IBM System Lab Services
PowerHA Tools for IBM i

Full System Replication Manager Users Guide

September 10, 2026
Version 5.3

The most recent version of this document and updated information can be found online:

http://ibm.biz/PowerHATools

Introduction 5

What's new in 4.1 5

What's new in 4.2 5

What's new in 4.3 6

What's new in 4.4 7

What's new in 4.5 7

What's new in 4.6 7

What's new in 5.1 8

What's new in 5.2 8

What's new in 5.3 8

Planning 9

Requirements 9

Planning Considerations 9

FSR Subsystem Considerations 9

Running FSR in another subsystem 10

Clustering Security Requirements 10

Requirement to coordinate QTIME between managing partitions 11

Migrating from Replication Consistency Groups to Policy-Based Replication 11

Setting up Full System Replication 12

Considerations 12

System name vs. LPAR name 12

System serial number and software licensing 12

IP addresses 12

Clustering 12

Production LPAR Communication Agent 13

Credentials 13

LPAR date and time 13

Communications ports 13

PowerVS Setup 14

Performing switches with Full System Replication 16

Scheduled switch 16

Unscheduled switch 17

Starting replication after unscheduled switch 17

Detach 19

Reattach 20

Displaying Full System Replication information 22

WRKCSE Option 5=Display 22

WRKCSE Option 12=Work with 22

WRKCSE Work with Option 10=Display replication 23

WRKCSEDTA Option 5 = Display 23

DSPCSEDTA *23*

Full System LUN Switch **26**

Configuring a Full System LUN Switch *26*

Performing a Full System LUN Switch *29*
Scheduled switch 29
Unscheduled switch 30

Full System Replication 3-Site Configuration with Copy Services Manager (CSM) **31**

Overview 31
Benefits of 3-Site FSR 31
Operational Considerations 31
Supported Storage and Requirements 31
Replication Flow 32
Role of Copy Services Manager (CSM) 32
Architectural Model 32
Metro-Metro Configuration 32
Metro-Global Configuration 33
Common Architectural Elements 34
Configuration Characteristics 34
Switching Behavior 34
Unplanned (Unscheduled) Switch 35

Commands **36**

ADDCSECRDE - Add CSE Credential Entry *36*

CFGSTRPRSC – Configure Startup Resources *36*

CHGBLDFLG – Change Build Flags *37*

CHGBRMSOBJ - Change BRMS Object Attributes *38*

CHGCLDVOL - Change Cloud Volume *41*

CHGCSECRDE - Change CSE Credential Entry *41*

CHGCSEDTA – Change CSE Data *41*

CHKCSE – Check Copy Services Environment *42*

CHKFSR3 – Check FSR 3-site *42*

CLEANLOGS – Clean Toolkit Logs *43*

CLRDDD – Clear Device Data Domain *44*

CPYCSEDTA – Copy CSE Data *44*

CRTCSEDTA - Create Copy Services Environment Data *44*

DLTCLDCS - Delete Cloud Consistency Set *47*

DLTCSEDTA - Delete CSE Data *47*

DSPCSEDTA - Display Copy Services Data *47*

DSPDDDLCK - Display DDD lock holder *47*

IPLLPAR – IPL Partition *48*

LOGMSG – Log Message *49*

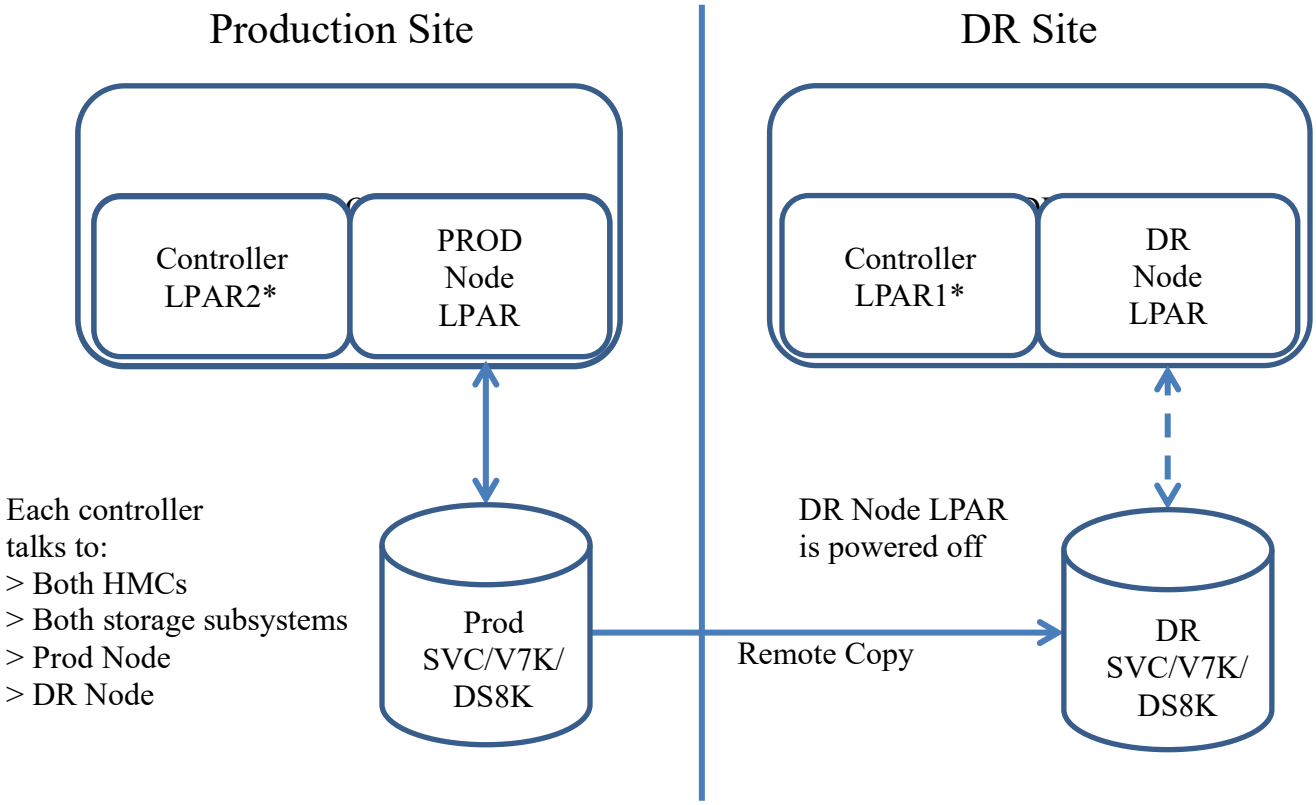
PRTCSE – Print CSE Information *50*

RLSCSELCK - Release Copy Services Lock *51*
RMVCSECRDE - Remove CSE Credential Entry *51*
RSTDDD – Restore Toolkit Device Data Domains *51*
RSTJOBSCDE – Restore Job Scheduler Entries *51*
RTVINF - Retrieve ICSM Information *52*
RTVSTRPRSC – Retrieve Startup Resources *59*
RUNDSCMD - Run DS Scripted Command *59*
RUNLPARCMD - Run command based on LPAR/SRLN *61*
RUNSVCCMD - Run SVC Command *62*
SAVDDD – Save Toolkit Device Data Domains *62*
SETCSELCK - Set Copy Services Lock *62*
SETDDDLCK - Set DDD lock holder *63*
SETUPFSR - Set up IBM Pwr HA tools – FSR *63*
STRCLDVSI - Start Cloud Power Virtual Machine *64*
SWCSE – Perform Full System Switch *64*
SWFSR3 – Switch FSR 3-site *65*
TSTCOMM – Test Communication Paths *66*
VIEWLOG - View Log File *67*
WRKCLDVOL - Work with Cloud Volumes *68*
WRKCSE - Work with Copy Services Environment *70*
WRKCSECRDL - Work with CSE Credential List *70*
WRKCSEDTA -Work with IBM i CSE Data *70*
WRKSTRPRSC -Work with Startup Resources *71*
ZAPDDDLCK - Zap DDD lock *73*

6

Introduction

The following picture shows an overview of the PowerHA Tools for IBM i - FSR relationships between various entities involved in a HA and/or DR architecture for an IBM i environment.



*It is recommended to have 2 control LPARs. If using only 1, then it should be at the DR site.

Note: Any reference to SVC in this document applies equally to V7000, V5000 and V3700.

What's new in 4.1

FSR 4.1 adds support for Full System Replication on the DS8000 storage subsystem using Metro or Global Mirror.

What's new in 4.2

FSR 4.2 primarily adds support for Full System FlashCopy, which is licensed separately. FSR benefits from several new features which were added for FSFC:

- All fixes to 4.1
- WRKCSEDTA - new command
- WRKCSECRDL enhancements:
    - Stores DSCLI credentials
        - No longer restricted to user QLPAR and password files
    - Option 2 displays the user profile
    - List is properly populated the first time
    - Password mismatch error message was incorrect
- WRKSTRPRSC enhancements:
    - Allow *VIRTUAL IP interfaces
    - Allow entries with duplicate LOCCODE
    - Data area LASTCONFIG removed when changes are made in WRKSTRPRSC
- FIXSTRPRSC enhancements:
    - Data area LASTCONFIG only created after successful completion of configuration
- CHKCSE / CHKPPRC enhancements:
    - Fails if QSYSOPR message queue delivery is not *HOLD
    - Does not crash QZRDIAEXT2
- Changed the product access codes:
    - 80 - FSR for Storwize
    - 82 - FSR+FSFC version 1 for Storwize
    - 87 - FSR+FSFC version 2 for Storwize
    - 88 - FSR+FSFC version 2 for DS8K
    - 89 - FSR for DS8K
- Source and target LPAR's are always IPL'd in B-mode.
    - Detach will always be in manual mode
    - Switch will use what is defined by the HMC (i.e. *PANEL)
- CHGASPACT(*FRCWRT) performed prior to detach
- WRKCSE enhancements:
    - Opt 10 on GMCV can fail - now fixed
    - DS8K MMIR errors have better messages
    - DS8K detach occasionally tried to re-attach
    - Changing GMIR environment handles prompting better
- SWCSE *UNSCHEDULED in batch will send inquiry message to QSYSOPR

What's new in 4.3

Version 4.3 was released with the intent of providing new enhancements for Full System Flash Copy. One of those enhancements was to remove the limit of how many CSE Data's were allowed, and that also benefits FSR. This was accomplished by re-structuring the PowerHA space where the FSR data was stored, therefore, all toolkit installations in the same PowerHA cluster must be at the same release level, and upgrading to version 4.3 requires recording, removing and recreating all the CSE Data's.

There are no other significant changes to 4.3 relevant for FSR. It is still recommended that customers upgrade to 4.3 as defects and new features will be delivered in version 4.3.

What's new in 4.4

Version 4.4 was primarily a release to add new Full System Flashcopy features but some of those features also benefitted Full System Replication. Primarily it increased the number of environments allowed in WRKCSE to nearly 1000. Most customers will skip this release and go directly to 4.5.

What's new in 4.5

Version brought many new features to Full System Replication and Flashcopy, including:

- Full System LUN Switch
- DS8K environments are changed to only support Host name and volume mappings. Volume groups are no longer supported.
- Support for cascading Flashcopies on DS8K. This means checking for multiple flashcopies at the target of GMIR replication prior to a switch have been disabled, as they are now allowed and will not prevent a switch.
- CSM Support, including dual-control and unique users
- Better FSR cleanup on failed detach
- Warn if IPL attribute set to start in restricted state
- CHKCSE verifies RCCG replicates the right volumes
- Support for aggregate communication lines
- Detach will invoke *FRCWRT on the source LPAR
- Better cleanup if the detach fails
- Prompt (F4) on the shutdown command in CSEDTA
- CHKCSE will warn if qzrdhasm.log is greater than 20 MB
- After a switch, replication is not restarted until the LPAR has finished it's IPL.
- IASP DS8K PPRC license sufficient to create full system replication environments (but no automation)

What's new in 4.6

Version 4.6 introduced PowerVS support for FSFC, but there were some enhancements for FSR:

- Command TSTCOMM added
- WRKCSEDTA filtering, position-to fields, and opt 15 = display log with auto refresh.
- Full System LUN Switch is supported for SVC
- WRKCSECRDL Validation (option 6)
- WRKCSEDTA Option 15 – Auto-refresh log
- Check that QLPAR on the source is neither disabled nor expired and has *ALLOBJ, *SECADM and *IOSYSCFG authorities
- SAVDDD gets new parameter *GEN
- SAVDDD and RSTDDD have more log entries

- SETUPFSR includes parameters for access codes.
- WRKCSEDTA F7 = WRKACTJOB SBS(QZRDFSR)
- Support for BRMS restricted state IP addresses
- WRKSTRPRSC has one port per location code for SRVIOS compatibility
- WRKSTRPRSC can resolve port numbers automatically
- *SGxx entries in WRKSTRPRSC

What's new in 5.1

The big announcement in 5.1 for FSR is compatibility with PowerVS Global Replication Service, which includes:

- WRKCLDVOL and DLTCLDVOL to work with Cloud volumes
- Managing iSCSI interfaces with WRKSTRPRSC
- STRCLDPVM for starting Cloud virtual machines
- PowerVS licensing is by subscription

Other improvements include:

- SETUPFSR migration parameter
- CHKCSE checks that STRTCP is *NO

What's new in 5.2

Version 5.2 added support for FlashSystem Policy-Based Global Mirror Replication (PCYGM).

Other improvements include:

- Support for FLASH and PPRC environments with the same name
- New command RTVSTRPRSC
- CHKCSE updates for LUN switching
- Confirmation panels on detach/re-attach
- Miscellaneous bug fixes
- References to PVM changed to VSI

What's new in 5.3

Version 5.3 added support for FSR 3-site using CSM over DS8K.

Other improvements include:

- HMC communications use REST APIs instead of ssh

- The Jsch open-source jar file is no longer required (this will break RUNSVCCMD)
- WRKSTRPRSC:
    - *CMN: Restricted BRMS interfaces
    - *RTE: Duplicate Priority Route
- WRKCSE opt 16 host names can be up to 63 characters

Planning

Planning for FSR installation is part of the services engagement associated with purchasing this product. This includes ensuring that all requirements/restrictions are followed. An overview of the Requirements and Restrictions is included below.

Refer to the online resource for current toolkit and OS version support matrix:

https://www.ibm.com/support/pages/node/6560768

Requirements

Prior to the start of installation, the services representative must ensure the following requirements have been met. As this information can frequently change, please refer to the requirements published on our websites:

http://ibm.biz/PowerHAToolsCompatibility

http://ibm.biz/FSRManager

Planning Considerations

FSR Subsystem Considerations

FSR jobs are submitted by IBM i clustering to the *JOBQ specified by the *JOBD associated with the **QLPAR** user profile. The default *JOBD is QDFTJOBD, and its default *JOBQ is QBATCH. This *JOBQ must be configured to allow multiple clustering jobs to run without any delay. Two jobs are submitted for each Copy Services Environment (CSE).

