the work was selected by the Master Scheduler to run in this host of the Common Recover Queue but the host cannot determine if the work completed.

2.1.3 Operating procedures for DFSMShsm

2.1.3.1 Altering the priority of DFSMShsm queued requests

Add note:

For queued requests in the Common Dump Queue or the Common Recover Queue the priority of one or more requests can be changed for requests that originated in the host that issued the ALTERPRI command that also matches the REQUEST or USERID filter criteria.

2.2 DFSMShsm Storage Administration Reference

Updates needed by section in the DFSMShsm Storage Administration Guide and Reference sections: SC23-6871-04

2.2.1 ALTERPRI command: Altering the priority of queued requests

Replaced the 4th note with the following:

Changing the priority of queued requests running in the Common Dump Queue or the Common Recover Queue will be completed for queued requests that originated on the host issuing the ALTERPRI command for a request number or userid that matches the parameters on the ALTERPRI command. A single ALTERPRI by USERID will effect requests in both common queues.

The QUERY COMMONQUEUE(RECOVER) command or the QUERY COMMONQUEUE(DUMP) command can be used to determine the requests running in their respective common queues. This will identify the HOST id of the originating host for the request and can be used to determine the host from which a corresponding ALTERPRI should be issued.

For a host that is connected to a Common Dump or Recover Queue and is not the Master Scheduler, the local queue will be searched and the Master Scheduler member's group queue will be searched. For a host that is the Master Scheduler, only the local queue will be searched.

2.2.2 CANCEL command: Canceling a queued or active DFSMShsm request

Replaced the 4th note with the following:

Canceling requests running in the Common Dump Queue or the Common Recover Queue will be completed for queued requests that originated on the host issuing the CANCEL command for a request number or userid that matches the parameters on the cancel command. Canceling active requests with the TCBADDRESS must be done on the host that is processing the request. A single CANCEL by USERID will effect requests in both common queues.

The QUERY COMMONQUEUE(RECOVER) command or the QUERY COMMONQUEUE(DUMP) command can be used to determine the requests running in their respective common queues. This will identify the HOST id of the originating host for the request and can be used to determine the host from which a corresponding CANCEL should be issued.

For a host that is connected to a Common Dump or Recover Queue and is not the Master Scheduler, the local queue will be searched and the Master Scheduler member's group queue will be searched. For a host that is the Master Scheduler, only the local queue will be searched.

2.2.3 **FRRECOV command: Requesting a fast replication recovery**

2.2.3.1 **Syntax of the FRRECOV command**

Syntax of the FRRECOV command (showing additions only)

Added the MESSAGEDATASET option on FRRECOV

FRRECOV Command with COPYPOOL and TOVOLUME:

```
> _____ >
  |__MESSAGEDATASET (DATE(yyddd) TIME(hhmmss))__|
  |__MSGDS_____|
```

2.2.3.2 **Optional parameters of the FRRECOV command**

Added optional parameters on the FRRECOV command

MESSAGEDATASET: Specifying the date and time used in the message data set name

Explanation:

MESSAGEDATASET is an optional parameter that allows you to specify the date and time used in the name of the fast replication command message data set. Specifying the date and time ensures that you can identify the message data set that corresponds with the FRRECOV COPYPOOL or TOVOLUME command. A 5 digit numeric value must be specified as the DATE parameter, and a 6 digit numeric parameter must be specified as the TIME parameter.

When DATE and TIME keywords are specified on the FRRECOV COPYPOOL or TOVOLUME command, and a duplicate message data set already exists with the specified date and time, DFSMSHsm will issue the ARC1873I message and continue without capturing the fast replication messages in a message data set.

When the MESSAGEDATESET keyword is not specified on the FRRECOV command, DFSMSHsm will obtain the date and time from the system and use that for the message data set generated name.

When DATE and TIME are specified on a FRRECOV DSNAME command, they are ignored.

Use of the message data set is controlled with the MESSAGEDATASET parameter of the SETSYS FASTREPLICATION command. Refer to [SETSYS command: Establishing or changing the values of DFSMSHsm control parameters](#) for more information.

Defaults:

None.

2.2.4 HOLD command: Preventing DFSMSHsm functions from running

2.2.4.1 COMMONQUEUE: Preventing common queue functions

The following note was added:

Note: HOLD COMMONQUEUE without any parameters will not have an affect on the common dump or recover queue functions

2.2.4.2 FRRECOV: Preventing fast replication recovery processing

The following was added:

The description below applies to preventing fast replication recovery and recover processing when running in a Common Recover Queue.

When the HOLD command (ALL, RECOVER, FRRECOV, or FRRECOV(TAPE)) is issued from a host that is connected to a Common Recover Queue (CVQ), the effect is different than in a non-CVQ environment. The actions taken by this host include:

HOLD of any type (ALL, RECOVER, FRRECOV) will prevent a host being selected by the Master Scheduler as a processing host, even though the request may be for another type that the HOLD was specified for.

HOLD of any type (ALL, RECOVER, FRRECOV) on a specific host will not prevent that host from submitting requests, including WAIT type requests, to the CVQ. This is because this host's resources are those of the group. It just won't be selected as a processing host (previous item).

HOLD by the Master Scheduler host for a specific type will affect the assignment of that type by the Master Scheduler host. A HOLD on both (RECOVER, FRRECOV) on the CVQ MS will prevent it from performing its assignment duties. For example, a HOLD RECOVER on the Master Scheduler host will prevent the assignment of only RECOVER requests but will still assign FRR CP requests volumes to processing hosts. The Query Active command message ARC1560I response will indicate this status.

- Requests submitted to the CVQ from this host that are active or queued are allowed to complete as long as work did not get assigned to the host issuing the HOLD
- A HOLD on a CVQ host may cause the work currently on that host for processing to be failed, based on whether EOD or EOVS is specified on the HOLD

Consider using SETSYS MAXDUMPRECOVERTASKS(0) from the Master Scheduler, instead of A HOLD, if you want the Master Scheduler to schedule CVQ work but not use its resources for CVQ volume restore processing.

When the HOLD command (ALL, RECOVER, FRRECOV, or FRRECOV(TAPE)) is issued from all hosts connected to a Common Recover Queue, which changes the Master Scheduler status to DISABLED, the actions taken by each host include:

- New WAIT request types from this host are failed locally
- New NOWAIT request types from this host are queued locally as the Master Scheduler status is DISABLED

- If EOD is specified on the HOLD, all requests are failed. If EOVS is specified, the current volume in a stack is completed and the remaining are failed.
- Queued WAIT request types in the CDQ Master Scheduler host are failed

HOLD ALL, RECOVER (any variety), and FRRECOV(TAPE)) all will have an impact in a Common Recover Queue (CVQ) for the selection of a host used for volume restore processing. Any one of these being held prevents a host from being selected for volume restore processing. These holds also impact the Master Scheduler's ability to perform scheduling of CVQ work for the specific HOLD type. Instead of the HOLD, consider using the SETSYS MAXDUMPRECOVERTASKS(0) command from a host connected to a CVQ. This prevents the selection of a host's resources for CVQ volume restore processing but allow volume restore requests to be submitted. It also allows the Master Scheduler to schedule CVQ work but not select its own resources for processing. If a HOLD was used, any takeover of the Master Scheduler responsibilities by another system are impacted by a previously issued HOLD for RECOVER or FRRECOV(TAPE). In this case the HOLD should only be used where a host is defined to the group with a Master Scheduler candidate status of NO, with SETSYS COMMONQUEUE(RECOVER(MSC(NO))) command.

In addition, when any variety of the HOLD command is issued from a host that is connected to a Common Recover Queue, its host resources are prevented from being selected to process a request even if the requests did not originate from the system that issued the HOLD command.

