



EMC[®] z/OS Migrator
Version 3.0.0

Product Guide
P/N 300-009-864
REV A01

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748-9103
1-508-435-1000
www.EMC.com

Copyright © 2010 EMC Corporation. All rights reserved.

Published March, 2010

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Chapter 1	Overview	
	Data migration challenge.....	22
	Introducing z/OS Migrator.....	22
	Technical benefits of z/OS Migrator and logical migration.....	22
	Business benefits of z/OS Migrator logical migration.....	23
	Migration technology.....	24
	Volume Mirror and Volume Migrator.....	27
	Volume Migrator.....	27
	Volume Mirror.....	27
	How z/OS Migrator fits in your environment.....	28
	z/OS Migrator components.....	30
	Multi-LPAR considerations.....	32
	Control interfaces.....	33
	Supported storage subsystems and dataset types.....	34
	Operating system support.....	35
	Restrictions.....	35
Chapter 2	z/OS Migrator Installation	
	Preinstallation checklist.....	38
	Hardware and software requirements.....	38
	Reserve handling requirements.....	40
	z/OS Migrator storage requirements.....	40
	Downloading maintenance updates.....	41
	Installing/upgrading EMCSCF (Mainframe Enablers).....	43
	Gathering z/OS Migrator installation information.....	45
	z/OS Migrator installation steps.....	47
	Step 1. Load ZOSM <code>vr</code> m.XMITLIB to disk.....	48
	Step 2. Run ZOSM <code>vr</code> m.XMITLIB(#EXTRACT).....	50

Step 3. Customize the RIMLIB JCL	50
Step 4. Run the installation jobs	54
Step 5. Cleanup	54
Step 6: Enter z/OS Migrator licensed feature code	55
Step 7: Apply maintenance	55
Step 8: Shut down and restart Mainframe Enablers.....	56
Step 9: Prepare to execute z/OS Migrator	56
Step 10: Allocate the z/OS Migrator system information dump dataset.....	58
Step 11: Allocate the z/OS Migrator database	60
Step 12: Configuring the z/OS Migrator server parameters	61
Step 13: Customize and run the z/OS Migrator server	63
Step 14. Prior to first execution, specify parameters in SAMPLIB	68
Step 15: Starting and stopping the z/OS Migrator server....	69

Chapter 3 Logical Migration Planning

Plan the migration	74
Establish a migration team.....	74
Pre-migration tasks	75
Pre-migration planning checklist.....	77
Migration action item checklist	78
Migration considerations.....	79
z/OS Migrator dataset selection planning	79
Datasets allocated after activation	85
Dataset extent and multi-volume dataset handling	85
Merging volumes and VTOC sizing considerations	86
Multi-LPAR considerations	86
Reserve handling requirements	87
Restrictions on storage subsystems and dataset types	93
Systems managed storage (DFSMS) allocation.....	94
Migrating VSAM clusters.....	97
Migrating DB2 cluster datasets	97
Hard-coded volume references	97
Catalog structure	97
Small versus large migration groups.....	98
Suspending/resuming data migrations.....	98

Chapter 4 Volume Mirror/Migrator Planning

- Overview 102
- Volume Mirror 102
 - Volume group creation103
 - Activation, monitoring, and mirroring.....106
 - Suspend and resume108
 - Multi-server communication.....110
- Volume Migrator 112
 - Components.....112
 - Volume Migrator operation113

Chapter 5 Running z/OS Migrator

- Running the z/OS Migrator ISPF Monitor..... 118
 - To start and customize the z/OS Migrator ISPF Monitor...118
- Migrating datasets and volumes..... 120
 - The z/OS Migrator migration process defined120

Chapter 6 Using the z/OS Migrator ISPF Monitor

- Overview 134
 - z/OS Migrator Functions panel.....135
- Option S – Set User Session options 137
 - Starting the z/OS Migrator ISPF Monitor for the first time137
- Option 1 – Manage Groups..... 140
 - Primary commands141
 - Line commands143
 - Creating a new group definition144
 - Defining a Logical Migration Group.....145
 - Defining a Volume Migrate Group157
 - Defining a Volume Mirror Group160
 - Saving a group definition162
 - Verifying a group definition.....163
 - Promoting a group definition165
 - Modifying a group definition165
- Option 2 – Monitor Promoted Groups..... 168
 - Line commands for groups.....169
 - Mirror/Migrate volume configuration display.....171
 - Logical Migration dataset configuration display172
 - Activating a group175
 - Suspending a group.....176
 - Resuming a suspended group176

Diverting a Logical Migrator group	176
Completing a Logical Migrator group	178
Deactivating a group.....	178
Option 3 - Display Configuration Information	179
Option 4 - Operator Interface	180
Option 5 – Display Host Messages	181
Option 6 – System Change Summary	183
Option 7 – Message Help	184
Option 8 – Monitor Command Line Help.....	187
Option 9 – Display Installation Options	189
Option V– View or Browse Log Dataset.....	190

Chapter 7 **z/OS Migrator Command Line Interface**

Overview.....	192
General syntax.....	193
Help facility	194
HELP	194
z/OS Migrator active commands.....	195
ACTivate.....	195
COMplete	195
DEACTivate	196
DElete.....	196
DIVert	196
DROP	197
RESume	197
SHUTDOWN	197
SPIN	198
STOP	198
SUSpend	199
TERM	199
z/OS Migrator environmental commands	200
IDLE	200
MAXCIO	200
MAXDIO	201
REMOVE	201
SET	202
z/OS Migrator display commands	207
Display DEBUG.....	207
Display DEV	208
Display DSDDET	208
Display DSN	209
Display EXT	210

Display EXTMAP	212
Display GROUP	212
Display HOSTS	213
Display JOBNames	213
Display PATHS	213
Display PROC	214
Display TRACE	214
Display VOLUMes	215
z/OS Migrator diagnostic commands	216
COMPARE	216
DEBUG	217
DUMP	217
INTERFACE	218
RELEASE	218
TRACE	218
VERify	219

Chapter 8 Defining z/OS Migrator Groups via Text Editor

Overview	222
General syntax rules	222
Syntax for defining logical migration groups	223
GROUP	224
SRCVLIST	227
SET	227
RENAME Unconditional	232
Defining groups for a volume migration	234
Syntax for defining mirror groups	235
GROUP	235
PAIR	237
Syntax for defining migrator groups	240
GROUP	240
PAIR	241
SWAPOPT	243

Chapter 9 z/OS Migrator Examples

Overview	246
Setting session options	246
Migrating a single dataset, non-SMS	249
Storage environment	249
Migration objective	249
Defining a Migration group	249

Migrating a single multi-volume dataset, non-SMS.....	252
Storage environment.....	252
Migration objective.....	252
Defining a Migration group.....	253
Migrating multiple datasets, non-SMS.....	255
Storage environment.....	255
Migration objectives.....	255
Defining a Migration group.....	256
Migrating multiple datasets, SMS volumes.....	261
Storage environment.....	261
Migration objectives.....	261
Defining a Migration group.....	262
Volume Mirror example.....	265
Volume Migrator examples.....	275
Split offline example.....	275
Constant Copy example.....	285

Chapter 10 z/OS Migrator Batch Utility

Overview.....	296
Execution parameters.....	297
Group definition dataset.....	298
SYSIN commands.....	299
<i>VERIFY</i> group-member-name.....	299
<i>PROMOTE</i> group-member-name.....	299
<i>PROMOTE DLM=xx</i>	299
<i>IF RC</i> operator value.....	299
<i>IF GROUP</i> groupname [<i>NOT</i>] <i>EXIST</i>	299
<i>IF GROUP</i> groupname <i>STATUS</i> [<i>NOT</i>] groupstatus.....	300
<i>WAIT nnnnn</i>	300
<i>WAIT GROUP</i> groupname <i>STATUS</i> [<i>NOT</i>] groupstatus.....	300
<i>MAXWAIT nnnnn</i>	301
<i>QUIT</i>	301
Usage examples.....	302
Logical sample batch jobs.....	303
ZOSM5070.....	303
ZOSM5071.....	307
Volume sample batch jobs.....	309
Sample SYSPRINT output.....	310

Appendix A Determining the CPU Serial Number

Determining the CPU serial number..... 312

Appendix B Detailed Parameter Reference

Execution parameters 314

- CFG314
- CPFX314
- CONGROUP.....316
- DB316
- DEBUG317
- DIVSUB317
- PARMLIB318
- REFRESH318
- SSID319
- START319
- SUBSYS320
- TRACE321

Global configuration parameters 322

- CONSTANT_COPY_INT322
- DWA_ADD_PAGE323
- HEARTBEAT_INT324
- IDLE325
- LMPFX326
- MAX_CHANNEL_IO327
- MAX_DEVICE_IO328
- MAXIO328
- MAXTRK329
- OPTIONS 330
- SSUB#331
- SUBNAME331
- Choosing a subsystem identifier.....332
- SWAPOPT333

Appendix C Authorization Return Codes

Authority checking 336

Reading/writing the z/OS Migrator security record 337

Appendix D Catalog Structure

Overview 340

Catalog management considerations 342

- BCS record types 342

	VSAM volume record (VVR)	342
	Catalog diagnostic recommendations	343
	BCS/VVDS/VTOC synchronization	343
	Dataset alias synchronization	346
	BCS/VVDS structure check	348
	BCS integrity check.....	349
	Sample Catalog Solution DUMP JCL	349
Appendix E	Diagnostic Information and Utilities	
	Reporting results.....	352
	z/OS Migrator Group syntax (for Logical Migrations only) ...	352
	Requested information for diagnostics	352
	z/OS Migrator unresponsive, no apparent activity	357
	Interrogating the dataset allocation status.....	357
	Requesting an address space dump	357
	SYSMDUMP dump	358
	z/OS Migrator diagnostic facilities.....	360
Appendix F	Interactions with EMC Software Products	
	Introduction	362
	AutoSwap	362
	Catalog Solution.....	362
	Compatible Flash	362
	Concurrent Copy	362
	Consistency Groups	363
	Controller-based data movers.....	365
	ResourcePak Base	365
	TimeFinder/Clone Mainframe Snap and Concurrent Sessions	365
	TimeFinder Clone and z/OS Migrator operations	366
	TimeFinder and z/OS Migrator operations.....	369
	SRDF and z/OS Migrator operations	372
Appendix G	EMCSAFI Security Interface	
	Overview.....	376
	Class names and resource names used.....	377
	ACF2 Requirements	378
	Disabling the security interface	378
	The EMCSAFD routine.....	378
	Disabling the EMCSAF interface.....	379
	Customizing the security interface	380

The EMCSAFI routine	380
EMCSAFRB structure	382
Enhanced SAF device number checking	387
Specifying the SAF resource name	387

Glossary

	Title	Page
1	z/OS Migrator business value	24
2	Metadata challenges	25
3	z/OS Migrator host-based data migration	28
4	The z/OS Migrator components	30
5	EMC JCL Customization Utility panel	52
6	Sample of z/OS Migrator SYSMDUMP GDG creation JCL	59
7	Sample of z/OS Migrator database creation JCL	60
8	Sample z/OS Migrator server startup configuration file	62
9	z/OS Migrator server JCL	64
10	Reserve handling	88
11	Volume group creation	103
12	Activation, monitoring, and mirroring	106
13	Suspend and resume	108
14	Specify the Session parameters panel	119
15	z/OS Migrator functions panel topology	134
16	z/OS Migrator Functions panel	135
17	Specify the Session parameters panel	137
18	Manage Group Members panel	140
19	Build Type pop-up	144
20	Define z/OS Migrator Group panel	145
21	Define z/OS Migrator Group - Source Dataset Options panel	146
22	Define Source Volume List panel	151
23	Specify Exclude Masks panel	151
24	Specify Rename Unconditional Masks panel	152
25	Target Specifications panel	152
26	Generated group definition	155
27	Build Type panel	157
28	Specify Migrate Group parameters panel	157
29	Build Type panel	160
30	Specify Migrate Group parameters panel	160

	Title	Page
31	Define z/OS Migrator Group panel	162
32	Verify group definition	164
33	Monitor Promoted Groups panel	168
34	Device List panel	171
35	Dataset List panel	172
36	Confirm Activate Request panel	175
37	z/OS Migrator Display Group parameters	177
38	Message pop-up panel	182
39	Host system messages	182
40	System Change Summary option	183
41	z/OS Migrator message display	184
42	z/OS Migrator message display example	185
43	Monitor command line help	188
44	Installation options	189
45	Dataset logs	190
46	Sample EMCSAFD	379
47	EMCSAFRB macro	383

	Title	Page
1	z/OS data components	26
2	Symmetrix system requirements	38
3	Mainframe hardware and software requirements	39
4	RIMLIB contents	50
5	SYSOPTN defaults and options	57
6	z/OS Migrator system options and defaults	66
7	Pre-migration planning checklist	77
8	Migration action item checklist	78
9	DFDSS masking rules	84
10	z/OS Migrator data migration process	120
11	Configuration and control parameters	138
12	Define z/OS Migrator Group panel - Primary commands	141
13	Manage Group Members panel - Line commands	143
14	Kinds of build types	144
15	Group options	145
16	Source dataset options	146
17	Target parameters	153
18	Migrate Group parameters	158
19	Mirror group parameters	161
20	Modes of groups	168
21	Monitor Promoted Groups - line commands for groups	169
22	Dataset List line commands	173
23	Dataset List filters	174
24	z/OS Migrator message severity table	185
25	Deciphering D M=CPU output	312
26	z/OS Migrator authority checking return codes	336
27	z/OS Migrator security record return codes	337
28	Dataset entries	341
29	Consistency Group migration requirements	364
30	TF Clone and Migrator operations before migration promotion	366

	Title	Page
31	TF Clone and z/OS Migrator operations between migrator source and target volumes after migration activation	367
32	TF Clone and z/OS Migrator operations between migrator and non-migrator volumes after migration activation	368
33	TimeFinder and z/OS Migrator operations inside the Migrator Group before promotion	369
34	TimeFinder and z/OS Migrator operations inside the Migrator Group during activation	370
35	TimeFinder and z/O Migrator operations outside of the Migrator Group	371
36	SRDF and z/OS Migrator operations before promotion	372
37	SRDF and z/OS Migrator operations after activation	373
38	Resource validation requests for z/OS Migrator	377
39	EMCSAFI routine: Return code zero	381
40	EMCSAFI routine: Return code non zero	381
41	Authority level	385
42	Requestor information	385
43	Resource information	385
44	Return codes	386

As part of an effort to improve and enhance the performance and capabilities of its product lines, EMC® periodically releases revisions of its hardware and software. Therefore, some functions described in this document may not be supported by all versions of the software or hardware currently in use. Your product release notes contain the most up-to-date information on product features.

If a product does not function properly or does not function as described in this document, please contact your EMC representative.

Note: This document was accurate as of the time of publication. However, as information is added, new versions of this document may be released to the EMC Powerlink website. Check the Powerlink website to ensure that you are using the latest version of this document.

Audience

This guide is intended for EMC customers to perform online migrations of their active mainframe datasets between logical volumes, and to perform online mirroring and migration of full volumes.

This guide provides installation, configuration, implementation, and usage information for the EMC z/OS Migrator.

Related documentation

Related documents include:

- ◆ *EMC z/OS Migrator Message and Code Guide*
- ◆ *EMC z/OS Migrator Release Notes*

For additional information about EMC documentation, contact your EMC sales representative or go to the EMC Powerlink website at:

<http://Powerlink.EMC.com>

Conventions used in this document

EMC uses the following conventions for special notices.

Note: A note presents information that is important, but not hazard-related.



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment.



IMPORTANT

An important notice contains information essential to operation of the software.

Typographical conventions

EMC uses the following type style conventions in this document:

Normal

Used in running (nonprocedural) text for:

- Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus)
- Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, filenames, functions, utilities
- URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, notifications

Bold:

Used in running (nonprocedural) text for:

- Names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, man pages

Used in procedures for:

- Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus)
- What user specifically selects, clicks, presses, or types

<i>Italic:</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> • Full titles of publications referenced in text • Emphasis (for example a new term) • Variables
<code>Courier:</code>	Used for: <ul style="list-style-type: none"> • System output, such as an error message or script • URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
Courier bold:	Used for: <ul style="list-style-type: none"> • Specific user input (such as commands)
<i>Courier italic:</i>	Used in procedures for: <ul style="list-style-type: none"> • Variables on command line • User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces indicate content that you must specify (that is, x or y or z)
...	Ellipses indicate nonessential information omitted from the example

Where to get help

EMC support, product, and licensing information can be obtained as follows.

Product information — For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to the EMC Powerlink® website (registration required) at:

<http://Powerlink.EMC.com>

Technical support — For technical support, go to EMC Customer Service on Powerlink. To open a service request through Powerlink, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send your opinion of this document to:

techpubcomments@EMC.com

This chapter provides an overview of the EMC z/OS Migrator, including functions and components. Topics are:

- ◆ Data migration challenge 22
- ◆ Introducing z/OS Migrator 22
- ◆ Volume Mirror and Volume Migrator 27
- ◆ How z/OS Migrator fits in your environment 28
- ◆ Supported storage subsystems and dataset types 34

Data migration challenge

As data storage capacity requirements continue to grow, the demand for 24/7 data availability increases. At the same time, there is a strong need to control costs. Data centers face a dilemma: the introduction of new storage technology is traditionally disruptive. The need for upgrades and reconfigurations conflicts with the need to ensure maximum availability of the data and, in turn, application availability.

Traditionally, customers have required tools to support the events that drive data migrations, including such needs as:

- ◆ **Technology refreshes**, at which time new storage devices are brought in to replace existing storage devices.
- ◆ **Disk maintenance**, where storage devices must be off-loaded to allow maintenance activities to be performed on a storage device.
- ◆ **Performance related issues**, where data must be migrated to better balance activity within or across storage devices.

Introducing z/OS Migrator

The EMC z/OS Migrator is a host-based data migration facility that provides the ability to perform traditional volume migrations as well as host-based volume mirroring. Together, these capabilities are referred to as the Volume Mirror and Migrator functions of z/OS Migrator. Volume Mirror and Migrator specific capabilities are described on [page 27](#).

In addition, z/OS Migrator can also migrate active mainframe datasets at the logical dataset (extent) level from one set of volumes to another, without any application downtime. This capability is referred to as Logical Migration, and z/OS Migrator can perform this migration in a non-disruptive manner to applications and business continuity.

Technical benefits of z/OS Migrator and logical migration

With z/OS Migrator, you can:

- ◆ Introduce new storage subsystem technologies with minimal disruption of service.

- ◆ Allow users to easily reclaim z/OS UCBs by simplifying the migration of datasets to larger volumes (combining volumes).
- ◆ Facilitate data migration while applications continue to run and fully access data being migrated, eliminating application downtime usually required when migrating data.
- ◆ Eliminate the need to coordinate application downtime across the business, and eliminate the costly impact of such downtime on the business.
- ◆ Improve application performance by facilitating the relocation of poor performing datasets to lesser used volumes/storage arrays.
- ◆ Ensure all metadata always accurately reflects the location and status of datasets being migrated.

Business benefits of z/OS Migrator logical migration

Traditionally, when performing data migrations at the dataset level, users have been forced to deal with application downtime during the actual relocation process. Though excellent facilities exist for moving data, they require the exclusive usage of the dataset during the migration process.

z/OS Migrator differs from these other migration techniques in a major way: it does not require the exclusive usage of any dataset during the migration process.

Due to this, the following business benefits can be gained through the implementation of z/OS Migrator:

- ◆ **Elimination of limited migration windows**

Migration windows typically take place on weekends. Time allotted is usually very small, possibly several hours in duration. During this time, applications are taken down, and then data can be moved. In contrast, z/OS Migrator allows you to move the data at leisure at any time of the day, turning a once-a-week migration window from 2 hours per week to the full week, or a 168-hour migration window.

◆ **Eliminating downtime**

By eliminating the migration window, businesses no longer have to plan to have extended periods of time where their applications are down. When applications are down, business is not performed as normal. This can be costly to a business. z/OS Migrator eliminates the need to have business outages.

◆ **Eliminating off-hour work**

Allowing migrations to be performed at any time allows storage administrators, operations staff, and business application support personnel the ability to work normal business hours. Additionally, this lowers the overall cost for the business by eliminating off-hours work.

◆ **Reduced hardware expenditures**

Since data can be migrated much more efficiently and quickly with z/OS Migrator, storage resources are freed up much faster. Whether the source devices are coming off-lease, or being redeployed, this results in savings.

Figure 1 illustrates the business value of z/OS Migrator.




Figure 1 z/OS Migrator business value

Migration technology

Volume level data migration facilities move logical volumes in their entirety. z/OS Migrator volume migration is performed on a track for track basis without regard to the logical contents of the volumes involved. Volume migrations end in a volume Swap which is entirely non-disruptive to any applications using the data on the volumes.

Frequently, however, it is advantageous to relocate datasets at a more granular level. So, in addition to volume migration, z/OS Migrator provides for logical migration, that is, the migration of individual

datasets. In contrast to volume migration functions, z/OS Migrator performs dataset migrations with full awareness of the contents of the volume, and the metadata in the z/OS system that describe the datasets on the logical volume.

Figure 2 describes the metadata that must be dealt with when dynamically relocating a dataset.




Figure 2 Metadata challenges

Table 1 describes the z/OS data components.

Table 1 z/OS data components

z/OS component	Description
VTOC	Contains information describing volume contents, including information about specific datasets (number of extents, size of extents, starting locations, and so forth).
VVDS	Introduced with ICF Catalogs, this contains similar information about a dataset as is contained within the VTOC. Additionally, information such as details on VSAM datasets such as RBA (relative block addresses) is also kept here.
Catalog	There can be and usually are multiple catalogs in a z/OS environment. A dataset is most typically cataloged, and this tells the system where the dataset is located (volume or volumes on which it resides). Catalog information is not only kept on a DASD volume, but is also kept in memory in the catalog address space.
Coupling facility	If Enhanced Catalog Sharing is in use, then VVR entries from the VVDS may be cached in the Coupling Facility.
Application address spaces	When an application opens a dataset, control blocks are built in memory that describe the dataset. These control blocks describe the location of the dataset and its extents (physical location on disk). When the application issues a read or write request, the access method being used (such as QSAM or VSAM) will build a channel program, which refers to the device and the physical cylinder and track containing the application data to be accessed.

To facilitate data migration at the logical data level, the data migration facility must understand the z/OS metadata structures that describe the logical data being migrated, as well as be able to manipulate and update the metadata dynamically while applications and the system continue to access and modify the data.

z/OS Migrator's dataset migration functions automatically manage all metadata during the migration process with the result that there is no disruption to active applications using the migrating datasets.

Volume Mirror and Volume Migrator

Volume Mirror and Volume Migrator are part of z/OS Migrator. You define a mirror or migrator member when you set up a new volume migration group. After that the z/OS Migrator panels contain volume specific selections.

Volume Migrator

Volume migration provides host-based services for data migration at the volume level on mainframe systems. It provides migration from third-party devices to Symmetrix devices as well as migration between Symmetrix devices.

Volume Mirror

Volume mirroring provides mainframe installations with volume-level mirroring from one Symmetrix device to another. It uses host resources (UCBs, CPU, and channels) to monitor channel programs scheduled to write to a specified primary volume and clones them to also write to a specified target volume (called a mirror volume).

After achieving a state of synchronization between the primary and mirror volumes, Volume Mirror maintains the volumes in a fully synchronized state indefinitely, unless interrupted by an operator command or by an I/O failure to a Volume Mirror device. Mirroring is controlled by the volume group. Mirroring may be suspended consistently for all volumes in the group.

How z/OS Migrator fits in your environment

Figure 3 outlines how z/OS Migrator fits into your z/OS environment.