*Tip: Create a new *JOBD for clustering that specifies Job Queue QSYSNOMAX. The JOBD parameter of the QLPAR User Profile must be changed to use this new Job Description. This will ensure that all clustering jobs can run as required in the QSYSWRK subsystem.*

Running FSR in another subsystem

If running in QBATCH or QSYSNOMAX is not a valid option, then it is possible for the environment to be changed so that the jobs run in an independent subsystem. The standard work management objects must be created: subsystem description, job description, job queue and class. Change user profile qlpar to use these objects.

Clustering Security Requirements

*Note: Some of the TCP/IP servers used by clustering require that the QUSER user profile's STATUS = *ENABLED and that it does NOT have *SECADM or *ALLOBJ special authority. It must also NOT be expired. If this is not possible, the file /QIBM/ProdData/OS400/INETD/inetd.conf must be changed to use a different profile that matches these requirements.*

```
Edit File: /QIBM/ProdData/OS400/INETD/inetd.conf
 Record :      10   of      20 by  10                    Column :   1     76 by 126
 Control : ____________________________________________________________          .

CMD ....+....1....+....2....+....3....+....4....+....5....+....6....+....7....+....8..
    # ##### DO NOT MODIFY THIS FILE #####
    #
    # Any changes made to this file will be lost during release upgrade.
    # User-defined services may be defined in file:
    #  /QIBM/UserData/OS400/INETD/inetd.conf
    #
    #
    # Clustering
    as400-cluster  stream  tcp  wait    QUSER  /QSYS.LIB/QCSTINETD.PGM QCSTINETD
    #
     ************End of Data*********************
```

The ALWADDCLU (Allow Add to Cluster) network attribute must be appropriately set on the target node if trying to start a remote node. This should be set to *ANY or *RQSAUT depending on the environment. If set to *RQSAUT, then -- Digital Certificate Manager (57xxSS1 Option 34) on IBM i 6.1, or later release must be installed.

To change the ALWADDCLU (Allow Add to Cluster) network attribute, use the following green screen command:

CHGNETA (Change Network Attributes)
Specify **ALWADDCLU = *ANY** or ***RQSAUT**

Note: This applies only to the control nodes.

Requirement to coordinate QTIME between managing partitions

To prevent simultaneous operations on the same environment, a cluster wide lock per environment has been added, set at the start of SWPPRC, SWCSE, CHKPPRC, CHKCSE and STRFLASH operations and released at the end. The default time for automatic release of the lock is 15 minutes, and the time-of-day for the timeout is calculated and communicated to the other nodes in the cluster. Therefore, it is preferable to make sure that the QTIME system value on all systems contain the same time-of-day. You should consider use of the Simple Network Time Protocol (SNTP) TCP server to automate keeping the time synchronized.

Note: If systems are in different time zones or the system times are significantly different, then this locking will work only on the local system.

Migrating from Replication Consistency Groups to Policy-Based Replication

An Expert Labs services engagement is highly recommended for migrating from RCCG to PCYGM replication. If there is sufficient bandwidth and storage capacity available, the duration of an exposure (i.e. no DR available) can be eliminated.

Setting up Full System Replication

Considerations

At setup time, the production system will be running, and all resources required for that partition will be known to the system. However some information about DR node resources may not be known until the DR node is activated for the first time. The IP addresses for both the source and target should be known prior to starting the setup.

The Controller partitions do not need to be dedicated to this purpose. They can be any other partition available, i.e., dev, QA, etc. They need to be able to communicate with the production partitions they will manage via TCP with the Full System Replication tools.

System name vs. LPAR name

The system name is stored in *SYSBAS via the CHGNETA command. Thus when the DR node is active, it sees its system name as that of the Production node. The LPAR name is stored in the hosting HMC.

System serial number and software licensing

The serial number is stored in the hardware so it will be different between the Production and DR nodes. Licenses for 3rd party software that depend upon a specific serial number will need to be adjusted after a switch. The new RTVLPARINF command can be used to determine the current LPAR for the system. The new RUNLPARCMD command can also be used in the startup program to make changes. Licenses for FSR allow multiple serial numbers and will not require adjustment. You should check with the vendors of other licensed software to determine how to install serial-number-based licenses for their products.

IP addresses

The IP addresses and subnetwork of a node may need to change as part of a switchover. The IP addresses for when the partition is active at the production site and at the DR site must be configured on the node. The replication tool will modify the “online at IPL” attributes of these addresses as part of each switch.

Clustering

One or two control nodes are required in a cluster.

The cluster function allows the CSE data to be mirrored between the control nodes so that either control node can run functions at any time.

Production LPAR Communication Agent

The Controller(s) must be able to communicate with the Production LPARs. To do so, subsystem QZRDFSR must be started on the Production LPAR, with pre-start program QZRDIAEXT2 handling service requests from the Controllers. The subsystem can be started by using command STRSBS /QZRDHASM/QZRDFSR . Use of the WRKSTRPRSC or FIXSTRPRSC on the production node will also automatically start the subsystem

The Communication Agent listens on a specified port, and default port is 55920. On the Production LPAR, use the command SETUPFSR to change the port. This information is stored in file /QIBM/Qzrdhasm/FSRPort. On the Controlling LPAR, the port is configured using CRTCSEDTA or CHGCSEDTA. The port specified must be the same on both the Controller and Production LPAR. Sensitive communications between the Production and Controller are encrypted.

Credentials

The Controller(s) must be able to communicate with all participants in the Full System Replication environment:

1. Production Power System HMC
2. DR Power System HMC
3. Production SVC/DS HMC
4. DR SVC/DS HMC
5. Production node
6. DR node

Configuring the user profiles and passwords for the SVCs, DS HMCs and Power Systems HMCs in the configuration is required as part of setup.

LPAR date and time

The date and time of an LPAR is stored in the HMC with offsets stored in IBM i *SYSBAS. Since only *SYSBAS, not the HMC data, is replicated to the DR site, the first time an LPAR is IPLed the date and time must be configured. Subsequent IPLs will retain the correct date/time.

Communications ports

The following communications ports are used by FSR Full System Replication and access must be allowed through the firewall:

- Managing partitions to SVC on port 22 (if applicable)
- Managing partitions to DS HMC on ports 1751/1750 (if applicable)
- Managing partition to Power System HMC on port 22
- DR and Production Controllers to Production and DR LPARs on port 55920 (or the port you specify)

PowerVS Setup

When using FSR in a PowerVS environment there are number of differences which need to be considered. Access to the VM's and the storage is indirect via REST API's. The configuration is performed using the PowerVS web gui (https//cloud.ibm.com) and specific toolkit commands. Once configured, detaches and switches occur in the same manner as on-prem operations.

It is expected what the client or cloud migration team will create and install the production and controller VM's (the DR VM need only be deployed). Also, a network for communication between the VMs should be configured.

Once these have been created, the setup process is:

1. Create or obtain API KEYs owned by the client
2. Log into DR VM to retrieve:
    a. LOCCOD's for public and private IP interfaces
    b. Serial number
3. Shut down the DR VM
4. Log into the controlling LPARs:
    a. Create the PowerHA cluster and device data domain
    b. Restore toolkit, SETUPFSR, ADDPRDACS
    c. Add credentials (WRKCSECRDL *CLOUD)
        i. Enter the API KEYs
    d. Start volume replication using WRKCLDVOL
        i. Toggle replication on production volumes
        ii. Create a consistency set (F7)
        iii. Add production volumes to the new consistency set
        iv. Rename auxiliary volumes
5. Attach new volumes to DR VM
    a. Mark one as boot
    b. Un-mark original deployment volume to not be boot
    c. Detach original deployment volume
    d. Delete original deployment volume
6. Create the Copy Services Environment (WRKCSE opt 1)
7. Create the CSE Data (WRKCSEDTA opt 1)
8. Log into production VM and:
    a. Create the IP interfaces for the DR VM (CFGTCP opt 1)
    b. Install toolkit (restore, SETUPFSR, WRKSTRPRSC)
    c. Modify startup program
    d. CHGIPLA STRTCP(*NO)
    e. CLEANLOGS
9. CHKCSE
10. WRKCSE opt 12
    a. Detach
    b. Reattach
    c. Switch to DR, don't start replication

d. Start replication after switch
e. Switch back to production

It will be noted through this manual where PowerVS varies from the standard on-prem installation.

Performing switches with Full System Replication

The process of switching between Production and DR nodes is performed on either the Production or DR Controller, and the first step is to verify that all is ready for a switch.

Enter command Check Copy Services Environ. (CHKCSE) and press F4=Prompt.

```
                    Check Copy Services Environ. (CHKCSE)

Type choices, press Enter.

Environment name . . . . . .   __________    Name
```

Key the environment name and press Enter. The following conditions are checked to determine readiness for a switch:

- The control node can communicate to the Production node.
- The environment must be *PPRC only (no FlashCopy)
- At least one of the HMCs for the Production LPAR must be available
- At least one of the HMCs for the DR LPAR must be available
- CSE data must show USE(*SYSTEM)
- Both storage units are available
- The *PPRC direction in the CSE data must match the replication direction
- QSYSOPR message queue delivery must be *HOLD

Scheduled switch

System maintenance on the Production node may require a planned switch to the DR node. Enter command Switch Copy Services Environ. (SWCSE) and press F4=Prompt.

```
                    Switch Copy Services Environ. (SWCSE)

Type choices, press Enter.

Environment name . . . . . .   __________    Name
Switch type  . . . . . .       *SCHEDULED    *SCHEDULED, *UNSCHEDULED
Type . . . . . .               *             *, *GMIR, *MMIR
Auto replicate . . . . . .     *DFT          *DFT, *YES, *NO
```

Key the environment name and type and Press Enter. The following steps are performed:

- If the Production node is running, send an *INQ message (IAS0029) to QSYSOPR requesting permission to proceed
- Call the customer-specified power down command
  - *STATUS message “Current SRC for <Production node LPAR> is nnnnnnnn” indicates progress toward shutdown
- Wait until the Production node LPAR is deactivated
- Pause replication
  - *STATUS message “Starting switchover of node <Production node> from LPAR <Production node> to LPAR <DR node>
- IPL the DR node
  - *STATUS message “Current SRC for <DR node LPAR> is nnnnnnnn” indicates progress of the IPL. Message “The PPRC switch for the *SYSTEM CSE data complete.” is sent to the joblog of the Controller where the SWCSE was entered
  - The LPAR is always IPLd to the B side
- After the DR node has IPL’d start replication if ‘Auto Replicate’ is not *NO
- Call the customer startup program
  - FIXSTRPRSC processes the information for *CMN, STG, *BRM, and *RTE differences entered with WRKSTRPRSC
  - Change license keys for other programs dependent upon the system serial number
  - Other customer-specified processing
- Start clustering; when clustering was set up, IP addresses for both the Production and DR nodes were specified

Unscheduled switch

If the Production node becomes inoperative, an unscheduled switch is necessary to activate the DR node. The only differences in the process described above are:

- CHKCSE cannot be used
- If running in batch, *INQ message is sent to QSYSOPR on the Controlling node requesting permission to proceed.
- No *INQ message is sent to QSYSOPR on the Production node requesting permission to proceed.
- The default is not to start replication automatically

Starting replication after unscheduled switch

To start replication after an unscheduled switch, use WRKCSE Option 12=Work with.
Note: screens show SVC but functionality is identical with DS8K however a Metro Mirror will not show Detach and Reattach options

```
                        Work with SVC PPRC Environment

  Environment  . . . . . . . :   <name>
  MMIR Status  . . . . . . . :   Idling
  Direction  . . . . . . . . :

Select one of the following:
     2. Pause
     3. Resume

     5. Switch
     6. Start Replication after Switch

     8. Detach
     9. Reattach
    10. Display replication
```

Select Option 6 (Start replication after switch). A Warning window opens:

```
.............................................................
:                 Confirm Start of Replication              :
:                                                           :
:  Warning; this option may be hazardous to the health      :
:  data of your production data. It is possible to start    :
:  replication in either direction. BEFORE CONTINUING,      :
:  CONFIRM THE NEW NODE ROLES BELOW.                        :
:                                                           :
:                                                           :
:  Press F10 to continue, F8 to reverse, F12 to cancel.     :
:                                                           :
:    Source LPAR  . :   ctciha4e                            :
:    Target LPAR  . :   ctciha9e                            :
:                                                           :
:                                                  Bottom   :
:  F1=Help        F3=Exit   F8=Reverse the Direction        :
:  F10=Continue   F12=Cancel                                :
:                                                           :
:...........................................................:
```

Be very careful here. The Source and Target LPARs are determined by FSR from the CSE exit data status of *NORMAL or *REVERSED. Be sure to confirm that FSR determined correctly before proceeding. Press F8=Reverse the Direction if you decide otherwise. Press F10 when ready to start replication.

Next, select Option 10 if available (Display replication) and press F11=Show status to monitor progress. For DS8K Metro Mirror, use option 12 (Work with Volumes) and Option 13 (Display out of sync sectors) instead.

```
                          Display Replication

Environment . . . :   <name>             Type  . . . . . :   MMIR
Consistency group :   <name>
Cycle period  . . :   *NONE
Primary . . . . . :   Master (or Aux)
State . . . . . :   Inconsistent_copying

  Relationship    State / in sync?         Freeze time      Progress
  <name>          inconsistent_copying                            96
  <name>          inconsistent_copying                            97
  <name>          inconsistent_copying                            96
  <name>          inconsistent_copying                            97
```

Eventually you will see that replication is running (Consistent synchronized).

```
                        Display Replication

Environment . . . :   <name>              Type  . . . . . :   MMIR
Consistency group :   <name>
Cycle period  . . :   *NONE
Primary . . . . . :   Master (or Aux)
State . . . . . :   Consistent_synchronized

  Relationship    State / in sync?          Freeze time     Progress
  <name>          consistent_synchronized
  <name>          consistent_synchronized
  <name>          consistent_synchronized
  <name>          consistent_synchronized
```

Detach

The detach function stops replication and IPLs the target node in manual mode and restricted state. The source node continues to run. On one of the Controllers, use the WRKCSE command and option 12 (Work with) to access this function.