Notes:

The MASTERSCHEDULER status for the Common Recover Queue on this host is not changed by this command. It will continue to have a MASTERSCHEDULER status of ENABLED as long as one host in the group is not held. If all of the group hosts have been held, then no host is eligible to do the work and the MASTERSCHEDULER status will be changed to be DISABLED. In this state, the wait-type requests are failed and no-wait type requests are placed on the Common Recover Queue.

You can further refine the capabilities of a Common Recover Queue host, as follows:

- If you wish a Common Recover Queue host to be able to submit command work, but not process the work, issue the SETSYS MAXDUMPRECOVERTASKS(0) command on that host. This can be either a Master Scheduler host or a non-Master Scheduler host. In the case of a Master Scheduler host, this does not impact the Master Scheduler responsibilities

2.2.4.3 RECOVER: Preventing recovery and restore


The following was added:

When the HOLD command (ALL, RECOVER, FRRECOV, or FRRECOV(TAPE)) is issued from a host that is connected to a Common Recover Queue (CVQ), the effect is different than in a non-CVQ environment. See the preventing fast replication recover processing description for these differences.

2.2.5 QUERY command: Displaying the status of parameters, statistics, and requests

2.2.5.1 Syntax of the QUERY command

Added optional parameter, RECOVER, on the QUERY COMMONQUEUE command



2.2.5.2 **COMMONQUEUE: Displaying common queue information**

Added COMMONQUEUE(RECOVER) optional parameter description

COMMONQUEUE(RECOVER) is a sub-parameter that you can specify to display the status of recover requests that are being processed and queued in the Common Recover Queue that the host issuing the QUERY command is a member of. This query provides a snapshot of all fast replication recover and volume recovery activity across all Common Recover Queue group members.

2.2.6 **RELEASE command: Releasing all or some DFSMSHsm functions for processing**

2.2.6.1 **DUMP: Releasing dump**

'Any one of these DUMP options being held prevents a CDQ host from being selected to process any CDQ dump' was replaced with the following' replaced with the following:

The RELEASE DUMP command will allow dump processing.

If any variant of RELEASE for dump is issued on a host running a Common Dump Queue (CDQ) that host will be eligible to be selected for CDQ volume dump processing if the following conditions are met:

- No other HOLD for DUMP is in place for the host.
- SETSYS MAXDUMPTASKS(x) is specified with a value greater than 0

If the host is the CDQ MS, it will also be eligible to perform MS assignment responsibilities duties for the dump function.

2.2.6.2 **FRRECOV: Releasing the fast replication recover function**

The following as the last paragraph was added:

The RELEASE FRRECOV command will have the same effect if the system issuing the command is connected to a Common Recover Queue (CVQ). However, this host's volume recover tasks in the CVQ will only be used by the Master Scheduler(MS) host if the following conditions are met:

- No other HOLD RECOVER is in place for the host. SETSYS MAXDUMPRECOVERTASKS(x) is specified with a value greater than 0

If the host is the CVQ Master Scheduler(MS), it will also be eligible to perform MS assignment responsibilities for fast replication recovery requests.

2.2.7.3 Optional parameters of the SETSYS command

Added the following optional parameters to the SETSYS COMMONQUEUE command:

RECOVER(CONNECT):

Explanation: RECOVER(CONNECT(*basename*)) is an optional sub-parameter of the COMMONQUEUE parameter. This sub-parameter specifies the XCF functional group name suffix that this host uses for volume restores from dumps. This includes the FRRECOV COPYPOOL FROMDUMP, FRRECOV TOVOLUME FROMDUMP, and RECOVER * FROMDUMP volume restore commands. When specified this host is allowed to send volume restores to the Common Recover Queue for processing in the HSMplex that was started with the same *basename*. For *basename*, substitute four characters. The four characters can be numeric, upper-case alphabetic, and national characters. The XCF group name is *ARCVbasename*. For example, if you specify REC1 for *basename*, the resulting XCF group name that this host will become a member of is ARCVREC1.

This host becomes a candidate for the Common Recover Queue Master Scheduler when the MSC keyword is not specified or specified as MSC with the CONNECT command or MSC(YES) on another separate command. If this Master Scheduler candidate system is selected during HSM startup or during recovery to be the Master Scheduler of this group, this system is considered to be the Master Scheduler for the group.

Subsequent SETSYS commands after a CONNECT can be issued to change the MASTERSCHEDULERCANDIDATE state of this member of the recover group. A specification of MSC(NO) will remove this member as a Master Scheduler candidate and if the member was currently active as the Master Scheduler it will be stopped. The system will re-assign the Master Scheduler to another recover group member that was defined as a candidate, if one is available. New incoming work will be assigned to run in this new Master Scheduler member. Active and queued work will be allowed to complete. A specification of MSC(YES) will allow this member to become the Master Scheduler if needed by the recover group. This stand-alone MSC specification will have no effect when issued on a host that is not connected to the group.

If it is known at startup that a host is not a candidate for being the Master Scheduler but the host needs to be a member of the recover group, the MSC(NO) and CONNECT(*basename*) sub-parameters should be specified together in the RECOVER keyword.

Remote recover processing in the group will not be performed if all of the group members are not allowed to be a Master Scheduler with the MSC(NO) specification.

Use the HOLD RECOVER command to prevent additional recover processing tasks being started by this system. See HOLD RECOVER command for additional details.

When the SETSYS COMMONQUEUE(RECOVER(CONNECT(*basename*))) is specified, the preference by the system will be to utilize the available recover tasks as defined by MAXDUMPRECOVERTASKS on all the recover group members in the HSMplex. When this SETSYS is not specified only local recover tasks where the command was issued will be utilized. This reduced parallelism might lengthen the total time needed for the volume recovery to complete. In either case, the number of concurrent tasks is further limited by the stacking of the volumes on the dump tapes. For example, if 100 volumes are being recovered are all on a single tape, then only a single recover task will be used.

SMS relationship: Not applicable.

SETSYS default: None.

DFSMSHsm default: If the RECOVER(CONNECT) sub-parameter is not specified, no XCF recover groups will be defined. Local processing of volume restores will be performed.

Note: A system must disconnect itself from using XCF groups before it attempts to connect to a new XCF

group.

The host will respond with message ARC1566I when the host connects to the common recover group and with message ARC1567I if it took on the Master Scheduler responsibilities for it.

RECOVER(MASTERSCHEDULERCANDIDATE):

Explanation: RECOVER(MASTERSCHEDULERCANDIDATE) is an optional sub-parameter of the COMMONQUEUE parameter. This specifies whether the host is a Master Scheduler candidate for the common recover group.

Subsequent SETSYS commands after a CONNECT can be issued to change the MASTERSCHEDULERCANDIDATE state of this member of the recover group. A specification of MSC(NO) will remove this member as a Master Scheduler candidate and if the member was currently active as the Master Scheduler it will be stopped. The system will re-assign the Master Scheduler to another recover group member that was defined as a candidate, if one is available. New incoming work will be assigned to run in this new Master Scheduler member. Active and queued work will be allowed to complete. A specification of MSC(YES) will allow this member to become the Master Scheduler if needed by the recover group. This stand-alone MSC specification will have no effect until the host is connected to the group.

If it is known at startup that a system is not a candidate for being the Master Scheduler but the system needs to be a member of the recover group, the MSC(NO) and CONNECT(*basename*) sub-parameters should be specified together in the RECOVER keyword.

SMS relationship: Not applicable.

SETSYS default: None.

DFSMSshm default: MASTERSCHEDULERCANDIDATE(Y), but will have no effect until a SETSYS CQUEUE(RECOVER(CONNECT(groupname))) is issued.

The host will respond with message ARC15668I if the host lost its Master Scheduler responsibilities and with message ARC1569I if it took them on.