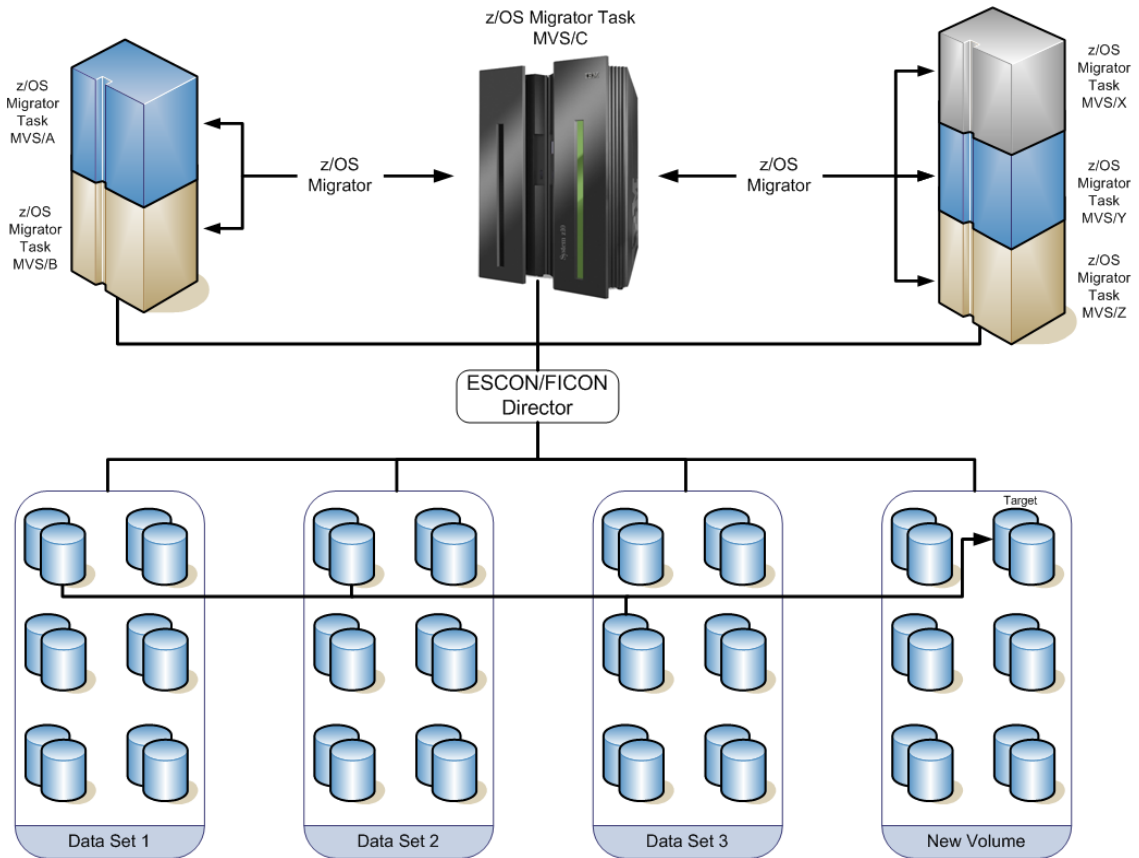


Figure 3 z/OS Migrator host-based data migration

Within your environment, z/OS Migrator offers the following benefits:

- ◆ Full target and source flexibility - All 3380 and 3390 Count Key Data/Extended (CKD/E) capable control units in the data center can participate in a migration.
- ◆ Full system sharing throughout the data center.

- ◆ Full access to the data at any point during a migration operation.
- ◆ Installs dynamically; no IPL is required.
- ◆ Works with pairs of datasets, including the original “source” dataset and the new “target” dataset.
- ◆ Uses both asynchronous and synchronous techniques to copy data from source locations to target locations, thus minimizing performance overhead.
- ◆ Is invisible to applications. Applications continue to access and update the source dataset as normal. z/OS Migrator ensures all updates are mirrored to the target dataset. Eventually, under user direction, z/OS Migrator can *divert* all I/O activity to the target dataset and simultaneously modify all metadata to reflect the new location of the dataset.
- ◆ Supports multiple system data-sharing environments.
- ◆ Guarantees physical data integrity with ongoing internal heartbeat monitoring, error detection, and recovery capability.

z/OS Migrator components

Figure 4 outlines the primary z/OS Migrator components and their functions.

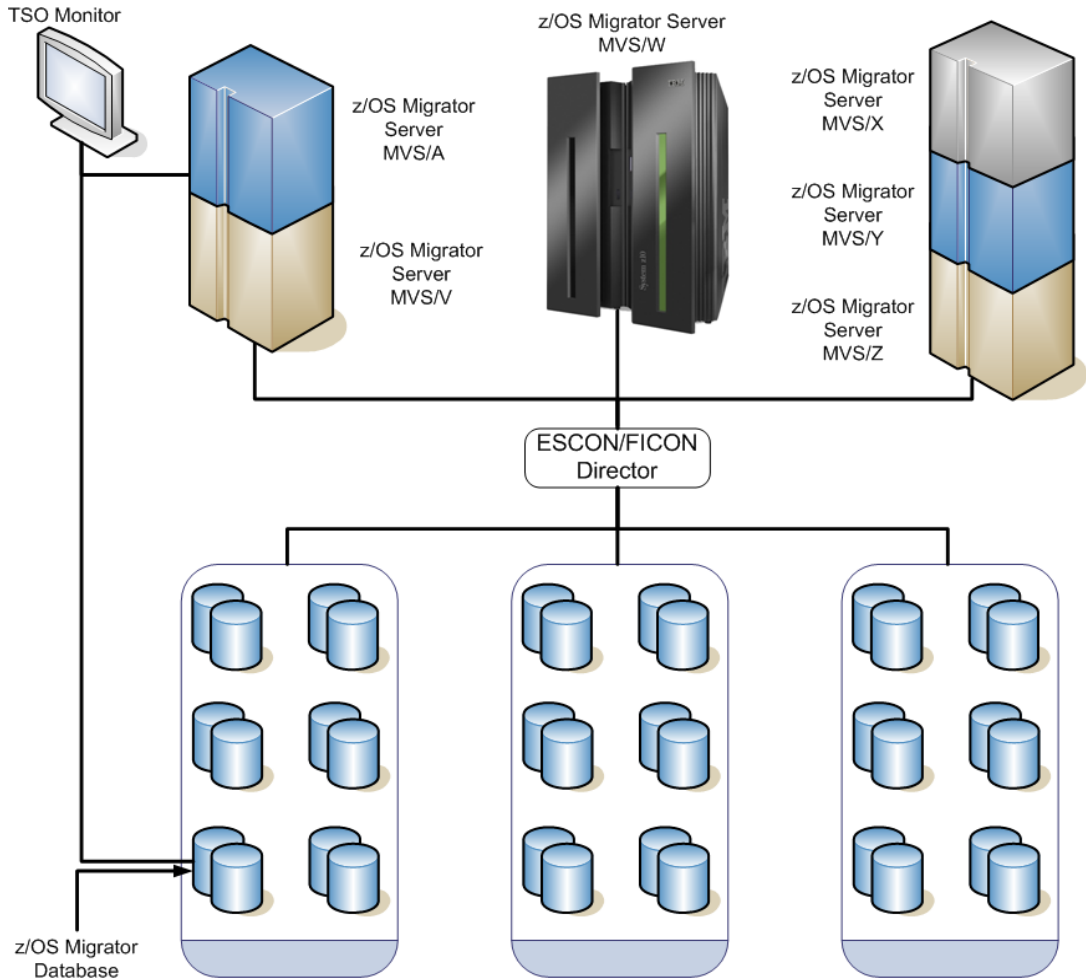


Figure 4 The z/OS Migrator components

z/OS Migrator server

The z/OS Migrator server is the primary component of the z/OS Migrator product framework. The z/OS Migrator server is a z/OS Started Task that must run on each system that has access to the data that will be migrated.

The z/OS Migrator server provides the following functionality:

- ◆ Fulfills data migration requests from other components such as the z/OS Migrator ISPF Monitor. For this reason, the z/OS Migrator server must be up and running before any data migration activities can take place.
- ◆ During a migration, the z/OS Migrator server initiates, coordinates, and performs the following actions:
 - Once z/OS Migrator groups are defined, they are promoted to the z/OS Migrator server, which validates them and stores them in the z/OS Migrator database.
 - Manages migration groups through all stages of a migration across all systems in the migration environment, including *Promotion, Activation, Copy, Synchronization, Diversion, and Completion*. A detailed description is in [“The z/OS Migrator migration process defined” on page 120](#).
 - Copies and synchronizes all migration groups during a migration.

**CAUTION**

The z/OS Migrator server must be installed and operational on all z/OS images that can access the data being migrated. If it is not, then you may be exposed to possible data integrity issues.

z/OS Migrator database

The z/OS Migrator database is the file used to store and share information about z/OS Migrator data migration activity. The z/OS Migrator database is used by the z/OS Migrator server to store and communicate migration information across all participating systems.

z/OS Migrator I/O monitor

The z/OS Migrator I/O monitor is a sub-component of the z/OS Migrator server. The z/OS Migrator I/O monitor is responsible for the following functions:

- ◆ Monitors all I/O activity to source datasets (extents).
- ◆ When in *Mirroring* phase, mirrors all write I/O activity to the target datasets to ensure that they are always synchronized.

- ◆ When in *Diversion* phase, re-directs (diverts) all I/O activity, including reads and writes, to the target dataset until completion has been achieved.
- ◆ Prevents all non-z/OS Migrator I/O to the target datasets until diversion.
- ◆ Monitors source volume VTOCs to ensure that all changes to extents, including additions, deletions, or modifications to such, are dynamically reflected on the target volumes.

Multi-LPAR considerations

z/OS Migrator is designed to be used in a multi-LPAR environment with shared DASD. Generally, a z/OS Migrator server must be active on all LPARs that have host connectivity to the volume pairs involved in the migration or mirroring operation. The one exception to this rule is if you are doing only volume migrations between Symmetrix devices. Otherwise, for all migrations from third-party devices, any volume mirroring operations, and all dataset ("logical") migrations, the z/OS Migrator server must be active on all servers. The reason is that the z/OS Migrator interfaces must be installed and operational to correctly monitor any application I/O that may affect the mirror or migration operation. Only with the interfaces in place, can the data integrity of any volume mirroring or dataset migration be guaranteed.

To a great extent, the interfaces, once installed, operate independently of the server. However, the server provides management of certain storage resources for the interfaces and if a storage request can not be satisfied, the interfaces have no choice but to suspend mirroring for a device pair or even a whole group. In some cases, if the z/OS Migrator server is not available to service a particular request, an application I/O may be blocked until the server is restarted.

Even if there is no current application I/O on an LPAR in the "z/OS Migrator complex," a number of state changing commands require that all participating servers be actively accessing the database. In some cases, you may receive a message indicating that the command system is waiting for another LPAR where the z/OS Migrator server appears to be idle. ["IDLE" on page 325](#) describes the IDLE parameter. If the "idle" LPAR is actually out of service, there is a command which allows processing to continue. ["REMOVE" on page 201](#) describes the REMOVE parameter. Otherwise, corrective action should be taken if required.

- The owner server** The z/OS Migrator Server to which the Activate command is issued is known as the 'owner' host. Certain z/OS Migrator actions for a group, such as Divert for a dataset group, can only be performed from the owner.
- Multi-server communication** In addition to running on all LPARs, all of the z/OS Migrator servers must share the same z/OS Migrator database. The database is used to communicate group status and other cross-system functional requests.
- Note that certain z/OS Migrator actions are serialized through the use of device reserves issued against the device on which the z/OS Migrator database resides. For this reason, you should place the z/OS Migrator database on a lightly-used volume against which no other reserves are issued. You can perform mirror or migrate functions against the database volume or any datasets on that volume.
- Missing volumes** For Volume migration functions only, it may be possible to proceed even if one or more devices are not currently accessible on agent systems. Note that all devices must still be accessible to the activation or owning system. Volumes on devices that are not available to a given LPAR are referred to as "missing" and may appear in certain device oriented displays indicated with '????' where the device address or volume serial number would normally appear.

Control interfaces

You interact with the z/OS Migrator server through three alternative control interfaces:

- ◆ The ISPF monitor
- ◆ Command line interface
- ◆ Batch interface

z/OS Migrator ISPF monitor

The z/OS Migrator ISPF monitor is the primary interface to the z/OS Migrator product. From the ISPF monitor, you can perform the following functions:

- ◆ Manage migration groups by adding, modifying, deleting, and promoting migration groups to the z/OS Migrator server.
- ◆ Interact with the z/OS Migrator server to monitor and actively manage active migrations from *Activation* to *Completion* phase. Moreover, you can use the interface to perform configuration modifications.

[Chapter 6 on page 133](#) describes the ISPF monitor and [Chapter 9 on page 245](#) describes how to use the ISPF screens for z/OS Migrator.

Command line interface

z/OS Migrator also has an operator command interface. Although the ISPF interface is the primary means of communicating with the z/OS Migrator server, the operator commands allow you to initiate essentially all server actions.

Note: [Chapter 7 on page 191](#) describes the command line interface.

The batch interface

You can use the batch utility to generate and promote configuration elements as well as to submit z/OS Migrator commands to the server and retrieve the response. The z/OS Migrator Batch Utility provides simple automation of routine migration tasks.

Note: [Chapter 10 on page 295](#) describes the batch utility usage.

Supported storage subsystems and dataset types

z/OS Migrator only supports 3380 and 3390 architecture, like geometry devices, and provides support for the following dataset types:

- ◆ Direct access (DA) datasets
- ◆ Extended format sequential datasets
- ◆ Extended partitioned datasets (PDSE)
- ◆ GDG base names and GDG datasets
- ◆ Extended format VSAM keyed sequential datasets (KSDS)
- ◆ Partitioned (PO) datasets (TYPE=HFS is not supported)
- ◆ Physical sequential (PS) datasets
- ◆ Striped sequential datasets
- ◆ BDAM datasets
- ◆ VSAM datasets:
 - Alternate Index (AIX)
 - ESDS
 - KSDS
 - Linear

- RRDS
- Spheres (KSDS | ESDS + PATHS + AIX)
- VRRDS
- ◆ Large sequential datasets indicated by DSNTYPE=LARGE (z/OS 1.7 and above)

Operating system support

z/OS Migrator is compatible with all IBM supported versions of the z/OS operating system.

Restrictions

z/OS Migrator has the following restrictions in regards to storage architecture and dataset types.

- ◆ The target controller must be at an equal or higher technology level and be downward compatible. In other terms, you cannot migrate data from a 2105 model controller to a 3990, but you can migrate data from a 3990 model controller to a 2105.

Note: Migration between different device geometries is *not* supported. You can only migrate data between similar geometry devices, for example: 3380 to 3380, or 3390 to 3390.

- ◆ Datasets cataloged on System Resident (SYSRES) volumes: Migration of system datasets, that is, datasets used by the operating system, typically located on system residence volumes, including linklisted datasets, is a complex operation. Be careful.
- ◆ Dataset types *not* currently supported in z/OS Migrator include:
 - VSAM datasets with the IMBED, KEYRANGE and REPLICATE parameters. These are obsolete VSAM parameters. Target dataset extents cannot be allocated using these parameters because z/OS does not support new allocations using them. The physical layout of extents of datasets defined with these parameters is different from the physical layout of datasets without them. Copying extents from a dataset containing these to a newly defined dataset that does not have them yields an unusable dataset.

If you copy an IMBED cluster to NOIMBED (the only type you can currently define), the target data component would have a track full of sequence set records at the start of each CA that VSAM would think was actual data.

- Catalogs
- ISAM
- Individual PDS members (z/OS Migrator does support migrating an entire PDS)
- Page and Swap datasets
- HFS/zFS datasets
- Datasets that are designated as "Unmoveable" are not supported, that is, DSORG=U | PSU
- VTOC, VVDS, and VTOCIX datasets
- Temporary (&&) datasets
- Uncataloged datasets
- EAV volumes/datasets

Note: Uncataloged datasets will fail activation.

- ◆ Volumes containing the following datasets cannot be migrated:
 - z/OS Migrator database dataset
 - Page datasets
 - CA-OPS/MVS datasets
 - Sysplex Couple datasets
- ◆ Volumes not supported as source or targets of a migration:
 - Extended Address Volumes
 - FBA volumes
 - Virtual devices (VDEV)
 - CSC gatekeeper volumes

This chapter describes how to install z/OS Migrator. Topics include:

- ◆ [Preinstallation checklist.....](#) 38
- ◆ [z/OS Migrator installation steps](#) 47

Preinstallation checklist

Before you install z/OS Migrator, take the following steps:

- ❑ Make sure that your system meets the hardware and software requirements listed in this section and in the *z/OS Migrator Release Notes*.
- ❑ Check the EMC Powerlink website for any z/OS Migrator updates or current release notes and download any available updates or release notes.
- ❑ Install/upgrade the EMCSCF component of ResourcePak® Base (if you do not already have the latest version on your host).

Note: z/OS Migrator requires ResourcePak Base (EMCSCF) 7.0 or higher, distributed in the Mainframe Enablers V7.0 (MFE700) Product Kit.

- ❑ Gather preinstallation information as listed in [“Gathering z/OS Migrator installation information”](#) on page 45.

Hardware and software requirements

[Table 2](#) lists Symmetrix hardware and Enginuity™ requirements. Before you install the z/OS Migrator software, make sure your Symmetrix system meets these requirements.

Table 2 Symmetrix system requirements

Item	Requirements
Hardware	All currently-supported Symmetrix systems.
Enginuity release	Enginuity revision level 5874 as well as many prior levels. The minimum supported level is 5x68. (For Enginuity revision level 5x74, patch 49947 is required.)



IMPORTANT

This minimum supported release level is accurate at the time of publication and is subject to change. Please check the Release and End of Life Service Dates on Powerlink.EMC.com for the most current information. Contact your EMC Customer Support Engineer to verify that your system meets these requirements.

Table 3 lists the mainframe hardware and software requirements for z/OS Migrator. Before you install the z/OS Migrator software, make sure your mainframe system meets these requirements.

Table 3 Mainframe hardware and software requirements

Item	Requirements
Hardware	<ul style="list-style-type: none"> • Any system that supports versions of the z/OS operating system currently supported by IBM. • FTP or TSO connection to an open systems host with a CD-ROM drive (if you are installing by download or CD). • Open systems host with a CD-ROM drive (if you are installing by download or CD).
Storage devices and arrays	<p>All 3380 and 3390 z/OS formatted volumes supporting Count Key Data/Extended (CKD/E) format.</p> <hr/> <p>Note: FBA formatted devices are not supported.</p> <hr/>
Software	<ul style="list-style-type: none"> • All z/OS-based operating systems that are currently supported by IBM. ^a z/OS Migrator is not suitable for use with native z/VM and VM/ESA, native VM/XA, native VSE/ESA, native VSE/XA or DOS, and z/VSE. <hr/> <p>Note: z/OS Migrator is not supported in native VM. However, z/OS Migrator can run as a z/OS guest under VM. VM does not allow volumes defined as unsupported to be attached to SYSTEM, or used to IPL a virtual machine. When running as a guest under VM, z/OS Migrator requires special consideration. You must define volumes to VM (SET RDEV) as TYpe UNSUPorted DEVClass DASD DPS Yes RESERVE_RELse Yes, and you must attach volumes to the guest.</p> <hr/> <ul style="list-style-type: none"> • EMC ResourcePak Base (EMCSCF) Version 7.0 or higher, distributed in the Mainframe Enablers V7.0 (MFE700) Product Kit. • TCP/IP protocol stack from IBM Version 3.1 or later. • RACF 1.9 or higher, or an equivalent SAF compliant security product.

a. z/OS Migrator supports the JES2 environment. It does not support the JES3 environment.

Reserve handling requirements

z/OS Migrator requires that all hardware reserves issued by z/OS address spaces, including system, application, and jobs, be converted to software reserves.

Place the following sample entry in the GRSRNLxx member of SYS1.PARMLIB; this would be the only reserve conversion entry required to convert all hardware reserves.

```
RNLDEF RNL(CON) TYPE(PATTERN) QNAME(*) /* do not include in excl list */
```

Refer to IBM Global Resource Serialization documentation for details.

In the event all hardware reserves cannot be converted, the following list of sample entries are a guideline for the required changes. These changes must be evaluated by systems programming staff to ensure compatibility with the current environment.

```
/* - - - - - z/OS Migrator Change SYSTEM to SYSTEMS - - - - - */
RNLDEF RNL(INCL) TYPE(GENERIC) QNAME(SYSDSN)
/* - - - - - z/OS Migrator Convert RES to ENQ - - - - - */
RNLDEF RNL(CON) TYPE(GENERIC) QNAME(SPFDSN) /* do not include in excl list */
RNLDEF RNL(CON) TYPE(GENERIC) QNAME(SPFEDIT) /* do not include in excl list */
RNLDEF RNL(CON) TYPE(GENERIC) QNAME(SYSIGGV2) /* do not include in excl list */
RNLDEF RNL(CON) TYPE(GENERIC) QNAME(SYSVTOC) /* do not include in excl list */
RNLDEF RNL(CON) TYPE(GENERIC) QNAME(SYSZVVDS) /* do not include in excl list */
RNLDEF RNL(CON) TYPE(GENERIC) QNAME(EMCTF) /* do not include in excl list */
/* - - - - - z/OS Migrator Exclude Datasets - - - - - */
RNLDEF RNL(EXCL) TYPE(SPECIFIC) QNAME(EMC$FM) RNAME(ZOSM.DATABASE.DSNAME)
```

“Reserve handling requirements” on page 87 provides complete details on reserve handling.

z/OS Migrator storage requirements

The basic ECSA, CSA, and Extended Private storage requirements for z/OS Migrator are as follows:

- ◆ A global area is acquired: X'698' (1,688) bytes of fixed ECSA.
- ◆ Operating system interface modules are loaded into fixed ECSA, and require at least X'DAA08' bytes (874.5 KB).
- ◆ The PC routine is loaded into pageable ECSA, currently X'B3F00' bytes (720 KB).

- ◆ The storage pools (all fixed ECSA) are initialized with X'480000' bytes (4.5 MB). When groups are mirroring or diverting, the storage pools will expand depending on the complexity of the channel programs being intercepted.

Note: At least 0.85 MB will be resident in fixed ECSA until the next IPL for all components, except for the PC routine and possibly the storage pools.

- ◆ When a group is promoted and activated, the parsing routine currently uses ECSA to store images.
- ◆ For each active group, there are (in fixed ECSA):
 - Control blocks for which space is not freed:
 - One DWGRP block, X'A8' (168) bytes, until the next IPL, or z/OS Migrator is shut down.
 - One DWGDEV block for each source and target device, X'F8' (248) bytes.
 - A sparse array for extent look-up from each device block. Minimum size is X'400' bytes (1 KB).
 - Control blocks for which space is freed when the group completes:
 - One QCDSN block for every dataset and each individual VSAM component or AIX, X'2B8' (696) bytes.
 - One DWEXT block for each dataset extent, X'68' (104) bytes.

Downloading maintenance updates

If there is no current maintenance, keep these instructions for future use when you do need to download maintenance.

You can download the latest maintenance updates and current release or service notes (identical to release notes) from the EMC Powerlink website:

<http://powerlink.emc.com>

Note: Keep in mind that you must register on Powerlink as a valid EMC customer before you can access Powerlink.

On the page for your product, you will see files for different product versions. For your version, you may see the following types of files:

- ◆ **ReadMe_*prod*vr*m*.txt** - contains starter instructions for the electronic kit.
- ◆ **ReadMe_*prod*vr*m*_Fixes.txt** - contains information about the release.
- ◆ **Service_Notes_*prod*vr*m*.pdf** (or **.txt**) - contains information discovered after initial release of the product.
- ◆ ***prod*vr*m*_fixes.zip** - contains the previous two documents as well as a software patch file and instructions about how to apply this maintenance.

where:

prod = The product prefix.

*vr*m** = The version, revision level, and modification level of the software you want.

Take the following steps to download these files:

1. Log in to:

<http://powerlink.emc.com>

2. Navigate to the **Downloads and Patches** pages for your product.

Result: On the page is a table of zip files and document files.

3. Take either of the following steps:

- To download a copy of a document, click either **ReadMe_*prod*vr*m*_Fixes.txt** or **Service_Notes_*prod*vr*m*.pdf** or **.txt**.

Note: **Service_Notes_*prod*vr*m*.pdf** is an Adobe Acrobat document. To view documents on Powerlink, you must have Adobe Reader installed. If you do not have Adobe Reader, you can install it from Powerlink at no charge. Just follow the path displayed on Powerlink.

- To download the zip file, click ***prod*vr*m*_fixes.zip**. Download the zip file to your home system, unpack the zip file, and follow the instructions it contains.

4. When you are finished reading or copying, return to previous pages to view other products and services, or choose logout from the menu bar at the top of any Powerlink page.