- This function is not supported for GMIR replication without change volumes for SVC.
- This function is not supported for MMIR replication with DS8K
- This function is supported for PCYGM with FlashSystem with the following caveats:
    - Environment variable QZ_ALLOW_PCYGM_DETACH exists on the job performing the detach
    - During the attended IPL, the user must ensure the LPAR will start to restricted state
    - When a command line is available but before the controlling is started the user must issue these commands on the target LPAR:
        - CHGIPLA STRRSTD(*YES)
        - CHGSYSVAL QSTRUPPGM VALUE(*NONE)
        - CRTDTAARA DTAARA(QZRDHASM/QDETACHED) TYPE(*CHAR)

```
                        Work with SVC PPRC Environment

  Environment  . . . . . . . :   <name>
  MMIR Status  . . . . . . . :   Consistent synchronized
  Direction  . . . . . . . . :   Normal (or Reversed)

Select one of the following:

     2. Pause
     3. Resume

     5. Switch
     6. Start Replication after Switch

     8. Detach
     9. Reattach
    10. Display replication
```

Select option 8 (Detach). The following steps are performed:

- The IPL attributes are set to Restricted=Yes if QLPAR has *SECADM authority
- Data area QZRDHASM/QDETACHED is created
- LPAR memory is flushed
- Replication is stopped
- The target node is IPLed in Manual mode while the production node continues to run
- The IPL attributes are set back to Restricted=No, if necessary
- Data area QZRDHASM/QDETACHED is deleted

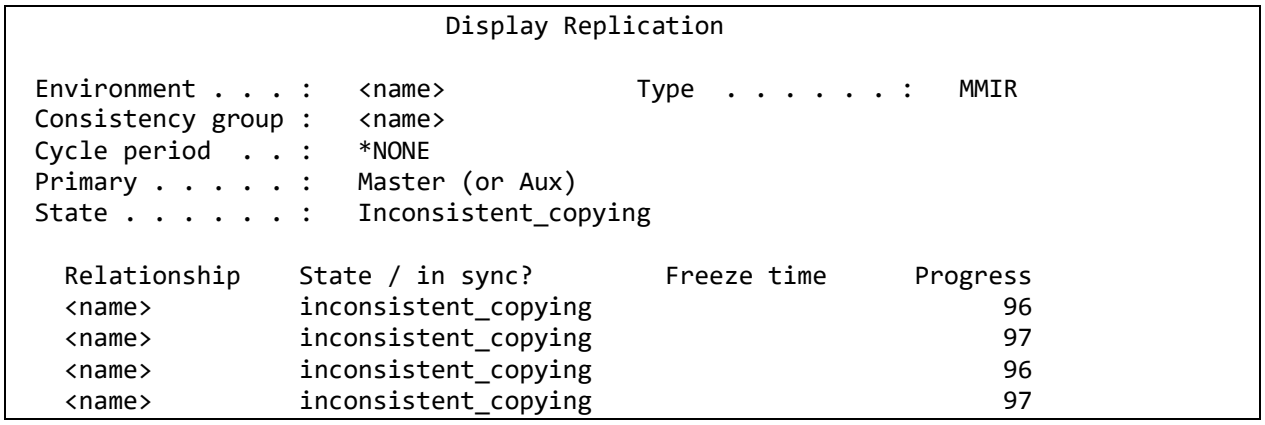
Wait for *STATUS message IAS00E4 (Current SRC for <LPAR name of target node> is C600 4301) indicating the console is available for signon. Sign on the console and provide responses to complete the IPL. Any data changes on the target node will be lost when reattach restarts replication from the Production node.

When the target LPAR has reach SRC 00000000 and is in status “Running” the detach is considered complete.

Reattach

When you are ready to resume normal operations, select option 9 (Reattach). The following steps are performed:

- The target node is powered down, if not already Inactive
- Replication is resumed in the same direction as before the detach. You can monitor progress toward “Consistent_synchronized” by selecting option 10 (Display replication) and pressing F11=Show status.



```
                          Display Replication

Environment . . . :   <name>             Type  . . . . . :   MMIR
Consistency group :   <name>
Cycle period  . . :   *NONE
Primary . . . . . :   Master (or Aux)
State . . . . . :   Inconsistent_copying

  Relationship    State / in sync?         Freeze time      Progress
  <name>          inconsistent_copying                            96
  <name>          inconsistent_copying                            97
  <name>          inconsistent_copying                            96
  <name>          inconsistent_copying                            97
```

Displaying Full System Replication information

WRKCSE Option 5=Display

```
                          Display a PPRC Environment
Press Enter to continue.

Environment . . . . . . . . . . . :   <name>
Type  . . . . . . . . . . . . . . :   MMIR (or GMIR)
Storage Type  . . . . . . . . . . :   SVC
ASP Device name   . . . . . . . . :   *SYSTEM
Source Copy Description . . . . . :   *NONE
Target Copy Description . . . . . :   *NONE
Production node . . . . . . . . . :   <PROD node name>
Preferred source SVC Information:
  IP Address  . . . . . . . . . . :   nnn.nnn.nnn.nnn
  Remote copy consistency group Id    n
Preferred target SVC Information:
  IP Address  . . . . . . . . . . :   nnn.nnn.nnn.nnn
  Remote copy consistency group Id    n
```

WRKCSE Option 12=Work with

```
                       Work with SVC PPRC Environment

  Environment  . . . . . . . . :   <name>
  MMIR Status  . . . . . . . . :   Consistent synchronized
  Direction  . . . . . . . . . :   Normal (or Reversed)

Select one of the following:

     2. Pause
     3. Resume

     5. Switch
     6. Start Replication after Switch

     8. Detach
     9. Reattach
    10. Display replication
```

Note: Direction shows what is actually occurring between the storage units, not what is in the CSE data.

WRKCSE Work with Option 10=Display replication

```
                              Display Replication

Environment . . . :   <name>            Type  :   MMIR (or GMIR)
Consistency group :   <RCCG ID and name>
Cycle period  . . :   *NONE
Primary . . . . . :   Master (or Aux)
State . . . . . . :   Consistent_synchronized

  Relationship    Master volume         Auxiliary volume
  <name>          <volume name>         <volume name>
  <name>          <volume name>         <volume name>
  <name>          <volume name>         <volume name>
  <name>          <volume name>         <volume name>
```

Primary = Master if replicating in the normal direction or Aux if replicating in the reversed direction.

WRKCSEDTA Option 5 = Display

```
                              Work with CSE Data

Type options, press Enter.
  1=Create   2=Change   3=Copy   4=Delete   5=Display

                        CRG
Opt     CSE Data        type       Text

        BRMS9M9P1       FSFC
        CTCHAFS         FSR
        EXIT9M9P1       FSFC
        FSFCDS8K        FSFC
        FSFC9F9G        FSFC
        FSFC9J9K        FSFC
        FSFC9M9N        FSFC       GMCV Flash
        FSFC9M9P1       FSFC
        FSFC9M9P1Q      FSFC
        FSFC9M9P2       FSFC
        FSR9J2          FSR
```

DSPCSEDTA

```
                              Display CSE CRG Data

Press Enter to continue.

CRG Name . . . . . . . . . . . :   <name>
```

```
Use  . . . . . . . . . . . . . . :   *SYSTEM
Copy type  . . . . . . . . . . . :   *PPRC

Environment name . . . . . . . . :   <name>
Preferred production node name :   <PROD node>
FSR communications port  . . . . :   55920
PPRC status  . . . . . . . . . . :   *READY
PPRC direction . . . . . . . . . :   *REVERSED
Controlling node name  . . . . . :   <DR Ctl>
Secondary controlling node name:   <PROD Ctl>

                                                More...
```

Note: Direction is what is recorded in the CSE exit data based upon FSR processing and may not match the replication direction that is actually occurring between the SVCs.

```
                              Display CSE CRG Data

Press Enter to continue.

Preferred production details:
  HMC LPAR name  . . . . . :   <name>
  HMC Profile name . . . . . :   <name>
  HMC managed system . . . . . :   <name>
  Primary HMC IP . . . . . :   nnn.nnn.nnn.nnn
  Secondary HMC IP . . . . . :

Preferred target details:
  HMC LPAR name  . . . . . :   <name>
  HMC Profile name . . . . . :   <name>
  HMC managed system . . . . . :   <name>
  Primary HMC IP . . . . . :   nnn.nnn.nnn.nnn
  Secondary HMC IP . . . . . :

                                                        More...
```

```
                              Display CSE CRG Data

Press Enter to continue.

Power down command . . . . . :   PWRDWNSYS OPTION(*IMMED)

                                                         Bottom
```

Full System LUN Switch

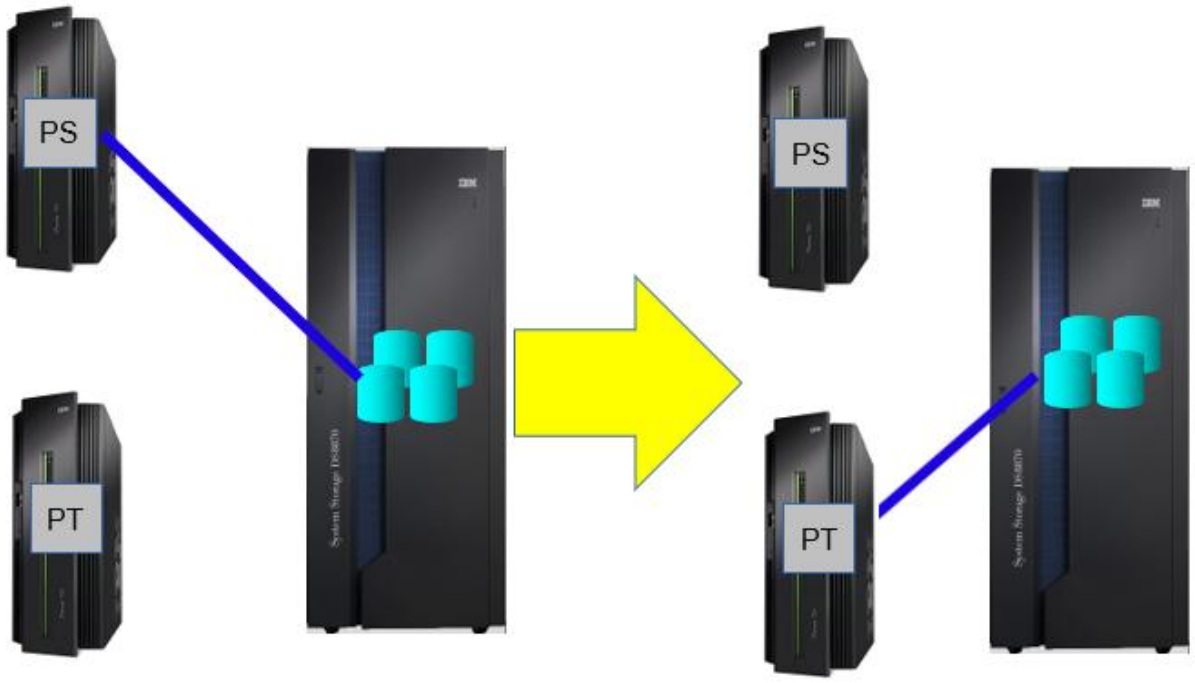
The LUN switch is very similar to Full System Replication, but instead of managing the replication it is only the LPAR host connections which are changed. This can be useful when performing a scheduled or unscheduled switch from one LPAR to another LPAR, and can be combined with Full System Replication.

Compared to Logical Partition Movement (LPM):

- An outage (shutdown and IPL) is required
- An LPAR must exist on each system
- The switch is performed and managed from a controlling LPAR running IBM i
- The switch can be executed with an active or deactivated LPAR (i.e. after a failure)
- Resources (communication, storage, BRMS details etc.) on the LPAR are managed by the toolkit same as with Full System Replication

The Full System LUN Switch can be combined with Full System Replication but there is limited support for this configuration and requires manually updating the WRKCSEDTA prior to switch. This will be better addressed in a future release.

Configuring a Full System LUN Switch



To configure Full System LUN Switch, first create the environment (WRKCSE)

```
                              Add an Environment

Enter Copy Services and ASP information

Environment name . . . . :   TESTLS
Copy Service Type  . . . :   LUN
Storage Type . . . . . . :   SVC or DS8K

Primary ASP  . . . . . . .   *SYSTEM                  33 - 255, *SYSTEM
```

Use option 16 to enter the Production and HA Node host connection information (note that the panels may differ slightly depending on whether they are using Storwize or DS8K units).

Use F6 to import hosts directly from the storage unit:

```
                     Select production and target hosts

  Select 2 hosts, press Enter.
    1=Select Production Host   2=Select HA Host

  Opt     Host name
          ctciha9l
```

```
          ctciha9k
          ctciha9j
   1      ctciha9e
   2      ctciha4p
          ctciha4q
          dummyhost
                                                                  More...
  F1=Help   F12=Cancel
```

```
                          Change Host Volume List

SVC IP Address . . . . . :   9.5.167.87
Production Host name . . . :   ctciha9e
HA Host name . . . . . :   ctciha4p



                                                                  Bottom
Type choices, press Enter

Volume ranges:
  0002-0005  0035-0035









F1=Help   F12=Cancel

Volume Ids imported successfully...press Enter to accept.
```

```
                           Work with Host Mappings

   Environment  . :   LUNSW
   Type . . . . . :   LUN
   SVC IP Address :   9.5.167.87



 Type option, press Enter.
   1=Add   2=Change   4=Remove   6=Change host volume list

           Production Node          HA Node                  Volume
   Opt     Host Name                Host Name                 Count

           ctciha9e                 ctciha4p                      5
```

```
                                                                    Bottom
F1=Help   F3=Exit   F6=Import host   F8=Validate   F12=Cancel
```

Then use WRKCSEDTA to create a Full System Replication entry using the LUN switch environment.

Continue to set up Full System LUN Switch similarly to FSR (WRKSTRPRSC etc).

Performing a Full System LUN Switch

Use the commands CHKCSE and SWCSE to perform the switches, specifying the CSE Data.

The process of switching between Production and DR nodes is performed on either the Production or DR Controller, and the first step is to verify that all is ready for a switch.

Enter command Check Copy Services Environ. (CHKCSE) and press F4=Prompt.

```
                    Check Copy Services Environ. (CHKCSE)

Type choices, press Enter.

Environment name . . . . . .   __________    Name
```

Key the environment name and press Enter. The following conditions are checked to determine readiness for a switch:

- The control node can communicate to the Production node.
- At least one of the HMCs for the Production LPAR must be available
- At least one of the HMCs for the DR LPAR must be available
- CSE data must show USE(*SYSTEM)
- The storage unit is available
- QSYSOPR message queue delivery must be *HOLD

Scheduled switch

System maintenance on the Production node may require a planned switch to the DR node. Enter command Switch Copy Services Environ. (SWCSE) and press F4=Prompt.

```
                    Switch Copy Services Environ. (SWCSE)

Type choices, press Enter.

Environment name . . . . . .   __________    Name
Switch type  . . . . . .       *SCHEDULED    *SCHEDULED, *UNSCHEDULED
Type . . . . . .               *             *, *GMIR, *MMIR
Auto replicate . . . . . .     *DFT          *DFT, *YES, *NO
```

Key the environment name and type and Press Enter. The following steps are performed:

- If the Production node is running, send an *INQ message (IAS0029) to QSYSOPR requesting permission to proceed
- Call the customer-specified power down command
  - *STATUS message “Current SRC for <Production node LPAR> is nnnnnnnn” indicates progress toward shutdown
- Wait until the Production node LPAR is deactivated
- Change the host connections to the DR node LPAR.
- IPL the DR node
  - *STATUS message “Current SRC for <DR node LPAR> is nnnnnnnn” indicates progress of the IPL. Message “The switch for the *SYSTEM is complete.” is sent to the joblog of the Controller where the SWCSE z entered
  - The LPAR is always IPLd to the B side
- Call the customer startup program
  - FIXSTRPRSC processes the information for *CMN, STG, *BRM, and *RTE differences entered with WRKSTRPRSC
  - Change license keys for other programs dependent upon the system serial number
  - Other customer-specified processing
- Start clustering; when clustering was set up, IP addresses for both the Production and DR nodes were specified

Unscheduled switch

If the Production node becomes inoperative, an unscheduled switch is necessary to activate the DR node. The only differences in the process described above are:

- CHKCSE cannot be used
- If running in batch, *INQ message is sent to QSYSOPR on the Controlling node requesting permission to proceed.
- No *INQ message is sent to QSYSOPR on the Production node requesting permission to proceed.