RECOVER(DISCONNECT):

Explanation: RECOVER(DISCONNECT) is an optional sub-parameter of the COMMONQUEUE parameter which specifies that a host disconnect itself from using the currently connected Common Recover Queue (CVQ) for processing. The following considerations apply to this disconnect process:

- All CVQ requests being processed on this host, and requests that originated from this host and are selected to be processed in the CVQ, are waited upon for completion before the host is disconnected from the CVQ. This host will no longer be selected to process new volume restore work from the CVQ.
- A HOLD RECOVER issued prior to the disconnect command prevents queued requests from being selected to be processed. After the disconnect has completed, a RELEASE RECOVER allows the queued requests to be processed locally.
- When the disconnect is for the host that is currently the Master Scheduler, all CVQ active and queued work is waited upon for completion before the host is disconnected. Once the host is disconnected, another member in the CVQ group that has been previously defined as a Master Scheduler candidate automatically takes over the Master Scheduler responsibilities. A HOLD RECOVER prior to a disconnection prevents queued requests from being selected for processing. A subsequent RELEASE

RECOVER after the disconnecting allows these requests to be processed locally.

- To verify the disconnection, use the QUERY ACTIVE command and look for message ARC1560I.
- Recover resources available on this host after the disconnection has completed are not known or available to the remaining CVQ members. All new recover requests received are placed on this host's local queue and processed locally.
- For shutdown procedures, a DISCONNECT should first be issued on each host being removed from the HSMplex to allow CVQ requests to complete first.
- The DISCONNECT FORCE option should be reserved for emergency situations. With FORCE specified, requests that originated from this host and are selected to be processed in the CVQ are not waited upon for completion before disconnecting. Messages are issued to indicate that a request has been failed but still may complete in the CVQ.

SMS relationship: Not applicable.

SETSYS default: None.

DFSMSHsm default: None.

The host will respond with message ARC1567I when the host disconnects from the group and with message ARC1568I if it lost the Master Scheduler responsibilities.

RECOVER:

Explanation: SETSYS CQUEUE(RECOVER) with no options specified will parse out with an error.

2.2.7.4 MAXDUMPRECOVERTASKS: Specifying the maximum number of concurrent volume recovery from dump tasks

Changed the description for 'nn': Specifying the maximum number of concurrent volume recovery from dump tasks

For nn, specify a decimal number from 0 to 64 to represent the number of volume recovery from dump tasks to be processed concurrently.

Added second note:

For Common Recover Queue (CVQ) processing consider using SETSYS MAXDUMPRECOVERTASKS(0) from the Master Scheduler (MS), instead of a HOLD RECOVER or FRRECOV, if you want the MS to assign CVQ work but not use its resources for CVQ fast replication or volume recovery processing. To end this state, use a SETSYS MAXDUMPRECOVERTASKS command with a non zero value.

2.2.7.5 Examples of how to code the SETSYS command

Added the following section

Specifying the SETSYS parameters for recover processing:

In this example, the SETSYS parameters associated with recover processing are specified. Where

appropriate, default values or recommended values are used in the example. These parameters are in effect until another SETSYS command is issued to change them.

- A maximum of two dump recover tasks will be used for FRRECOV and RECOVER FROMDUMP volume recovers
- The DFSMSHsm host is connected to the XCF group ARCVREC1 for Common Recover Queue processing. This host is eligible to be the Master Scheduler for the Common Recover Queue. All available dump recover tasks in the HSMplex hosts that are connected to the group as members can be used to process volume restores as they are received on this local host. The parameter used is COMMONQUEUE(RECOVER(CONNECT(REC1) MSC)).

```
SETSYS MAXDUMPRECOVERTASKS(2) +
COMMONQUEUE (RECOVER (CONNECT (REC1) MSC))
```

2.2.8 Using the QUERY command

The following messages were changed for the specified QUERY:

SETSYS:

Modified: ARC1500I

```
ARC1500I PLEXNAME=HSMplex_name, PROMOTE PRIMARYHOST={YES | NO}, PROMOTE SSM={YES | NO}, COMMON RECALL QUEUE BASE NAME={basename | *****}, COMMON RECALL QUEUE TAPEDATASETORDER={PRIORITY | FBID | *****}, COMMON DUMP QUEUE BASE NAME={basename | *****} AND MSC={YES | NO | ***}, COMMON RECOVER QUEUE BASE NAME={basename | *****} AND MSC={YES | NO | ***}
```

WAITING:

Modified: ARC1542I

```
ARC1542I WAITING MWES ON COMMON QUEUES: COMMON RECALL QUEUE=type_mwes, COMMON DUMP QUEUE=type_mwes, COMMON RECOVER QUEUE=type_mwes, TOTAL=total_mwes
```

2.3 z/OS DFSMS Using the new functions

Publication SC23-6857-04 updated with the following information:

2.3.1 Using the DFSMSHsm Enhancements in z/OS V2R2

Message Simplification:

Changed Message simplification to indicate FRRECOV TOVOLUME FROMDUMP supported:

To simplify the analysis of fast replication errors during FRBACKUP, FRRECOV COPYPOOL and copy pool

auto dump, and `FRRECOV TOVOLUME FROMDUMP`, DFSMSHsm now collects all related DFSMSHsm and DFSMSdss messages and records them in a unique fast replication message data set. . . .

The message simplification function also supports the [Common Dump Queue](#) and [Common Recover Queue](#) processing.

Common Recover Queue:

The following new section was added for the Common Recover Queue:

With this enhancement, storage administrators have the ability to distribute the processing of volume restores (`RECOVER *`, `FRRECOV COPYPOOL`, and `FRRECOV TOVOLUME` with `FROMDUMP`) to all DFSMSHsm tape resources available in a group of DFSMSHsm hosts.

New parameters on the `SETSYS COMMONQUEUE` command define the common recover group and its members (hosts) in the HSMplex. The `HOLD`, `RELEASE`, `QUERY`, `CANCEL`, and `ALTERPRI` commands all now support this function. XCF system facilities are used to provide this support. Message simplification supports a host running a Common Recover Queue

2.3.2 Administrating

Added the following new section for the Common Recover Recover

Administrating Common Recover Queue:

This topic describes the administration tasks for using the Common Recover Queue enhancements in V2R2.

Using the `SETSYS COMMONQUEUE RECOVER(CONNECT(basename))` command to connect hosts as members of the Common Recover Queue:

Use the `SETSYS COMMONQUEUE(RECOVER(CONNECT(basename)))` command to connect hosts as members of the Common Recover Queue group of hosts in an HSMplex. When specified, this system will be allowed to send and receive volume restores to and from other systems in the HSMplex that was started with the same basename.

The DFSMSHsm host that manages all of the volume restore requests in the Common Recover Queue (CVQ) is called the Master Scheduler. All volume restore requests are sent to the Master Scheduler, and it assigns work to all of the other hosts eligible in the group to process requests. Each group needs a Master Scheduler to manage CVQ processing for the group; otherwise, volume restore requests are processed locally.

A system becomes a candidate for being the Common Recover Queue Master Scheduler when the `MASTERSCHEDULERCANDIDATE(YES)` sub-parameter is specified with `SETSYS COMMONQUEUE (RECOVER(CONNECT(basename)))`. Note that this is the default if not specified. You should specify `MASTERSCHEDULERCANDIDATE(NO)` if you wish to prevent this host from becoming a Master Scheduler for the group.

One of the hosts that connects to the group as a Master Scheduler candidate becomes the Master Scheduler.

For more information, refer to the description of the `COMMONQUEUE` parameter for the `SETSYS` command in z/OS DFSMSHsm Storage Administration.

Using the `QUERY COMMONQUEUE` command to display requests queued or active in the Common Recover Queue (plus other `QUERY` options):

Use the QUERY COMMONQUEUE with the RECOVER sub-parameter to display the status of RECOVER and FRRECOV requests that are being processed and queued in the Common Recover Queue (CVQ) of which the host issuing the command is a member. This query will provide a snapshot of all CVQ activity across all group members.