Note: Do not apply any maintenance until after z/OS Migrator is accepted. [“Step 7: Apply maintenance” on page 55](#) describes maintenance procedures.

Installing/upgrading EMCSCF (Mainframe Enablers)

If you are using the volume migration feature, you need to have SCF active and an implied or explicit association between z/OS Migrator and SCF, either by default or via the `/SCF$xxx DD` statement.

z/OS Migrator requires that you install Mainframe Enablers and any currently available maintenance before you install z/OS Migrator.

Note: SCFINI must contain the statement: `SCF.DAS.ACTIVE=YES`. This is required for volume migrations for the SWAP to take place at the conclusion of a volume migration.

If the required version of Mainframe Enablers is not installed before you install z/OS Migrator, the SMP/E APPLY job fails with a missing PREREQ. If this happens, you must install the required version of Mainframe Enablers before you can rerun the SMP/E APPLY job.

Note: Only the most recent version of EMCSCF will install on your system.

If you do not have Mainframe Enablers running on the system, or do not have the required version running on your system, you can download the Mainframe Enablers kit from EMC Powerlink.

Note: The *EMC z/OS Migrator Release Notes* contain current additional requirements.

Shutting down an existing version of EMCSCF

An older version of EMCSCF may already be running on the systems on which you are installing z/OS Migrator, so it is important that you perform an orderly shutdown of that version of EMCSCF and remove it before you install a new version of Mainframe Enablers.

You can remove the existing EMCSCF from memory by entering:

```
F server,INI,SHUTDOWN
```

If MSC¹ is active, before you can shutdown EMCSCF, you must enter:

```
F server,MSC DISABLE
```

-
1. In SRDF/A environments, the Multi-Session Consistency (MSC) task provides consistency across multiple Symmetrix systems for SRDF/A groups. Multi-Session Consistency is enabled by a Licensed Feature Code.

Installing the new versions of EMCSCF

Take the following steps to download and install Mainframe Enablers from Powerlink:

1. Log in to a privileged account on an open systems host (root on UNIX or administrator on Windows).
2. Allocate a working directory on the open system for installation.
3. Using your browser, open:
<http://powerlink.emc.com>
4. From the Powerlink home page, navigate to the **Downloads and Patches** page for Mainframe Enablers.

Note: If you are not able to access this location, you may not have registered your software or registered it incorrectly. Follow the prompts to register your software, correct your registration, or contact EMC in the event of a problem.

5. Click the version of Mainframe Enablers you want to download. The product version consists of a zip file that contains the installation kit and the installation instructions.
6. After you download the installation kit and instructions, choose logout from the menu bar at the top of any page.
7. Unpack the zip file and follow the instructions within the zip file to install Mainframe Enablers.
8. After z/OS Migrator is installed and customized, restart the most current version of EMCSCF.

Running multiple EMCSCF copies

You can run multiple instances of EMCSCF as separate subsystems. You may find this desirable when you are testing new versions of EMCSCF or EMCSCF-enabled products.

To run multiple instances of EMCSCF as separate subsystems, add the following DD statement to the EMCSCF test procedure.

```
//SCF$nnnn DD DUMMY
```

where:

nnnn

Defines this instance of EMCSCF as a unique z/OS subsystem. The DD statement would then be used in any task where you want to use this copy of EMCSCF.

For example:

Test version of EMCSCF

```
//EMCSCF EXEC PGM=SCFMAIN,TIME=1440,REGION=0M
//STEPLIB DD DISP=SHR,DSN=test.load_library
//SCFINI DD DISP=SHR,DSN=init_dataset
//SYSABEND DD SYSOUT=*
//SCF$0100 DD DUMMY
```

Any task needing to use this instance of EMCSCF would add a `//SCF$nnnn DD DUMMY` statement. If a version of EMC z/OS Migrator needed to use this version of EMCSCF, the ZOSMPROC JCL for EMC z/OS Migrator would use the matching DD statement.

Gathering z/OS Migrator installation information

Before beginning the z/OS Migrator installation, identify, or decide upon, the following items:

CLIST library and EDIT macro

You need to determine the name of a CLIST library where you can store the edit macro. You also need to determine a name for the edit macro created by the installation dialog.

Product dataset name prefix

You need to choose the dataset prefix you will use to install z/OS Migrator. Names for the product datasets consist of a final qualifier, such as LINKLIB, and a dataset prefix. For example, if you choose a dataset prefix of EMC.ZOSMvrm, the LINKLIB dataset is named:

```
EMC.ZOSMvrm.LINKLIB
```

EMC recommends that you use EMC.*product* (that is, EMC.ZOSMvrm) if it agrees with your site standards.

Ensure that you have RACF ALTER authority (or the equivalent from another security manager) for the datasets created with this dataset prefix.

Note: Hereafter, datasets created using this dataset prefix are referred to as if they had been created with the suggested value. The actual *fmid* for your installation may be different.

Mainframe Enablers dataset prefix

You need to specify the dataset prefix you used when you installed Mainframe Enablers. EMC recommends that you use *EMC.fmid* if it agrees with your site standards.

SMP/E dataset name prefix

You need to choose the prefix for the SMP/E datasets into which you install z/OS Migrator. If you have installed another EMC product using SMP/E, you should install z/OS Migrator into the same CSI.

For example, if you already have SMP/E maintained EMC products and the SMPLOG dataset is called "EMC.SMPE.SMPLOG," the SMP/E dataset prefix should be "EMC.SMPE."

If you are installing an EMC SMP/E maintained product for the first time, EMC recommends using "EMC.SMPE."

SMP/E datasets volser

You need to choose the disk volume onto which you will install the distribution libraries (required by SMP/E). This may be the same volume as you use for the product libraries. However, many customer sites prefer to keep SMP/E-related datasets on separate volumes from product libraries. An amount of space similar to that needed for the product libraries is required.

Install to disk volser

You need to determine the disk volume onto which you will install the target (that is, runtime) datasets. The space required is nominal.

Disk unit name

Decide upon a disk unit name for the above volumes. For many users, "SYSDA" will suffice. However, use whatever generic or esoteric name your local standards require.

z/OS Migrator installation steps

The z/OS Migrator installation kit consists of a PDS containing TSO TRANSMIT images of files needed to perform an SMP/E indirect-library installation on the product. This PDS is packaged as a TSO TRANSMIT file on a CD or as an electronic download from EMC Powerlink.

To install z/OS Migrator on an IBM mainframe:

- ◆ “Step 1. Load ZOSMvrm.XMITLIB to disk”
- ◆ “Step 2. Run ZOSMvrm.XMITLIB(#EXTRACT)”
- ◆ “Step 3. Customize the RIMLIB JCL”
- ◆ “Step 4. Run the installation jobs”
- ◆ “Step 5. Cleanup”
- ◆ “Step 6: Enter z/OS Migrator licensed feature code”
- ◆ “Step 7: Apply maintenance”
- ◆ “Step 8: Shut down and restart Mainframe Enablers”
- ◆ “Step 9: Prepare to execute z/OS Migrator”
- ◆ “Step 10: Allocate the z/OS Migrator system information dump dataset”
- ◆ “Step 11: Allocate the z/OS Migrator database”
- ◆ “Step 12: Configuring the z/OS Migrator server parameters”
- ◆ “Step 13: Customize and run the z/OS Migrator server”
- ◆ “Step 14. Prior to first execution, specify parameters in SAMPLIB”
- ◆ “Step 15: Starting and stopping the z/OS Migrator server”

The following sections describe these steps in more detail.

Step 1. Load ZOSMvrm.XMITLIB to disk

1. Take one of the following steps:
 - If you are installing z/OS Migrator from a CD, mount the CD on an open systems host. Allocate a working directory on the open system for the installation. Copy the contents of the CD to the working directory.
 - If you are installing z/OS Migrator from an EMC Powerlink download, take the following steps:
 - a. Log in to a privileged account on an open systems host (root on UNIX or administrator on Windows).
 - b. Allocate a working directory on the open system for the installation.
 - c. Log on to the Powerlink website.
 - d. Navigate to **Downloads and Patches**. Then, click on your product.

Note: If you are not able to access this location, you may not have registered your software or registered it incorrectly. Follow the prompts to register you software, correct your registration, or contact EMC in the event of a problem.

Result: You see a page for the product you selected.

- e. Click the product version you want to download. The product version consists of a zip file that contains the installation kit and the installation instructions.
 - f. Download the installation kit into the working directory on the open system.
2. If your current host is a Windows system, unzip the file in the working directory. If your current host is a UNIX system, unzip and untar the file in the working directory.
3. Locate the ZOSMvrm.XMITFILE. This file is in TSO TRANSMIT format and contains a flattened copy of ZOSMvrm.XMITLIB, a PDS that holds other TRANSMIT images, the JCL to extract them, and necessary SMP/E installation files.
4. On the target mainframe, allocate a file to which you can FTP ZOSMvrm.XMITFILE. Use the dataset name prefix you intend to use for product installation.

For example, if you intend to install the product with a dataset name prefix of EMC.ZOSMvrm, as recommended, name the file EMC.ZOSMvrm.XMITFILE.

5. Allocate the dataset with the following characteristics:

```
LRECL=80
BLKSIZE=3120
DSORG=PS
SPACE=(CYL,(36,2))
```

Note: The SPACE parameter here and elsewhere assumes that you are allocating the dataset on a 3390 device.

6. FTP the file to the mainframe in binary format. Your FTP session may look something like the following:

```
ftp hostname
(username and password prompts)
cd ..
25 "'"' is working directory name prefix
binary
200 Representation type is image
put ZOSMvrm.XMITFILE 'EMC.ZOSMvrm.XMITFILE'
```

7. Use TSO RECEIVE to receive the file into a PDS. The PDS is created by the RECEIVE command and does not have to be preallocated. However, you *must* specify a dataset name using the DA[taset] parameter or the file will be allocated using your TSO prefix (usually your logonid). The dataset name specified must have the final qualifier of XMITLIB.

For example:

```
receive indataset('EMC.ZOSMvrm.XMITFILE')
INMR901I Dataset EMC.ZOSMvrm.XMITLIB from userid on nodename
INMR906A Enter restore parameters or 'DELETE' or 'END' +
da ('EMC.ZOSMvrm.XMITLIB')
```

If you did not specify "DA(...)" as above, the dataset would be allocated as *userid.XMITLIB*.

Step 2. Run ZOSMvrm.XMITLIB(#EXTRACT)

Now run ZOSMvrm.XMITLIB(#EXTRACT) to extract *ds-prefix*.RIMLIB and the SMP/E indirect libraries. Take the following steps:

1. Edit the #EXTRACT member of the newly RECEIVED library. You can edit the #EXTRACT job by running the SETUP REXX program you can find in the XMITLIB dataset. The SETUP REXX program prompts you for all of the information needed to edit the JOB.

If you wish to edit the JOB manually, make the following changes:

- a. Change the job card to one that conforms to your standards.
 - b. Globally change *ds-prefix* to the dataset prefix of this library (which will be the dataset prefix for the product libraries).
 - c. Globally change DVOL to a disk volser that can be used to hold the extracted libraries.
 - d. Globally change DISK_UNIT to a site-appropriate unit name.
2. Submit #EXTRACT. Step completion codes should be 0, except for the DELETE step, which will have a step completion code of 8 unless the job is a rerun.

Step 3. Customize the RIMLIB JCL

The RIMLIB library (*ds-prefix*.RIMLIB) is a PDS containing JCL to install the product. After you extract the RIMLIB PDS, you find that RIMLIB has the contents shown in [Table 4](#).

Table 4 RIMLIB contents (1 of 2)

File	Contents
#01ALLOC	Allocate target and distribution libraries
#02DDDEF	Add or replace product library DDDEFS to SMP/E CSI
#03RECEV	SMP/E RECEIVE z/OS Migrator function into global zone
#04APPLY	SMP/E APPLY z/OS Migrator function into target zone
#05ACCPT	SMP/E ACCEPT product sysmods into distribution zone
#06CLEAN	Deletes indirect libraries and DDDEFs used for them

Table 4 RIMLIB contents (2 of 2)

File	Contents
#90SAFJB	JCL to assemble and link z/OS Migrator with EMCSAFD ^a
#92SAFJB	JCL to re-enable the use of the EMC supplied EMCSAFI security module.
#99MAINT	SMP/E RECEIVE and APPLY service
SETUP	JCL customization dialog launcher
ZOSJCL	JCL customization dialog driver
ZOSWIN1	JCL customization dialog data entry panel

a. Member EMCSAFI is the security interface. Member EMCSAFD contains assembler source code that you can use to disable the EMCSAFI security interface.

Take the following steps to customize the installation JCL using the automated dialog:

1. Edit the RIMLIB library (*ds-prefix.RIMLIB*).
2. Locate the member named SETUP on the member selection list, type **EX** in the selection column next to it and press **Enter**.

```

Menu Functions Confirm Utilities Help
-----
EDIT          EMC.ZOSMvrm.RIMLIB          Row 00001 of 00013
Command ==>          Scroll ==> CSR
      Name          Prompt  Size  Created      Changed      ID
-----
#01ALLOC          45  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
#02DDDEF          51  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
#03RECEV          22  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
#04APPLY          22  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
#05ACCP          22  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
#06CLEAN          53  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
#90SAFJB          28  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
#92SAFJB          60  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
#99MAINT          27  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
ex_____  SETUP          13  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
_____  ZOSJCL          206 yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
_____  ZOSWIN1         51  yyyy/mm/dd  yyyy/mm/dd  hh:mm:ss  idstring
**End**

```

The customization panel shown in [Figure 5](#) is displayed.

```

                                EMC JCL Customization Utility
COMMAND ==> _____

Type EXEC on the command line and press ENTER to proceed, or PF3 to exit.

CLIST library          ==> 'EMC.ZOSMvrm.RIMLIB'
Edit macro name       ==> ZOS
Product dsname prefix ==> EMC.ZOSMvrm
Mainframe Enablers
dsname prefix         ==> EMC.MFE700
SMP/E dsname prefix  ==> EMC.SMPE
SMP/E datasets volser ==> _____
Install-to disk volser==> _____      Disk unit name ==> SYSDA

Enter your job card below ('%MEMBER%' will be replaced by member name):
=>

```

Figure 5 EMC JCL Customization Utility panel

3. Enter or change the following information on the panel shown in [Figure 5 on page 52](#) to customize your installation:
 - a. The CLIST library field (shown in [Figure 5](#)) is set by default to the name of the RIMLIB library. This field should contain the name of a library in which you want the edit macro created by this dialog to be stored.

The default value is fine for most users and need not be changed.
 - b. In the Edit macro name field (shown in [Figure 5 on page 52](#)), either:
 - Accept the default name displayed.
 - If necessary, change the name of the edit macro.

Note: Normally, you should not have to change the name.

Result: The edit macro is created in the CLIST or EXEC library from the data entered on this panel and applied to all members of RIMLIB that start with a # character.

- c. In the Product dsname prefix field (shown in [Figure 5 on page 52](#)), enter the dataset name prefix you want to use for the z/OS Migrator target datasets. EMC suggests EMC.*fmid* (for example, EMC.ZOSM*vrn*).
- d. In the Mainframe Enablers dsname prefix field (shown in [Figure 5 on page 52](#)), enter the dataset name prefix that was used to install Mainframe Enablers.
- e. In the SMP/E dsname prefix field (shown in [Figure 5 on page 52](#)), enter the dataset name prefix of the SMP/E datasets into which you installed Mainframe Enablers (EMCSCF).

For example, if you called the SMPSCDS dataset EMC.SMPE.SMPSCDS, enter EMC.SMPE.

- f. In the SMP/E datasets volser field (shown in [Figure 5 on page 52](#)), enter the six-character volume serial number of the disk volume on which you want to allocate the SMP/E distribution libraries for z/OS Migrator.

This volume may be the same as the volume you specify in step [g](#), or you may elect to keep these datasets on a separate volume.

- g. In the Install-to disk volser field (shown in [Figure 5 on page 52](#)), enter the six-character volume serial number of the disk volume to which you want to install the z/OS Migrator libraries.
- h. In the Disk unit name field, you can specify an esoteric disk name that is appropriate to your site. SYSDA is the default, but you can overwrite it with another esoteric disk name.
- i. Enter a site-appropriate job card.

The job card is initially set to a value which may be suitable to many users. The first seven characters of the job name is set to your TSO userid, plus "X."

You can set the job name to %MEMBER%. This causes the edit macro to set each job name equal to the JCL member name (that is, #01ALLOC, #02DDDEF, and so forth).

Do not use any parameter that contains an ampersand (&), such as NOTIFY=&SYSUID. An ampersand in the job card can cause edit macro errors.

4. When you are satisfied with your entries, type **exec** on the command line and press **Enter**.

Result: If the dialog completes successfully, you see something similar to the following:

```
BUILDING AN EDIT MACRO(z/OS Migrator) IN
'EMC.ZOSMvzm.RIMLIB'
PROCESSING MEMBER: #01ALLOC
PROCESSING MEMBER: #02DDDEF
PROCESSING MEMBER: #03RECEV
PROCESSING MEMBER: #04APPLY
PROCESSING MEMBER: #05ACCPT
PROCESSING MEMBER: #06CLEAN
PROCESSING MEMBER: #90SAFJB
PROCESSING MEMBER: #92SAFJB
PROCESSING MEMBER: #99MAINT
***
```

Step 4. Run the installation jobs

Carefully examine each job before you submit it to make sure that it was customized the way you intended.

Submit the customized jobs in the following order, making sure that each job completes successfully before submitting the next one:

1. #01ALLOC
2. #02DDDEF
3. #03RECEV
4. #04APPLY
5. #05ACCPT

You should expect completion codes of 0 (zero) for all jobs except for #02DDDEF and #05ACCPT, where 04 is acceptable.

SMP/E installation is now complete.

Step 5. Cleanup

After you are satisfied that z/OS Migrator is correctly installed and functioning properly, run the #06CLEAN job to delete datasets and DDDEFS used during the installation process that are no longer needed.

Step 6: Enter z/OS Migrator licensed feature code

The z/OS Migrator software is enabled by entering a licensed feature code (LFC).

The LFC is a 16-digit number on the EMC license key you received when your site received z/OS Migrator. Enter the LFC into the ResourcePak Base SCFINI initialization parameter table. Once you have done so, the LFC enables that feature.

You can find directions about how to apply the LFCs with the shipping documents. You can also find information about the license key installation process in the *EMC Mainframe Enablers Installation and Customization Guide*.

Step 7: Apply maintenance

If there is no current maintenance, keep these instructions for future use when you do need to download maintenance.

You need to install any available maintenance for z/OS Migrator before you start running. You can obtain the latest maintenance updates and current release or service notes (identical to release notes) from the **Software Downloads and Licensing** option on the Support menu at the EMC Powerlink website.

Note: You must register on Powerlink as a valid EMC customer before you can access Powerlink. Confirm that your license for this software is registered. Otherwise you will not be able to access the **Software Downloads and Licensing** option of the Powerlink website.

On the product page there may be more than one version listed. For your version, you may see the following types of files:

- ◆ ReadMe_ZOSMvrm_Fixes.txt — lists the fixes included in the release.
- ◆ Service_Notes_ZOSMvrm.txt — updates the release notes with information discovered after initial product release.
- ◆ ZOSMvrm_Fixes.zip — contains the two previous files as well as a software patch file (EZ300Fix.bin) and a sample job (smpjob.txt) with instructions about how to apply the maintenance.

Downloading files from Powerlink

1. Log on to:
<http://powerlink.emc.com>
2. Navigate to the **Software Downloads and Licensing** page for your product.
Result: You see a page with zip and document files for the product you selected.
3. Take either of the following steps:
 - To download a copy of a document, click `ReadMe_ZOSMvrm_Fixes.txt`, or `Service_Notes_ZOSMvrm.txt`.
 - To download the zip file, click `ZOSMvrm_fixes.zip`. Download the zip file to your home system, unpack the zip file, and follow the instructions it contains.
4. When you are finished reading or copying from Powerlink, return to the previous pages to view other products and services, or choose logout from the menu bar at the top of any page.

Step 8: Shut down and restart Mainframe Enablers

After you install z/OS Migrator, you need to shut down and restart the latest versions of EMCSCF (Mainframe Enablers).

Note: The *Mainframe Enablers Installation and Customization Guide* provides more information about starting EMCSCF.

Step 9: Prepare to execute z/OS Migrator

Before you can start utilizing z/OS Migrator to perform dataset and volume migrations, you must customize z/OS Migrator to run in your environment. The following steps must be followed to perform the initial setup of the z/OS Migrator components for execution after installation.

- ◆ Review [“SYSOPTN1 system defaults” on page 57](#) for information on setting up and running the z/OS Migrator server.

- ◆ Review “[Running the z/OS Migrator ISPF Monitor](#)” on page 118 for information on setting up and running the z/OS Migrator ISPF Monitor.

SYSOPTN1 sample job

Member SYSOPTN1 contains the input statements for the z/OS Migrator system defaults.

The initial z/OS Migrator display contains the following default dataset values:

- ◆ Company: *YOUR COMPANY*
- ◆ Site: *YOUR SITE*
- ◆ WTO/WTOR for AutoOps (messages for automated operations required): *NO*
- ◆ WTO AutoOps Route Codes (MVS routecodes): *2, 4*
- ◆ Local Time on Messages (display time): *LOCAL*

You can change these values by running the SYSOPTN1 batch job, with the new values you want displayed for your company.

Review and select the z/OS Migrator system defaults listed in [Table 5 on page 57](#).

SYSOPTN1 system defaults

[Table 5](#) describes the system defaults (SYSOPTN1 batch job). For examples of the SYSOPTN1 batch job, please refer to member SYSOPTN1 in SAMPLIB.

Table 5 **SYSOPTN defaults and options (1 of 2)**

Parameter	Function	Option	Default
z/OS Migrator VERSION	Must be the first non-comment entry. Specifies what version of z/OS Migrator is being executed.	3	The option must be specified.
Company Name =	Customer corporation name. Used in the Monitor display and SYSOUT listings. Up to 64 characters allowed.	Customer name	YOUR COMPANY
Site =	Local customer site name for a specific site, for example. Up to 64 characters allowed.	Name	YOUR SITE

Table 5 SYSOPTN defaults and options (2 of 2)

Parameter	Function	Option	Default
WTO/WTOR for AutoOps =	Specifies whether messages for automated operations are required.	YES NO	NO
WTO AutoOps Route Codes =	Specifies the MVS routing codes.	1-28	2, 4
Display Time As =	z/OS Migrator defaults to Local time on SYSOUT. This can be set to another value, such as GMT (Greenwich Mean Time). ("Option S – Set User Session options" on page 137 provides information on setting displays.)	GMT LOCAL Blank	LOCAL

Step 10: Allocate the z/OS Migrator system information dump dataset

z/OS Migrator writes comprehensive system diagnostic information to the datasets defined by the SYSMDUMP DD statement. To aid in timely analysis and diagnosis, EMC recommends that all z/OS Migrator server started tasks be created with the capability to write dump output to a Generation Data Group (GDG) dataset, as shown in [Figure 6 on page 59](#).

hlq.ZOSM300.SAMPLIB contains a member named 'DEF@GDG' for use in defining the SYSMDUMP dataset GDG base. Alternative methods exist for administering the SYSMDUMP dataset. For more information, please contact EMC Technical Support.

One GDG must be defined for each LPAR that z/OS Migrator runs on. The GDG should have the following format, since the same values will have to be specified in the z/OS Migrator server JCL:

```
DHLQ = YOURHLQ      High level qualifier for dumps
VER  = ZOSM300      z/OS Migrator version 3 release 0.0
SMF  = SMFID        SMFID
```

Figure 6 Sample of z/OS Migrator SYSDUMP GDG creation JCL

```

JOB CARD1
JOB CARD2
/*
/*      SAMPLE JCL TO DEFINE THE SYSDUMP GDG BASE
/*      RECOMMENDED NUMBER OF GDG ENTRIES IS 3
/*      ONE SUCH GDG MUST BE DEFINED FOR EACH LPAR THAT Z/OS MIGRATOR RUNS ON.
/*      THE GDG SHOULD HAVE THE FOLLOWING FORMAT, SINCE THE SAME VALUES
/*      WILL HAVE TO BE SPECIFIED IN THE Z/OS MIGRATOR SERVER JCL
/*      DHLQ = YOURHLQ           HIGH LEVEL QUALIFIER FOR DUMPS
/*      VER = ZOSM300           Z/OS MIGRATOR VERSION 3 RELEASE 0.0
/*      SMF = SMFID            SMFID
/*
//DEFGDG EXEC PGM=IDCAMS,REGION=1024K
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DEFINE GDG (
    NAME(YOURHLQ.ZOSM300.SMFID.SYSDUMP) /* NAME           */ -
    LIMIT(3)                          /* MAX NR OF GENERATIONS */ -
    OWNER(ZOSM300)                     /* OWNER MAX 8 CHARS     */ -
    SCRATCH                             /* DELETE VERSION IF UNCATALOGED */ -
)
LISTCAT ENTRIES(YOURHLQ.ZOSM300.SMFID.SYSDUMP) ALL
/*

```

Step 11: Allocate the z/OS Migrator database

The z/OS Migrator database is used to pass information between z/OS Migrator systems. It contains all information related to data migration operations. The z/OS Migrator database is automatically initialized by the z/OS Migrator server when it first starts.