Full System Replication 3-Site Configuration with Copy Services Manager (CSM)

Overview

A **3-site Full System Replication (FSR) configuration with Copy Services Manager (CSM)** provides enhanced disaster recovery by maintaining **two independent target sites** for a single production IBM i system. This configuration allows protection against both local site failures and regional outages while preserving automation, consistency, and operational control through PowerHA Tools for IBM i.

In an FSR 3-site architecture, IBM i Copy Services Manager (CSM) coordinates multiple storage replication relationships so that data changes on the production system are replicated to **two separate disaster recovery (DR) sites** using a combination of **Metro Mirror and Global Mirror** technologies.

While the CSM uses DS8K storage, there is no toolkit interaction with the DS8K.

Benefits of 3-Site FSR

- Protection against both site-level and regional outages
- Lower RPO at the Metro site with long-distance DR coverage
- Automated role changes through PowerHA Tools commands
- Centralized replication management via CSM
- Flexible recovery options

Operational Considerations

- Detach operations may be restricted depending on replication type.
  Specifically, when replicated is cascaded then detach is not possible.
- Only **one active production LPAR** is supported at any time.
- CHKFSR3 validates replication consistency across all targets.
- Startup resource differences (WRKSTRPRSC) must be defined for **each target site**.
- Licensing and serial-number-based applications must be validated for all sites.

Supported Storage and Requirements

3-site FSR with CSM requires:

- IBM DS8000 with CSM support
- PPRC licenses (Metro Mirror and/or Global Mirror, depending on topology)
- PowerHA Tools for IBM i version supporting CSM
- One or (ideally) two controlling LPARs
- IBM PowerHA for i (5770-HAS) on the controlling LPARs
- Correctly configured Copy Services environments for MMIR and optionally GMIR
- Consistent volume mappings across all sites
- Journal/Change volumes for GMIR operations

Replication Flow

The replication flow in a 3-site FSR environment is as follows:

1. All production I/O is written to the **source volumes** at the production site.
2. Data is synchronously replicated to the **Metro Mirror target** (secondary site).
3. Data is replicated to the **Metro** or **Global Mirror target** (tertiary site).
4. Copy Services Manager monitors the health, state, and direction of both relationships.
5. PowerHA Tools for IBM i uses CSM state information during switches and recovery operations.

This design allows very low RPO at the Metro site while providing geographical separation via the Global site.

Role of Copy Services Manager (CSM)

CSM is a required component for 3-site FSR configurations and performs the following functions:

- Coordinates **multi-target replication relationships**
- Maintains replication consistency across all volumes
- Controls replication direction during role changes
- Reports replication state information to PowerHA Tools
- Prevents conflicting operations across replication paths

FSR relies on CSM-provided status rather than directly managing individual PPRC relationships when CSM replication is enabled.

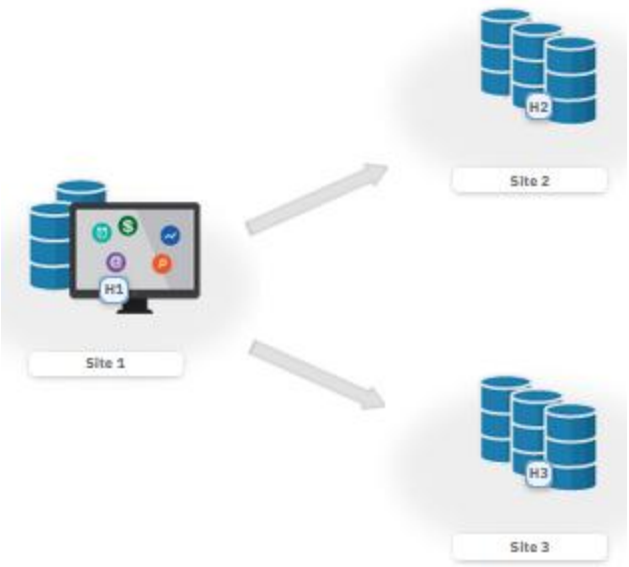
Architectural Model

A 3-site Full System Replication (FSR) configuration with Copy Services Manager (CSM) supports multiple replication topologies to address different availability and distance requirements. The architecture consists of three IBM i sites, each with its own storage subsystem, with replication relationships coordinated by CSM.

CSM manages replication in a **multi-target topology**, ensuring that the same source volumes are replicated to both target sites while maintaining write-order consistency.

The following 3-site models are supported:

Metro-Metro Configuration

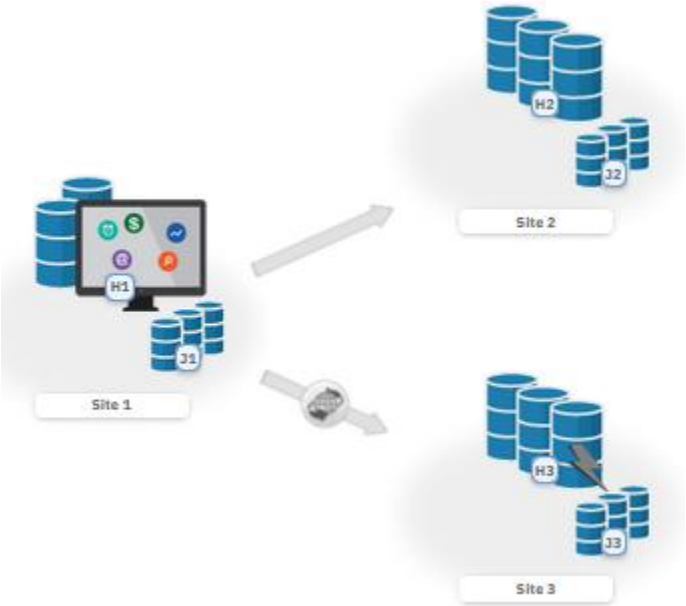


In a **metro-metro** configuration:

- The **Primary Site (Production)** replicates synchronously to **two metro-distance target sites**
- Both replication relationships use **Metro Mirror (MMIR)**
- CSM manages a **multi-target synchronous replication topology**
- Either metro site can be activated during a failover or switchover
- Provides the **lowest possible RPO**, typically zero, for both target sites

This configuration is designed for high availability across multiple nearby data centers and enables protection against the loss of an entire metro site while maintaining synchronous data consistency.

Metro-Global Configuration



In a **metro-global** configuration:

- The **Primary Site (Production)** replicates synchronously to a **Metro Site** using **Metro Mirror**
- The same production data is replicated asynchronously to a **Global Site** using **Global Mirror**
- CSM coordinates the mixed synchronous and asynchronous relationships
- The Metro Site is typically the first choice for recovery, with the Global Site providing protection against regional outages

This configuration balances **low RPO and rapid recovery** at the metro site with **geographic separation** at the global site.

Common Architectural Elements

Regardless of topology, all 3-site FSR configurations with CSM share the following characteristics:

- A single active **production IBM i LPAR**
- Two independent **target IBM i LPARs**
- CSM-controlled multi-target replication
- Three separate FSR environments (CSEs) for each replication relationship
    - MMIR, MMIR2 and MMIR3 or
    - MMIR, GMIR and GMIR2
- A common configuration profile (CSE Data)
- Only **one target LPAR is activated at a time**
- PowerHA Tools for IBM i orchestrates system-level switching using replication state provided by CSM

Configuration Characteristics

When using 3-site FSR with CSM, the following characteristics apply:

- Each node is labeled H1, H2 or H3
- For Metro-Metro architectures:
    - Three separate **FSR environments** are defined:
        - One Metro Mirror environment (MMIR) between H1 and H2
        - One Metro Mirror environment (MMIR2) between H1 and H3
        - One Metro Mirror environment (MMIR3) between H2 and H3
- For Metro-Global architectures:
    - Three separate **FSR environments** are defined:
        - One Metro Mirror environment (MMIR) between H1 and H2
        - One Global Mirror environment (GMIR) between H1 and H3
        - One Global Mirror environment (GMIR2) between H2 and H3
- All environments reference the **same production LPAR**
- Each environment has its own CSE definition
- One CSEDTA exists, which details the H1, H2 and H3 HMC and IP information.
- Only **one target is activated at a time** except for detach operations.
- Replication to remaining targets is paused or reversed as required during a switch

Switching Behavior

Switches are performed using the command SWFSR3.

Planned (Scheduled) Switch

During a scheduled switch:

1. The administrator determines which DR site becomes active.
2. FSR verifies availability of controllers, HMCs, and storage systems.
3. A confirmation panel is presented to verify the switch can proceed
4. Production LPAR is shut down in a controlled manner.
    1. Alternately, the user can specify a manual shutdown is required.
5. PPRC relationship is paused or reversed.
6. Selected target LPAR is IPLed.
7. Replication can be automatically restarted based on configuration.
8. The unused site remains dormant but consistent.

Unplanned (Unscheduled) Switch

In an unscheduled outage:

- The administrator determines which DR site becomes active.
- Replication is **not auto-started** by default
- Manual confirmation is required before resuming replication

Commands

ADDCSECRDE - Add CSE Credential Entry

ADDCSECRDE allows the user to add credential entries. This information is used by the toolkit whenever it needs to communicate with the device specified on the Host IP address parameter, and is used when communicating with the DS8K, HMC and SVC's. This information is encrypted and stored in the PowerHA device data domain, and is available to all the nodes in the cluster.

```
                    Add CSE Credential Entry (ADDCSECRDE)

 Type choices, press Enter.

 Host IP address  . . . . . .  _________________
 User ID  . . . . . .  _________________
 Password . . . . . .
 Confirm password . . . . . .
 Host description . . . . . .  _________________
```

Host IP address: Enter the IP address of the host which the credentials are for. A valid IP address is required; name resolution is not performed by the toolkit.

User ID: Enter the user ID which exists on the device specified by the Host IP address.

Password: Enter the password for the user ID which exists on the device specified by the Host IP address.

Confirm password: Re-enter the password to ensure that it is correct.

Host description: Enter a description to make it easy for humans to see which host the IP address is for.

CFGSTRPRSC – Configure Startup Resources

This command is to be placed at the beginning of the client startup program on the production LPAR. It will use the information in WRKSTRPRSC to determine the role (usage) of the LPAR at IPL, and then enact the changes specified in WRKSTRPRSC. When finished, it will return the LPAR usage as a variable.

```
                    Configure Startup Resources (CFGSTRPRSC)

 Type choices, press Enter.
```

```
 Type of IPL  . . . . . .                  Variable
```

Type of IPL: *CHAR 5 variable returned which contained the usage (*PS, *PT, *SGxx etc).

This is how it is used in the startup program:

```
DCL VAR(&IPLTYPE) TYPE(*CHAR) LEN(5) VALUE('?????')

/* BEGIN POWERHA TOOLS FOR IBM I - FSFC */
QZRDHASM/LOGMSG MSG('STARTING CFGSTRPRSC')
QZRDHASM/CFGSTRPRSC RTNVAR(&IPLTYPE)
MONMSG     MSGID(CPF0000)
IF COND(%SST(&IPLTYPE 1 2) *NE '*P' *AND +
        %SST(&IPLTYPE 1 2) *NE '*H') +
   THEN(RETURN)
STRTCP
MONMSG     MSGID(CPF0000)
/* END POWERHA TOOLS FOR IBM I - FSFC */
```

Note that if the program detects it is a flashcopy or detach the startup program will be terminated.

Starting TCP should be delayed until after CFGSTRPRSC has run, which also means CHGIPLA should be set to STRTCP(*NO).

CHGBLDFLG – Change Build Flags

This command alters flags held in the BUILD data area that are used for debugging and other purposes within PowerHA Toolkit for IBM i. The command should only be used when so recommended by an IBM consultant or support personnel.

```
                     Change ICSM Build Flag (CHGBLDFLG)

 Type choices, press Enter.

 Specific build flag  . . . . . .   *DEBUGEXIT    *DEBUGEXIT, *DEBUGFLASH...
 Setting  . . . . . .   *OFF          *ON, *OFF
```

Specific build flag: Determines which flag to toggle. The possible values are:

- ***DEBUGCRG:** Change the DEBUGCRG flag used with the toolkit. Setting it on will result in more detailed messages related to the CRG DDD space.
- ***DEBUGEXIT**: Change the DEBUGEXIT flag used with the toolkit for debugging the CRG exit program.

- ***DEBUGFLASH:** Change the DEBUGFLASH flag used with the toolkit. Setting it on will result in more detailed messages related to flash operations.

- ***DEBUGTPC:** Change the DEBUGTPC flag used within the toolkit for debugging CSM (formerly TPC-R).

- ***DEBUGDDD:** Change the DEBUGDDD flag used within the toolkit. Setting it on will result in more detailed messages related to the DDD space.

- ***TEST:** Change the TEST flag used within the toolkit. This should only be used for toolkit development testing.

- ***XTRALOGS:** Change the XTRALOGS flag used within the toolkit. This will result in additional logging messages generated and retained.

Setting: Used to toggle the specified build flag. The possible values are:

- ***OFF:** Change the above flag to off.

- ***ON:** Change the above flag to on.

CHGBRMSOBJ - Change BRMS Object Attributes

This command allows you to change an attribute of a BRMS object. It is particularly useful for programmatically changing values which do not have a BRMS command line interface.

```
                  Change BRMS Object Attributes (CHGBRMSOBJ)

 Type choices, press Enter.

 BRMS Object type to change . . .  __________     *DEVICE, *MEDPCY, *CTLGATTR
 Name of BRMS object to change  .  __________     Object name
 BRMS Attribute to change . . . .  __________     *LOC, *MEDCLS, *MOVPCY...
 New attribute value  . . . . . .
_____________________________________________
_____________________________________________________________________________
_______
```

Object type to change: Specify the type of object to change. The objects which can be changed are:

- ***CTLGATTR:** The following control group attributes (WRKCTLGBRM opt 8) can be changed:
    - *DEVICE

- ***DEVICE:** The following device attributes (WRKDEVBRM) can be changed:

- *LOC
- *TEXT

- ***MEDPCY:** The following media policy attributes (WRKPCYBRM) can be changed:
  - *LOC
  - *MEDCLS
  - *MOVPCY
  - *MARKDUP
  - *MARKHST
  - *MINVOL
  - *TEXT
  - *VOLSEC

Name of BRMS object to change: Specify the name of the BRMS object to change.

BRMS Attribute to change: Specify the attribute of the object to change. The valid combination of object to change and the attribute to change varies.