The QUERY SETSYS command displays the current settings for the Common Recover Queue. For QUERY ACTIVE, QUERY WAITING, QUERY REQUEST, and QUERY USER, in a Common Recover Queue, only requests that originated from the host that issued the QUERY command are returned.

For more information, refer to QUERY command in z/OS DFSMSHsm Storage Administration.

Using HOLD for volume restores in a Common Recover Queue:

To prevent the Master Scheduler host from selecting a CVQ-connected host to process volume restores from tapes, use SETSYS MAXDUMPRECOVERTASKS(0) on hosts that should not be selected for CVQ processing. An alternative is to use the HOLD RECOVER and FRRECOV(TAPE) commands, but this also holds all Master Scheduler assignment duties on a host that is the Master Scheduler or a candidate for one. It is recommended that SETSYS MAXDUMPRECOVERTASKS(0) be used to configure hosts to not process volume restores, while HOLD RECOVER or FRRECOV(TAPE) be reserved for situations where ongoing volume restore processing needs to be held for environmental reasons.

Using ALTERPRI to change the priority of queued requests in the Common Recover Queue:

For environments where a host is connected to a Common Recover Queue, you can alter the priority of queued requests by the request number and by the userid that issued the request. This will change the priority of all requests that are queued in the Common Recover Queue of which the system issuing the command is a member.

For more information, see the ALTERPRI command description in z/OS DFSMSHsm Storage Administration.

Using the DISPLAY command to monitor XCF for the Common Recover Queue:

The DISPLAY XCF command can be used to monitor the XCF facilities utilized by the Common Recover Queue. Tracking and monitoring RECOVER, and FRRECOV functions remain unchanged.

Message simplification (changed):

Using MESSAGEDATASET with the FRBACKUP and FRRECOV COPYPOOL/TOVOLUME commands to specify the date and time used in the message data set name

MESSAGEDATASET is an optional parameter that allows you to specify the date and time used in the name of the fast replication command message data set. Specifying the date and time ensures that you can identify the message data set that corresponds with the FRBACKUP request, FRRECOV COPYPOOL request, or FRRECOV TOVOLUME FROMDUMP requests. If you don't specify the date and time with the MESSAGEDATASET parameter, DFSMSHsm uses the date and time when the message data set name is generated.

For more information, refer to the FRRECOV command description in z/OS DFSMSHsm Storage Administration.

2.4 z/OS DFSMSHsm implementation and customization guide

Publication SC23-6869-02 updated with the following information:

2.4.1 Specifying commands that define your environment

Added the following new section:

Defining the Common Recover Queue (CVQ) environment

DFSMSHsm supports an HSMplex-wide Common Recover Queue (CVQ). With CVQ, volume restore requests are distributed to a group of hosts for processing. This increases the number of available tasks to perform the work and improves throughput by distributing the workload instead of concentrating it on a single host's address space.

For an overview of the CVQ environment and how to define it, refer to the z/OS DFSMSHsm Storage Administration guide.

2.4.2 Customizing in a sysplex environment

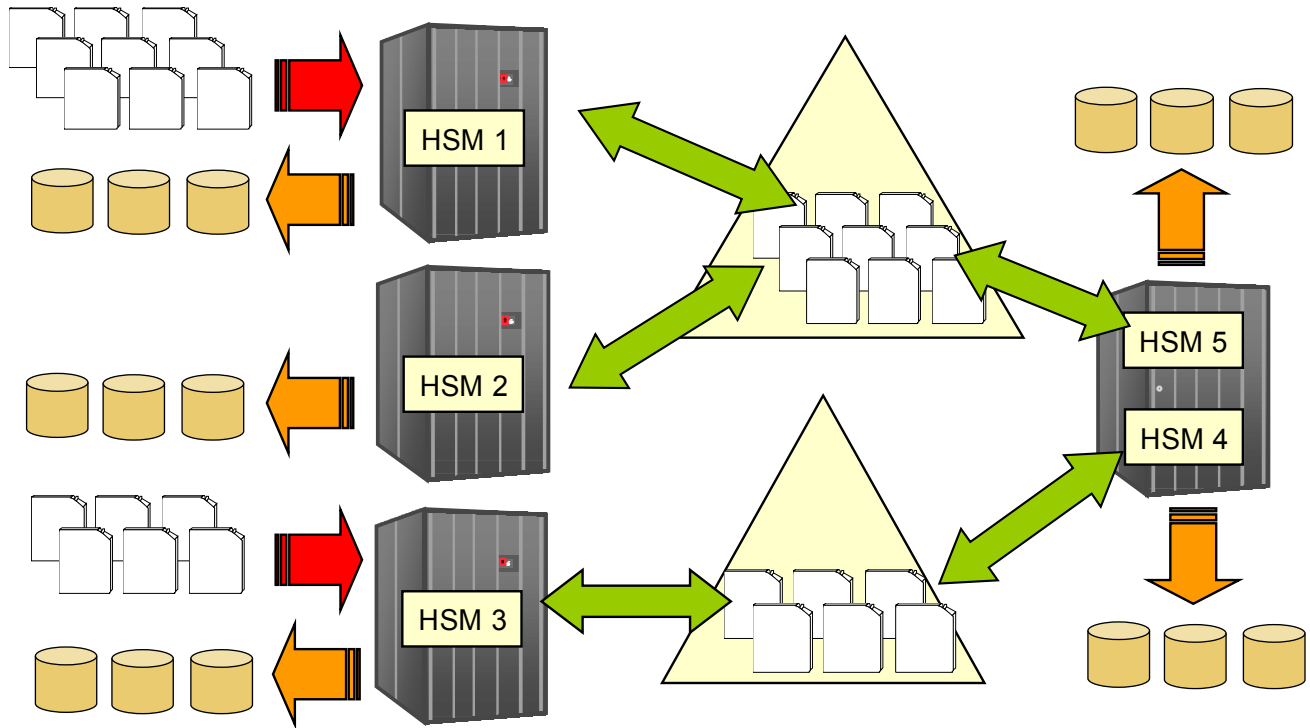
Added the following new section:

Common recover queue (CVQ) configurations

For an overview of the CVQ environment, refer to the z/OS DFSMSHsm Storage Administration guide.

A standard HSMplex configuration is one where all hosts are connected to the same CVQ and all hosts are eligible to process volume restores regardless of what host was used to submit the requests. The CVQ is a queue of volume restore requests that is shared by these hosts, managed by a *Master Scheduler (MS)* host and implemented through the use of the Cross-system Coupling Facility (XCF) for host-to-host communication between an XCF defined group and its members. The purpose of the CVQ is to balance volume restore processing across the resources available in all the hosts and return results back to the host where the request originated to post the user complete.

As shown in the figure, the CVQ group allows for flexible configurations, providing the capability to define multiple queues in the same HSMplex, allow group members to both receive and process requests, just process requests and receive requests but not process them.



In CVQ environments

- A CVQ can not span an HSMplex
- A host can only be connected to one CVQ
- Can coexist with a CDQ environment with the same or subset of hosts
- Multiple CVQ groups can be defined

Results are returned to the submitting host but progress and status messages are recorded on the host processing the request.

The *Submitting Host* receives requests from commands and sends it to the *MS host* and when the command is completed by the group it is notified to post the user that the command has completed.

The *Master Scheduler (MS)* is the DFSMSHsm host that manages all of the volume restore requests in the CVQ. It accepts requests from a *submitting host* and itself. The MS assigns the requests to eligible hosts (*processing hosts*), including itself, that have available tasks to process the work while balancing the utilization of the dump tasks in the group. The MS also manages the interaction between the *processing host* for tape optimization and the *submitting host* for the command completes notifications.

The *Processing Host* receives assigned work requests from the MS and completes the work and interacts with the MS to manage stacking.

Any given host in the CVQ could be any or all of the host types depending on how your environment is setup.