The z/OS Migrator database must be allocated before you start the z/OS Migrator server. [Figure 7 on page 60](#) provides sample allocation details. hlq.ZOSM300.SAMPLIB contains a member named 'DEF@DB' for use in defining the z/OS Migrator database.



CAUTION

It is imperative that you properly allocate the z/OS Migrator database as described in the following example, and not preformat it in any way. When the z/OS Migrator server starts, it automatically formats the database. Additionally, both the primary and secondary allocations MUST be of the same size.

```

Jobcard1
Jobcard2
//*
//* Sample JCL to define the z/OS Migrator Server Data Base
//*
//*      Set the primary and secondary space allocation
//*      definitions to the same value.
//*      Example: SPACE=(CYL,(15,15),CONTIG),
//*
//DEF      EXEC PGM=IEFBR14
//SYSPRINT DD  SYSOUT=*
//ZOSMDB   DD  DSN=yourhlq.ZOSM300.DB,
//          DCB=(RECFM=F,LRECL=4096,BLKSIZE=4096),
//          UNIT=SYSALLDA,
//          VOL=SER=xxxxxxx,
//          SPACE=(CYL,(xx,xx),CONTIG),
//          DISP=(,CATLG)
//
/*

```

Figure 7 Sample of z/OS Migrator database creation JCL

Note: Multiple servers that are expected to cooperate MUST share the same database file. If there are multiple database files, you must NOT activate the same group using different databases. Using multiple databases is NOT recommended.

Placing the z/OS Migrator database

The z/OS Migrator database should be placed on a volume with low activity. The volume must also not contain data that you will wish to migrate from or to. It is recommended that the z/OS Migrator database not be placed on a volume where other control datasets reside (such as, MIM, the StorageTek LSM Control dataset, and so on).



CAUTION

Short-term reserves will be issued against the volume containing the z/OS Migrator database by the z/OS Migrator servers. You should carefully consider the placement of the z/OS Migrator database. It is strongly recommended that you choose a volume against which no other reserves are generally issued and which is not heavily active.

z/OS Migrator will not migrate datasets to or from the volume containing the active z/OS Migrator database; keep this in mind when considering placement of the dataset.

The database uses a RESERVE with the following major/minor name and may be converted to a global ENQ only if all LPARs involved in the migration are in the same GRS/MIM environment.

Note: If you are using CA-MIM, the database file and the SECCOM library must NOT be protected by CA-MIM. The MIM4006 and MIM4008 messages in the *CA-Multi-Image Manager Statements and Commands* manual provide additional detail.

- ◆ Major name (qname): EMC\$FM
- ◆ Minor name (rname): <44-byte database name>

z/OS Migrator will periodically issue RESERVE macros to the volume containing the z/OS Migrator database and this might affect other applications if this RESERVE is not being converted to a global ENQ by GRS.

Step 12: Configuring the z/OS Migrator server parameters

The z/OS Migrator server started task JCL specifies the Configuration PDS member that contains the startup parameters. The z/OS Migrator server configuration PDS member is located in

hlq.ZOSM300.SAMPLIB(CONFIG). These parameters are described in detail in Appendix B, “Detailed Parameter Reference” on page 313.

Figure 8, provides an example of what comprises the z/OS Migrator server configuration PDS.

```

* Sample Start Up Configuration File
* -----
*   4 Digit Server ID
*
SUBNAME=
*   1-8 Digit Command Prefix
CPFX=
*   The z/OS Migrator Data Base File Name (whole name)
*
DB=yourhlq.ZOSM300.DB
*
* DIVERSION SUBSET PARAMETER SETS THE NUMBER OF DATASETS PROCESSED
* BETWEEN RESERVE/RELEASE ON SOURCE/TARGET CATALOG PAIRS DURING
* DIVERSION.  DEFAULT IS ZERO AND THIS ALLOWS THE RESERVE TO BE HELD
* FOR THE TIME IT TAKES TO RENAME/RECATALG ALL DATASETS IN THE GROUP
* BEING PROCESSED.
*
DIVERSION_SUBSET=0
*   High Level Qualifiers for Target Extent Control Files (HLQ only)
*
LMPFX=yourhlq.EXTMAP
*   Additional 4k Pages to add to Working Set Size
*
DWA_ADD_PAGE=20
*   Maximum Overall I/O Concurrency
*
MAXIO=25
*   Maximum I/O Concurrency per Channel Path
*
MAX_CHANNEL_IO=15
*   Maximum I/O Concurrency per Individual Device
*
MAX_DEVICE_IO=3
*   Indicates the minimum number of seconds another LPAR must be
*   unresponsive before it is considered IDLE (Default is 45 seconds)
*
IDLE=0045
*   Optional, size of each copy I/O in Tracks.  Default/Max is 15
*   Suggestion - 'n' should be a factor of 15, e.g. 3
*
*MAXTRK=n
*** End Of ParmS

```

Figure 8 Sample z/OS Migrator server startup configuration file

Step 13: Customize and run the z/OS Migrator server

This section provides information and procedures for setting up your z/OS Migrator migration environment.

Before you can begin using z/OS Migrator to migrate data, you must customize and initialize the z/OS Migrator server. The z/OS Migrator server is the integral part of the z/OS Migrator product and is required to run on all z/OS images that will be accessing the data being migrated.

Please review all the information on setting up and operating the z/OS Migrator server before attempting to perform any data migrations.

The information on customizing, initializing, and operating the z/OS Migrator server is contained in the following sections:

- ◆ [“Step 10: Allocate the z/OS Migrator system information dump dataset” on page 58](#)
- ◆ [“Step 10: Allocate the z/OS Migrator system information dump dataset” on page 58](#)
- ◆ [“Step 11: Allocate the z/OS Migrator database” on page 60](#)
- ◆ [“Step 12: Configuring the z/OS Migrator server parameters” on page 61](#)
- ◆ [“Shutting down the z/OS Migrator server” on page 70](#)
- ◆ [“z/OS Migrator running in a shared storage environment” on page 70](#)

Customizing the z/OS Migrator server

Running the z/OS Migrator server as a started task

The z/OS Migrator server is intended to be run as a started task. In this way, z/OS Migrator will always be available to ensure migrations can be performed when required. This also ensures that processing continues for active migrations if one or more cooperating LPARS or servers is restarted.

You should implement the z/OS Migrator started task to run with performance attributes that are equivalent to that used by system started tasks, generally above batch processing. The z/OS Migrator server is an I/O oriented task but should have appropriate priority to schedule internal events in a timely manner.

Figure 9 below and on the following pages provides an example of z/OS Migrator server JCL.

Note: The JCL located in *hlq.ZOSM300.SAMPLIB(ZOSMPROC)* might have been updated and therefore may differ from the example provided here.

Figure 9 z/OS Migrator server JCL

```

/**      SAMPLE JCL to start the z/OS Migrator server as a started task.
/**
/**      For detailed assistance refer to the z/OS Migrator
/**      Product Guide, Chapter 2, z/OS Migrator Installation,
/**      Customizing the z/OS Migrator server.
/**
/**
//ZOSMPROC PROC CHLQ=hlq,           High Level Qualifier for z/OS Migrator
//                MEM=member,       Configuration Member Name
//                VER=ZOSM300,       z/OS Migrator Version 1 Release 3.0
//                START=WARM,        Warm start the z/OS Migrator server
//                DHLQ=hlq,          High Level Qualifier for DUMPS
//                SMF=smfid          SMFID
/**
//IEFBR14 EXEC PGM=IEFBR14
//SYSUT1 DD DSN=&DHLQ..&VER..&SMF..SYSMDUMP(+1),DISP=(,CATLG),
//                DCB=(LRECL=4160,RECFM=FBS),DSORG=PS,
//                UNIT=SYSALLDA,SPACE=(CYL,(500,500),RLSE)
/**
//ZOSM300 EXEC PGM=EFMMMAIN,COND=(0,LT),REGION=0M
//                PARM=('CFG=&CHLQ..&VER..SAMPLIB(&MEM),
//                START=&START')
/**
/**      Following is the z/OS Migrator Load Library
/**      It is recommended this remain a steplib and not
/**      added to the link list because of possible conflicts
/**      with other EMC products.
/**
/**      IF Z/OS MIGRATOR WAS INSTALLED WITH SMP/E THE MAINFRAME
/**      ENABLERS LINKLIB MUST BE ADDED TO THE STEPLIB CONCATENATION
/**      TO AVOID S806 ABENDS.
/**
//STEPLIB DD DISP=SHR,DSN=&CHLQ..&VER..LINKLIB
/**      DD DISP=SHR,DSN=DS-SCF-PREFIX.LINKLIB
/**
/**      Following is the z/OS Migrator Security Library
/**
//SECCOM DD DISP=SHR,DSN=&CHLQ..&VER..SECCOM
/**

```

Figure 9 (continued) ZOSMPROC

```

/**          *****
/**          * Required DD statements. *
/**          *****
/**
//EFMMSYLM DD  SYSOUT=* ----> PROCESS / DIVERSION TASK MESSAGES
//EFMMSYNC DD  SYSOUT=* ----> PROCESS / VOLUME MIGRATION MESSAGES
//SYSPRINT DD  SYSOUT=* ----> Z/OS MIGRATOR MESSAGES
/**
/**          GROUP ACTIVATION MESSAGE LOGS (ALLOWS FOR CONCURRENCY)
/**
//QCOUTPUT DD  SYSOUT=* ----> Snap API Messages
//QCOUT002 DD  SYSOUT=* ----> Snap API Messages
//QCOUT003 DD  SYSOUT=* ----> Snap API Messages
//QCOUT004 DD  SYSOUT=* ----> Snap API Messages
//QCOUT005 DD  SYSOUT=* ----> Snap API Messages
//QCOUT006 DD  SYSOUT=* ----> Snap API Messages
//QCOUT007 DD  SYSOUT=* ----> Snap API Messages
//QCOUT008 DD  SYSOUT=* ----> Snap API Messages
//QCOUT009 DD  SYSOUT=* ----> Snap API Messages
//QCOUT010 DD  SYSOUT=* ----> Snap API Messages
/**
/**
/**          *****
/**          * Contact Technical Support for assistance with *
/**          * SYSDUMP dataset definition and administration. *
/**          *****
/**          NOTE - DISP=MOD IS IMPORTANT
/**
//SYSDUMP DD  DISP=MOD,DSN=&DHLQ..&VER..&SMF..SYSDUMP(+1)
/**
/**          Z/OS MIGRATOR CONTAINS ITS OWN EMC SCF (RESOURCEPAK BASE)
/**          COMPONENTS. THE SCF$DMY DD STATEMENT ENSURES THAT ANY
/**          EXISTING SCF INTERFACE IS IGNORED.
/**
/**
//SCF$xxxx DD DUMMY
/**
/**          *****
/**          * OPTIONAL DD STATEMENTS *
/**          *****
/**
/**          OUTPUT IS CONTROLLED BY DEBUGGING SETTINGS - UNCOMMENT
/**          AS DIRECTED BY EMC TECHNICAL SUPPORT.
/**
/** EFMMCMPR DD  SYSOUT=* ----> COMPARE OUTPUT MESSAGES
/** EFMMCTLM DD  SYSOUT=* ----> COPY TASK MESSAGES
/** EFMMDBHB DD  SYSOUT=* ----> HEARTBEAT TASK MESSAGES
/** EFMMUTIL DD  SYSOUT=* ----> USED FOR DIAGNOSTIC PURPOSES

```

Figure 9 (continued) ZOSMPROC

```

//* EFMMRMAN DD   SYSOUT=* ----> MESSAGES FROM DIVERT VERIFICATION
//* EFMMVUCB DD   SYSOUT=* ----> USED FOR DIAGNOSTIC PURPOSES
//* EFMMXMAP DD   SYSOUT=* ----> EXTENT MAP DIAGNOSTICS
//* SYSSNAP  DD   SYSOUT=* ----> Z/OS MIGRATOR SNAP DUMPS
//*
//*          DISABLE COMPUWARE'S ABEND-AID
//*
//ABNLIGNR DD   DUMMY      ----> DISABLE ABEND-AID PROCESSING
//*
```

Note: The SCF\$xxxx statement is only necessary if z/OS Migrator uses a non-default instance of EMCSCF (ResourcePak Base).

z/OS Migrator system options and defaults

Table 6 defines the JCL substitution variables and their meaning:



CAUTION

The keyword **REGION=0MB** may be specified on jobcards in order to avoid virtual storage issues within the z/OS Migrator server address space. For more information on the **REGION** parameter, please refer to the *z/OS MVS JCL Reference Manual (SA22-7597)*.

Table 6 z/OS Migrator system options and defaults (1 of 2)

Parameter	Function	Option	Default
CHLQ=	High-level qualifier specified for the z/OS Migrator product when it was installed.		None. Must be specified.
MEM=	Name of the configuration member. The configuration file defines the operational parameters for the z/OS Migrator server. “Step 12: Configuring the z/OS Migrator server parameters” on page 61 provides more information.		None. Must be specified.
START=	Specifies the type of startup for z/OS Migrator to perform. Refer to “z/OS Migrator server START options” on page 67 for more details.	HOT WARM COLD	WARM

Table 6 z/OS Migrator system options and defaults (2 of 2)

Parameter	Function	Option	Default
VER=	Specifies the version of z/OS Migrator.		ZOSM300
DHLQ=	High Level Qualifier specified for the system information dump dataset, created in “Step 10: Allocate the z/OS Migrator system information dump dataset” on page 58.		None. Must be specified.
SMF=	SMFID for the LPAR on which this started task will run.		SMFID Must be specified.

If z/OS Migrator was installed with SMP/E, the Mainframe Enablers LINKLIB must be added to the Steplib concatenation to avoid s806 Abends.

In statement STEPLIB DD, change "DS-SCF-PREFIX" to the hlq specified for Mainframe Enabler when it was installed.

z/OS Migrator server START options

The z/OS Migrator server stores its operating parameters and maintains migration control and status information. This data provides key information on the data migration activities being managed by z/OS Migrator. The START option, specified when starting the z/OS Migrator server, controls how z/OS Migrator will treat previous information it finds in memory. There are two START options:

- ◆ **“WARM start”** (interchangeable with HOT start)
- ◆ **“COLD start”**

WARM start

WARM start is the default option. It is also interchangeable with HOT start.

During a WARM start, all existing common data structures and the data space are simply reused as-is. All prior I/O interface routines remain in place. There should be no impact on any active work, although if an application I/O request for a mirrored dataset could not be serviced during the hiatus, the group may have been suspended and should be resumed.

If there was an intervening IPL, then there are no existing common structures or data space. The group configuration objects and target extent files are read to re-establish a working environment and the

groups should come up in whatever state they were in before the IPL, with the exception that a group that has not reached mirror state will not be resumed if the restart is done on the owning system.

HOT start

A HOT start can be used interchangeably with a WARM start.

COLD start

During a COLD start, the previous data structures are abandoned and the server will come up as though there had been an IPL. Thus, the I/O interfaces will not have been continuously active, resulting in a data integrity exposure if there were intervening application I/O. A HOT or WARM start after an IPL is actually no different from a COLD start except for the assumption that there were no “interesting” application I/O from the “failure” of the old environment until z/OS Migrator was restarted after the IPL. A COLD start with mirroring groups will cause the groups to suspend, and is automatically disallowed if any groups were in *Diversion* phase.

Step 14. Prior to first execution, specify parameters in SAMPLIB

Prior to the first execution of z/OS Migrator, EFMXEXEC in the Mainframe Enablers SAMPLIB needs to be edited with specify installation specific parameters. The following values need to be replaced with your installation specific data. Values that need to be replaced include:

- ◆ DS_PREFIX — The Mainframe Enablers dataset prefix for ISPF libraries
- ◆ INSTALLED_LIB — HLQ used when installing z/OS Migrator
- ◆ YOUR_MIRROR_COMMAND_PREFIX — the command prefix used for communicating with the z/OS Migrator subsystem, which should match the value specified in the z/OS Migrator configuration/execution parameter CPEX.
- ◆ YOUR_MIRROR_LINKLIB — the fully qualified data set name of the z/OS Migrator linklib created during installation.
- ◆ YOUR_MIRROR_SECURITY_LIBRARY — the fully qualified data set name of the SECCOM library create during installation.

This is an example from the EFMXEXEC showing the values that need replaced.

```

DS_PREFIX = "INSTALLED.LIB" /* specify the ISPF DS-PREFIX here */
TYPEM     = ''
IDM       = ''
APPID     = "EFM#"
/*****
/* Note: The code below allows you to set variables for all users using */
/* this REXX clist. While the users can modify these variables in */
/* their current session, these variables will override them the */
/* next time this exec is invoked. Set these variables to a value */
/* or '' otherwise unwanted variable values will result. Setting a */
/* variable to '' allows the user to specify their own value and */
/* retain it until the variable is set in this procedure. */
*****/
UNIT      = "SYSALLDA"
VOLUME    = ''
UNITSW    = "3390"
MIR_PREFIX = "YOUR_MIRROR_COMMAND_PREFIX"
ZOSM_LOAD_LIBRARY = "'YOUR_MIRROR_LINKLIB'"
ZOSM_SECURITY_LIBRARY = "'YOUR_MIRROR_SECURITY_LIBRARY'"
/*****

```

Note: The z/OS Migrator ISPF Monitor is described in detail in [Chapter 6](#) on page 133.

Step 15: Starting and stopping the z/OS Migrator server

Starting the z/OS Migrator server

You are now ready to start the z/OS Migrator server. To start the server with the default parameters, issue the start command as follows:

S procname

Substitute the procname you have created. This will start the z/OS Migrator server started task with the default parameters that you specified during customization. (“[Option S – Set User Session options](#)” on page 137 describes the customization process.)

**CAUTION**

The z/OS Migrator server must always be active and able to immediately process requests. z/OS Migrator requires a high execution priority and should be put in a response-oriented performance group.

Shutting down the z/OS Migrator server

When you stop the z/OS Migrator server, the z/OS Migrator I/O Monitor routines will remain active within the system.

**CAUTION**

In the extreme situation where the z/OS Migrator server will not exit the system, the MVS FORCE ARM command may be used. It is highly recommended that this situation be avoided and that you follow IBM guidelines for the FORCE ARM command. Before attempting this, however, we recommend that you contact EMC Customer Support.

The commands that you enter can be through a z/OS Modify command or by using the z/OS Migrator server command prefix you specified when configuring the z/OS Migrator server. The valid options are:

- ◆ SHUTDOWN
- ◆ STOP

z/OS Migrator running in a shared storage environment

z/OS Migrator supports running in a shared storage environment. This might include multiple LPARs running z/OS, separate processors, single or multiple sysplexes, or a combination of all.

**CAUTION**

When running in a shared storage environment, the z/OS Migrator server must be running on all active z/OS systems and utilizing the same z/OS Migrator database. If you do not comply with this requirement, you will be exposing yourself to serious data integrity issues.

All processors active in the shared storage environment MUST be running as part of the same z/OS Migrator environment.

To accomplish this, you simply need to ensure that the z/OS Migrator started task is:

- ◆ Active on all z/OS images

Utilizing the same z/OS Migrator database across all z/OS Migrator servers.

Logical Migration Planning

This chapter contains a migration methodology, and best practice guideline to be used in conjunction with a logical data migration project. Additionally, because z/OS Migrator is a dataset-level migration tool that includes updating the information in the ICF catalog and interacts with SMS, there are sections that discuss their role in the migration process. Topics include:

- ◆ [Plan the migration](#)..... 74
- ◆ [Migration considerations](#)..... 79

Plan the migration

Planning for migration prior to execution contributes to having a successful migration. Planning includes establishing a migration management team, and performing pre-migration tasks.

Frequently you have to work with the local CE and the Z/OS sysprog that performs the I/O gen of the new box being added that will contain the target devices. Together you establish when the EMC box will be ready (task:CE), when the I/O gen is complete and the new UCB's along with volser's (task:Sysprog). With this knowledge you can feel comfortable that the target devices are ready for the migration.

Then check to see if the source volumes are in ConGroups and make sure the CE and Sysprog have done their part to ensure the new storage array is in place and ready to go.

Establish a migration team

You should create a migration management team to manage the migration process and a migration team to perform the pre-migration tasks. When establishing a migration management team, make sure that the team consists of at least the following people:

- ◆ Primary Migration Manager
- ◆ Alternate Migration Manager
- ◆ Application Coordinator
- ◆ Security Coordinator
- ◆ Technical Lead Coordinator

The migration team should include a technical data migration team, whose members have been trained to use z/OS Migrator and can also act as migration mentors.

Pre-migration tasks

Before performing dataset migrations, complete the pre-migration tasks to ensure that the migration will go smoothly and to minimize the risk of problems.

Perhaps the most basic preparation is to make sure the prospective target devices have been genned, and volumes labeled. It is also important to know what features have been specified and whether source volumes are currently participating in ConGroups. See [Appendix F on page 361](#) for additional considerations in this area.

Use the following list to help with the pre-migration tasks:

1. Identify the application owners.
2. Notify the migration to the entire application community about the upcoming migration at least 30 days prior to the intended migration date. Include an invitation to attend the initial migration kickoff meeting.
3. Gather information about the SMS storage environment and applications.
4. Inform Security and Compliance groups about the migration.
5. Schedule a pre-migration rehearsal that includes all the members of the migration team and define a data sampling that will enable the application groups to appropriately conduct the pre- and post-migration verification process.
6. Establish a weekly meeting schedule that includes remote participants.
7. Identify and follow the required change control process.
8. Establish a migration status call-in process.
9. Use the [“Pre-migration planning checklist” on page 77](#) to ensure that all of the pre-migration planning steps have been executed.
10. Perform a catalog diagnostic.
11. Create a new target volume SMS Storage Group Configuration.
12. Add the volume to the appropriate storage group in DISNEW status.
13. Change the status of the source configuration to DISNEW.

14. Create a list of source volumes that contain data to be moved and the associated target volumes or storage group.
15. Establish a naming standard for the temporary names of the new target datasets.

Note: The Source name becomes the Target name. Naming is discussed in [“Target and Source dataset names”](#) on page 121.

16. Establish a naming standard for the migration groups.
17. Identify the time slot for the migration process to transition from mirroring to diversion, and then to completion, based on DSN or group activity, and fallback requirements. Mirroring, diversion, and completion are explained in [“The z/OS Migrator migration process defined”](#) on page 120.

Note: Based on internal policies associated with a data migration process, some of the aforementioned list items may not be applicable.

Pre-migration planning checklist

Table 7 Pre-migration planning checklist (1 of 2)

Action item	Assigned to	Status
Create a migration management team consisting of the following people: <ul style="list-style-type: none"> • Primary migration manager • Alternate migration manager • Application coordinator • Security coordinator • Technical lead coordinator 		
Create a migration team to perform the pre-migration tasks; ensure that the team members have been trained to use z/OS Migrator and can also act as migration mentors.		
Identify the application owners.		
Notify the entire application community about the upcoming migration at least 30 days prior to the intended migration date. Include an invitation to attend the initial migration kickoff meeting.		
Gather information about the SMS storage environment and applications.		
Inform the security and compliance groups about the migration.		
Schedule a pre-migration rehearsal that includes all the members on the migration team and define a data sampling that will enable the application groups to appropriately conduct the pre- and post-migration verification process.		
Establish a weekly meeting schedule that includes remote participation dial in.		
Follow the required Change Control process.		
Establish a migration status call-in process.		
Use the Migration Planning Checklist to ensure that all of the Pre-migration planning steps have been executed.		
Perform a catalog diagnostic.		
Create the new target volume SMS storage group configuration.		
Add the volume to the appropriate storage group in DISNEW status.		
Change the status of the source configuration to DISNEW.		