- ***DEVICE:** Specifies the device(s) to be associated with the object. Only one device is currently supported. The possible values are:
  - *BKUPCY**:** The value for the backup device field in this control group uses the default value for this field from the backup policy.
  - *SYSPCY**:** The value for the backup device field in this control group uses the default value for this field from the system policy.
  - device-name: Specify the names of the devices that are used in processing this backup control group. Only one device is currently supported.
  - *NONE: There is no device for this save operation. Save files are used to store the saved data.
  - *MEDCLS: Devices for this policy or control group are selected based on device types that support the density for the media class specified in the media policy. The *MEDCLS special value is used for devices that are part of a device pool, such as several systems that share a single or set of devices.

- ***LOC:** Specifies the storage location to be associated with the object. The possible values are:
  - *ANY: Any device in any location capable of saving to the specified media class can be utilized for the save operation.
  - location-name: Specifies the location name of the device that is to be used with this policy. Only devices assigned this location name can be used for the save operation.

- ***MEDCLS:** Specifies the name of the media class to be used for any volume created using this media policy.The possible values are:
  - *NONE: No media class is assigned for this media policy.
  - *ADSM: No media class is assigned for this media policy. The media is managed by the TSM server.

- media-class-name: Specify the name of the media class that is to be used with this media policy.

- ***MOVPCY:** Specifies the name of the move policy to be associated with this object. The possible values are:
  - *NONE: Specifies that no move policy is associated with this media policy.
  - *ADSM: Specifies no move policy is associated with this media policy. A TSM server is responsible for movement and storage of the media.
  - move-policy-name: Specifies the name of the move policy that is to be assigned to volumes used by this media policy.

- ***MARKDUP:** Specifies whether history items created by the save will be marked for duplication. You can use DUPMEDBRM VOL(*SRCHHST) command to duplicate the marked saved items. The possible values are:
  - *NO: History items created during a save operation that use this media policy will not be marked for duplication.
  - *YES: History items created during a save operation that use this media policy will be marked for duplication.

- ***MARKHST:**Specifies whether media volumes will be marked for duplication. If a volume is a member of a media set and it is marked for duplication, all members of the media set will be marked for duplication. When the DUPMEDBRM VOL(*SEARCH) command is used, active volumes that have been marked will be selected and duplicated. The possible values are:
  - *NO: Volumes written to during a save operation that use this media policy will not be marked for duplication.
  - *YES: Volumes written to during a save operation that use this media policy will be marked for duplication.

- ***MINVOL:** Specifies the minimum number of expired volumes that must be present before any save can begin. The possible values are:
  - *NONE: There is no check done to determine the minimum number of required volumes before a save operation begins.
  - number-of-required-volumes: Specify the number of expired media volumes that must be available before any BRMS save operation will begin. The number of volumes can range from 1 to 9,999.
- ***TEXT:** Specifies text that describes the object. The possible values are:
  - *NONE: No text will be used to describe the object.
  - text-description: Specifies the text description to be used.

- ***VOLSEC:** The secure volume attribute will be changed. Value values are:
  - *NO: Volume security will not be applied to this media class. Volumes that do not have volume security can be read by anyone.
  - *YES: Only users with special authorities *ALLOBJ or *SECADM can read media volumes in this media class.
  - *ADSM: Volume security is managed by the TSM server.

New attribute value: Specify the new value value.

CHGCLDVOL - Change Cloud Volume

This command allows you to change an attributes of cloud volumes.

```
                                        Change Cloud Volumes (CHGCLDVOL)

Type choices, press Enter.

API Key Name . . . . . .  _________________
Volume ID  . . . . . .  _____________________________________________
               + for more values  _____________________________________________

Action to perform  . . . . . .  _________ *STOPREPL
```

API key name: Specify the key name (from WRKCSECRDL *CLOUD) containing the cloud instance id and PowerVS credentials.

Volume ID: Specify the IDs of the volumes to be changed.

Action to perform: Specify the actin to perform on each volume.

- ***STOPREPL:** Remove the volume(s) from the consistency set and stop replication.

CHGCSECRDE - Change CSE Credential Entry

This command can be used to change a user credential entry in the Device Domain Data associated with IBM i Copy Services Manager. The information is encrypted using a 128 bit AES key before it is retained.

For a description of the parameters, see the command ADDCSECRDE.

CHGCSEDTA – Change CSE Data

CHGCSEDTA is used to change the CSE Data as it relates to the operational characteristics of the toolkit. The parameters of CRTCSEDTA and DSPCSEDTA are similar and these descriptions apply to those commands as well. Please refer to CRTCSEDTA for parameter details.

CHKCSE – Check Copy Services Environment

The Check Copy Services Environment (CHKCSE) command verifies that all of the requirements are in place to perform a scheduled CSE switch. It is recommended to call this command after making changes to the CSE or CSEDTA, and to schedule it to run during the day to catch any changes in configuration which may prevent a switch.

An error in the configuration will result in an escape message being sent.

CHKCSE performs many more checks than SWCSE, including whether the LUNs attached to the source LPAR are the ones replicated. If some LUNs are missing from the replication then CHKCSE will fail; if LUNs are replicated but are not configured to the LPAR then CHKCSE will issue a warning. To prevent CHKCSE from checking any LUNs create a system environment variable QZ_DISABLE_LUN_CHECK. Note: Disk mappings using vSCSI are not supported and they will result in the LUN check failing.

CHKCSE will warn if qzrdhasm.log is greater than 20 MB

```
                    Check Copy Services Environ. (CHKCSE)

Type choices, press Enter.

Environment name . . . . . .   __________    Name

                           Additional Parameters

Message queue  . . . . . .   *NONE         Name, *QSYSOPR, *NONE
  Library  . . . . . . . . .     *LIBL       Name, *LIBL
```

Environment name: Specifies the name of the Environment which is to be checked. This name also refers to the CSEDTA name.

Message queue: The message queue and library name where inquiry messages will be routed. The default is *NONE. If a user message queue name is provided, a library name is required.

CHKFSR3 – Check FSR 3-site

The Check FSR 3-site (CHKFSR3) command verifies that all of the requirements are in place to perform a scheduled CSE switch in a 3-site environment. It is recommended to call this command after making changes to the CSE or CSEDTA, and to schedule it to run during the day to catch any changes in configuration which may prevent a switch.

An error in the configuration will result in an escape message being sent.

CHKFSR3 performs many more checks than SWFSR3, including whether the LUNs attached to the source LPAR are the ones replicated. If some LUNs are missing from the configuration then CHKFSR3 will fail; if LUNs are replicated but are not configured to the LPAR then CHKFSR3 will issue a warning. To prevent CHKFSR3 from checking any LUNs create a system environment variable QZ_DISABLE_LUN_CHECK. Note: Disk mappings using vSCSI are not supported and they will result in the LUN check failing.

CHKFSR3 will warn if qzrdhasm.log is greater than 20 MB

```
                    Check FSR 3-site. (CHKFSR3)

Type choices, press Enter.

Environment name . . . . . .   __________      F4 to prompt
Replication type . . . . . .   ________      *GMIR, *MMIR, *GMIR2…

                           Additional Parameters

Message queue  . . . . . . . .   *NONE         Name, *QSYSOPR, *NONE
  Library  . . . . . . . . . .     *LIBL       Name, *LIBL
```

Environment name: Specifies the name of the Environment which is to be checked. This name also refers to the CSEDTA name. Use F4 to see a list of available environments.

Replication type: The type of Copy Service environment. This required parameter specifies the type of the copy services environment to be checked.

The possible values are:

- ***GMIR -** The Global Mirroring environment with this environment name will be checked. This is for the H1-H3 leg.
- ***MMIR -** The Global Mirroring environment with this environment name will be checked. This is for the H1-H2 leg.
- ***GMIR2** - The Global Mirroring 2 environment with this environment name will be checked. This is for the H2-H3 leg.
- ***MMIR2** - The Metro Mirroring 2 environment with this environment name will be checked. This is for the H1-H3 leg
- ***MMIR3** - The Metro Mirroring 3 environment with this environment name will be checked. This is for the H2-H3 leg

Message queue: The message queue and library name where inquiry messages will be routed. The default is *NONE. If a user message queue name is provided, a library name is required.

CLEANLOGS – Clean Toolkit Logs

In prior releases this command was named CLNICSMLOG.

This command will clean the logs related to the IASP Copy Services Manager, Full System Replication Manager, or Full System FlashCopy Manager. CLEANLOGS removes specified entries from the /QIBM/Qzrdhasm/Qzrdhasm.log, removes any temporary directories created by DMPINF, dmpinf zip files and toolkit joblogs in /QIBM/Qzrdhasm/joblogs/ which are older than the number of days specified on this command. This command should be run periodically to prevent log files from consuming excessive space.

Days of information to retain: How many days of information to retain.

FSFC environment: The name of the FSFC environment to be cleaned.

- ***NONE:** Do not clean up the logs for any environment.
- ***ALL**: Clean up the logs for all environments.
- **environment-name:** Clean up the logs for the named environment.

CLRDDD – Clear Device Data Domain

CLRDDD is a utility command used to clear data out of PowerHA device data domains. It should only be used when instructed by IBM personnel and may remove toolkit configuration data.

CPYCSEDTA – Copy CSE Data

This interactive command will create a new Cluster Resource Group and write configuration data to the CSE Data device domain based on information from an existing CSE data. The configuration data may be modified before the CSE Data is created.

Please refer to the command CRTCSEDTA for a detailed description of the parameters.

CSE Data name: The name that identifies the CSE data.

To CRG: Specifies the name of the CRG device which is to be prepared for backup. This name identifies a CRG device description on the owning cluster node. This is also the Environment name referenced elsewhere within ICSM. The environment is configured using WRKCSE.

CRTCSEDTA - Create Copy Services Environment Data

CRTCSEDTA is used to enter the CSE Data as it relates to the operational characteristics of the toolkit. The parameters of CHGCSEDTA and DSPCSEDTA are similar and these descriptions apply to those commands as well. Flash and Replication, IASP and Full system all use CRTCSEDTA, but the parameters differ based on the initial selections.

```
                                Create CSE CRG

 Supply all required values, press Enter.

 CRG Name . . . . . . . . . . . :   TEST

 Use  . . . . . . . . . . . . . :   *SYSTEM

 Copy type  . . . . . . . . . . .   *PPRC        *PPRC, *FLASH
```

For FSR operations the CSE data name must be the same as the environment name, as only one CSEDTA is expected for FSR (this differs from FSFC). Entering use *SYSTEM and type *PPRC indicates this is going to be a CSE Data for Full System Replication. Press Enter and the remaining FSR-specific parameters are displayed.

```
                                Create CSE Data

 Supply all required values, press Enter.

 CSE Data Name  . . . . . . . . :   TEST
Use  . . . . . . . . . . . . . :   *SYSTEM
Copy type  . . . . . . . . . . :   *PPRC

Environment name . . . . . . . .   __________
Production node name . . . . . .   __________
FSR communications port  . . . .   55920
Primary controlling node . . . .   __________   Name
Secondary controlling node . . .   __________   Name

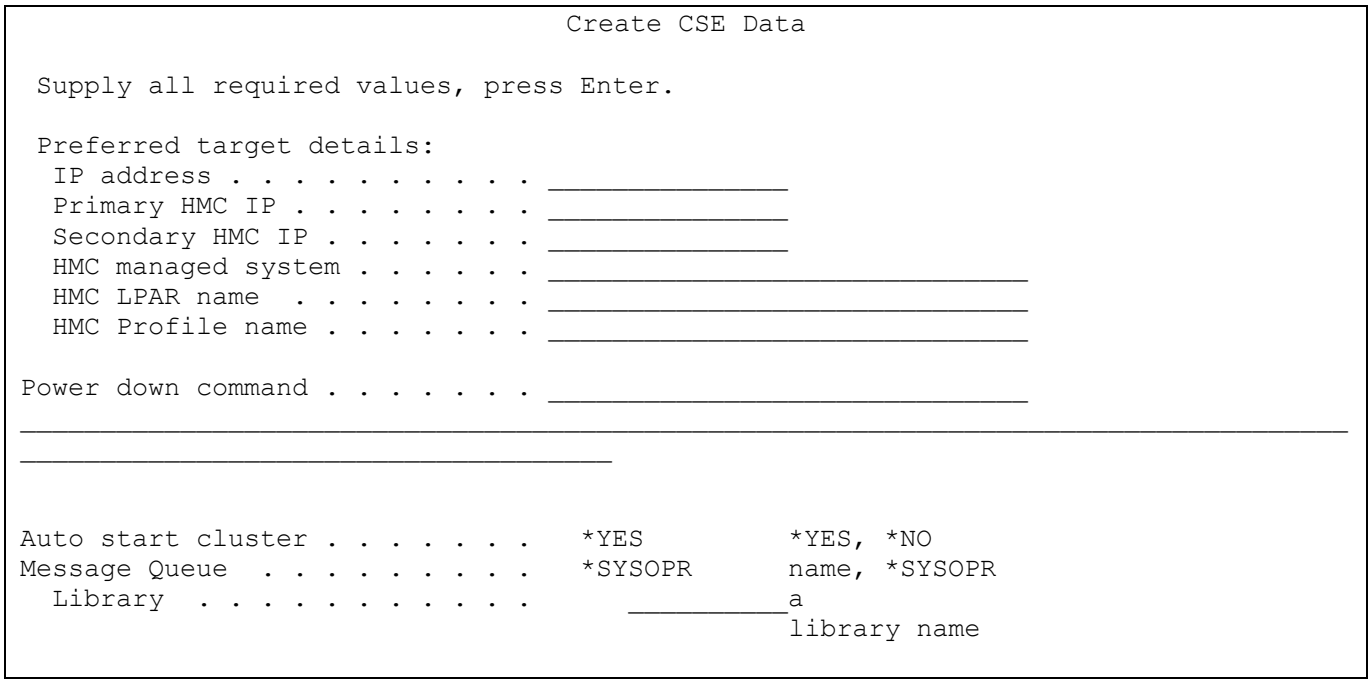
Preferred source details:
  IP address . . . . . . . . . . _________________
  Primary HMC IP . . . . . . . . _________________
  Secondary HMC IP . . . . . . . _________________
  HMC managed system . . . . . . _____________________________________
  HMC LPAR name  . . . . . . . . _____________________________________
  HMC Profile name . . . . . . . _____________________________________
                                                                  More...
F1=Help   F3=Exit   F4=Prompt   F6=Query HMC   F12=Cancel
```

Environment: Refers to the environment previously created in WRKCSE. This must be an FSR environment of the same name as the CSESDTA, and determines the external storage details.

Primary and Secondary controlling nodes: Enter the cluster node names of the controllers. There is no practical distinction between primary and secondary nodes. If there is only one node, do not enter a secondary node.

Communications port: The port on the source LPAR the controllers will use to communicate with the source and target LPARs. Program QZRDIAEXT2 in subsystem QZRDFSR is expected to be listening to this port when the FlashCopy operations are performed.

Preferred source details: Specify the details for the preferred source LPAR. The *IP Address* is used by the controller to reach the preferred source host. The *Primary HMC IP* is required and *Secondary HMC IP* is optional. *The HMC LPAR name***,** *HMC profile name***,** and *HMC managed system* parameters are case sensitive and can be queried directly from the HMC by pressing F6 after the HMC IP addresses are entered. If the toolkit receives an unexpected response from the currently used HMC it will automatically retry the response on the other HMC, if it is configured.