For cases where you do not want volume restore tasks on a host to be used by the CVQ group avoid using HOLD RECOVER or FRRECOV(TAPE). Instead use a SETSYS MAXDUMPRECOVERTASKS(0) to prevent its volume restore tasks from being used. It has the same effect of HOLD RECOVER or FRRECOV(TAPE) without the risk of affecting MS responsibilities. A HOLD on the MS will prevent it from assigning and processing requests for the CVQ.

2.4.3 Patch for FRRECOV COPYPOOL FROMDUMP performance to bypass the EXCLUSIVE NONSPEC ENQ

Add the following to the description of DFSMSshm PATCH command,

PATCH .MCVT.+3D5 BITS(.....1):

(Original description)

To improve FRRECOV COPYPOOL FROMDUMP performance, you may want to bypass the EXCLUSIVE NONSPEC ENQ that is performed when DSS allocates its temporary data set or temporary work area for ICKDSF to use when building the VTOC index. The enqueue is designed to prevent a lockout that may occur involving SYSZTIOT and SYSVTOC if the temporary data set allocation is directed to a volume that DFSMSshm is dumping or restoring and an EOVS is issued.

Users who wish to bypass this enqueue process and are willing to accept the possibility of the enqueue lockout between SYSZTIOT and SYSVTOC can do so using the following DFSMSshm PATCH command:

To activate the feature, issue the following:
 PATCH .MCVT.+3D5 BITS(.....1)

To deactivate the feature, issue the following:
 PATCH .MCVT.+3D5 BITS(.....0)

(Added description)

Given that the Common Recovery Queue's (CVQ) primary goal is to increase throughput of parallel volume restores, this patch will be enabled for all CVQ-submitted FRR CP FROMDUMP requests without the need for the user to set the enabling patch.

2.5 z/OS MVS System Messages, Vol 2 (ARC-ASA)

Publication SA38-0669-05 updated with the following information:

2.5.1 ARC1500I

Message ARC1500I modified:

```
ARC1500I PLEXNAME=HSMplex_name,PROMOTE PRIMARYHOST={YES | NO},PROMOTE SSM={YES | NO},COMMON RECALL
QUEUE BASE NAME={basename | *****},COMMON RECALL QUEUE TAPEDATASETORDER={PRIORITY | FBID |*****},
COMMON DUMP QUEUE BASE NAME={basename | *****} AND MSC={YES | NO | ***}, COMMON RECOVER QUEUE BASE
NAME={basename | *****} AND MSC={YES | NO | ***}
```

Explanation: A QUERY command was issued with the SETSYS parameter. This messages describes the current implementation.

Parameter Meaning:

COMMON RECOVER QUEUE

BASE NAME *basename* is the name of the Common Recover Queue specified on SETSYS CQUEUE(RECOVER(CONNECT(*basename*))) and running in the PLEXNAME HSMplex_name.

**** indicates that SETSYS CONNECT has not been specified for the Common Recover Queue.

MSC

=YES, this host can be the Master Scheduler for the Common Recover Queue. The system enabled to be the Master Scheduler is assigned at DFSMSHsm startup or during recovery.

=NO, this host cannot be the Master Scheduler for the Common Recover Queue

=***, a SETSYS command with the MSC keyword was not issued for this host. By default this host can be the Master Scheduler for the Common Recover Queue. The system enabled to be the Master Scheduler is assigned at DFSMSHsm startup or during recovery.

2.5.2 **ARC1542I**

Message ARC1542I modified:

```
ARC1542I WAITING MWES ON COMMON QUEUES: COMMON RECALL QUEUE=type_mwes, COMMON
DUMP QUEUE=type_mwes, COMMON RECOVER QUEUE=type_mwes, TOTAL=total_mwes
```

2.5.3 **ARC1560I**

Message ARC1560I modified:

```
ARC1560I COMMON type QUEUE FACTORS: GROUPNAME={groupname | *}, CONNECTION
STATUS={CONNECTING | CONNECTED | DISCONNECTING | UNCONNECTED | QUIESCED},
MASTERSCHEDULER HOSTID={hostid | NONE | *}, MASTER SCHEDULER STATUS={DISABLED |
ENABLED | QUIESCED | *}
```

Explanation: A QUERY command was issued with the ACTIVE parameter. DFSMSHsm issues this message, to describe the the current group status and related common queue factors of the host issuing the QUERY command.

The message variables are explained:

Variable Meaning

type

DUMP, common queue active status is described.

RECOVER, common queue active status is described.

groupname

XCF group name defined to the functional group running in the HSMplex.

CONNECTION STATUS describes the current connection state of this host with the *groupname*. NONE is an invalid status and should be reported to IBM support.

When the MASTER SCHEDULER STATUS is set to:

ENABLED the Master Scheduler *hostid* of the DFSMSHsm Common *type* Queue function has tasks available in the group to perform the work submitted by this host. For the work submitted, the Master Scheduler host will assign the work to a processing host in the Common *type* Queue if the function is not held on the Master Scheduler host.

ENABLED(RECOVER HELD) The host, *hostid*, that issued the QUERY is the Master Scheduler of the DFSMSHsm Common Recover Queue function. It has tasks available in the group to perform the work submitted by a Common Recover Queue host, but may not be able to because of a HOLD. A HOLD RECOVER was issued on this host to receive this status. For this Master Scheduler host, RECOVER requests will not be assigned but queued, RECOVER WAIT type requests received will be failed, and FRRECOV work received will be assigned to a processing host.

ENABLED(FRRECOV HELD) The host, *hostid*, that issued the QUERY is the Master Scheduler of the DFSMSHsm Common Recover Queue function. It has tasks available in the group to perform the work submitted by a Common Recover Queue host, but may not be able to because of a HOLD. A HOLD FRRECOV was issued on this host to receive this status. For this Master Scheduler host, FRRECOV requests will not be assigned but queued, FRRECOV WAIT type requests received will be failed, and RECOVER work received will be assigned to a processing host.

ENABLED(HELD) The host, *hostid*, that issued the QUERY is the Master Scheduler of the DFSMSHsm Common *type* Queue function. It has tasks available in the group to perform the work submitted by a Common *type* Queue host, but is not be able to because of a HOLD.

For a Common Recover Queue Master Scheduler host, FRRECOV and RECOVER requests will not be assigned, FRRECOV and RECOVER WAIT type requests will be failed, and NOWAIT type requests will be queued and not assigned to a processing host. An HOLD FRRECOV and HOLD RECOVER or a HOLD ALL was issued on this host to receive this status.

For a Common Dump Queue Master Scheduler host, FRBACKUP, BACKVOL and AUTO dump requests will not be assigned, FRBACKUP and BACKVOL dump WAIT type requests will be failed, and NOWAIT type requests will be queued and not assigned to a processing host. A HOLD on dump of any type was issued on this host to receive this status.

QUIESCED the Master Scheduler *hostid* of the DFSMSHsm Common *type* Queue function has no tasks available in the group to perform the work for this host as it is transitioning out of its roll of being the Master Scheduler with a new one being assigned by the system if available. The SETSYS command COMMONQUEUE(*type*(MSC(N)) or COMMONQUEUE(*type*(DISCONNECT)) was issued on the Master Scheduler *hostid*.

DISABLED(HELD)

the Master Scheduler *hostid* of the DFSMSHsm Common *type* Queue function has no tasks available in the group to perform the work for this host. The Common *type* Queue function is HELD on all members of the group or the maximum number of tasks for the function on all the members have been set to zero.

*

, the system was not defined to perform common queue processing for DFSMSHsm function, *type*. There is not Master Scheduler status to report.