Table 7 Pre-migration planning checklist (2 of 2)

Action item	Assigned to	Status
Create a list of source volumes that contain data to be moved and the associated target volumes or storage group.		
Establish a naming standard for the temporary names of the new target datasets.		
Establish a naming standard for the migration groups.		
Identify the time slot you want to move from mirroring to diversion to completion based on DSN or group activity, and fallback requirements.		

Migration action item checklist

Note: Populated as part of the pre-planning process.

Table 8 Migration action item checklist

Action Item	Assigned to	Status

Migration considerations

Before using z/OS Migrator to perform data migrations, you should review this section for details that may affect your data migration environment. The following topics are included in this section:

- ◆ “z/OS Migrator dataset selection planning” on page 79
- ◆ “Datasets allocated after activation” on page 85
- ◆ “Dataset extent and multi-volume dataset handling” on page 85
- ◆ “Merging volumes and VTOC sizing considerations” on page 86
- ◆ “Reserve handling requirements” on page 87
- ◆ “Restrictions on storage subsystems and dataset types” on page 93
- ◆ “Systems managed storage (DFSMS) allocation” on page 94
- ◆ “Migrating VSAM clusters” on page 97
- ◆ “Migrating DB2 cluster datasets” on page 97
- ◆ “Hard-coded volume references” on page 97
- ◆ “Catalog structure” on page 97
- ◆ “Small versus large migration groups” on page 98
- ◆ “Suspending/resuming data migrations” on page 98

z/OS Migrator dataset selection planning

z/OS Migrator includes many dataset selection criteria, from a single dataset to thousands of datasets and multiple groups. When using selection criteria that go beyond a fully qualified dataset name, it is recommended that a listing utility such as SCAN in EMC Catalog Solution® be used.

Note: Under a licensing agreement with EMC, Rocket Software has taken over sales and support of EMC’s Catalog Solution. Please contact Rocket to report an issue. You can call the Rocket Software Support Hotline at (617) 614-4323, or visit the Rocket Mainframe Productivity Suite website: <http://www.rocketsoftware.com/mainframe-productivity/rcp-access.html>

The dataset selection classifications using Catalog Solution provide extensive categories that can be used as selection criteria for z/OS Migrator candidates and to identify datasets that should be excluded.

The following sections outline selection factors, volume checklist, how to change SMS volume selection, and how to select datasets.

Selection factors

The following are important factors related to z/OS Migrator dataset selection planning:

- ◆ VSAM CANDIDATE volumes must be identified because these volumes can change from CANDIDATE to PRIMARY. Candidate volume status can be changed using IDCAMS – ALTER REMOVE VOLUME.
- ◆ The Catalog Solution SCAN exit CS#MULTI can be used to generate a Multi Volume VSAM Cluster Report. “[Catalog Solution](#)” on page 362 describes the z/OS Migrator interaction with EMC Catalog Solution.
- ◆ Verify that the DSN length is less than 44 characters because a target DSN could include an additional qualifier. Can be identified using Catalog Solution.
- ◆ To prevent any new allocation in a SMS configuration, the status of the source volume must be DISNEW.
- ◆ Datasets that will be deleted and recreated within the z/OS Migrator migration cycle should be excluded and handled by SMS Data Redirection.
- ◆ GDG’s datasets should be reviewed for delete activity cycle. Once again, if the cycle is shorter than the z/OS Migrator migration cycle, let them move themselves.
- ◆ Non-SMS managed volumes that are mounted with a STORAGE status should be changed to PRIVATE.
- ◆ Space utilization: It is important that the storage subsystem have sufficient space to accommodate normal processing. Although in most cases there will be new storage available for a technology refresh, this might not be the case for dataset relocation or work load balancing.
- ◆ Go through a test/rehearsal type run to gain familiarity and confidence.
- ◆ Security: The security profile associated to a particular dataset might prevent moving the dataset.
- ◆ When moving a dataset for performance reasons, be sure that the new target volume does not already have a performance problem.

- ◆ When doing volume consolidation, be sure that you are not selecting multiple source volumes that have performance problems.
- ◆ Try to schedule the completion stage to take place as part of a preplanned outage.
- ◆ Identify Link List datasets or any other datasets that have a specific volser dependency. Since Link List datasets have a volser dependency, moving them without making the necessary changes that include the new volser could cause a problem.
- ◆ Identify all the volumes associated to a multi-volume dataset. The same number of target volumes will be required.
- ◆ Use the [“Migration action item checklist” on page 78](#) to establish tasks, assignments, and status.

Volume checklist

- ◆ Make sure that the status of all the source volumes is DISNEW.

Note: If DISNEW is enabled using VARY, the next ACS TRANSLATE and VALIDATE could reverse that status.

- ◆ Make sure that all source and target volumes contain an INDEX VTOC.
- ◆ Make sure that the target configuration is added to the appropriate storage group and the status of all volumes is DISALL. When you are ready to start the migration, change the status to ENABLE.
- ◆ Make sure that all the target volumes that have been initialized for SMS (ICKDSF STORAGEGROUP) contain a configuration with a large enough VVDS (10 cylinders) and contain an INDEX VTOC.
- ◆ Use ISMF to verify that the volumes in the target configuration have INDEX VTOC enabled. Additionally, a procedure should be implemented to verify that the status of all INDEX VTOCs is ENABLED because it is required on an SMS managed volume.

Changing the SMS volume selection

The following procedure is an example of a quick way to alter the SMS allocation status for a range of volumes in an SMS storage group using ISMF. The objective of the procedure is to migrate from three MOD3s to one empty MOD9, and control allocation; the volume range is PROD01 to PROD03.

To alter the SMS allocation status:

1. Change the SMS volume selection for the three 3390 MOD3s to **DISNEW**.
2. Create new 3390 MOD9 PRODA1.
3. Using **ISMF OPTION 6**:
 - a. Select **Enter**.
 - b. Type the **STORAGE GROUP** name and select **OPTION 4 VOLUME**.
 - c. Select **Enter**.
 - d. Type the volume prefix **PROD** and the range **FROM 01 TO 03** and select **OPTION 3 ALTER**.
 - e. Select **Enter**.
 - f. Change the SMS volume status for PROD01 through PROD03 to **DISNEW**.
4. Type **SETSMS SCDS(XX.XXXXXX.SCDS)**. Where XX.XXXXXX.SCDS is the SCDS.
5. Enable PRODA1 as **ENABLE**.
6. Complete the allocation phase.
7. Repeat [Step 3](#) and change the volume selection for PROD01 through PROD03 to **DISALL**.

```
SETSMS SCDS(XX.XXXXXX.SCDS)
```

Where XX.XXXXXX.SCDS is the SCDS.

Selecting datasets

Datasets are selected during the *Activation* phase by performing a catalog look-up using the Source DSN values provided in the Define Logical Migration Group panel. (“[Step 1. Group definition](#)” on [page 121](#) provides detailed information.)

Note: Uncataloged datasets will fail activation.

The Owning z/OS Migrator system (the host where the Group *Activate* command was submitted) performs the catalog look-up. When choosing the Owning system, select the system that possesses the most paths and that performs most of the I/O operations to the data that you are moving. Datasets cataloged in multiple catalogs will be migrated only from the perspective of the Owning system's master catalog.



CAUTION

Source Dataset Cataloged in Multiple Catalogs:

All master catalogs (MCATs) should have equivalent alias definitions, with each alias entry pointing to the same user catalog (usercat) across all MCATs within the SYSPLEX or shared storage complex. The risk to a migration is that z/OS Migrator identifies the datasets to be migrated through a catalog search performed on the z/OS Migrator Owning system, the system where the z/OS Migrator Monitor is used to perform the Group ACTIVATE. If the MCATs do not have synchronized aliases, then there is a risk of having more than one catalog entry for a given dataset. The z/OS Migrator ACTIVATE phase recognizes only the dataset associated with the catalog known to the Owning system.

MCAT alias synchronization errors should be resolved as part of your migration planning and preparation.

DFDSS masking rules

z/OS Migrator supports dataset masking for the SOURCE DSN, EXCLUDE DSN, and TARGET DSN parameters following the DFDSS masking rules. Table 9 explains the DFDSS masking rules.

Table 9 DFDSS masking rules

Character	Meaning
* (single asterisk)	Takes the place of one qualifier or indicates that you are specifying one part of a qualifier.
** (double asterisk)	Used with other qualifiers, denotes one of the following: <ul style="list-style-type: none"> The non-existence of leading, trailing, and middle qualifiers. They play no role in the selection.
% (percentage)	Indicates a single-character mask.

Rules for using asterisks in a qualifier

- ◆ The maximum number of permissible asterisks in a qualifier is two.
- ◆ When two asterisks are present in a qualifier, they must be the first and last characters. For example:
 - Permissible qualifiers:
 - **
 - *A*
 - Not permissible qualifiers:
 - **A*
 - *A*B*
 - *A*B
 - A*B*C

Masking rules

- ◆ SOURCE DSN, EXCLUDE DSN and TARGET DSN dataset name mask rules are different. The SOURCE DSN and EXCLUDE DSN masks allow partial index level masking.

For example, while EMC.DATA*.ABC** is valid for SOURCE DSN or EXCLUDE DSN, it is not valid as a TARGET DSN mask.
- ◆ The SOURCE DSN and EXCLUDE DSN dataset name masks follow DFDSS rules for dataset selection.
- ◆ The TARGET DSN dataset mask follows DFDSS rules for RENAME processing.

Examples

Some valid TARGET DSN masks are:

- A.** - Replaces the first index level of the source dataset name with "A" in the target dataset name.
- A.B.** - Replaces the first two index levels of the source dataset name with "A.B" in the target dataset name.

Some invalid TARGET masks are:

- **.DEF.** - Does not know which level to replace.
- A.*BCD* - The entire index level must be wild, or not at all.

Datasets allocated after activation

Datasets matching the source DSN masks that are cataloged after *Activation* are not included in the datasets population being migrated. If the migration objective is to depopulate the 'source' volumes, you should prevent allocations of datasets on these volumes.

In DFSMS-managed environments, the source volumes must be changed to volume status 'DISNEW' early in the migration project. This will cause these volumes to be excluded from SMS volume selection and the 'target' volumes within the storage group to be selected.

Dataset extent and multi-volume dataset handling

When migrating datasets that have multiple extents, z/OS Migrator images the target dataset to be an exact reflection of the source dataset. This means that if the source dataset in a migration currently is in six extents, then the target dataset is allocated with six identically sized extents.

Similarly, z/OS Migrator handles multi-volume datasets in a very similar manner. If a dataset over five volumes, then z/OS Migrator ensures that the target dataset being migrated is five volumes with the same number of extents on each volume as on the corresponding source volume.

z/OS allows a single dataset to have extents on up to 59 volumes. Under SMS, you only need to supply one volume in the Target Volume field. SMS will select the actual volumes based on your ACS routines.

When consolidating a group of datasets from multiple volumes to a smaller pool of new volumes, give consideration to those datasets that have been migrated into HSM. While z/OS Migrator does not recall and migrate those datasets currently HSM migrated, make sure that the new storage pool has sufficient volumes to re-stage any previously HSM migrated multi-volume dataset.

Note that the z/OS Migrator Define Logical Migration Group panel only allows five target volumes to be specified. If more volumes are needed, you can edit the group definition using the Modify option to add additional volumes to the group definition. Instructions on how to add more target volumes are included in [“Example 2: Modifying a group definition to add more target volumes”](#) on page 166.

Merging volumes and VTOC sizing considerations

Although z/OS Migrator facilitates the consolidation of datasets from multiple smaller volumes to larger volumes, z/OS Migrator does not resize VTOCs for the extra metadata that will reside on a larger volume.

For this reason, ensure that your larger volume VTOCs are appropriately sized to contain the data you plan on migrating to the volume.

For example, when many small datasets are consolidated onto a larger target volume, a larger VTOC may be required to hold all the dataset DSCB records. There should also be sufficient unused VTOC capacity on any volume to accommodate future additions. Otherwise, there could be unusable space on the volume due to the VTOC being full.

Multi-LPAR considerations

In general, z/OS Migrator should be active on all LPARS sharing the data to be migrated. [“Multi-LPAR considerations”](#) on page 32 provides more information. In a complex environment where DASD is only partially shared among multiple LPARs, it may be useful to develop a matrix that identifies the volumes to be migrated and the LPARs involved. This situation is best handled by having specific z/OS Migrator server complexes for each unique LPAR combination. If run concurrently, of course, each will have to have its own z/OS subsystem ID, command prefix, and most importantly, there will

need to be a unique shared z/OS Migrator database for each z/OS Migrator complex.

Reserve handling requirements

z/OS Migrator has explicit requirements in the area of reserve handling. More specifically, z/OS Migrator requires that all hardware reserves issued by z/OS address spaces, including system, application, and jobs, be converted to software reserves.

Understanding global resource serialization and hardware reserves

In a shared storage environment, in which multiple z/OS images utilize the same storage resources, the z/OS operating system will issue a *hardware reserve* to the physical device in order to control exclusive usage of a resource. This ensures that no other z/OS image in the shared storage environment can access that device, during the time that the reserve is held.

The Global Resource Serialization (GRS) and Computer Associates Multiple Image Manager (MIM) of the z/OS operating system serialize usage of resources by converting hardware reserves to globally propagated ENQ requests. Through enqueue/dequeue facilities, the integrity of shared resources can be guaranteed. GRS and MIM communicate the need for exclusive usage across all systems in the shared storage complex and ensure exclusive usage by issuing a “software enqueue” on all systems in the shared storage environment. This converts the coarse physical protection of a hardware reserve to the more focused logical protection of a global ENQ.

Since z/OS Migrator migrates datasets from one disk volume to another, there are implications for the use of serialization facilities. Considerations for converting hardware reserves to globally propagated ENQ requests include:

- ◆ Ensuring the data integrity of data for datasets being migrated
- ◆ The avoidance of deadlocks in a multi-system environment

Ensuring data integrity of datasets being migrated

To ensure the data integrity of datasets being migrated, it is imperative that hardware reserves be converted to globally propagated ENQ requests. In Figure 10, the dataset DS1 is being migrated from source volume A to target volume B, and dataset DS2 is being migrated from source volume B to target volume A.

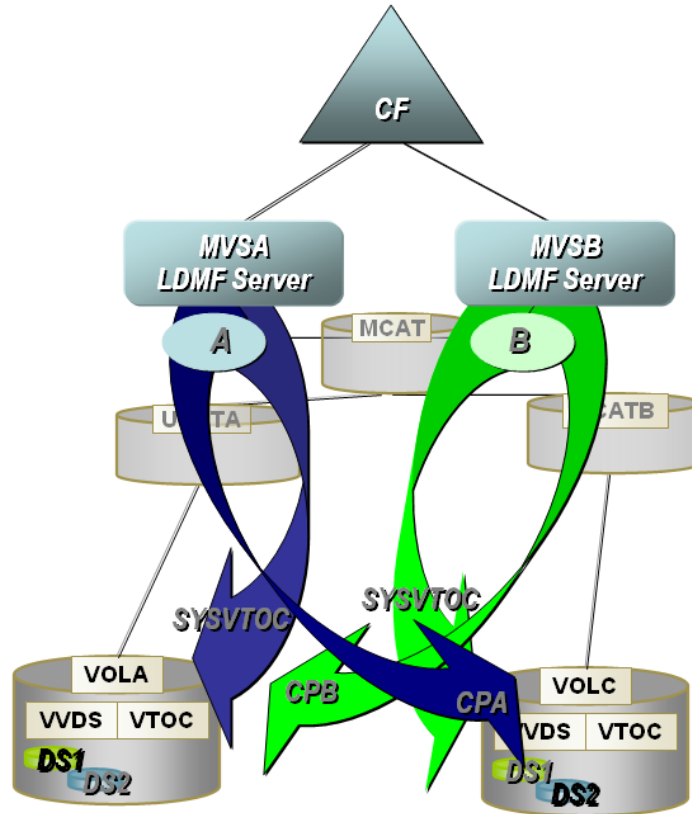


Figure 10 Reserve handling

If we assume that applications are running on both z/OS server MVS A and MVS B, and both require exclusive serialization of the resource, then the question becomes, *How does z/OS Migrator impact normal serialization processing resulting in the requirement to convert hardware reserves to ENQ requests?*

z/OS Migrator impact on serialization — z/OS Migrator requires resource serialization of catalog, VTOC, and VVDS resources during *Activation*¹ and *Diversion*, as well as during *Synchronization* and *Mirroring*, whenever new extents are added. With regard to dataset serialization, because long-running applications continue to read and write data from the source location, all serialization will be targeted to the source location for resources on that volume. In this example, if reserves are *not* converted, all reserves related to dataset DS1 located on VOLA will be issued to that location across all applications and systems for serialization of the resource. Indirectly, all I/O activity mirrored or diverted to the target dataset located on VOLC will also be serialized.

z/OS Migrator impact on serialization after diversion but before completion — Once z/OS Migrator has entered the *Diversion* phase, it is absolutely imperative that hardware reserves are converted to global ENQ requests to ensure data integrity of dataset(s) being migrated.

Applications that were active before *Diversion* and that remain active during *Diversion* will continue to serialize usage of the resource residing on the source volume, as with DS1 on VOLA in [Figure 10 on page 88](#).

If serialization is through hardware reserve, then the reserve will be issued against VOLA. However, a subsequent allocation to DS1 by a job starting after *Diversion* begins will be satisfied using the catalog entry that has been updated to refer to VOLC as the location of DS1.

Reserve serialization on behalf of this application will target VOLC. Consequently, neither access to the data will be properly serialized; changes may be lost as a consequence of the latest allocation having serialized on a different volume than that of the initial allocation.

Examples

An example follows:

- ◆ MVSA reserves VOLA for dataset DS1 prior to the time when z/OS Migrator diverts the catalog entry to VOLC. MVSA reads data, updates the data, and writes it back. z/OS Migrator re-directs the update to VOLC.

1. Activation, Diversion, Synchronization, and Mirroring are discussed in “The z/OS Migrator migration process defined” on page 120.

- ◆ On MVSB, a new allocation of DS1 reserves VOLC, reads data from dataset DS1, modifies the data and writes it back to dataset DS1.
- ◆ During this time, the z/OS Migrator channel program for MVSA will have to wait because VOLC is reserved to MVSB. Once MVSB releases the reserve on VOLC, MVSA will overwrite the data that was changed by the user that allocated DS1 on MVSB.

In this scenario, a classic data integrity exposure exists. The data has been overlaid and there is no indication that this has occurred.

As long as all dataset-related reserves are converted to global enqueues, the resource is protected by name and remains secure. Conversely, if the reserve is not converted, the benefit of reserve protection is lost between such pre- and post-diversion allocations and the preceding scenarios can occur.

Prevention of SYSVTOC/SYSZVVDS deadlocks in a multi-system environment

The SYSVTOC and SYSZVVDS resource names are the QNAMES for resource serialization of each volume's VTOC and VVDS.

Datasets grow and shrink with usage, and as they do, the VTOC is updated to reflect the dataset's space usage on a volume. The VVDS is updated to reflect the change in extents for VSAM datasets. When extent changes occur, z/OS Migrator mirrors the VTOC changes on the target volume. For example, if DS1 on VOLA in [Figure 10 on page 88](#) adds a new extent, then z/OS Migrator must also add a new extent on VOLC.

If SYSVTOC and SYSZVVDS QNAMES are not converted, then hardware reserves must be issued to VOLA and VOLC. A reserve on a resource, followed by an access request on another resource, greatly increases the likelihood of a deadlock. If this type of deadlock occurs, the group can be *Suspended* and *Deactivated* prior to *Diversion*; otherwise, once in *Diversion*, the application address space must be cancelled to eliminate the deadlock. This possibility is further accentuated if there is any overlap between the source and target volume sets across active migration groups, as illustrated by the second dataset in the previous example. The conversion of SYSVTOC and SYSZVVDS hardware reserves to globally propagated ENQ requests essentially eliminates such deadlocks. z/OS Migrator still issues hardware reserves for certain functions, but special handling code in z/OS Migrator eliminates the opportunity for multiple

systems to request the same volume's resources exclusively at the same time.

Note that when there is reserve contention, the IOS Start Pending message may be seen on the console or system log. Use commands like the following to investigate the processes involved:

```
D GRS, C
D GRS, DEV=ccuu
D GRS, ANALYZE, BLOCKER
```

Converting hardware reserves

To accomplish reserve conversion, you can use IBM's Global Resource Serialization (GRS) or Computer Associates Multiple Image Manager (MIM).

You can find the IBM publication (SA22-7600) *z/OS MVS Planning: Global Resource Serialization* at the following Web page:

<http://publibz.boulder.ibm.com/epubs/pdf/iea2g440.pdf>

z/OS Migrator RESERVE CONVERSION REQUIREMENTS:

```
Qname: EMCTF
Rname: *
```

The above reserve request needs to be added to the reserve conversion list for GRS or MIM as a generic conversion. This needs to be evaluated before being implemented in the customer environment to ensure it doesn't conflict with other customer requirements.

z/OS Migrator Reserve Monitor

The z/OS Migrator Reserve Monitor is a program that uses the z/OS ENQ/DEQ Batched Exit (ISGNQXITBATCH) to record Reserve request statistics. When correctly shut down, or in response to a Modify Display operator command, the program will display a table of Reserve resources with the following elements:

- ◆ QNAME: Reserved resource major name.
- ◆ RESERVES: The number of times a Reserve resulted in a request to physically reserve a device.
- ◆ CONVERTED: The number of times a Reserve was converted to a global ENQ request.

Sample reserve monitor output:

```

EFM9803I SWH10RMN RESERVE counts:
  _QNAME_          RESERVES  CONVERTED
  EMC$FM           155         0
  SYSZJES2         120         0
  SYSZRACF          5         0
  SYSIGGV2          0         52
  SYSZVVDS         56         0
  IGDCDSXS          4         0
  SYSVTOC           0         5

```

Note: If reserves for individual datasets are not converted to software reserves, data integrity cannot be assured once z/OS Migrator has started diverting I/O requests for datasets in the migration group. If SYSVTOC and SYSVVDS reserves are not converted, deadlock exposures exist.

To execute the Reserve Monitor, use a JCL based on the following sample:

```

//jobname      JOB
//RMON         EXEC PGM=EFMMRMON
//STEPLIB     DD DISP=SHR,DSN=<authorized load library>
//SYSUDUMP    DD SYSOUT=*

```

Note: Run EFMMRMON from the z/OS Migrator load library or move it to any APF authorized library.

When the Reserve Monitor starts, a message is issued to the console displaying the name that must be used in operator Modify and Stop commands. When the monitor is submitted as a batch job, this will always be the job name:

```
EFM9802I "name" RESERVE Monitor program ready.
```

To produce an instantaneous report of the Reserve requests monitored since the program became active:

- ◆ Enter the modify command:

```
F jobname, DISPLAY
```

To produce the final report and terminate the monitor:

- ◆ Enter the Stop command:

```
P jobname
```

Restrictions on storage subsystems and dataset types

z/OS Migrator only supports 3380 and 3390 architecture, like geometry devices. “Supported storage subsystems and dataset types” on page 34 lists the supported dataset types.

z/OS Migrator has the following restrictions in regards to storage architecture and dataset types

The target controller must be at an equal or higher technology level and be downward compatible. In other terms, you cannot migrate data from a 2105 model controller to a 3990 model controller, but you can migrate data from a 3990 model controller to a 2105 model controller.

When you want to verify that the SOURCE and TARGET devices have equal device characteristics issue the following command against the SOURCE and TARGET UCBs:

```
DS QD,nnnn,RDC
```

Compare the output. If they have equal characteristics the output will be identical (except for the UCB and Volume labels).

Note: Migration between different device geometries is NOT supported. You can only migrate data between similar geometry devices, for example: 3380 to 3380, or 3390 to 3390.

Datasets cataloged on System Resident (SYSRES) volumes

Migration of system datasets, that is, datasets used by the operating system, typically located on system residence volumes, including linklisted datasets, is a complex operation. We recommend that you contact Technical Support personnel before migrating such datasets.