Preferred target details: Specify the details for the preferred targe LPAR. The *IP Address* is used by the controller to reach the preferred target host, it can be the same as the source LPAR. The *Primary HMC IP* is required and *Secondary HMC IP* is optional. *The HMC LPAR name***,** *HMC profile name***,** and *HMC managed system* parameters are case sensitive and can be queried directly from the HMC by pressing F6 after the HMC IP addresses are entered. If the toolkit receives an unexpected response from the currently used HMC it will automatically retry the response on the other HMC, if it is configured.

Power down command: This is the command that will be issued on the running LPAR prior to the switch. It can be a PWRDWNSYS or a program call that will perform more graceful shutdown procedures. The power down command should NOT perform a restart; the toolkit will do that. If no shutdown is desired, enter *NONE. Prompting is available using F4, but note that the controller is used for prompting despite the command will actually run on the source LPAR.

Auto start cluster: Specify *YES or *NO depending on whether the CHKCSE/SWCSE command should issue the STRCLUNOD if the cluster is not started.

Message queue: The message queue and library name where inquiry messages will be routed. The default is *SYSOPR. If a user message queue name is provided, a library name is required.

DLTCLDCS - Delete Cloud Consistency Set

The Delete Cloud Consistency Set (DLTCLDCS) command deletes the specific consistency set. The consistency set must be empty.

API Key Name: This is a required parameter. Specifies the name of the API Key to use to find the cloud instance.

Consistency Set Name: The name that identifies the consistency set to delete.

- ***EMPTY:** Delete any empty consistency sets found.
- **Name:** Delete the consistency set specified.

DLTCSEDTA - Delete CSE Data

This command will delete the specific CSE Data from the device domain data.

CSE Data name: The name that identifies the CSE Data.

DSPCSEDTA - Display Copy Services Data

This interactive command displays the specific CSE Data configuration information in the device domain data space.

CSE Data name: Specifies the name of the CSE Data and the exit data that is to be displayed.

DSPDDDLCK - Display DDD lock holder

This is a utility command for display the job and user who holds a cluster-wide lock on the specified artifact. This command should be used if a log message indicates a lock is preventing an operation.

Name: Enter the name of the object to display a lock for.

Type: Enter the type of object to display a lock for.

IPLLPAR – IPL Partition

This command will activate the specified LPAR.

HMC partition name: Specify the name of the LPAR to activate. This parameter is case sensitive.

Partition profile: Specify the partition profile to use when activating the partition. This parameter is case sensitive.

- ***LPARNAME:** The name of the profile is the same as the partition name.

Managing system: The managed system name must be specified. This parameter is case sensitive.

Primary HMC IP address: Enter the IP address of the primary HMC. This address must have credentials listed in WRKCSECRDL.

Secondary HMC IP address: Enter the IP address of the secondary HMC. This address must have credentials listed in WRKCSECRDL. This IP address is optional and will only be used if an unexpected response is received from the primary HMC.
IPL Source: Specifies whether an initial-program-load (IPL) is started from the A-source, B-source or D-source of the system. This parameter allows you to control which Licensed Internal Code (LIC) storage source of the system to IPL. Also, the source of the system determines where LIC program temporary fixes (PTFs) are applied. This parameter also allows the system to be upgraded to a new release from an install image on DASD.

LIC has three storage areas known as the A-source, the B-source and the D-source. The D-source is the install media. The A- and B-sources are part of the system memory. Initially, the A- and B-sources are identical, but when Licensed Internal Code fixes are performed temporarily (PTF), the temporary fixes are stored on the B-source. When the same fixes become permanent, they are copied from the B-source to the A-source; therefore, the fixes reside on both the A-source and the B-source.

When you want to send temporary fixes to the B-source, you must start the system from the A-source, which causes the fixes to be sent to the B-source.

When you start the system from the A-source, you are running the system from the permanent fixes. When you start the system from the B-source, you are running the system from a mixture of temporary and permanent fixes. When you start the system from the D-source, you are using the Licensed Internal Code loaded from the install media.

It is recommended that you specify RESTART(*YES); otherwise, you cannot be assured which source of the system is actually started. This precaution can save you some time.

- ***PANEL:** The partition is started from the source that is currently shown on the operator's panel, the A-source, the B-source, or the D-source.

- **A:** The partition is started from the A-source.

- **B:** The partition is started from the B-source.
- **D:** The system is started from the D-source, the install media.

Keylock position: Specifies whether the LPAR will be activated for attended (manual) or unattended (auto) mode.

- ***PANEL:** The partition is started in the mode that is currently shown on the operator's panel.
- ***AUTO:** The partition is started in unattended mode.
- ***MANUAL:** The partition is started in attended mode.

Block until LPAR is active: This parameter determines whether IPLLPAR will return control to the caller immediately or after the LPAR has reached the "Running" state.

- ***YES:** IPLLPAR will remain active until the LPAR has reached a "Running" state.
- ***NO:** IPLLPAR will terminate immediately after activating the LPAR.

Max minutes to block: Specifies how long IPLLPAR will wait for the partition to reach a "Running" state. This is ignored if **Block until LPAR is active** is *NO. If the partition does not enter the "Running" state before the specified number of minutes has elapsed, IPLLPAR will issue and escape message.

Confirm activation: Configures IPLLPAR to wait on an inquiry message before activating the partition.

- ***YES:** IPLLPAR will issue an inquiry message and will wait for a positive response before activating the partition.
- ***NO:** IPLLPAR will activate the partition immediately.

LOGMSG – Log Message

The default log is in /QIBM/Qzrdhasm/qzrdhasm.log but some toolkit processes will change to a different log. LOGMSG will insert the specified message to the current log in use by the job.

Message: Enter the message to be inserted, from 20-120 characters.

Message format: Determines how the message will look in the log.

- ***STD:** Use the standard message format. An example:

  2017-07-27 12:46:00 standard message (logmsg)

- ***CMDSTR:** Use the 'command start' message format. An example:

  682401 2017-07-27 12:51:10 Start command start message starting from job 682401/AASLAND/QPADEV001Q

Message type: Indicates the message status tag. The possible values are:

- ***INFO:** The message does not have a tag.
- ***ERROR:** The message is tagged with 'Error'
- ***WARNING:** The message is tagged with 'Warning'

PRTCSE – Print CSE Information

The PRTCSE command prints Copy Services Environment (CSE) information, that was entered by the WRKCSE command and stored in cluster Device Domain Data (DDD). The information is printed to a QPRINT spool file. Unlike the environment list generated within WRKCSE, this command can access defined enviromnents from any node within the cluster. Any combination of ENV and TYPE parameter values may be used.

Environment name: Specifies the name of the environments for which all information is to be dumped. The possible values are:

- ***ALL:** All types of environments with the Copy Service Type below or all environments may be dumped.
- **Environment name:** All types of environments with the supplied environment name may be dumped.

Copy Service Type: The type of Advanced Copy Service enviromnent. This parameter is required. The possible values are:

- ***ALL:** All types of environments with the environment name or all environments may be dumped.
- ***FLASH:** The FlashCopy environment with the environment name or all FlashCopy environments may be dumped.

- ***GMIR:** The PPRC Global Mirroring environment with this environment name or all GMIR environments may be dumped.

- ***LUN:** The LUN Level connection switching environment with this environment name or all LUN environments may be dumped.

- ***MMIR:** The PPRC Metro Mirroring environment with this environment name or all MMIR environments may be dumped.

RLSCSELCK - Release Copy Services Lock

This command releases a cluster wide environment lock used within IASP Copy Services Manager.

Environment name: Specifies the name of the Environment (which may be an Independent ASP). A special value of *ALL may be used with RLSTYPE(*OVR) to override all ICSM locks for the cluster.

Release type: The possible values are:

- ***JOB:** Releae the lock that is held by this process.

- ***OVR:** Release a cluster wide environment lock which may be held by another process or when used with ENV(*ALL), release all ICSM locks for the cluster. Use this only to override locks held by processes which are no longer active.

RMVCSECRDE - Remove CSE Credential Entry

This command can be used to remove a user credential entry from the Device Domain Data associated with IBM i Copy Services Manager.

Host IP address: The address of the host to be removed.

RSTDDD – Restore Toolkit Device Data Domains

The Toolkits uses Device Data Domains (DDD) to store WRKCSE, WRKCSEDTA and WRKCSECRDL information. Use RSTDDD to restore the DDD from files created by SAVDDD.

The sole parameter is the path to the directory which contains the DDD backup files generated by SAVDDD.

RSTJOBSCDE – Restore Job Scheduler Entries

CFGSTRPRSC will save the existing job scheduler entries to a userspace then hold them. Use this command to display or restore and release the stored job entries if CFGSTRPRSC does not do so automatically. This command can also be used after recovering from a flashcopy backup.

RTVINF - Retrieve ICSM Information

The Retrieve Information (RTVINF) command retrieves specific information from IASP Copy Services Manager. This command is only valid when executed within a CL program.

Environment name: This is a required parameter. Specifies the name of the Environment (which may be an Independent ASP) for which information is to be retreived.

Information: This is a required parameter and specified the type of information to be returned. The possible values are:

- ***CLUNODENAME:** The node name of this system within the cluster. The length of the first CL return variable below must be at least eight characters.

- ***FLASHNODES:** The backup Flash Copy node names for the Flash configurations. Up to six FlashCopy environments may be configured. All six CL return variables must be specified and each of their lengths must be at least eight characters.

- ***FLASHSTATUS:** First Flashcopy status. The length of the first CL return variable below must be at least three characters.
    - 0 = Ready
    - 100 = Flashed
    - other number = Flash in process
    - blank = undetermined.

- ***FLASH2STATUS:** Second Flashcopy status. The length of the first CL return variable below must be at least three characters.
    - 0 = Ready
    - 100 = Flashed
    - other number = Flash in process
    - blank = undetermined.

- ***FLASH3STATUS:** Third Flashcopy status. The length of the first CL return variable below must be at least three characters.
    - 0 = Ready
    - 100 = Flashed
    - other number = Flash in process
    - blank = undetermined.

- ***FLASH4STATUS:** Fourth Flashcopy status. The length of the first CL return variable below must be at least three characters.

  - 0 = Ready
  - 100 = Flashed
  - other number = Flash in process
  - blank = undetermined.

- ***FLASH5STATUS:** Fifth Flashcopy status. The length of the first CL return variable below must be at least three characters.
  - 0 = Ready
  - 100 = Flashed
  - other number = Flash in process
  - blank = undetermined.

- ***FLASH6STATUS:** Sixth Flashcopy status. The length of the first CL return variable below must be at least three characters.
  - 0 = Ready
  - 100 = Flashed
  - other number = Flash in process
  - blank = undetermined.

- ***GMIRDIRECTION:** The length of the first CL return variable below must be at least one character.
  - N = Normal
  - R = Reversed
  - Blank = undetermined.

- ***GMIRNODEROLE:** This key value MUST replace any use of *PPRCNODEROLE for GMIR type environments, else unpredictable results may occur. The length of the first CL return variable below must be at least one character.
  - S = Source
  - T = Target
  - blank = undetermined

- ***GMIRSTATE:** The length of the first CL return variable below must be at least one character.
  - 0 = Stopped
  - 1 = Running
  - 2 = Failover
  - 3 = Suspended
  - 4 = Lagging
  - 5 = Extreme Lagging
  - 6 = Paused
  - 9 = Other
  - M = Multi-target Incremental Resync Pair
  - blank = undetermined. RTVLPARINF

- ***GMIRSTATEDIRECT:** State is returned in RTNVALUE and the direction is returned in RTNVALUE2. The length of the first and second variables must be at least one character.

    - State (RTNVALUE):
        - 0 = Stopped
        - 1 = Running
        - 2 = Failover
        - 3 = Suspended
        - 4 = Lagging
        - 5 = Extreme lagging
        - 6 = Paused
        - 9 = Other
    - Direction: (RTNVALUE2):
        - M = Multi-target Incremental Resync Pair
        - N = Normal
        - R = Reversed
        - blank = undetermined

- ***GMIR2DIRECTION:** The length of the first CL return variable below must be at least one character:
    - N = Normal
    - R = Reversed
    - Blank = undetermined

- ***GMIR2NODEROLE:** This key value MUST replace any use of *PPRCNODEROLE for GMIR2 type environments, else unpredictable results may occur. The length of the first CL return variable below must be at least one character
    - S = Source
    - T = Target
    - blank = undetermined.

- ***GMIR2STATE:** The length of the first CL return variable below must be at least one character
    - 0 = Stopped
    - 1 = Running
    - 2 = Failover
    - 3 = Suspended
    - 4 = Lagging
    - 5 = Extreme Lagging
    - 6 = Paused
    - 9 = Other
    - M = Multi-target Incremental Resync Pair
    - blank = undetermined.

- ***GMIR2STATEDIRECT:** State is returned in RTNVALUE and direction is returned in RTNVALUE2. The length of first and second CL return variables below must be at least one character.
    - State (RTNVALUE):
        - 0 = Stopped

        - 1 = Running
        - 2 = Failover
        - 3 = Suspended
        - 4 = Lagging
        - 5 = Extreme lagging
        - 6 = Paused
        - 9 = Other
    - Direction: (RTNVALUE2):
        - M = Multi-target Incremental Resync Pair
        - N = Normal
        - R = Reversed
        - blank = undetermined

- ***IASPNAME:** The name of the Independent ASP used for the VRYCFG commands. The length of the first CL return variable below must be at least ten characters.

- ***LUNCONNECTION:** The length of the first CL return variable below must be at least one character.
    - P = Connected to the Production Node
    - H = Connected to the HA node
    - Blank = undetermined

- ***LUNNODEROLE:** The values indicate the normal role of this node. The length of the first CL return variable below must be at least one character.
    - P = Production Node
    - H = HA node
    - Blank = undetermined

- ***MMIRDIRECTION:** The length of the first CL return variable below much be at least one character.
    - N = Normal
    - R = Reversed
    - Blank = undetermined

- ***MMIRNODEROLE:** This key value should replace any use of *PPRCNODEROLE for MMIR type environments. The length of the first CL return variable below must be at least one character.
    - S = Source
    - T = Target
    - Blank = undetermined

- ***MMIRSTATE:** The length of the first CL return variable below must be at least one character.
    - 0 = Failed
    - 1 = Running
    - 2 = Failover

- 3 = Suspended
- 4 = Resuming
- 6 = Paused
- M = Multi-target Incremental Resync Pair
- blank = undetermined.

- ***MMIRSTATEDIRECT:** State is returned in RTNVALUE and direction is returned in RTNVALUE2. The length of first and second CL return variables below must be at least one character.
  - States:
    - 0 = Failed
    - 1 = Running
    - 2 = Failover
    - 3 = Suspended
    - 6 = Paused
    - M = Multi-target Incremental Resync Pair
    - blank = undetermined.
  - Directions:
    - N = Normal
    - R = Reversed
    - blank = undetermined.