2.5.4 **ARC1562I**

Message ARC1562I GROUPNAME text changed to GROUP NAME to match what gets issued by the system:

ARC1562I VOLUME {*volser* | *volser*, SGROUP=*sg* | *volser*, COPY POOL = *cpname*, } {DUMPING | FRBACKUP | FRBACKUP DUMP | FRBACKUP DUMPOONLY | RESTORING | FRRECOV | FRRECOV FROMDUMP}, FOR USER {*userid* | **AUTO** | *} REQUEST {*request-number* [ON HOST=*hostid*] | NONE }, BEING PROCESSED FROM A COMMON QUEUE, GROUP NAME=*groupname*

2.5.5 **ARC1563I**

Message ARC1563I GROUPNAME text changed to GROUP NAME to match what gets issued by the system:

ARC1563I VOLUME {*volser* | *volser*, SGROUP=*sg* | *volser*, COPY POOL = *cpname*, } {DUMPING | FRBACKUP | FRBACKUP DUMP | FRBACKUP DUMPOONLY | RESTORING | FRRECOV | FRRECOV FROMDUMP} FOR USER {*userid* | **AUTO** | *} REQUEST {*request-number* [FROM HOST=*hostid*] | NONE }, WAITING TO BE PROCESSED ON A COMMON QUEUE, nmwe MWES AHEAD OF THIS ONE, GROUP NAME=*groupname*

Variable Meaning

volser is specified for full volume dump and recover functions. An '*****' may appear for FRBACKUP or FRRECOV requests that have yet to be converted to their individual requests for each volume described by *volser*.

2.5.6 **ARC1564E**

Message ARC1564I in publication changed to ARC1564E and modified with the following:

ARC1564E DFSMSHsm CANNOT PROCESS THE SETSYS REQUEST TO {CONNECT TO | DISCONNECT FROM} THE {DUMP | RECOVER} GROUPNAME *group_name*, REASON = *reason description*

Explanation: A SETSYS COMMONQUEUE command failed to complete as expected. *group_name* indicates the name of the XCF group that was being connected to or disconnected from. The *reason description* could be one of the following:

2.5.7 **ARC1565E**

Message ARC1565E modified:

```
ARC1565E HOST=hostid HAS LOST ITS CONNECTION TO ITS {DUMP | RECOVER} GROUP
group_name DUE TO A {CONNECTIVITY LOSS | DFSMSHSM ABEND}
```

Explanation: This host unexpectedly lost its connection to the specified group name.

2.5.8 **ARC1566I**

Message ARC1566I modified:

```
ARC1566I CONNECTION TO THE {DUMP | RECOVER} GROUP group_name WAS {SUCCESSFUL |
UNSUCCESSFUL, RC=retcode, REASON=reascode }
```

Explanation: DFSMSHsm attempted to connect to the function group (DUMP or RECOVER) that was specified by the SETSYS COMMONQUEUE(function (CONNECT(basename))) command. For an UNSUCCESSFUL attempt the return code and reason code are those that were returned by XCF services. For an explanation of the return and reason codes, see the IXCJOIN macro in z/OS MVS Programming: Sysplex Services Reference.

2.5.9 **ARC1567I**

Message ARC1567I modified:

```
ARC1567I DISCONNECTION FROM THE {DUMP | RECOVER} GROUP group_name WAS
{SUCCESSFUL | UNSUCCESSFUL, RC=retcode, REASON=reascode}
```

Explanation: DFSMSHsm attempted to disconnect to the function group (DUMP or RECOVER) that was previously connected to this system by the SETSYS COMMONQUEUE(function (DISCONNECT))command. For an UNSUCCESSFUL attempt the return code and reason code are those that were returned by XCF services. For an explanation of the return and reason codes, see the IXCLEAVE macro in z/OS MVS Programming: Sysplex Services Reference.

2.5.10 **ARC1568I**

Message ARC1568I modified:

```
ARC1568I HOST=hostid HAS LOST ITS MASTER SCHEDULER RESPONSIBILITIES FOR THE
COMMON {DUMP | RECOVER} QUEUE GROUP
groupname
```

Explanation: The HOST identified by *hostid* was the Master Scheduler host for the indicated functional group (DUMP or RECOVER) group, *groupname*. This could have been triggered by the issuance of the SETSYS command that set its Master Scheduler candidate status to NO on a host that was the active Master

Scheduler at the time. For example, SETSYS CQUEUE(DUMP(MSC(N))). The DFSMSHsm host may have also been disconnected from the common functional group because DFSMSHsm was stopped or the SETSYS DISCONNECT command was issued. For example, SETSYS CQUEUE(DUMP(DISCONNECT)).

2.5.11 **ARC1569I**

Message ARC1569I modified:

ARC1569I HOST=hostid HAS TAKEN OVER THE MASTER SCHEDULER RESPONSIBILITIES FOR THE COMMON {DUMP | RECOVER} QUEUE GROUP groupname

Explanation: The HOST identified by *hostid* was assigned by the system to be the Master Scheduler host for the indicated functional group (DUMP or RECOVER) group, *groupname*. This message is normally issued during the common function queue CONNECT processing to indicate who was assigned this responsibilities. However, it could be triggered when the system detected that the current Master Scheduler could no longer perform its responsibilities for the common function group. See message ARC1568I.

2.5.12 **ARC1127I**

Message ARC1127I added:

ARC1127I RECOVER FAILED DUE TO AN ERROR WITH THE COMMON RECOVER QUEUE

Explanation: DFSMSHsm attempted to process a recover operation using its common queue. The volume, command, or operation and the reason for the failure are described by message ARC1001I. In message ARC1001I, the *reascode* identifies the results. Summary message will also be issued with the same *reascode*.

The message variables are explained:

Variable Meaning

reascode

100 - A SETSYS COMMONQUEUE(*type*(DISCONNECT FORCE)) command was processed. This return code indicates that active requests from the common queue for this function running in this system were allowed to finish. However, the remote requests submitted to the common queue for this function from this system are failed but may still complete on another system in the functional group.

102 - A SETSYS COMMONQUEUE(*type*(DISCONNECT)) command was processed for a host member that is the current Master Scheduler. This return code indicates that active requests from the common queue for this function running in the group were allowed to finish. However, any request received from the remote members in the group on this Master Scheduler system that are waiting to be processed will be failed.

105 - A host in the common queue group became not defined. This return code indicates that the work was assigned by the Master Scheduler to run in this host of the common queue but the host cannot determine if the work completed

106 - The current Master Scheduler host for the Common Recover Queue was quiesced with a SETSYS CQ(RECOVER(MSC(N))) so that the Master Scheduler responsibilities could be taken over by

another member in the group. This return code indicates that other members in the group sent recover work to the quiescing Master Scheduler and that this work was removed from its queue. The recover work will get re-queued on the submitting host to be processed by the new Master Scheduler or processed locally if no other host took over the Master Scheduler responsibilities for the Common Recover Queue.

System Action: The system will take the action taken as described by each return code.

Application Programmer Response: Notify the storage administrator

Source: DFSMSHsm

2.5.13 **ARC1132I**

Message ARC1132I added:

ARC1132I RECOVER FROM DUMP FAILED IN THE COMMON RECOVER QUEUE

Explanation: A RECOVER or FRRECOV FROMDUMP command failed when processed in a common queue. One or more dump volumes were not processed. This is a summary message where the highest reason code for a volume not processed is provided in the reason code for ARC1001I return code 32.

System Action: The RECOVER or FRRECOV FROMDUMP operation ends. Some volumes may have been completed. DFSMSHsm processing continues.

Application Programmer Response: Evaluate the reason codes for the individual volume(s) that have failed. See the Programmer Response for the specific messages that have been issued prior to this ARC1001I RC01

Source: DFSMSHsm

2.5.14 **ARC1829I**

Message ARC1829I modified and added reason code 106:

ARC1829I FAST REPLICATION FAILED DUE TO AN ERROR WITH THE COMMON DUMP OR RECOVER QUEUE

Explanation: DFSMSHsm attempted to process a FRBACKUP or FRRECOV operation using its common queue. The volume, command, or operation and the reason for the failure are described by message ARC1001I. In message ARC1001I, the *reascode* identifies the results.