Dataset types NOT currently supported in z/OS Migrator

- VSAM datasets with the IMBED, KEYRANGE and REPLICATE options specified. IMBED, REPLICATE and KEYRANGE are obsolete VSAM parameters. Target dataset extents cannot be allocated using these parameters because z/OS does not support new allocations using them. The physical layout of extents of datasets defined with these parameters is different from the physical layout of datasets

without them. Copying extents from a dataset containing these to a newly defined dataset that does not have them would yield an unusable dataset.

If you copied an IMBED cluster to NOIMBED (the only type you can currently define), the target data component would have a track full of sequence set records at the start of each CA that VSAM would think was actual data.

- Catalogs
- ISAM
- IMS OSAM databases
- Individual PDS members (z/OS Migrator does support migrating an entire PDS)
- Page and Swap datasets
- HFS/zFS datasets
- Datasets that are designated as "Unmoveable" are not supported, that is, DSORG=U | PSU
- VTOC, VVDS, and VTOCIX datasets
- Temporary (&&) datasets

Note: Uncataloged datasets will fail activation.

Systems managed storage (DFSMS) allocation

DFSMS is a software suite that automatically manages data from the time that it is created until it expires. DFSMS provides allocation control for availability and performance, backup/restore, and disaster recovery services, space management, and tape management. DFSMS consists of DFSMSdfp, an element of z/OS; and DFSMSdss, DFSMShsm, and DFSMSrmm, features of z/OS. Using the ISMF or a VARY command, you can create a criterion to eliminate allocation of new datasets on any volume in a SMS subsystem. This status is known as DISNEW and prevents all new allocation.

When starting a migration that includes moving all the datasets that populate a volume, the status of the source volumes must be DISNEW. It is important to verify, using ISMF, that the storage pool containing these volumes has sufficient amount of free space to allow allocation restrictions to be enabled.

The following are the components that address dataset allocation

Data Class — Addresses the attributes of a dataset.

Storage Class — The required performance level required for the dataset that is in direct alignment with a specific service level agreement (SLA).

Management Class — The required management that addresses the appropriate backup and migration cycle of the dataset.

Storage Group — The storage group is made up of a volume or group of volumes that are designated for specific datasets based on a pre-determined criteria (for example, size, use, type).

SMS Data Redirection

A storage subsystem that is under SMS control has the capabilities to direct and disable allocation to specific storage pools. In an environment that has a high delete and reallocation activity rate, data can be moved by simply disabling one volume and enabling a different volume in the same subsystem. This process is known as "SMS Data Redirection."

As SMS Data Redirection is done on a dataset level, z/OS Migrator becomes a perfect complementary solution for those datasets that are allocated and not deleted on a scheduled basis. The first phase when migrating all the datasets in an SMS environment would be to have SMS move that data using SMS Data Redirection. z/OS Migrator will enable a scheduled dataset movement process, taking less time and enabling the utilization of performance enhancements from the new device for those datasets not moved using SMS.

When allocations are performed within an SMS environment, the SMS ACS routines are invoked. These routines are customized locally. z/OS Migrator allows you to specify several parameters as part of a migration group that influence the volumes selected for allocation of the migration group's target dataset(s).

Source dataset name

The source dataset name determines the volume that is selected for the destination of the migrated dataset. z/OS Migrator drives the allocation of target datasets through the site's ACS routine by substituting each source dataset's attributes (dataset name, physical characteristics, management criteria) at *Activation*.

Target Storage Class name

z/OS Migrator allows you to specify a target Storage Class name to be used during *Activation*.

The target Storage Class name in the migration group definition (if present), is substituted for the Storage Class of the source dataset when z/OS Migrator passes the allocation request to SMS.

You must edit the relevant ACS routines to permit allocations to be re-directed to the target Storage Class prior to Group Activation.

Note: The source dataset characteristics are passed to SMS in lieu of user-specified characteristics normally required for new dataset allocations.

Alternatively, specifying a target Storage Class name to be used during *Activation* without enabling the ACS selection logic to direct source dataset allocations to the specified Storage Class first, may result in the specified target Storage Class being ignored or may result in an allocation error.

Target volumes

z/OS Migrator allows you to specify a volume or list of volumes to be used during allocation for all target datasets. Depending on how the ACS routines have been customized, and how the SMS volume statuses of QUINEW or DISNEW are applied, specific volumes requested during *Activation* may or may not be honored.

In cases where the target dataset takes extents that cannot be accommodated on volumes selected at *Activation*, the selection of additional target volumes comes from the storage group containing current target volumes.

Migrating data to a different storage group

ACS routines determining Storage Class must be customized to use z/OS Migrator to migrate datasets from one Storage Group to another Storage Group. In the case where the resulting data must remain associated with the original Storage Class, that is, no target Storage Class is specified, the ACS routines must be customized to direct new allocations to the desired Storage Group prior to *Activation*.

In cases where the target dataset takes extents that cannot be accommodated on volumes selected at *Activation*, the selection of additional target volumes comes from the storage group containing current target volumes.

Migrating to new volumes

If you are migrating data within a Storage Group, and the volumes you are migrating to are newly installed, you must ensure that you initialize and then add them to the Storage Group definition prior to migration. z/OS Migrator does not take any action to dynamically add volumes to a Storage Group.

Migrating VSAM clusters

You cannot currently migrate individual components of a VSAM cluster. This is because the metadata for VSAM clusters must be managed in its entirety. Therefore, make sure you specify the base cluster name, or a dsname mask that will include it, when setting up VSAM-oriented migration groups. Likewise, the target must be given an appropriate target cluster name. [Table 17, “Target parameters,” on page 153](#) provides more information.

Migrating DB2 cluster datasets

Since DB2 datasets are VSAM clusters, refer to the information mentioned above in [“Migrating VSAM clusters.”](#) If you are migrating the entire database, make sure that masks matching all database clusters are included.

Hard-coded volume references

You must take into consideration instances where the ICF catalog is not the only mechanism used to locate datasets, as part of your migration planning.

Storage environments with comprehensive storage management policies implemented by DFSMS or another storage management facility are heavily reliant on the catalog environment to access datasets.

z/OS Migrator will migrate datasets and update their catalog entries without regard to the user or application design considerations that necessitate hard-coded volume serial numbers for accessing datasets. Hard-coded volume references must be updated with the new volumes selected during the migration *Activation* phase. This task may be performed during the *Diversion* phase, prior to the next restart of affected applications.

Catalog structure

Logical Migration is a dataset-level migration tool that includes updating the information in the ICF catalog. It is very important that you have a high-level outline of the ICF catalog structure to help you perform diagnostics for the pre-migration process.

A healthy catalog environment is essential to a successful migration. It is important to resolve all catalog problems before starting the migration.

Note: [Appendix D on page 339](#) shows details of the Catalog Structure.

Small versus large migration groups

It is recommended that when defining migration groups, you do not migrate too many source datasets at one time. Smaller migration groups:

- ◆ Are easier to manage during the z/OS Migrator data migration life-cycle
- ◆ Allow quicker movement through the z/OS Migrator migration phases (because a single large dataset will not delay other smaller datasets)
- ◆ Are easier to track and clean up multiple groups

Suspending/resuming data migrations

With z/OS Migrator, it is possible to halt a data migration while the group is in *Copy* or *Mirror* state, before you reach the *Pending Diversion (P-Divert)* or *Diversion* phase. This is accomplished by using the Suspend command through the z/OS Migrator ISPF Monitor, Option 2 - **Monitor Promoted Groups**.

Note: You cannot suspend a group while it is in *Active* state, that is, during initial processing. [“Active, Copy, and Mirror states” on page 175](#) and [“Migrating datasets and volumes” on page 120](#) provide more information.

When the Suspend command (P line command) is issued, the migration group is suspended indefinitely. The group can be Resumed (R line command) at any time. When resumed, z/OS Migrator revalidates the migration group and restarts migration. In general, migration restarts in *Pending Mirror* phase (*P-Mirror*) as described in [“Resuming a suspended group” on page 176](#).

[“Migrating datasets and volumes” on page 120](#) describes the migration steps.

Note: Suspended groups may be unable to resume if a system failure or IPL intervenes. Therefore, it is not recommended that groups be left suspended for extended periods.

This chapter provides an overview of the Volume Mirror and Volume Migrator components of z/OS Migrator, including their features and requirements. Topics are:

- ◆ Overview 102
- ◆ Volume Mirror 102
- ◆ Volume Migrator 112

Note: The words mirror and migrator in z/OS Migrator should not be confused with the Symmetrix concepts of mirror positions or with RAID protection mechanisms.

Overview

In addition to logical migration, z/OS Migrator provides for volume mirroring and migration.

To use them, type **1** in the z/OS Migrator main window, and press **Enter**. When the Manage Group Members panel appears, type **CR** and press **Enter**. The Build Type pop-up appears.

```

+----- Build Type -----+
|                               |
| Select with an S a build type. |
|   s Build Migrate Member      |
|   _ Build Mirror Member       |
|   _ Build Logical Migrate Member |
|                               |
| F3=Cancel                     |
+-----+

```

To select a Mirror Group or a Migrate Group type an **S** in the blank before your choice.

Volume Mirror

Volume mirroring provides mainframe installations with host-based, volume-level mirroring. This ability is particularly useful for performance-sensitive volumes. In such situations, it may provide a lower response overhead than other tools.

Note: The integrity of mirror (and migrate) operations depends on the cooperation of all systems that can access the devices being mirrored (or potentially migrated). Your installation needs to include the activation of the z/OS Migrator server in the IPL procedures on all of your systems.

To use Volume Mirror effectively, you need to understand how it operates. Volume Mirror operation can be divided into the following stages:

- ◆ Volume group creation
- ◆ Activation, monitoring, and mirroring
- ◆ Multi-server communication

Volume group creation

Figure 11 shows the stages of volume group creation:

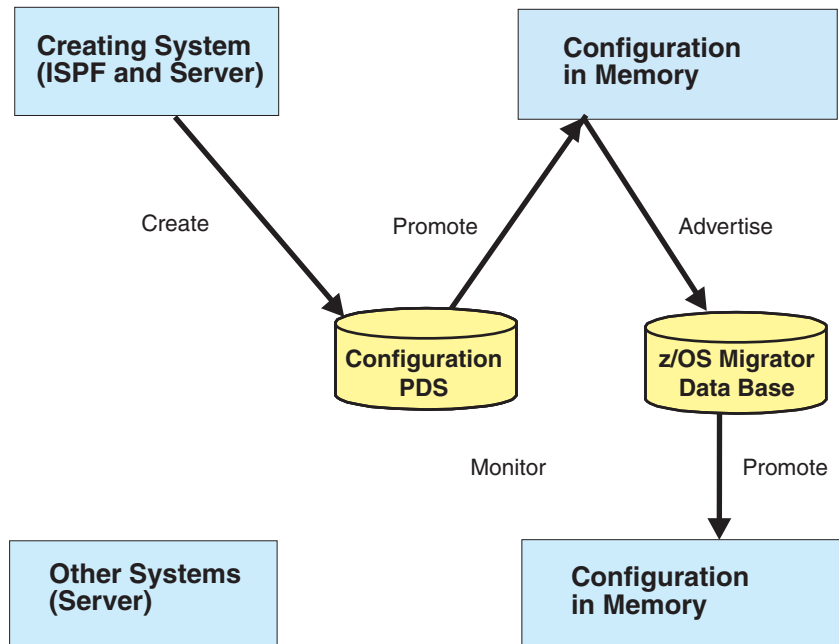


Figure 11 Volume group creation

When you initially start up the z/OS Migrator server, with no mirror volume groups defined, the z/OS Migrator server takes the following steps:

- ◆ Builds its basic global area in common storage
- ◆ Installs its subsystem I/O interface for command processing

At this point, you can employ the ISPF interface to define volume groups. Each group is:

- ◆ A named list of device pairs
- ◆ A specification of various processing options

Device pairs Each volume group can contain multiple device pairs. Each device pair in the group can be:

- ◆ A single pair of volumes
- ◆ A list of pairs
- ◆ A list of pair ranges

As described previously, a device pair consists of a primary and a mirror volume. The primary volume is the device to which your application is writing data. The mirror volume is the volume to which you want z/OS Migrator to mirror the data.

The mirror volume needs to be a Symmetrix volume for consistent migrations. Other restrictions can be found in [“Supported storage subsystems and dataset types”](#) on page 34.

Processing options The processing options are a series of configuration parameters you set at configuration time. There are three categories of configuration parameters:

- ◆ Execution parameters — that set values for execution.
- ◆ Volume group parameters — that specify how to process a particular volume group.
- ◆ Global parameters — that set certain values that apply to all volume groups.

You can define configuration parameters in three ways:

- ◆ Customize the startup JCL with execution parameters you want to use at z/OS Migrator startup.
- ◆ Create a partitioned dataset to hold global configuration parameters.
- ◆ (For volume group parameters) Use the ISPF dialog to define Mirror or Migrator volume groups.

Group definition storage After you create a volume group definition, Volume Mirror stores an image of that definition in a partitioned dataset referred to as the Configuration PDS. The Configuration PDS is used to cache the volume group definitions and other status information.

Once defined, the volume group may be “promoted.” Promotion causes the definition to be loaded into the z/OS Migrator database, where you can access it by various Mirror control functions using the global storage area.

All z/OS Migrator servers on multiple LPARs monitor the contents of the z/OS Migrator database continually. After a volume group has been promoted on one z/OS Migrator server, that volume group is automatically promoted on all z/OS Migrator servers sharing the z/OS Migrator database.

Note: Because short-term reserves are issued against the volume containing the z/OS Migrator database, you should place the z/OS Migrator database on a volume against which other reserves are not generally issued, which is not generally active, and which is shared by all participating LPARs.

Server activity

To maintain the integrity of Volume Mirror control, the z/OS Migrator server needs to be active on all systems that share any of the devices involved in mirroring, either as primary or mirror devices. This is because, as a host-based product, Volume Mirror needs to monitor the I/O from all systems.

All write I/O to the primary device needs to be propagated to the mirror on the system on which it occurs. Additionally, each system is responsible for preventing spurious I/O from other sources from being written to the mirror device.

For the same reason, it is not possible to use any of the primary or mirror devices for any operation that involves the transfer of data to the device entirely within the controller. The host must be able to see all writes to the devices. For example, the primary volume cannot be the target of a snap operation.

Activation, monitoring, and mirroring

Figure 12 shows the activation, monitoring, and mirroring stage:

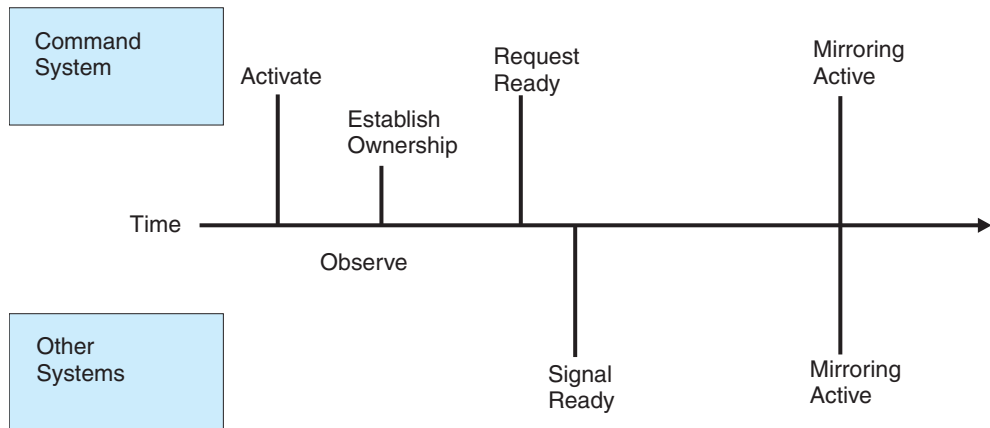


Figure 12 Activation, monitoring, and mirroring

After you have stored a volume group, any z/OS Migrator server can activate that group through an ISPF request, an operator command, or a batch command. That z/OS Migrator server becomes the owner of the volume group.

The activation process involves the following steps:

1. Validate the current state of the volume group on the activation host.
2. Activate the Volume Mirror I/O interface for these devices.
3. Communicate the activation request to other hosts within the complex.

The group owner

The host that performs the activation is referred to as the “owner.” The owner must be able to address all the devices in the volume group. After the other z/OS Migrator servers have signaled that they have also activated, the owner proceeds with the synchronization phase.

Synchronization During synchronization, all data is copied from the primary to the mirror volumes. During synchronization, additional track modification is monitored. When the number of tracks remaining to be copied falls below a certain threshold, I/O is briefly suspended and the last tracks copied. At this point, the primary and mirror are identical, and the I/O interface takes over to maintain that state.

Monitoring changed tracks At this point all participating systems are essentially performing the same task – monitoring I/O and mirroring writes. The owner continues to provide one additional service. The owner monitors changed tracks for the possibility that mirroring on one or more pairs will be suspended, and will need to be resumed.

The role of the server Technically, the z/OS Migrator server could be stopped on any system, including the owner, and mirroring would continue, carried on in the I/O interface. However, new activations and other command-oriented processes could not take place.

The server also provides special storage management as an ongoing service. In particular circumstances, the I/O interfaces may need to duplicate the channel program, or portions of the data from the channel program, issued to the primary device. The special storage management facilities in the z/OS Migrator server provide a number of storage pools for usage by the I/O interfaces.

Absence of the z/OS Migrator server in this circumstance may result in one or more volumes in the group being suspended. The exact impact of this is difficult to predict because it is highly dependent on the nature of the channel programs used by the specific application access method. However, the server could then be restarted with a warm start and it would immediately rejoin the network of Servers in the role it had before it was stopped: as owner or merely as a participating system.

In the same way, the temporary absence of a server, even the volume group owner, due to system IPL, does not disrupt mirror function for a volume group on other systems.

Suspend and resume

Figure 13 shows the suspend and resume stage:

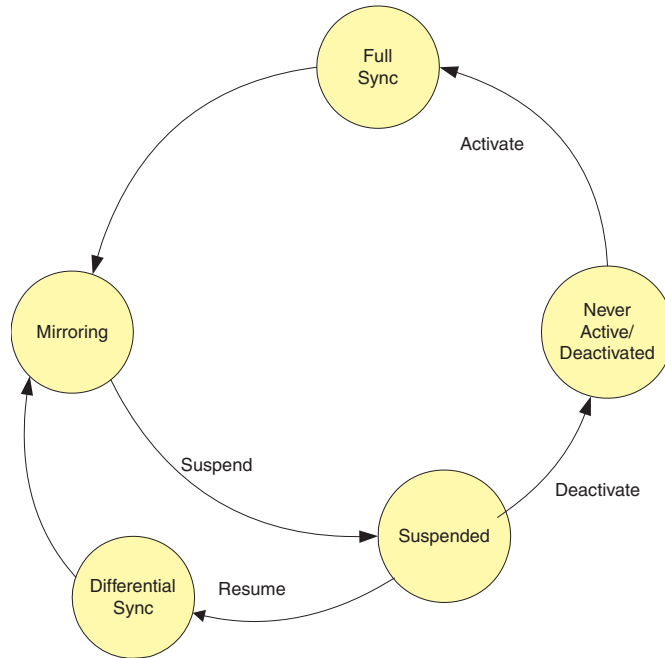


Figure 13 Suspend and resume

Mirroring can be suspended for a variety of causes. One possible cause of suspension is an error on the replicated I/O to the mirror.

In any case, if mirroring is suspended, there is no impact on the primary application I/O; however, mirroring is discontinued for the affected device pair. Naturally, console messages document the event and status displays report the suspended state.

A number of options are available within Volume Mirror, depending on the user and application requirements.

Default behavior option

The default behavior of the system is to extend the suspension of one mirroring pair to all the pairs in the volume group automatically and consistently. This guarantees that the set of mirror volumes remains in a consistent state with regard to *dependent I/O*. For more information refer to *dependent write* in the Glossary.

One type of suspension does not result in a consistent suspend for the volume group. If a primary volume is initialized so that a format write is performed on cylinder 0, track 0, for example ICKDSF INIT, then only the pair of which it is a part is suspended.

You may subsequently direct a RESUME command at the volume group to resynchronize this volume pair. Mirroring of the affected device pair resumes without having any impact on other mirroring. Non-formatting volume label operations, such as ICKDSF REFORMAT, do not result in suspend processing.

Deactivation option

Another possibility is to convert any suspend event automatically into a deactivation event. Deactivation differs from suspension in that a full resynchronization is required to resume mirroring on the volume group. In addition, write processing is no longer prevented to the mirror volume while the volume group is deactivated.

AutoVary

Another Volume Mirror option is the “AutoVary” option. During the synchronization process, the primary volume would generally be online while the mirror must be offline.

After the volumes are synchronized, Volume Mirror automatically varies the mirror volume to match the state of the primary. Thereafter, Volume Mirror continues to match the online/offline state of the mirror to that of the primary at any time.

One good reason to do this is to bind parallel access volumes (PAV) aliases to the mirror. To maintain performance levels, if the primary device is PAV, then the mirror should be PAV as well. Otherwise, longer I/O queues at the mirror slow down I/O to the primary.

Apparent volsers

Physically, the mirror has the same volser as the primary since it is an exact copy. To allow the mirror to be online, Volume Mirror automatically responds to any request to read the volume label by substituting a volser altered according to a specification you make during configuration. This altered volser is referred to as the *apparent volser*.

After setting the apparent volser, you can perform read operations on the mirror. Simply make sure that your applications that logically refer to the volser use the apparent volser.

An exception is made for applications, such as backup operations, which physically read the first track. These will see the true volume label with the volser identical to that of the primary. This allows you to:

- ◆ Perform a consistent suspend of the volume group
- ◆ Run backups against the mirror volumes to achieve a consistent backup copy of the data
- ◆ Perform a resume of the volume group

Remember that many volume accesses, which can be thought of as strictly reading from a volume, may actually update the VTOC, specifically, the last accessed date. Such accesses fail as though an I/O error has occurred because of Volume Mirror blocking the write. This condition is additionally externalized in the system log.

Multi-server communication

As described previously, the volume group promotion and the activation process relies on the communication of status information through the z/OS Migrator database. Volume group deletion is also communicated across z/OS Migrator servers in that way.

Currently, if you want to change a volume group definition, you must deactivate and delete that volume group definition. You can then recreate the volume group definition and promote the group again. Activation causes the volume group to be resynchronized and mirroring to resume.

Naturally, the mirroring process itself requires much tighter integration of multiple z/OS Migrator servers on different LPARs. The I/O interfaces accomplish this through a number of low-level techniques that physically communicate the current mirror state to all LPARs at the moment a write is performed.

In this way, the exact moment at which full synchronization occurs is communicated, as well as suspend and resume events. This communication is managed entirely within the Volume Mirror I/O interface for maximum integrity.

On LPARs where writes are performed infrequently, the mirror status may appear to be delayed. In this case, the status is communicated through the z/OS Migrator database during its normal polling period.

Device pair status

Keep in mind that a participating system may not detect the current state of a device pair until that system actually attempts application I/O. In this case, the state may be thought of as “pending” and will be fully recognized at the start of the next I/O.

Because this can make status displays somewhat confusing, status is secondarily communicated through the z/OS Migrator database. This allows the displays to reflect the current status in a user friendly way. However, because the z/OS Migrator server is only polling the z/OS Migrator database for periods of some seconds, other systems do not report the expected status instantaneously.

One important piece of information recorded in the z/OS Migrator database is the time of the last “heartbeat” for each participating z/OS Migrator server. By monitoring these values, each z/OS Migrator server can report on, and make decisions regarding, whether other z/OS Migrator servers are live. As with other status within the z/OS Migrator server, you can use an operator, ISPF, or batch command to report the z/OS Migrator server status.

Volume Migrator

Volume Migrator provides host-based services for data migration on mainframe systems.

Migrator provides installations with a volume level migration facility:

- ◆ From third-party devices to Symmetrix devices
- ◆ Between Symmetrix devices

Volume Migrator lends itself to workload balancing applications or to upgrading a site to use newer Symmetrix models, such as Symmetrix V-MAX™. Volume Migrator performs volume-level migration non-disruptively, and without requiring that EMC SRDF® be installed.

Components

As with Volume Mirror, Volume Migrator uses z/OS Migrator server and the I/O interface. The z/OS Migrator server performs the following tasks:

- ◆ Manages the overall Volume Migrator configuration
- ◆ Loads configured volume groups
- ◆ Services user commands
- ◆ Coordinates with the actions of z/OS Migrator servers on other systems

As with Volume Mirror, you can control the z/OS Migrator server's Volume Migrator functions through the identical set of ISPF panels¹, batch commands, and operator commands you use to control Volume Mirror.