- ***MMIR2DIRECTION:** The length of the first CL return variable below must be at least one character.
  - N = Normal
  - R = Reversed
  - blank = undetermined.

- ***MMIR2NODEROLE:** This key value should replace any use of *PPRCNODEROLE for MMIR2 type environments. The length of the first CL return variable below must be at least one character.
  - S = Source
  - T = Target
  - blank = undetermined.

- ***MMIR2STATE:** The length of the first CL return variable below must be at least one character.
  - 0 = Failed
  - 1 = Running
  - 2 = Failover
  - 3 = Suspended
  - 4 = Resuming
  - 6 = Paused
  - M = Multi-target Incremental Resync Pair
  - blank = undetermined

- ***MMIR2STATEDIRECT:** State is returned in RTNVALUE and direction is returned in RTNVALUE2. The length of first and second CL return variables below must be at least one character.
    - States:
        - 0 = Failed
        - 1 = Running
        - 2 = Failover
        - 3 = Suspended
        - 6 = Paused
        - M = Multi-target Incremental Resync Pair
        - blank = undetermined.
    - Directions:
        - N = Normal
        - R = Reversed
        - blank = undetermined.

- ***MMIR3DIRECTION:** The length of the first CL return variable below must be at least one character.
    - N = Normal
    - R = Reversed
    - blank = undetermined.

- ***MMIR3NODEROLE:** This key value should replace any use of *PPRCNODEROLE for MMIR3 type environments. The length of the first CL return variable below must be at least one character.
    - S = Source
    - T = Target
    - blank = undetermined.

- ***MMIR3STATE:** The length of the first CL return variable below must be at least one character.
    - 0 = Failed
    - 1 = Running
    - 2 = Failover
    - 3 = Suspended
    - 4 = Resuming
    - 6 = Paused
    - M = Multi-target Incremental Resync Pair
    - blank = undetermined

- ***MMIR3STATEDIRECT:** State is returned in RTNVALUE and direction is returned in RTNVALUE2. The length of first and second CL return variables below must be at least one character.
    - States:
        - 0 = Failed
        - 1 = Running

        - 2 = Failover
        - 3 = Suspended
        - 6 = Paused
        - M = Multi-target Incremental Resync Pair
        - blank = undetermined.
    - Directions:
        - N = Normal
        - R = Reversed
        - blank = undetermined.

- ***MULTITARGET**: Multi-target configuration
    - N = None
    - M = Multi-target using two Metro Mirror environments
    - G = Multi-target using one Metro Mirror environment and one Global Mirror environment.

- ***PPRCNODEROLE:** The length of the first CL return variable below must be at least one character.
    - S = Source
    - T = Target
    - blank = undetermined

- ***PPRCNODES:** The backup PPRC node names for the up to three PPRC configurations. RTNVALUE will contain the MMIR configuration node name. RTNVALUE2 will contain the GMIR configuration node name. RTNVALUE3 will contain the LUN configuration node name. All three CL return variables must be specified and each of their lengths must be at least eight characters.

- **PPRCNODE:** The Production node name. RTNVALUE will contain the node name. The CL return variable must be specified with a length of at least eight characters.

Return value: This is a required parameter. Specifies a CL character variable name for returned value.

Return value 2: This is a required parameter for *MMIRSTATEDIRECT, *GMIRSTATEDIRECT, *FLASHNODES and *PPRCNODES. For *MMIRSTATEDIRECT and *GMIRSTATEDIRECT, the PPRC direction ('N' or 'R') is returned here. For *FLASHNODES, the second FlashCopy node will be returned here. For *PPRCNODES, the GMIR target node will be returned here. Specifies a CL character variable name for the second returned value.

Return value 3: This is a required parameter for *FLASHNODES and *PPRCNODES. Specifies a CL character variable name for the third returned value. For *FLASHNODES, the third FlashCopy node will be returned here. For *PPRCNODES, the LUN HA node will be returned here.

Return value 4: This is a required parameter for *FLASHNODES if a forth FlashCopy environment is configured. Specifies a CL character variable name for the fourth returned value. The fourth FlashCopy node will be returned here.

Return value 5: This is a required parameter for *FLASHNODES if a fifth FlashCopy environment is configured. Specifies a CL character variable name for the fifth returned value. The fifth FlashCopy node will be returned here.

Return value 6: This is a required parameter for *FLASHNODES if a sixth FlashCopy environment is configured. Specifies a CL character variable name for the sixth returned value. The sixth FlashCopy node will be returned here.

RTVSTRPRSC – Retrieve Startup Resources

This command is used to determine the role of the LPAR. It will use the information in WRKSTRPRSC to determine the role (usage) of the LPAR at IPL, and then return the usage as a variable.

```
                        Retrieve Startup Resources (RTVSTRPRSC)

Type choices, press Enter.

Type of IPL  . . . . . .                Variable
```

Type of IPL: *CHAR 5 variable returned which contained the usage (*PS, *PT, *SGxx etc).

RUNDSCMD - Run DS Scripted Command

The RUNDSCMD command is used in a CL program to run a scripted command through DSCLI and validate the results from the expected result list provided in the command parameters. The command optionally returns the total of numeric values in a specified column of the results.

Three exception messages may be issued:

- IAS0301 - Parameter error detected.
- IAS0302 - Result file error detected.
- IAS0303 - Results not as expected.

Script input file: The complete path and name of the script stream file. A value of '*' is allowed when no script is to be executed, and only prior results are to be validated.

Profile input file: The complete path and name of the profile stream file. A value of '*' is allowed when that value is also used for the script parameter.

Results output file: The complete path and name of the results stream file. This parameter is always required.

User: Specify the user profile name to be used to run the DSCLI scripted command.

Result validation list: A list of from one to ten validation entries to be matched against the results returned by the scripted command. The columns in the result file must be comma separated (refer to the profile that is in use).

Each validation list entry contains two or three elements:

- **Column position:** The comma separated column position for this element. This required value must be from one through 20.

- **Expected value:** The case sensitive character string that is expected. This value is required, must be enclosed in apostrophes, and may contain alphanumeric characters, blanks and punctuation marks except commas. Strings that are not to be found may be specified by placing a minus sign (-) as the first character of the string.

- **Logic to next in list:** The locical operator (*AND or *OR) to the next expected value in the list. This value is required on all list elements except the last. Note that the results of each file row are evaluated from the first to the last expected value. Careful consideration must be made when this value is mixed (both *AND and *OR used on a single RUNDSCMD command). If more complex result checking is required, the RUNDSCMD command may be run again specifying '*' for the SCRIPT parameter and combined results may be evaluated within the user's CL program.

Result file rows: The rows that are to be validated. The possible values are:

- ***ALL**: The results specified in the validation list are expected to be found in at all of the result file rows.

- ***ONE:** The results specified in the validation list are expected to be found in at least one of the result file rows.

Summation column: The comma separated column positional value that is to be summed and returned in the TOTAL parameter. Numeric values from one through 20 are allowed. The default is *NONE. If a numeric value is specified, the content of that column in the result file may only contain numeric data.

CL variable for returned total: Specifies the name of the CL program variable that receives the total value being returned. The type and length for the CL variable must be TYPE(*DEC) LEN(9 0). This parameter is required when the value of SUMCOLUMN is other than *NONE.

Return column: The comma separated column positional value that is to be returned in the RTNVALUE parameter. Numeric values from one through 20 are allowed. The default is *NONE. Refer to the RTNKEY parameter below. This parameter is requires that the RTNCOLUMN parameter is also specified.

Return key value: Specifies the key string used to locate the row in the result file where the return value (RTNVALUE) will be extracted. The default is *NONE.If a numeric value is specified for RTNCOLUMN and this parameter is *NONE, the value in the first row of the result file is returned in RTNVALUE. If a numeric value is specified for RTNCOLUMN and this parameter specifies a key value, the value in the first row of the result file that contains that key is returned in RTNVALUE. This parameter is requires that the RTNCOLUMN and RTNVALUE parameters are also specified.

CL variable for returned value: Specifies the name of the CL program variable that receives the character value being returned. The type and length for the CL variable must be TYPE(*CHAR) LEN(80). This parameter requires that the RTNCOLUMN parameter is also specified.

RUNLPARCMD - Run command based on LPAR/SRLN

This command will run a specific command based on the specified LPAR and serial numbers detected. This command is usually used in the startup program to perform different tasks based on the serial and LPAR numbers.

NOTE: It is possible to test the behavior of this command on different serial and LPAR numbers by creating the following data areas:

```
CRTDTAARA DTAARA(QZRDHASM/FAKELPAR) TYPE(*CHAR) LEN(3)
VALUE('123')

CRTDTAARA DTAARA(QZRDHASM/FAKESRLN) TYPE(*CHAR) LEN(8)
VALUE('1234567')
```

The command RUNLPARCMD will look for these data areas and use them instead of detecting the actual serial or LPAR numbers.

Serial number: Specify the serial number required to run this command. NOTE: Prompt (F4) to see the current serial number.

- ***ANY:** Run this command on any LPAR, regardless of serial number.
- **Serial-number:** Run the command on the LPAR with the specified serial number.

LPAR number: Specify the LPAR number required to run this command. NOTE: Prompt (F4) to see the current LPAR number.

- ***ANY:** Run this command on any LPAR regardless of the LPAR number.
- **lpar-number:** Run the command on the LPAR with the specified number.

Command to execute: Run this command if the LPAR and serial numbers match those of the current LPAR. NOTE: Prompt (F4) to assist with building the command.

RUNSVCCMD - Run SVC Command

This command establishes an ssh session to an SVC, runs the user entered command, and stores the results in a stream file. NOTE: This command can establish an ssh session with any type of host, not just SVCs.

Command: Specifies the command to be run.

SVC IP address: Specifies the IPv4 formatted IP address of the SVC to be used.

Results output file: Specifies the IFS path name of the output file containing the results of the SVC command.

Display results: Specifies whether the results of the SVC command should be displayed in addition to be written to the output file.

SAVDDD – Save Toolkit Device Data Domains

The Device Data Domains (DDD) are where the Toolkit stores WRKCSE, WRKCSEDTA and WRKCSECRDL information. The command SAVDDD will extract that information and place it into IFS files. Any sensitive information including passwords will be encrypted.

Use the RSTDDD command to restore / recreate the DDD information.

Directory path: Where the DDD elements will be saved to.

- ***GEN:** A directory using the current timestamp will be created
- **Path:** The path to the directory in which to place the files. The directory must exist.

SETCSELCK - Set Copy Services Lock

This command sets a cluster wide environment lock used within IASP Copy Services Manager.

Environment name: This is a required parameter. Specifies the name of the Environment (which may be an Independent ASP).

Timeout minutes: The number of minutes before the environment lock will self time out. The default is fifteen minutes.

SETDDDLCK - Set DDD lock holder

This is a utility command for setting the job and user who holds a cluster-wide lock on a specified artifact.

Name: Enter the name of the object to display a lock for.

Type: Enter the type of object to display a lock for.

SETUPFSR - Set up IBM Pwr HA tools – FSR

This command prepares IBM Power HA toolkit for IBM i for Full System replication operations after has operational library (QZRDHASM) has been restored. If a version change has been detected this command will make any changes required. The QZRDFSR subsystem will be started.

```
                  Setup IBM i Copy Services Mana (SETUPFSR)

Type choices, press Enter.

Node role  . . . . . . . . . . . .   __________      *CTL, *PRD
FSR communications port  . . . .   *SAME           1-65535, *SAME, *DFT
Toolkit access code for *CTL . .   *SAME
Toolkit access code for *PS  . .   *SAME
Toolkit access code for *PT  . .   *SAME
```

Node Role: Special values indicating which role the system will be acting. The possible values are:

- ***CTL:** The system is a controlling node in Full System Replication.
- ***PRD:** The system is the source LPAR in Full System Replication.

FSR Communications Port: The port to which the production node will be using for communications. This value must match the FSR communications port value specified in either CRTCSEDTA or CHGCSEDTA that is run on the controlling node. The possible values are:

- ***SAME:** The value does not change. This is the only valid value for NODETYPE(*CTL)
- ***DFT:** The default port (55920) will be used.
- **1 – 65535:** Specify the port number that is to be used.

Toolkit access code for *CTL: Enter the access code when configuring the controlling LPAR.

Toolkit access code for *PS: Enter the preferred target access code when configuring the production LPAR.

Toolkit access code for *PT: Enter the preferred source access code when configuring the production LPAR.

STRCLDVSI - Start Cloud Power Virtual Machine

The Start Cloud VSI (STRCLDVSI) command will cause a VSI (Power Cloud virtual server instance) to start by issuing a request to the cloud instance mapped to the API Key name.

VSI: Specifies the id of the Power Cloud virtual server instance (VSI) that is to be started.

Block until VSI is active: Specifies whether the command should block until the VSI has finished the IPL and is active. The possible values are:

- ***YES:** The command will block until the IPL is finished.
- ***NO:** The command will not block.

Max minutes to block: Specifies the maximum anount of time in minutes that the command should wait for the IPL to finish and the LPAR to become active.

SWCSE – Perform Full System Switch

The Switch Copy Service Environment (SWCSE) command causes the Production Node to be powered down and the DR node to be powered up, and replication in the mirroring environment to be paused or reversed. If the source storage is not available then replication will be idled and secondary access granted to the target LPAR.

```
                      Switch Copy Services Environ. (SWCSE)

Type choices, press Enter.

Environment name . . . . . .   __________     Name
Switch type  . . . . . . . .   *SCHEDULED     *SCHEDULED, *UNSCHEDULED
Type . . . . . . . . . . . .   *              *, *GMIR, *MMIR
Auto replicate . . . . . . .   *DFT           *DFT, *YES, *NO
```

Environment name: Specifies the name of the Environment which is to be checked. This name also refers to the CSEDTA name.

Switch type: Indicates whether this is a *scheduled* or *unscheduled* switch. The possible values are:

- ***SCHEDULED**: This CSE switch is planned. The Production system is available and subsystem QZRDFSR is running and listening on the port specified in the CSEDTA.
- ***UNSCHEDULED:** This CSE switch is not planned. The Production system is not available.

Type: The type of the IBM i Copy Services environment. This optional parameter may be used to specify the type of the copy services environment to be switched. The possible values are:

- *: The specific type (GMIR or MMIR) will be resolved by the command processing program. Only one of these types may be configured with this environment name.
- ***GMIR:** The Global Mirroring environment with this name will be switched.
- ***MMIR:** The Metro Mirroring environment with this name will be switched.

Auto replicate: This optional parameter specifies whether replication is to be automatically started after the switch. The possible values are

- ***DFT:** Use the default setting for the specific switch type. For a scheduled switch the default is *YES. For an unscheduled switch the the default is *NO.
- ***YES:** A switch will auto replicate. The replication will not start until the DR node has finished it's IPL.
- ***NO:** A switch will not auto replicate.