The message variables are explained:

Variable Meaning

reascode

106 - The current Master Scheduler host for the Common Recover Queue was quiesced with a SETSYS CQ(RECOVER(MSC(N))) so that the Master Scheduler responsibilities could be taken over by another member in the group. This return code indicates that other members in the group sent FRR work to the quiescing Master Scheduler and that this work was removed from its queue. The FRR work will get re-queued on the submitting host to be processed by the new Master Scheduler or processed locally if no other host took over the Master Scheduler responsibilities for the Common Recover Queue.

2.5.15 **ARC1870I**

Message ARC1870I modified for TOVOLUME requests:

```
ARC1870I MESSAGE DATA SET msgdsname CREATED FOR {FAST REPLICATION {BACKUP |
BACKUP DUMP | BACKUP DUMPONLY | PREPARE | RECOVERY | RECOVERY FROM DUMP} | AUTO
DUMP} OF {COPY POOL cpname | VOLUME volser }
```

Description changes

Explanation:

DFSMSHsm created message data set *msgdsname* for the specified function. Messages issued for the specified operation will be written to the message data set.

When the function completes, DFSMSHsm will rename the data set, indicating the status of the operation in the final qualifier. The message data set naming convention is as follows:

For FRBACKUP requests and ccopypool auto dump:

```
msgdsprefix.FB.Ccccccc.ippppppp.Dyyddd.Thhmmss.z
```

For FRRECOV COPYPOOL requests:

```
msgdsprefix.FR.Ccccccc.ippppppp.Dyyddd.Thhmmss.z
```

For FRRECOV TOVOLUME requests:

```
msgdsprefix.FR.Vvolser.Dyyddd.Thhmmss.z
```

where:

msgdsprefix

is the value of the HLQ parameter of the SETSYS FASTREPLICATION(MESSAGE DATASET) command, if specified. If the HLQ parameter is not specified, the default prefix of 'HSMMSG' will be used.

cccccc

is up to 7 characters to identify the copy pool that was processed. For a non-DB2 copy pool, this will be the first 7 characters of the copy pool name. For a DB2 copy pool, this will be the first 7 characters of the location qualifier in the DB2 copy pool name.

i

is an identifier, indicating whether the message data set contains messages for:

D – DB2 data

L – DB2 log

X – non-DB2 copy pool data

ppppppp

is up to 7 characters to identify the copy pool that was processed. For a non-DB2 copy pool, this will be the next 7 characters of the copy pool name (the characters following *cccccc*). For a DB2 copy pool, this will be the next 7 characters of the location qualifier in the DB2 copy pool name.

volser

is the *volser* of the volume that was recovered.

yyddd

is the year and day when the FRBACKUP, FRRECOV, or copy pool auto dump operation started.

hhmmss

is the time in hours, minutes, and seconds when the FRBACKUP, FRRECOV, or copy pool auto dump operation started.

z

indicates the status of the FRBACKUP, FRRECOV, or copy pool auto dump operation. This qualifier will appear as **** in the ARC1870I message, to represent one of the following values:

S – Success

The operation completed successfully.

F – Fail

One or more errors occurred while processing the operation.

I – In progress

DFSMSHsm will allocate the data set with this final qualifier, and rename the data set with one of the qualifiers above once the operation completes.

FB, FR, C, *V*, D, and T are constants.

System Action:

DFSMSHsm processing continues.

Application Programmer Response:

If the last qualifier of the message data set indicates that the operation failed (final qualifier is 'F'), see the message data set contents to determine the cause of the error. If the final qualifier is 'I', DFSMSHsm was unable to rename the data set at the end of the operation.

Source: DFSMSHsm

2.5.16 ARC1873I

Message ARC1873I modified for TOVOLUME requests:

```
ARC1873I MESSAGES NOT CAPTURED FOR {FAST REPLICATION {BACKUP | BACKUP DUMP |
BACKUP DUMPNLY | PREPARE | RECOVERY | RECOVERY FROM DUMP} | AUTO DUMP} OF {COPY
POOL cpname | VOLUME volser }
```

Explanation:

The specified DFSMSHsm function for copy pool cpname, or volume volser has started. DFSMSHsm was either unable to allocate a message data set for this request, or the message data set was allocated with incorrect attributes. DFSMSHsm will not capture messages issued for the specified function in a message data set.

2.6 z/OS Summary of Messages and Interface Changes

Publication SA23-2300-7 updated with the following information:

2.6.1 DFSMSHsm summary of interface changes

Added the following to the command interface changes:

| Command Name | Release | Description | Reason for change |
|--------------|-----------|---|--|
| SETSYS | z/OS V2R2 | New: New parameters have been added for use with COMMONQUEUE keyword to allow DFSMSHsm hosts to participate in a Common Recover Queue (CVQ) | Distributed tape processing enhancements |
| | | Changed: The MAXDUMPRECOVERTASKS parameter now accepts a value of 0, which allows a host running in a Common Recover Queue to submit requests, while preventing the Common Recover Queue from using its available recover tasks for processing, without using a HOLD which would prevent the Master Scheduler's host from assigning requests to run in the Common Recover Queue. | Distributed tape processing enhancements |
| QUERY | z/OS V2R2 | New: The RECOVER parameter has been added to the QUERY COMMONQUEUE command to display active and queued requests in the DFSMSHsm Common Recover | Distributed tape processing enhancements |

| | | | |
|--|--|--------|--|
| | | Queue. | |
|--|--|--------|--|

2.7 z/OS DFSMSHsm Data Areas

Publication GC14-7504-02 updated with the following information:

2.7.1 MWE - Management Work Element

Changed or new fields in the MWE:

| | | | | |
|----------|-----------|---|--------------------|--|
| 135(87) | BITSTRING | 1 | MWEFLG10 | Additional Flags |
| | 1... .. | | MWEF_CDQ_MS | CDQ submitted MWE |
| | .1.. .. | | MWEF_CVQ_MS | CVQ submitted MWE |
| | ..1. | | MWEF_SERVICE_CALL | Service call MWE from QUERY, CANCEL, ALTERPRI |
| | ...1 | | MWEF_CxQ_NoReturn | MWE purge should skip the return to the submitting host |
| | 1... | | MWEF_CxQ_Submitted | This MWE has gone to the CxQ |
| |111 | | * | Not used |
| ... | | | | |
| 640(280) | ADDRESS | 4 | MWE_CPELMP_MSGDS | CPELM ptr for msg ds abend processing |
| 640(280) | ADDRESS | 4 | MWE_MS_MWEP | Pointer to MS MWE |
| ... | | | | |
| 660(294) | ADDRESS | 4 | MWE_RCBP | PTR to RCB for task processing this MWE |
| 664(298) | BITSTRING | 8 | MWE_MS_TOKEN | Token of the CVQ MS |
| 66C(306) | SIGNED | 4 | * | Reserved |

2.7.2 MWE Control Block Cross-Reference

New fields added or changed in the MWE:

| Name | Hex Offset | HexValue | Struct Level |
|------------------------|---------------|---------------|--------------|
| MWEF_CGQ_MS | 87 | 40 | 4 |
| MWEF_CVQ_MS | 87 | 40 | 4 |
| MWEF_CxQ_NoReturn | 87 | 20 | 4 |
| MWEF_CxQ_Submitted | 87 | 10 | 4 |
| ... | | | |
| MWE_MS_MWEP | 280 | | 4 |
| ... | | | |