Note: [Chapter 6 on page 133](#) describes the ISPF interface and [Chapter 9 on page 245](#) provides examples of how to use the ISPF screens for both Volume Mirror and Volume Migrator. [Chapter 7 on page 191](#) provides information about the operator commands.

1. Some of the ISPF panels have slightly different fields for Volume Mirror and Volume Migrator.

Volume Migrator operation

In using the Volume Migrator, you must first define migration volume groups. Each migration volume group consists of one or multiple device pairs. The pairs can be:

- ◆ A pair of volumes
- ◆ A list of pairs
- ◆ A list of pair ranges

Pairs consists of a *source volume* and a *target volume*. The source volume is a data volume from which you want to copy data. The target volume is the volume to which you want to copy that data.

The mirror volume needs to be a Symmetrix volume for consistent migrations. Other restrictions can be found in [“Supported storage subsystems and dataset types” on page 34](#).

Group storage and activation

After you create a volume group, the volume group definition is stored as a member of the z/OS Migrator database. After you have stored a volume group definition, you can promote the group and, once promoted, the group is loaded into memory where you can access it through ISPF request, operator command, or batch command.

The group owner

As with Volume Mirror, the host that performs the activation is the “owner.” The owner must be able to address all the devices in the group. After the other z/OS Migrator servers have signaled that they have also activated the migrator for all the devices in the volume group known to them, the owner proceeds with the synchronization phase.

Synchronization

During synchronization, Volume Migrator also uses host resources (UCBs, CPU, and channels) to copy the contents of a source volume to a target volume. Additional track modification is monitored during the process. When the number of tracks remaining to be copied falls below a certain threshold, I/O is briefly suspended and the last tracks are copied. At that point, the source and target are identical.

Migration complete options

After achieving a state of synchronization between volumes, Volume Migrator gives you several completion options:

- ◆ Completing a group by swapping at completion of the volume copy (SWAP)
- ◆ Completing the migration consistently with CONSISTENT swap option (CONSISTENT)
- ◆ Creating a copy of a volume for testing (SPLIT)
- ◆ Doing intermittent synchronization (CONSTANTCOPY)

Note: SWAP is the only migration completion option that completes the migration. CONSISTENT is an option that leaves the group in MIRROR status, allowing the user to SWAP the volumes consistently at a later time of their choosing.

Each of these options is described below:

Swapping the source and target devices

If the SWAP option is selected, z/OS Migrator dynamically and non-disruptively swaps the source and target devices. This option performs the following tasks:

- ◆ Swaps contents of the UCBs for the source and target volume.
- ◆ Redirects all I/O to the target volume, transparently to the applications using the volume.
- ◆ Terminates the relationship between the source and target volumes.

Consistent Swapping

If the CONSISTENT option is selected, a volume group ending in a swap will now be managed in the following way:

- ◆ When the volume group is activated, each volume pair is added to a specialized z/OS Migrator AutoSwap group.
- ◆ Once all pairs are added to the group, the group is validated by AutoSwap.
- ◆ Migration of the volume pairs is initiated, pair by pair.
- ◆ As the initial copy operation for each pair completes, a final synchronization is performed with application I/O entering a brief inactive state, and the pair entering Mirror mode. In this

state, all application write I/O is synchronously performed to both the source and the target, thereby keeping each volume pair, and any replication pairs, in sync.

- ◆ Once all volume pairs are mirroring, the group will be reported to be in Mirror state and the user may enter the z/OS Migrator SWAP group command to initiate a group swap from all source volumes to all target volumes. Complete data consistency is maintained across such a swap.

Once the swap has completed, various cleanup operations are automatically performed and the group is then reported as Complete.

Splitting the source and target volumes

If the SPLIT option is selected, z/OS Migrator severs the relationship between the source and target volumes and leaves the target volume offline.

- ◆ An option is available to relabel the target volume with a user-specified volser and vary the volume online.
- ◆ Another leaves the target volume offline without changing the volume serial numbers.

Constant Copy mode

If the CONSTANTCOPY option is selected, z/OS Migrator continues to copy additional changed tracks periodically from the source to the target volumes.

This option maintains a state of near synchronization between the source and target volumes on a continuous basis. Volumes in constant copy mode can be subsequently SPLIT or SWAPed without waiting for a long synchronization period to complete.

Change tracking during the synchronization process and in CONSTANT COPY mode is achieved by using the Symmetrix Differential Data Facility for Symmetrix devices and by using proprietary system interfaces for third-party migration source volumes.

Running z/OS Migrator

This chapter provides information for running z/OS Migrator and includes starting the z/OS Migrator ISPF Monitor, as well as the basic steps for starting the migration process. Topics include:

- ◆ [Running the z/OS Migrator ISPF Monitor](#) 118
- ◆ [Migrating datasets and volumes](#) 120

Running the z/OS Migrator ISPF Monitor

Before you can begin using the z/OS Migrator ISPF Monitor, it must be customized to your environment. (“[Option S – Set User Session options](#)” on [page 137](#) describes the customization process.) The z/OS Migrator ISPF Monitor is a feature that is installed at the same time as the z/OS Migrator program. The z/OS Migrator ISPF Monitor utilizes REXX execs. As an ISPF application, it requires ISPF Version 5.2 and TSO/E Version 2.4 at the minimum.

Note: Refer to “[Step 14. Prior to first execution, specify parameters in SAMPLIB](#)” on [page 68](#), for details on the required editing of the Mainframe Enablers SAMPLIB file with specific installation parameters prior to the first execution of z/OS Migrator.

To start and customize the z/OS Migrator ISPF Monitor

1. Start the z/OS Migrator ISPF Monitor.

Execute the supplied REXX exec, replacing the *hlq* with the high-level qualifier you specified when installing the z/OS Migrator product.

Note: The z/OS Migrator ISPF Monitor components are generally installed as part of ResourcePak Base of Mainframe Enablers. The supplied REXX exec, EFMXEXEC, can be found in the Mainframe Enablers SAMPLIB.

```
EX 'hlq.prodid.SAMPLIB(EFMXEXEC)'
```

2. Customize User Session Options.

When entering the z/OS Migrator ISPF Monitor for the first time, the **Specify the Session parameters** panel displays.

```

                                Specify the Session parameters
Command ==>

Command Prefix . . . . . @$
Configuration PDS. . . . 'JDOE1.XYZXYZMM.DW001'
Console Message Timeout. 1 (0-10)

Use Log PDS. . . . . Y (Y/N) Days to save . . 7 (1-30)
  Use Browse or View. . V (B/V)
  Log PDS Name. . . . . 'JDOE1.EFMM300.TFMLLOG'

ZOSM Load Library . . . 'JDOE1.EFMM300.LINKLIB'
ZOSM Security Library .. 'JDOE1.EFMM300.SECCOM'
UNIT . . . SYSALLDA and VOLUME . . _____ (optional) for allocations

Stack Console output . . Y (Y/N)
Clear variables. . . . . N (Y/N)
Simulate Mode. . . . . N (Y/N)
Debug Mode . . . . . N (Y/N)

                                F1=Help   ENTER=Save Session   F3=Exit

```

Figure 14 Specify the Session parameters panel

You will be required to fill in the information so that your z/OS Migrator ISPF Monitor can interact with the z/OS Migrator server.

Each time the z/OS Migrator ISPF Monitor is started, it checks to see if it was started with a different version of the monitor. If a change was detected, a panel will pop up indicating this situation and informing you that you will be sent to the **User Session Options** panel to verify the settings.

On the **User Session Options** panel, you will be required to set the configuration and control parameters according to the settings specified during the z/OS Migrator server customization process. [“Option S – Set User Session options” on page 137](#) provides details on the customization process.

3. Save your User Session options.

You must save the User Session options you have defined by pressing **Enter** and then **PF5**, before exiting this panel by pressing **PF3**.

Migrating datasets and volumes

The following sections describe the general process and phases that are involved in migrating datasets and volumes.

All product features and capabilities discussed in the z/OS Migrator Migration Process, as well as how to use the z/OS Migrator ISPF Monitor, are explained in greater detail in [Chapter 6, “Using the z/OS Migrator ISPF Monitor”](#) on page 133.

The z/OS Migrator migration process defined

Before you begin using z/OS Migrator to migrate your data, you should become familiar with the way z/OS Migrator works. [Table 10](#) outlines the z/OS Migrator dataset migration process and phases.

Table 10 z/OS Migrator data migration process

Migration Phases	Description	Details
Step 1. Group definition	The dataset(s) to be migrated are defined in a migration group.	“Step 1. Group definition” on page 121 and “Selecting datasets” on page 83.
Step 2. Activate	Activating a migration group initiates the data migration process for the defined migration group.	“Step 2. Activate a migration group” on page 124.
Step 3. Copy	Data is asynchronously copied from the source datasets to the target datasets that are defined in a migration group.	“Step 3. Copy” on page 125.
Step 4. Synchronize	All final differences between source and target datasets in a migration group are synchronized and the migration group is prepared for <i>Mirroring</i> .	“Step 4. Synchronize” on page 126.
Step 5. Mirror	The migration group is put into a state of synchronous mirroring.	“Step 5. Mirror” on page 126.
Step 6. Diversion	In this phase, the actual logical relocation of datasets occurs. Source and target datasets metadata are modified and all I/O activity is redirected to the new location.	“Step 6. Divert” on page 127.
Step 7. Complete	Although the metadata has been modified, applications that were active before diversion will continue to have their I/O redirected until they de-allocate the dataset.	“Step 7. Complete” on page 128.
Step 8. Post-Completion	Once the migration has completed, you will want to clean up your migration group's source datasets and the storage resources they reside upon.	“Step 8. Post-completion” on page 130.

Step 1. Group definition

Before a data migration can begin, a group definition must first be created. A group definition consists of the following:

- ◆ Migration group name
- ◆ Migration options
- ◆ Source dataset(s)
- ◆ Target dataset(s) name and location

Target and Source dataset names

Referring to the affected datasets during a migration can be confusing, so we provide the following discussion of terminology.

As part of the group definition, you specify a "target dataset name." Of course, the objective of the dataset migration is to relocate the original dataset to a new target location under its own name.

By "target dataset name" we actually mean a name which is given to the dataset in the target location during the migration process, and then reassigned to the source dataset after the migrated dataset becomes active in the new location.

The key moment in the migration of any dataset is called Diversion, actually the start of the period during which existing application I/O will be redirected to the target location. As part of the Diversion process, the prior dataset names of the source (the "real" name) and the target are "swapped" with the result that the target becomes the "real" dataset and the source can be abandoned or deleted. In fact, depending on options you've selected (see ACTVTERM) you may have to wait until the migration is fully complete for the group as a whole before you can delete the renamed source datasets and their volumes.

Remember that once the migration is complete the dataset with the original source dataset name is the one you want to keep. The dataset with the original "target" dataset name is the one you want to delete.

The following procedure is a general outline for defining a group. [“Option 1 – Manage Groups” on page 140](#) provides more information.

To create a group definition:

1. In the z/OS Migrator main panel type **1** (for Manage Groups) and press **Enter**.

The Manage Group Members panel appears.

2. In the Manage Group panel type the **CREATE** (cr) command and press **Enter** to display the **Define z/OS Migrator Group** panel.

Define z/OS Migrator Group		Scroll ==> CSR
Command ==>		
Primary Commands	: EXit NExt	
Group Name _____	
Group Options	. . . N Replace Existing Datasets (Y/N)	
	Y Tolerate Allocation Failure (Y/N)	

3. Define the source dataset(s) to be moved.

When defining the source dataset(s), you must use:

- A minimum of a dataset name, either an absolute dataset name or a dataset name mask, or a comma-delimited list of either or both dataset name formats.
- And/or a valid DFSMS Storage Class name.
- EXCLUDE is optionally used to bypass selection of datasets that match the Source DSN, Volume or Storage Class criteria. EXCLUDE can be an absolute dataset name or a dataset name mask, or a comma-delimited list.

4. Select group options, and then type the **NEXT** command (or **NE**) and press **Enter** to proceed to the subsequent panels.

There are a number of options you may specify for a group definition. These focus on diagnostic information as well as options related to reusing target datasets. [Table 16, “Source dataset options,” on page 146](#) provides more information.

5. Define the target dataset(s) location(s).

Specify a combination of the following to determine the target location:

- You must specify a Target DSN as an absolute dataset name or dataset name mask. A Target dataset is cataloged and paired with each selected source dataset. It functions as a work dataset in a source-target dataset pair. The following requirements should be carefully considered when selecting a target dataset name:
 - The name (mask) must be easily identifiable. Once the migration has completed, you will, at some point, wish to delete the datasets from the source volumes. Following migration, the source volumes will contain the datasets cataloged as the target dataset names. Ensuring that these are easily identifiable is important.
 - In a DFSMS managed environment, the Target Storage Class field can be used to direct the allocations to a different Storage Group than the one that currently contains the volumes where the source datasets reside. To accomplish this, the site's ACS routines must also allow new allocations of the source datasets to be placed into the specified target Storage Class. [“Systems managed storage \(DFSMS\) allocation” on page 94](#) provides more information.

Note: In DFSMS managed environments, the source dataset name determines the volume that is selected for the destination of the migrated dataset because z/OS Migrator drives the allocation through the site's ACS routines using each source dataset's attributes (dataset name, physical characteristics, management criteria) at *Activation*.

- Specify one or a list of target volumes. In a DFSMS managed environment, specific volumes will have their target datasets allocated if they are within a Storage Group that is valid for the source datasets' Storage Class or the target Storage Class, if specified, and they have been assigned the “Guaranteed Space” attribute.



CAUTION

If you migrate target datasets that are not owned by z/OS Migrator, extra care must be taken since these datasets are not protected on the non-owning system.

6. Verify your new group definition.

Enter the `Verify Primary` command on the **Define Logical Migration Group** panel. This validates your input. [“Verifying a group definition” on page 163](#) provides additional information.

7. Save your group definition.

Once you have validated your input, press **Enter**. This will cause your input to be reflected back in the lower half of the panel in the format that becomes the group's member in the Group Definition PDS. You will be prompted for a name. The name will be prefixed with “LD”. The group definition will be saved to the Group Definition Dataset.

If the member name already exists, a new name will be generated. The new name is in the form of LDxxx001, where xxx is today's Julian date.

8. Promote the group definition.

You must now promote the group definition to the z/OS Migrator database. The *Promoted* group is then listed by the group name you provided during the creation phase and is identified by a “P” to the right of the name. [“Promoting a group definition” on page 165](#) provides additional information.

Step 2. Activate a migration group

During the *Activation* phase, z/OS Migrator allocates datasets on the target volumes. If any errors occur during allocation, the migration group may not activate. [“TOLERATE_ALLOCATION_FAILURE \(or TOLALLOF\)” on page 224](#) describes this parameter.

Note: HSM migrated datasets are skipped when you activate a migration group, and therefore are not migrated. As a result, these datasets are not visible when you attempt to display them using the `Display DSN groupname /S` command.

To activate a migration group:

1. Select **Option 2 - Monitor Promoted Groups** as described in [“Option 2 – Monitor Promoted Groups” on page 168](#).
2. Place an **A**, for *activate*, next to the migration group and press **Enter**.

The time to activate a migration group depends on the number of target datasets that must be allocated and cataloged. In this case, the migration group will temporarily go into a *P-Active* state.

Once the migration group has been activated, z/OS Migrator automatically moves the migration group to the *Copy* phase.

[“Activating a group” on page 175](#) provides more information on *Activation*. [“Datasets allocated after activation” on page 85](#) for more information on datasets allocated after *Activation*.

Step 3. Copy

During the *Copy* phase, z/OS Migrator initiates and performs these tasks:

- ◆ The z/OS Migrator server asynchronously copies the source datasets to the target datasets.
- ◆ The z/OS Migrator I/O Monitor Routines monitor all activity to all source datasets and track all modifications to the source dataset.
- ◆ Once having completed the initial copy of source datasets, the z/OS Migrator server refreshes the changed data repeatedly until it reaches a point where it can quickly *Synchronize* the source and target datasets with minimal disruption to applications or the system. Once this point is achieved, z/OS Migrator automatically moves to the *Synchronization* phase for this migration group.

In some instances, a migration group may take some time to complete the *Copy* phase across all datasets in the migration group. The *Copy* phase happens after the group is activated. [“Activating a group” on page 175](#) provides more information.

Managing performance during the copy phase

During the *Copy* phase, depending on the size of a migration group, a good deal of I/O could possibly be driven by z/OS Migrator. To control the pacing of I/O during the z/OS Migrator *Copy* phase, you can modify parameters that the z/OS Migrator server uses.

The parameters that can be modified include:

- ◆ MAXIO
- ◆ MAX_CHANNEL_IO
- ◆ MAX_DEVICE_IO

These parameters are located in the z/OS Migrator Configuration Member as documented in [“Running the z/OS Migrator ISPF Monitor” on page 118](#).

Step 4. Synchronize

A migration group enters the *Synchronization* phase when z/OS Migrator determines it can quickly synchronize all data between source and target datasets in the migration group without causing disruption to the systems and applications using this data.

During the *Synchronization* phase, z/OS Migrator initiates and performs the following tasks:

- ◆ The z/OS Migrator I/O Monitor routines *dynamically* hold all I/O to the source dataset so final synchronization can be achieved.
- ◆ Copies all remaining differences from the source datasets to the target datasets. At this point, there is very little difference between datasets and this operation occurs very quickly.
- ◆ Allows normal I/O operations to continue once synchronization has been achieved.

When the *Synchronization* phase has completed, the migration group indicates that it is in *Mirror* state.

Step 5. Mirror

Once the source and target datasets within the migration group have successfully synchronized, they enter the *Mirror* phase.

During the *Mirror* phase, updates to source and target datasets in the migration group are applied simultaneously. If an I/O error occurs at the target volume, the Group is *Suspended*.

Once the *Mirror* phase has been achieved, *mirroring* continues until one of the following occurs:

- ◆ A *Divert* command is issued against the migration group.
- ◆ A *Suspend* command is issued against the migration group.

Step 6. Divert

The *Diversion* phase executes in two sub-phases:

- ◆ [Sub-phase 1: Processing of the Divert command](#)
- ◆ [Sub-phase 2: Diversion of I/O from active applications to the target dataset](#)

Sub-phase 1: Processing of the Divert command

During this sub-phase, z/OS Migrator modifies all metadata for the source and target dataset pairs within the migration group, effectively swapping the identities of the source and target.

To accomplish this, z/OS Migrator:

- Serializes access to the metadata for collections of datasets cataloged in a particular source/target catalog pair.
- Updates all metadata to accomplish the identity swapping of the source and target datasets:
 - Modifies all volume-based metadata in the VTOC, VTOCIX, and VVDS.
 - Modifies the catalog entries for the source and target dataset pairs, and refreshes catalog data buffers for the catalogs involved across all z/OS images in a shared storage complex.

Sub-phase 2: Diversion of I/O from active applications to the target dataset

With the source and target dataset identities switched, I/O to any of the source datasets previously allocated by ongoing applications is *diverted* to the target dataset instead.

Note: Once the `Divert` command processing is completed, any new application allocations will automatically be directed to the target dataset.

[“Diverting a Logical Migrator group” on page 176](#) provides more information on the *Diversion* phase. The following outlines the process you must initiate to move a migration group to the *Divert* phase.

To divert a migration group:

1. Select **Option 2 - Monitor Promoted Groups** as described in [“Option 2 – Monitor Promoted Groups” on page 168](#).

2. Place a **V**, for *Divert*, next to the migration group and press **Enter**.

The migration group enters *Pending Diversion* state (*P-Divert*). Once all metadata changes have been made and the entire migration group has been *diverted*, the state of the migration group changes to *Diverted*.

The z/OS Migrator Monitor executing on a particular system (LPAR) will reflect the status of the group on that LPAR. Datasets that are allocated on their source volumes (persistent allocations) by jobs executing on this LPAR will be shown as *Diverted*. If there are group datasets with persistent allocations on other LPARs, but none are still being diverted on this LPAR, the z/OS Migrator Monitor will show the group as *Pending Complete*.

Note: While z/OS Migrator diverts active DB2 volumes, the I/Os are temporarily held until the renaming is complete.

Datasets allocated PRIOR TO diversion

Prior to *Diversion*, various z/OS address spaces may have source datasets allocated. Once the *Divert* command processing has completed, the z/OS Migrator I/O monitor ensures that all I/O requests for these address spaces are diverted (that is redirected) to the target dataset, which now has the original source dataset's name. This continues until the address space or application terminates, or the allocation is freed. (See note above.)

Datasets allocated AFTER diversion

After the *Divert* command processing has completed, the catalog correctly points to the target datasets, which have been given the original source dataset names. Therefore, any new allocations made using the source dataset name will be made directly to the target datasets on their new volumes. The source datasets may now be referred to using the original target dataset name, generally just to be deleted after *Completion* as no longer useful. Remember that, after *Diversion*, the source datasets become out-of-date as application I/O is no longer directed to them. (See note above.)

The note on [page 121](#) discusses the naming of target dataset names.

Step 7. Complete

Dataset migration groups will remain in the *Diversion* phase until all applications have relinquished their allocation of migrating datasets and z/OS Migrator no longer has to redirect I/O to the targets.

Group *Completion* is achieved LPAR by LPAR in the shared storage complex depending on the dataset usage of the applications on each server. Each server periodically examines allocation information to see if any active address space still has any source dataset allocated. Once all such allocations are freed through the normal action or completion of the system or application processing, then it is no longer necessary for z/OS Migrator to divert any I/O. Locally, the z/OS Migrator processing is complete for the group. However, the group cannot be marked fully *Complete* until all LPARs reach the same state.

The default operation of z/OS Migrator is to keep the source dataset and source volume allocated to the application during the diversion phase. This prevents the space occupied by the dataset from being reused and also prevents the source volume from being varied offline. Once all applications free the original source allocation then the source resources may be reused.

If the ACTVCMP option is set, then additional processing is performed by z/OS Migrator during the diversion phase. This processing redirects allocations for the source datasets/devices to the target datasets/devices. This allows for removal or reuse of the source devices while the group is still in diversion. I/O still needs to be redirected until the group moves to completed status. The group remains in diversion status even though the source device can now be varied offline.

There are events that occur in the usage of a dataset when z/OS Migrator determines diversion is no longer required. Some examples of these events are:

- ◆ Application reopens the dataset
- ◆ New extents are allocated to the dataset
- ◆ Applications are recycled

When these events occur the group goes to completed status and the processing for the group is terminated.

Depending on timing and the manner in which you display the status of a group, that is, ISPF Monitor versus operator command, you may see the group progress locally from *Divert* to *Pending Complete* to *Pending All Complete*. The group will remain in *Pending All Complete* status until z/OS Migrator servers see the group in that status. The status of the group will then display as *Complete*.

Identifying address spaces being diverted

To identify z/OS address spaces that are being *diverted* is a simple process with z/OS Migrator. By using the z/OS Migrator ISPF Monitor, **Option 2 - Monitor Promoted Groups**, you can display all address spaces that are currently diverted across all z/OS images in a shared storage environment.

Step 8. Post-completion

Data migrations in many cases are undertaken to free a storage resource. Plans for the storage resources may include freeing the physical storage array as it is coming off lease, or reusing the storage for other application data needs. Regardless of the reason, before the resources can be reused, you must take the following actions.

To initiate post-completion storage resource clean-up:

1. Ensure that the migration group has *Completed*.

In the z/OS Migrator ISPF Monitor, **Option 2 - Monitor Promoted Groups**, check that the migration group is in the *Complete* state.

2. Delete the renamed source datasets of the migration group. (“[Target and Source dataset names](#)” on page 121 discusses the naming of source and target datasets.)

Once you have identified these datasets, you may then delete them, which will make the space available for reuse or, once all data has been migrated, free the storage resource for disposition.



CAUTION

Never reallocate storage space from migration groups where the migration group is not in the *Complete* state. If you do, you may experience serious data integrity problems.

3. Delete the migration group from the z/OS Migrator database.

Once the migration group has *completed*, you must delete the migration group from the z/OS Migrator database. Simply go to the z/OS Migrator ISPF Monitor, **Option 2 - Monitor Promoted Groups**, to delete the group.