SWFSR3 – Switch FSR 3-site

The Switch FSR 3-site (SWFSR3)command causes the Production Node to be powered down and the selected DR node to be powered up, and replication in the mirroring environment to be paused or reversed. If the source storage is not available then replication will be idled and secondary access granted to the target LPAR.

```
                          Switch FSR 3-site (SWFSR3)

Type choices, press Enter.

Environment name . . . . . .   __________     Name
Switch type  . . . . . . . . .                *SCHEDULED, *UNSCHEDULED
Replication type . . . . . . .                *GMIR, *MMIR, *GMIR2…
Auto replicate . . . . . . . .   *DFT          *DFT, *YES, *NO
```

Environment name: Specifies the name of the Environment which is to be checked. This name also refers to the CSEDTA name.

Switch type: Indicates whether this is a *scheduled* or *unscheduled* switch. The possible values are:

- ***SCHEDULED**: This CSE switch is planned. The Production system is available and subsystem QZRDFSR is running and listening on the port specified in the CSEDTA.
- ***UNSCHEDULED:** This CSE switch is not planned. The Production system is not available.

Replication Type: The type of the IBM i Copy Services environment. This optional parameter may be used to specify the type of the copy services environment to be switched. The possible values are:

- ***GMIR -** The Global Mirroring environment with this environment name will be switched. This is for the H1-H3 leg.
- ***MMIR -** The Global Mirroring environment with this environment name will be switched. This is for the H1-H2 leg.
- ***GMIR2** - The Global Mirroring 2 environment with this environment name will be switched. This is for the H2-H3 leg.
- ***MMIR2** - The Metro Mirroring 2 environment with this environment name will be switched. This is for the H1-H3 leg
- ***MMIR3** - The Metro Mirroring 3 environment with this environment name will be switched. This is for the H2-H3 leg

Auto replicate: This optional parameter specifies whether replication is to be automatically started after the switch. The possible values are

- ***DFT:** Use the default setting for the specific switch type. For a scheduled switch the default is *YES. For an unscheduled switch the the default is *NO.
- ***YES:** A switch will auto replicate. The replication will not start until the DR node has finished it's IPL.
- ***NO:** A switch will not auto replicate.

TSTCOMM – Test Communication Paths

This interactive command will test the communication on the ports specified.

Role: This parameter specifies what type of communication role the partition is to take. The possible values are:

- ***SERVER:** The partition will act as a server.
- ***CLIENT:** The partition will act as a client.

IP Address: This parameter specifies the IP address of the other partition. For a client, this will be the server address. For a server, this will be the expected address received after the accept operation completes.

Port: This is the IP port that will be used for the test.

To use this command, log into the two LPARs to test the communication between and free the port to be tested (this may mean ending QZRDFSR). Designate one LPAR as the server, the other as the client. On the server, issue the following command

TSTCOMM ROLE(*SERVER) IPADDR('<client IP address>') PORT(55920)

When TSTCOMM is 'Listening on socket' issue the following command on the client:

TSTCOMM ROLE(*SERVER) IPADDR('<server IP address>') PORT(55920)

The results will indicate either success or failure (timeout) communicating on the specified port.

VIEWLOG - View Log File

This interactive command allows editing of the Advanced Copy Services log file on the local system and displaying the same on another system. This utility can also view any stream file on the systems.

System name: The system where the file is located. The possible values are:

- ***LOCAL:** View the main log file on the local system.
- ***SNMP:** View the SNMP log file on the local system.
- **System name:** View the main log file on the named system.

Stream file: The specific file to be viewed. The possible values are:

- ***SNMP**: View the SNMP log file. This option allows the SNMP log file on another system to be viewed.

- **Path and name of the stream file:** The default is /QIBM/Qzrdhasm/qzrdhasm.log

WRKCLDVOL - Work with Cloud Volumes

This command retrieves and displays a list of the PowerVS volumes associated with the provided API **k**ey.

API key name: Specify the key name (from WRKCSECRDL *CLOUD) containing the cloud instance id and PowerVS credentials.

Volume name: Provide the name of the volume(s) to display. Names with special characters, spaces etc. may need to be enclosed in single quotes. The possible values are:

- ***ALL**: Display the volumes regardless of volume name.
- **Name:** Display only the volume with the provided name.
- **Generic-name:** Specify the generic name of the volumes to display. A generic name is a character string that contains one or more characters followed by an asterisk (*). If a generic name is specified, all volumes that have names with the same prefix as the generic name are shown.
- **Substring:** Specify a substring to be included in the names of the volumes displayed. A substring is a character string that contains one or more characters followed by two asterisks (**). If a substring is specified, all volumes that have names including the subtring are shown.

Master volume name: Provide the name of the master volume(s) to display. Names with special characters, spaces etc may need to be enclosed in single quotes. The possible values are:

- ***ALL:** Display the volumes regardless of master volume name.
- **Name:** Display only the volume with the provided master volume name.
- **Generic-name:** Specify the generic master volume name of the volumes to display. A generic name is a character string that contains one or more characters followed by an asterisk (*). If a generic name is specified, all volumes that have master volume names with the same prefix as the generic name are shown.
- **Substring:** Specify a substring to be included in the master volume names of the volumes displayed. A substring is a character string that contains one or more characters followed by two asterisks (**). If a substring is specified, all volumes that have master volume names including the substring are shown.

Auxiliary volume name: Provide the name of the auxiliary volume(s) to display. Names with special characters, spaces etc may need to be enclosed in single quotes. The possible values are:

- ***ALL:** Display the volumes regardless of auxiliary volume name.
- **Name:** Display only the volume with the provided auxiliary volume name.
- **Generic-name:** Specify the generic auxiliary volume name of the volumes to display. A generic name is a character string that contains one or more characters followed by an asterisk

(*). If a generic name is specified, all volumes that have auxiliary volume names with the same prefix as the generic name are shown.

- **Substring:** Specify a substring to be included in the auxiliary volume names of the volumes displayed. A substring is a character string that contains one or more characters followed by two asterisks (**). If a substring is specified, all volumes that have auxiliary volume names including the substring are shown.

Replication-enabled volumes: Indicate whether replication-enabled volumes are to be included, exclusively included, or excluded. The possible values are:

- ***ALL:** Do not exclude based on replication-enabled status.
- ***ONLY:** Only include replication-enabled volumes.
- ***NONE:** Do not include replication-enabled volumes.

Master volumes: Indicate whether master volumes are to be included, exclusively included, or excluded. The possible values are:

- ***ALL:** Do not exclude master volumes.
- ***ONLY:** Only include master volumes.
- ***NONE:** Do not include master volumes.

Auxiliary volumes: Indicate whether auxiliary volumes are to be included, exclusively included, or excluded. The possible values are:

- ***ALL:** Do not exclude auxiliary volumes.
- ***ONLY:** Only include auxiliary volumes.
- ***NONE:** Do not include auxiliary volumes.

Volume ID: Specify the volume ID of the volume to display.

Virtual Server Instance ID: Display the volumes attached to the specified VSI.

The command lists volumes by consistency set (a consistency is a group of volumes replicating together). The following options and functions are available:

4=Delete: Deletes a Cloud Volume. You will be presented with a confirmation display prior to deleting the Cloud Volume.
5=Display: Display details about the volume or consistency set.
7=Rename: Rename a Cloud Volume. You will be prompted for a new name for the selected Cloud Volume.
8=Onboard: Onboard a Cloud Volume. Onboarding a Cloud Volume is a part of setting up replication.
10=Modify Consistency Set: Modify a Consistency Set. Add or remove volumes to/from a consistency set.
12=Toggle replication: Toggle replication. Enable or disable replication for a Cloud Volume. Enabling replication creates an auxiliary volume and a replication relationship between them. It is a

requirement prior to adding it to a volume group. Disabling replication removes the replication relationship and deletes the auxilary volume.
13=Toggle bootable: Toggle the boot status of a volume. Change the boot status of a volume. A VM must have at least one bootable volume.
14=Reset Consistency Set: Reset a consistency set in error status. Reset the status of a consistency set from error back to available.

F6=Create Volume: Press F6 to create a new volume.
F7=Create Consistency Set: Press F7 to create a new Consistency Set.
F10=View log: View the toolkit flight recorder log.
F11=Display replication state: Press F11 to show the replication state of the volumes.

WRKCSE - Work with Copy Services Environment

This interactive command allows creation and editing of Flash Copy, Metro-Mirroring, Global-Mirroring and LUN environments defined for use within IBM i Copy Services Manager. Menus are also provided for basic operations. The information is retained in cluster Device Domain Data associated with IBM i Copy Services Manager.

There are no command parameters.

WRKCSECRDL - Work with CSE Credential List

This interactive command allows addition, editing, and removal of user credentials needed to establish SSH sessions with a host. Normally the host will be either a SAN Volume Controller (SVC) or a Hardware Management Console (HMC). The information is encrypted using a cryptographic derived 128 bit AES key and retained in cluster Device Domain Data associated with IBM i Copy Services Manager.

There are no command parameters.

WRKCSEDTA -Work with IBM i CSE Data

This interactive command provides basic operation to manage CSE Data used for Metro-Mirroring, Global-Mirroring, and Flash environments defined for use within IBM i Copy Services Manager.

Select: Select which CSE Data entries are to be displayed. The possible values are:

- ***CRG:** Display only the CSE Data entries that are complete.
- ***ALL:** Display all CSE Data entries.

WRKSTRPRSC -Work with Startup Resources

This interactive command provides the information for the toolkit to use. It should be executed on the source LPAR. The information is used by commands FIXSTRPRSC, CFGSTRPRSC, and STRFSFLASH.

There are five main panels on WRKSTRPRSC which can be selected on the command or cycled through using F11. The common element that ties the records together is the LPAR Usage. Once the toolkit (FIXSTRPRSC, CFGSTRPRSC) determines the current LPAR Usage, the rules on the remaining panels apply to when the LPAR is used in that role. Some LPAR Usage identifiers allow for multiple variations; for example, you can have two preferred source LPARs (*PT01, *PT02) depending on which serial number it is running on.

If the toolkit cannot determine the current role of the LPAR it will exit without performing any modifications.

The panels are:

- ***SYS:** Determine the LPAR Usage based on the serial and LPAR numbers.
  - ***PSxx:** Preferred Source – where your production LPAR is normally running.
  - ***PTxx:** Preferred Target – where the target LPAR can run after a roleswap
  - ***SGxx:** Safeguarded Copy – used when recovering from a Safeguarded Copy.
  - ***FCxx:** Flashcopy Target – where the backups will occur after a Safeguarded Copy.
  - **Default CSEDTA**: Used for PVI configurations to determine which configuration to use for flashcopy backups.
- ***CMN**: Determines which communications resources to bring online. All other lines and IP addresses are left offline.
  - **Usage**: Indicate the role in which these resources will be used
    - ***PSxx:** Activated when the serial and LPAR numbers match the system usage for the preferred source.
    - ***PTxx:** Activated when the serial and LPAR numbers match the system usage for the preferred target.
    - ***SGxx:** Activated when the serial and LPAR numbers match the system usage for the Safeguarded Copy for recovery.
    - ***FCxx:** Activated when the serial and LPAR numbers match the system usage for Safeguarded Copy for backups.
    - ***CTL1:** The IP for the primary controller. Used by CFGSTRPRSC in a PVI configuration to reach the controllers. STRFSFLASH CTLR(*AUTO) will use this entry first to find a working controller. Hardware Resource Location is not used.
    - ***CTL2:** The IP for the secondary controller. Used by CFGSTRPRSC in a PVI configuration to reach the controllers. STRFSFLASH CTLR(*AUTO) will use this entry to find a working controller if the primary is not reached. Hardware Resource Location is not used.
  - **IP Interface:** Specifies the IP interface to be activated and set to autostart when the LPAR is in the specified role. The IP address must exist on the source LPAR.

    - **Line Description:** Specifies the line description to be activated and set to autostart when the LPAR is in the specified role. The line description must exist on the source LPAR.
        - ***IPADDR:** Use the line description currently assigned to the specified IP address.
        - ***VIRTUALIP:** The specified IP address is virtual and does not have a line description associated with it. Start this IP address after all preceding IP addresses have been started. The Hardware Resource Location must be *NONE.
    - **Hardware Resource Location:** Indicates the hardware resource location that will be associated with the specified line description.
        - ***LINE:** Use the resource location currently associated with the specified line description.
        - ***NONE:** No hardware resource location is specified. This is only valid for *VIRTUALIP
    - **Port:** Specify the port of the resource if necessary.
        - ***DFT:** Use the location code to determine the port number.
    - **BRMS:** Specify whether this IP is to be used for BRMS Restricted State Communications.
- ***ISCSI:** Specifies the iSCSI resources to configure based on the usage of the LPAR. All prior information will be removed during a switch. The system will only configure one iSCSI resource.
    - **Usage:** Configure the resource for the specified LPAR usage.
    - **IP Interface:** The IP interface of the remote iSCSI device.
    - **Port:** The port to use. The default 3260 is typical for FalconStor VTL's.
    - **Target Device IQN:** Enter the IQN on the remote iSCSI to connect to.
    - **Client Device IQN:** Enter the IQN assigned to the IBM i LPAR for the specified role.
- ***STG:** Determines the tape storage resources to configure based on the usage of the LPAR.
    - **Usage:** Configure the resource for the specified LPAR usage.
    - **Device Description:** What the toolkit will vary on and is expected to be used in the backups (WRKCTLGBRM etc).
    - **Device Type:** Indicate whether the device is media library (*MLB) or tape drive (*TAP). If *DEV is specified, the toolkit will use the type which the device description refers to.
    - **Serial Number:** Enter the serial number of the tape resource. The toolkit will find the resource name (TAPMLBxx) associated with this serial number (it must be operational) and change the specified device description to it.
- ***BRM:** Change the specified BRMS attributes. For example, you can change the location associated with a device, or the devices used by a control group, etc based on where the LPAR is IPL'ing. Note that only a handful of objects and attributes can be changed. See also the CHGBRMSOBJ command which has similar functionality.
    - **Usage:** Apply this rule to if the LPAR usage matches.
    - **Object Name:** The name of the BRMS object to change
    - **Object Type:** The type of BRMS object to change
    - **Attribute:** Which BRMS attribute to change
    - **New Value:** The new value for the specified attribute

- ***RTE:** Routing entry changes based on LPAR usage. If there are any new routes for a specific LPAR usage, then *all the existing routes (CFGTCP opt 2) will be removed*. The new routes will be created.
    - **Usage:** Apply the new route if the LPAR usage matches.
    - **Destination:** The route destination (IP or subnet)
    - **Subnet Mask:** The destination subnet mask.
    - **Next Hop:** The next hop in the network for the packet destined for the network
    - **Preferred Interface:** Which interface to use to send these packets
    - **Duplicate Route Priority:** Indicate the duplicate route priority.

The WRKSTRPRSC information is stored on the production LPAR in the following files in QZRDHASM, and should be backed up:

- QZRSTRPCMN
- QZRSTRPBRM
- QZRSTRPSTG
- QZRSTRPRTE
- QZRSTRPSYS

ZAPDDDLCK - Zap DDD lock

This is a utility command for releasing the job and user who holds a cluster-wide lock on a specified artifact. It is deprecated and has been replaced with RLSCSELCK.

Name: Enter the name of the object to display a lock for.

Type: Enter the type of object to display a lock for.