MWE_MS_TOKEN 298 3

2.7.3 FRGCB—Fast Replication Global Control Block

Changed or new fields in the FRGCB:

| | | | | |
|----------------------|-----------------------|------------------|------------------------------|--|
| 8 (8) | BITSTRING | 4 | FRGCB_FLAGS | Miscellaneous flags |
| ... | | | | |
| 11(B) | 1... .. | | FRGCBF_VVDSPERF | On = Enable VVDS performance enhancement. Support Data set VVR RBA (VOLVRBA). |
| | .1.. .. | | FRGCBF_ENABLE_MSGDS | On = Enable use of FR message data sets |
| | ..1. | | FRGCBF_DGEN_CLSMSGDS | On = DCTL needs to post DGEN to close a message data set |
| | ...1 | | FRGCBF_GCTL_CLSMSGDS | On = GCTL needs to close a msg data set |
| | 1... | | FRGCBF_GCTL_OPNMSGDS | On = GCTL needs to open a msg data set |
| |1.. | | FRGCBF_FRTCB_VALID | On = FRGCB_FRTCB_PTRs are valid |
| |1. | | FRGCBF_ROUTEALL_UNALLOC | On = User wants the CATALOG UNALLOCATE commands to be routed to all LPARS |
| |1 | | FRGCBF_GCTL_CLSMSGDS_NOCPREL | On = A task wants GCTL to close a message data set |
| | 1111 | * | | Not used |
| ... | | | | |
| 104 (68) | BIT(32) | 4 | FRGCB_FRRM_ECB | Volume pairs on FRRMQ ready for processing Also used for enhanced messaging support |
| ... | | | | |
| 816 (330) | SIGNED 2 | | FRGCB_MSGDS_COUNT | Count of message data sets that are open |
| 816 (330) | SIGNED 4 | | FRGCB_MSGDS_COUNT | Count of message data sets that are open |
| 818 (332) | UNSIGNED 1 | *(24) | Unused | |
| 818 (332) | UNSIGNED 1 | *(22) | Unused | |

2.7.4 FRGCB Data Area Cross-Reference

New fields added or changed in the FRGCB:

| Name | Hex Offset | HexValue | Struct Level |
|------------------------------|------------|----------|--------------|
| FRGCBF_GCTL_OPNMSGDS | B | 08 | 3 |
| FRGCBF_ROUTEALL_UNALLOC | B | 02 | 3 |
| FRGCBF_GCTL_CLSMSGDS_NOCPREL | B | 01 | 3 |

2.7.5 JGCB – HSMplex Global Control Block

Changed or new fields in the JGCB:

| | | | | |
|---|----------------------|----------------------|-----------------------|---|
| 24 (18) | CHARACTER | 3 | JGCB_PLEX_PREFIX | Prefix for HSMplex name |
| 27 (1B) | UNSIGNED | 1 | *(33) | Reserved |
| 32(20) | SIGNED | 4 | JGCB_TRBUFF_PTR | PDA trace buffer (CVQ) |
| 36 (24) | UNSIGNED | 1 | *(24) | Reserved |
| ... | | | | |
| 780 (30C) | CHARACTER | 8 | JGCB_GRP_GROUP_NAME | Function group name 4-char prefix is fixed for CDQ – 'ARCD', CVQ – 'ARCV' CGQ – 'ARCG' . 4-char suffix is set from SETSYS CONNECT command for the function |
| ... | | | | |
| 794 (31A) | BIT(8) | 1 | JGCBF_GRP_FLAGS2 | Processing flags |
| | 11.. | | * | Unused |
| | ..1. | | JGCBF_GRP_REMOTE_WORK | Remote work for this function running on this system |
| | ...1 | | JGCBF_GRP_TASK_WAIT | On = CxQ nonlocal task is waiting for MS |
| | ...1 1111 | | * | Unused |
| | 1111 | | * | Unused |
| ... | | | | |
| Note: Reference these indexed with CGQ, field(CGQ) if they ever need to be defined | | | | |
| 952 | (3B8) | CHARACTER | 32 | JGCB_GRP_AFMS_CGQ Auto func data for rec |
| Note: Reference these indexed with CVQ, field(CVQ) if they ever need to be defined | | | | |
| 952 | (3B8) | CHARACTER | 32 | JGCB_GRP_AFMS_CVQ Auto func data for rec |
| ... | | | | |
| 0 (0) | STRUCTURE | * | JGCB_TR_BUFF | GXMxCB trace buffer based on JGCB_TRBUFF_PTR |
| 0 (0) | SIGNED | 4 | JGCB_GXMxCB_CNT | Count of GXMxCB's found |
| 4(4) | STRUCTURE | *(100) | JGCB_GXMxCB | Copy of each GXMxCB |
| ... | | | | |

Indexes into the CxQ grouped function array, JGCB_GRP_ELEMENT

| Length | Type | Value | Name | Description |
|--------|---------|-------|------|--------------------------------|
| 1 | DECIMAL | 1 | CDQ | Index for Common Dump Queue |
| 1 | DECIMAL | 2 | CVQ | Index for Common Recover Queue |

2.7.6 JGCB Data Area Cross-Reference

New fields added or changed in the JGCB:

| Name | Hex Offset | HexValue | Struct Level |
|------------------------------|----------------|----------|--------------|
| JGCB_GRF_AFMS_CGQ | 3B8 | | 4 |
| JGCB_GRF_AFMS_CVQ | 3B8 | | 4 |
| JGCBF_GRF_TASK_WAIT | 31A | 10 | 4 |
| JGCB_TRBUFF_PTR | 20 | | 2 |

2.7.7 XCFP – XCF Implementation macro

Changed or new fields in the XCFP:

DFSMSHsm GROUPED FUNCTION XCF MEMBER CONTROL BLOCK LIST
 THIS BLOCK INCLUDES THE CONNECTION STATE, USER STATE
 INFORMATION AND DESCRIPTIVE INFORMATION OF A MEMBER
 CONNECTED TO A GROUPED HSM FUNCTION (XCF).
 FOR A MASTER SCHEDULER (MS) SYSTEM AN ELEMENT FOR
 EACH MEMBER OF THE GROUP WILL BE MAINTAINED. THIS ALLOWS
 THE MS TO KEEP TRACK OF THE CONNECTION STATUS AND THE USE
 INFORMATION FOR EACH MEMBER IN THE GROUP. THE HEAD OF THIS
 WILL BE THAT OF THE MASTER SCHEDULER SYSTEMS ELEMENT IN THE
 LIST.
 FOR A SYSTEM THAT IS NOT A MS ONLY A SINGLE ELEMENT FOR THIS
 SYSTEMS CONNECTION STATUS AND USE INFORMATION WILL BE
 CREATED. THE HEAD or PTR TO THIS WILL BE THAT OF THIS SYSTEM
 THESE LINKED LISTS WILL BE CREATED FOR EACH
 FUNCTION DEFINED TO USE XCF GROUPS. POINTERS TO EACH LISTS
 ARE:

HEAD - JGCB_GRF_GXMCB_HEAD(x)

PTR - JGCB_GRF_GXMCB_PTR(x)

TAIL - JGCB_GRF_GXMCB_TAIL(x)

where 'x' - is the index for the group function.

These include -

CDQ for DUMP

CVQ for RECOVER

...

~~146 (92) UNSIGNED 2 GXMCB_SUPPORT_LEVEL FUNCTIONAL SUPPORT LEVEL FOR THIS HSM MEMBER 1 = CDQ (SEE LEVEL CONSTANTS)~~

146 (92) UNSIGNED 2 GXMCB_SUPPORT_LEVEL FUNCTIONAL SUPPORT LEVEL FOR THIS HSM MEMBER 1 = CDQ or CVQ (SEE LEVEL CONSTANTS)

...

HSM FUNCTIONAL SUPPORT LEVEL CONSTANTS

| Length | Type | Value | Name | Description |
|--------------|--------------------|--------------|-----------------------|--------------------------|
| 1 | DECIMAL | 1 | GXMCB_LVL1 | CDQ SUPPORTED |
| 1 | DECIMAL | 1 | GXMCB_LVL1 | CDQ, CVQ SUPPORTED |
| 2 | CHARACTER | DQ | GXMCB_CDQ | Common Dump Queue |
| 2 | CHARACTER | VQ | GXMCB_CVQ | Common Recover Queue |

2.7.8 XCFP Data Area Cross-Reference

No changes.