When a migration group is deleted, z/OS Migrator performs these tasks:

- It ensures that the migration group is in a state that will allow it to be deleted.

- It removes the migration group from internal z/OS storage (memory). This frees valuable system memory.
- It deletes the migration group from the z/OS Migrator database.

Note: If you do not delete migration groups, there is a possibility that you may eventually run into z/OS memory issues with ECSA. Therefore, it is recommended that you delete the migration group once it reaches *completion*.

Using the z/OS Migrator ISPF Monitor

This chapter provides detailed descriptions and procedures for using the z/OS Migrator ISPF Monitor. Topics include:

- ◆ Overview 134
- ◆ Option S – Set User Session options 137
- ◆ Option 1 – Manage Groups 140
- ◆ Option 2 – Monitor Promoted Groups 168
- ◆ Option 3 - Display Configuration Information 179
- ◆ Option 4 - Operator Interface 180
- ◆ Option 5 – Display Host Messages 181
- ◆ Option 6 – System Change Summary 183
- ◆ Option 7 – Message Help 184
- ◆ Option 8 – Monitor Command Line Help 187
- ◆ Option 9 – Display Installation Options 189
- ◆ Option V– View or Browse Log Dataset 190

Overview

The z/OS Migrator ISPF Monitor is a component that is installed at the same time as the z/OS Migrator tool. The Monitor component consists of REXX execs, which require ISPF Version 5.2 and TSO/E Version 2.4 at the minimum. The ISPF Monitor is used to create group definitions, including relevant source/target pairs, and manage the dataset migration environment.

Your tool for controlling both logical and volume migration is the ISPF interface. You initiate the ISPF interface using the EFMXEXEC REXX script that you can associate with a menu item in an installation menu.

Note: To ensure that only one copy of the ISPF Monitor is run by each TSO user, the monitor is protected by the allocation of a dummy dataset. Should the ISPF monitor terminate abnormally for any reason, this dataset may be left allocated. If this situation occurs, attempting to restart the ISPF Monitor will fail with the message; `Exec Already Active`. To correct, make sure there isn't another copy of the ISPF Monitor being used by the current TSO user, and issue the TSO command, **FREE FI(LOCK)**. A restart of the ISPF Monitor should now be successful.

Figure 15 shows the general structure of the z/OS Migrator panels:

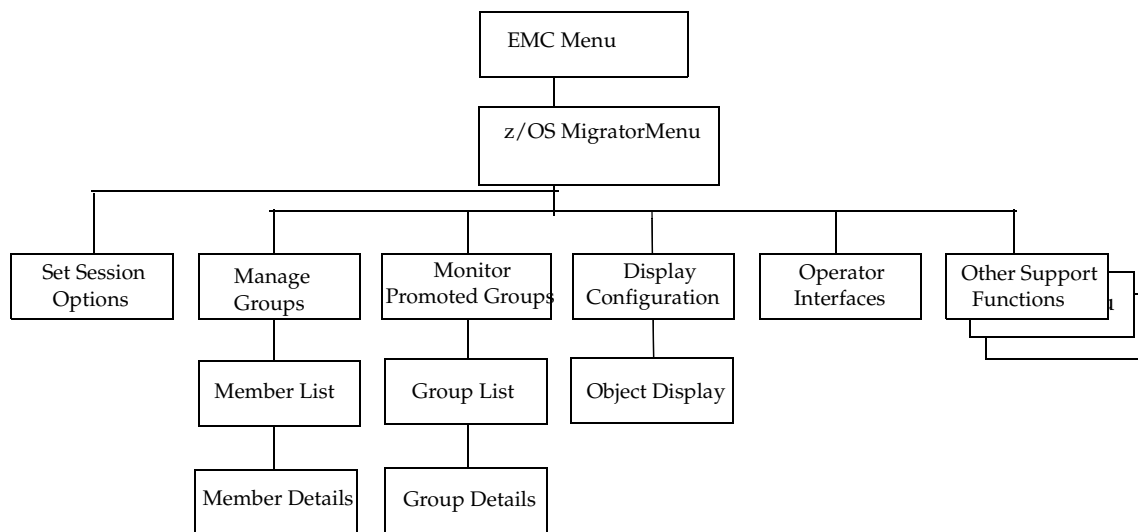


Figure 15 z/OS Migrator functions panel topology

Most of the infrastructure shown in [Figure 15](#) is common to both dataset and volume migration. They diverge in the appearance of specific panels for displaying the group unit list and for defining groups.

z/OS Migrator Functions panel

Once the monitor (EXEC 'hlq.prodid.SAMPLIB(EFMXEXEC)', located in the Mainframe Enabler SAMPLIB) has been started, it provides complete control over z/OS Migrator session(s).

The first panel displayed is the z/OS Migrator Functions panel, also referred to as the Main menu.

```

EMC Logical Data Migration Facility

Function Number                06/03/07 10:10:19
or Command ==>                SMFID: X04
                                CPFY: @$

                                Z/OS Migrator 3.0.0 FUNCTIONS

1 Manage Groups                6 System Change Summary
2 Monitor Promoted Groups      7 Message Help
3 Display Configuration Information 8 Monitor Command Line Help
4 Operator Interface           9 Display installation Options
5 Display Host Messages

S Set User Session Options
V View or Browse Log dataset

F1=HELP      F3=Exit

```

Figure 16 z/OS Migrator Functions panel

The z/OS Migrator Functions panel provides the following functions:

- ◆ “Option 5 – Set User Session options” on page 137
- ◆ “Option 1 – Manage Groups” on page 140
- ◆ “Option 2 – Monitor Promoted Groups” on page 168
- ◆ “Option 3 - Display Configuration Information” on page 179
- ◆ “Option 4 - Operator Interface” on page 180
- ◆ “Option 5 – Display Host Messages” on page 181
- ◆ “Option 6 – System Change Summary” on page 183
- ◆ “Option 7 – Message Help” on page 184
- ◆ “Option 8 – Monitor Command Line Help” on page 187

- ◆ [“Option 9 – Display Installation Options” on page 189](#)
- ◆ [“Option V– View or Browse Log Dataset” on page 190](#)

Note: Option V - View or Browse Log dataset is only displayed on the z/OS Migrator Functions panel if logging is set to active under the User Session Options. [“Option S – Set User Session options” on page 137](#) provides details.

Help panels are available for all actions within the z/OS Migrator ISPF Monitor. Some panels have multiple help panels depending upon the location of the cursor.

Option S – Set User Session options

This option allows you to configure your z/OS Migrator session options and control parameters, including setting the command prefix value and specifying the PDS configuration database.

Starting the z/OS Migrator ISPF Monitor for the first time

When the z/OS Migrator ISPF Monitor is started for the first time, the Set User Session Options panel is displayed. This is where you set your user environment options.

When invoked, Option S displays the Specify the Session Parameters panel.

```

                                Specify the Session parameters
Command ==>

Command Prefix . . . . . @ $
Configuration PDS. . . . 'JDOE1.XYZMM.DW001'
Console Message Timeout. 1 (0-10)

Use Log PDS. . . . . Y (Y/N) Days to save . . 7 (1-30)

    Use Browse or View. . V (B/V)
    Log PDS Name. . . . . 'JDOE1.ZOSMvrm.TFMLLOG'

ZOSM Load Library . . . 'JDOE1.ZOSMvrm.LINKLIB'
ZOSM Security Library .. 'JDOE1.ZOSMvrm.SECCOM'
UNIT . . . SYSALLDA and VOLUME . . _____ (optional) for allocations

Stack Console output . . Y (Y/N)
Clear variables. . . . . N (Y/N)
Simulate Mode. . . . . N (Y/N)
Debug Mode . . . . . N (Y/N)

                                F1=Help    ENTER=Save Session    F3=Exit

```

Figure 17 Specify the Session parameters panel

To modify or specify configuration and control parameters on the Specify the Session parameters' panel:

1. Type the required information next to the corresponding parameter(s).
2. Press **Enter**, followed by **PF5** to save your settings.

Table 11 describes the configuration and control parameters on the Specify the Session Parameters panel.

Table 11 Configuration and control parameters (1 of 2)

Parameter	Description
Command Prefix	Command prefix value used to identify the z/OS Migrator subsystem and communicate with the z/OS Migrator service task. This must match the prefix (CPFX=) specified at runtime in the execution parms.
Configuration PDS	<p>When initially created, group definitions and all their corresponding configuration information are stored in the Group Definition Dataset or PDS. Once a group definition is promoted, it is transferred to the z/OS Migrator database, and the original member remains.</p> <p>Enter a fully qualified PDS name in standard ISPF format. Automatic compression of the PDS occurs when a group configuration/member is deleted. If the specified dataset does not exist, you will be prompted for space allocation.</p> <p>This definition may be specified in multiple user configurations on the same LPAR or other LPARs, assuming the dataset is available on other LPARs.</p> <p>Note: Multiple users specifying the same dataset implies that the dataset could be reserved by one user and lock out another.</p>
Use Log PDS <ul style="list-style-type: none"> • Background logging • Use Browse or View • Log PDS Name 	<p>A value of Y enables the logging feature. When logging is enabled, the specified PDS is created and each day's worth of data is stored as a PDS member.</p> <p>If logging is enabled you must specify a Log Dataset Name.</p> <ul style="list-style-type: none"> • Indicates whether you want commands and responses logged to the Monitor message log PDS. • Specify B to browse or V to view log output (when logging is enabled). • Enter your log PDS name in standard ISPF format. You must specify a fully qualified dataset name.
z/OS Migrator Load Library	Specify the installed z/OS Migrator Load library dataset name. The name will be verified and must be a fully qualified dataset.
z/OS Migrator Security Library	Specify the installed z/OS Migrator Security library dataset name. The name will be verified and must be a fully qualified dataset.
Unit and Volume	Specify the unit and volume serial number if the Log PDS or other intermediate datasets must be allocated.
Stack Console Output	Indicates whether you want commands and responses to be "pushed down" in the Operator Interface dialog.

Table 11 Configuration and control parameters (2 of 2)

Parameter	Description
Clear Variables	Setting this value to Y will cause the panel variables to be cleared when moving between panels. A value of N will retain any panel variables that were entered.
Simulate Mode	This mode affects the Configure and Query commands. When this mode is enabled by specifying Y , the query and configure commands are built, echoed back to the terminal, and <i>not</i> processed.
Debug Mode	Setting this value to Y will enable the debug trace facility in this interface.

If one of the critical datasets is changed to a new dataset (as detected during the SAVE process), a popup appears, informing you that you need to close and restart the ISPF monitor. This message appears only once it has been verified that the dataset(s) are cataloged and that no error is present in the user settings.

Option 1 - Manage Groups

Select the Manage Groups option from the z/OS Migrator Functions panel to display the Manage Group Members panel. The Manage Group Members panel allows you to create and manage group definitions, edit an existing definition, validate/verify a new group definition for adherence to syntax and system parameters, delete a group definition, and promote a group definition.

You must have UPDATE, ALTER, or CONTROL authority, or the security must be inactive (SAF Not Active, or No Security) for you to be able to perform the following functions: Promote, Verify, and Restore. If not, you will get the following message:

```
This user does not have the required authority for the
z/OS Migrator server database. You may proceed with
limited functionality. For full functionality, you must
change your security setting to be able to update this
dataset.
```

```

                                Manage Group Members                                Row 1 of 40
Command=>                                                                Scroll=> CSR
Primary Cmds: CR = Create EX = Exit   F = Find  FIL = Filter SW = Switch
Line   Cmds:  C = Copy   D = Delete E = Restore M = Modify      SMFID X04
                P = Promote R = Rename S = Display V = Verify    CPFIX @$
*****
-  ABC          P
  * GROUP NAME:ABC                MEMBER GENERATED BY  JDOE1
  * ON THURSDAY, 25 OCT 2007 AT 18:20:04 FROM:  LOGMIGR CONTROL COMMANDS
-  BIPPY        P
  * GROUP NAME:BIPPY              MEMBER GENERATED BY  JDOE1
  * ON MONDAY, 17 DEC 2007 AT 17:30:32 FROM:  MIRROR CONTROL COMMANDS
-  BUILD        P
  * GROUP NAME:BUILD              MEMBER GENERATED BY  JDOE1
  * ON THURSDAY, 25 OCT 2007 AT 18:20:04 FROM:  LOGMIGR CONTROL COMMANDS
-  LDBACKG1
  * MIGRATION DIRECTION:          FROM MOD 9'S TO MOD 3'S
  * aLL aPPlication dATA          (1520-1527)   TO (1500-151F)
-  GAMMA        P
  * GROUP NAME:GAMMA              MEMBER GENERATED BY  JDOE1
  * ON THURSDAY, 25 OCT 2007 AT 18:20:04 FROM:  MIGRATE CONTROL COMMANDS
-  LARGE        P
  * GROUP NAME:LARGE              MEMBER GENERATED BY  PTN
  * ON FRIDAY, 16 NOV 2007 AT 15:26:25 FROM:  ZOSM CONTROL COMMANDS

```

Figure 18 Manage Group Members panel

The Manage Group Members panel lists both the members of the configuration PDS and those groups that have been promoted. Once a group is promoted, both its original PDS member and the promoted group will appear in the list. Group definitions are listed by the member names you assigned when they were created—all names are prefixed with “LD”. *Promoted* groups are listed by group name and are identified by a “P” to the right of the name in the list. Each entry in the list for group members created via the ISPF interface provides information as to when the group was defined. [“Promoting a group definition” on page 165](#) provides more information.

Primary commands

When creating a new group definition, you need to be familiar with the Primary commands, usually noted at the top of each panel. The commands that you can specify change depending on which panel is currently being viewed.

The Primary commands that you will encounter while using the panels of Option 1 are described in the following table. Notice that some are the same as the line commands, described in [Table 13 on page 143](#).

Note: The Primary commands listed in [Table 12](#) apply to several or all options of the z/OS Migrator ISPF Monitor. Refer to this table as you use the other options.

Table 12 Define z/OS Migrator Group panel - Primary commands (1 of 2)

Command	Description
CRreate	Creates a new migration group. A member with the following format will be created: LDxxxxxxx - Group Definition member where xxxxxxxx is a 6-character name specified by the user. This type of member can be modified and promoted.
EXit	Exits the group creation process. (Or use PF3.)
Find	Finds the next occurrence of the specified string. The row containing the string is moved to the top of the display.
FILter	Limits the list to show only the promoted groups (indicated with a 'P' next to the group name) or the unpromoted groups.

Table 12 Define z/OS Migrator Group panel - Primary commands (2 of 2)

Command	Description
IMport	<p>Note: This command is available if option 'Build Data Set Exclude List' is set to 'Y' when entering the source options for a Logical Migration group.</p> <p>Allows a dataset (up to 44 characters) and member (up to 8 characters) to be specified for import. This member should contain a list of dataset name exclusion masks.</p>
SWitch	Allows a direct switch to Option 2 - Monitor Groups from Option 1.
MOre	Returns to the source parameter panel to define additional selections for the group.
EDit	Enters ISPF edit for the group member being created. Can be used to make minor corrections or to extend the group definition manually.
SAve	<p>Saves the group definition that has been created. You will be prompted to enter the 6-character unique member name suffix.</p> <p>A suffix of LD will be used to complete the name. If the name provided isn't unique, a member name with the following format will be substituted. <i>LDjjjsss</i> - where <i>LD</i> is the normal suffix, <i>jjj</i> is the current Julian day and <i>sss</i> is a sequential number starting at 001.</p>
VErify	<p>Submits the group definition to the server for syntax scan. This option will cause the service task to analyze the selected definition's configuration parameters and syntax to ensure that they match system requirements.</p> <p>You cannot verify a promoted group.</p>
PRomote	<p>Loads the defined group into server memory. This makes the group available to be activated.</p> <p>The promoted group will be identified by a 'P', signifying its promoted status.</p>

Line commands

Table 13 lists the line commands you can enter on the Manage Group Members panel.

Table 13 Manage Group Members panel - Line commands

Command	Description
Copy (C)	Place a C next to a group definition to copy the selected definition. You will be prompted to enter a 1- to 6-character member name. You cannot copy a promoted group. Note: The time stamp for the copied group will reflect the time stamp of the original group. Edit the time stamp manually to reflect the time the new group was created.
Delete (D)	Place a D next to a group definition to delete the group definition member for the configuration PDS. This option is only valid for promoted groups. "Option 2 – Monitor Promoted Groups" on page 168 describes how to delete a promoted group.
Restore (E)	Place an E next to a promoted group to <i>copy</i> the selected group back to the configuration PDS. You will be prompted to enter a unique 6-character member name suffix; the new name will be prefixed with "LD". This option is only valid for promoted groups. Note: <i>Restoring</i> a promoted group does not in any way affect the group; z/OS Migrator makes a copy of the group in the configuration PDS.
Modify (M)	Place an M next to a group definition to enter ISPF edit for the specified group definition member. This option is not valid on promoted groups.
Promote (P)	Place a P next to a group definition to load the group into server memory and store it in the z/OS Migrator database. All cooperating servers will automatically promote the group as well. The promoted group will be identified by 'P' to the right of the group name in the menu.
Rename (R)	Place an R next to a group definition to rename the selected definition. This option is only valid for promoted groups. To rename a promoted group, first <i>Restore</i> the group to its original group definition state and rename the restored definition.
Display (S)	Place an S next to a group definition to view the selected definition. This option is not valid on promoted groups.
Verify (V)	Place a V next to a group definition to validate its syntax and configuration parameters. This option will cause the service task to analyze the selected definition's configuration parameters and syntax to ensure that they match system requirements. You cannot verify promoted groups.

Creating a new group definition

To create a new group definition:

Type the **CREATE** command (or **CR**) on the **Manage Group Members** panel and press **Enter**.

The Build Type pop-up appears.

```

+----- Build Type -----+
|                               |
|   Select with an S a build type. |
|   _ Build Migrate Member      |
|   _ Build Mirror Member      |
|   _ Build Logical Migrate Member |
|                               |
|   F3=Cancel                   |
+-----+

```

Figure 19 Build Type pop-up

The pop-up gives you the option of creating one of the three kinds of groups.

Table 14 Kinds of build types

Build Member Type	Action
Migrate Member	Migrating a set of volumes
Mirror Member	Host-based mirroring for a set of volumes
Logical Migrate Member	Logical migration of a collection of datasets

Type an **S** on the line of the group type you wish to define.

- ◆ [“Defining a Logical Migration Group” on page 145](#), shows you how to create a Logical Migrate, or dataset, group definition.
- ◆ [“Defining a Volume Mirror Group” on page 160](#) shows you how to create a Mirror Member.
- ◆ [“Defining a Volume Migrate Group” on page 157](#) shows you how to create a Migrate Member. Migrate and Mirror groups are volume oriented.

Defining a Logical Migration Group

1. Type an **S** at **Build Logical Migrate Member**.

The **Define z/OS Migrator Group** panel displays.

```

Define z/OS Migrator Group
Command ===>                               Scroll ===> DATA

Group Name . . . . _____

Group Options . . . N Replace Existing Data Sets (Y/N)
                   Y Tolerate Allocation Failure (Y/N)

                   PF3: Exit      Enter: Process

```

Figure 20 Define z/OS Migrator Group panel

Note: If you select **Build Mirror Member** or **Build Migrate Member**, an appropriate panel for mirroring or migrating a volume will start you on a path toward defining a volume group.

2. On the **Define z/OS Migrator Group** panel, type a name for the group and specify **Y** or **N** for the group options.

Table 15 Group options

Parameter	Description
Group Name	Specify a 1- to 8-character name by which this group will be known once it is promoted.
Replace Existing Datasets	<p>This parameter causes z/OS Migrator to replace an existing target dataset. Valid input characters are Y and N:</p> <ul style="list-style-type: none"> • Y specifies that an existing target dataset may be replaced. The existing target dataset is erased, and the new target dataset is allocated. • N specifies that an existing target dataset may not be replaced. The migration will fail if the target dataset name already exists. <p>Note: This affects all datasets defined for the group. N (or no) is the default setting.</p>
Tolerate Allocation Failure	<p>Valid input characters are Y and N:</p> <ul style="list-style-type: none"> • Y specifies that execution continues on successfully allocated targets (overrides normal target allocation failure procedures). This is useful when the target dataset name uses wildcard characters. • N specifies that group activation fails upon detection of allocation errors.

3. Press **Enter**.

Define z/OS Migrator Group – Source Data Set Options displays.

```

EFMPDEF3                Define z/OS Migrator Group
Command ==>>>                Scroll ==>>> DATA

Group Name . . . . . CAPTURE
Source Data Set Options . . N Trace (Y/N)      N AllocSeq (D/S/N)
                          . Y Sphere (Y/N)      N Rename UnConditional (Y/N)
                          . y Build Data Set Exclude list (Y/N)
Source Data Set Name/Mask . xyztest.**_____
Source Volume List Name . . svlcap__ Saved SVLCAP
Storage Class . . . . . _____
Source Volume(s) . . _____
                          _____
                          _____
                          _____

                          PF3: Exit      Enter: Process

```

Figure 21 Define z/OS Migrator Group - Source Dataset Options panel

4. On **Define z/OS Migrator Group – Source Data Set Options**, enter the source dataset mask and optional source volumes.

Instead of individual source volumes, you may enter the name of a source volume list to be created in a subsequent panel. Select other options as described in [Table 16](#).

Table 16 Source dataset options (1 of 4)

Parameter	Description
Trace	<p>If selected, z/OS Migrator will produce diagnostic output during the migration cycle. No output is the default setting.</p> <p>Note: Each dataset defined for the group may use a different setting.</p>
AllocSeq	<p>Dataset Allocation Sequence. Valid input characters for this option are:</p> <ul style="list-style-type: none"> D (DATASET) specifies to process datasets in ascending name sequence. S (SIZE) specifies to process datasets in descending size sequence. The largest datasets are processed first and the smallest are processed last. N (NONE) specifies to process datasets in the order that they are selected for processing. This may appear random. <p>Note: N is the default setting.</p>

Table 16 Source dataset options (2 of 4)

Parameter	Description
Sphere	<p>This parameter specifies whether, for any VSAM cluster copied, all associated AIX clusters and paths are to be copied. Individual names of sphere components do not need to be specified. Only the base cluster is eligible to be migrated if SPHERE(N) is specified, individual AIX clusters are ineligible. Each SET statement comprises a complete selection specification, and each may have a different SPHERE value.</p> <p>Valid input characters are Y and N:</p> <ul style="list-style-type: none"> • Y specifies that all associated AIX clusters and PATHs are to be copied. • N specifies that only the selected base cluster will be copied. <hr/> <p>Note: Y is the default setting; however, when migrating non-VSAM datasets, z/OS Migrator ignores the SPHERE parameter.</p> <hr/> <p>If SPHERE(Y) is specified, then the selection of the base cluster is determined by the source selection mask, however, multiple related components (KSDSIESDS + AIX + PATH) will be migrated in a single action, and they may not have a common naming structure. Unless there is a single AIX whose name is base-cluster-name.AIX, then the RENAMEUnconditional parameter must be used to specify the naming rule for the associated target PATH and AIX components.</p>
Rename UnConditional	<p>Builds a RENAMEUnconditional statement.</p> <p>A method to provide alternate naming conventions to components being copied as part of a sphere.</p> <p>Valid input characters are Y and N:</p> <ul style="list-style-type: none"> • Y specifies to allow the creation of RENAMEUnconditional statements. • N specifies not to create RENAMEUnconditional statements. <hr/> <p>Note: SPHERE=Y must also be specified.</p> <hr/> <p>"RENAME Unconditional" on page 232 provides more information.</p>
Build Dataset Exclude list	<p>Builds the dataset exclusion list.</p> <p>A method to eliminate datasets from the selection list.</p> <p>Valid input characters are Y and N:</p> <ul style="list-style-type: none"> • Y specifies to allow the creation of the Exclude list. • N specifies not to create the Exclude list. <hr/> <p>Note: You may specify multiple dataset names or dataset masks to exclude by <i>modifying</i> the group definition. "Example 1: Modifying a group definition to exclude multiple datasets" on page 166 provides more information.</p>

Table 16 Source dataset options (3 of 4)

Parameter	Description
Source Dataset Name/Mask	<p>Enter the source dataset name for this required parameter. To specify multiple datasets as source, use dataset name masking. “Selecting datasets” on page 83 provides detailed information on dataset masking r</p> <p>Both the source and target datasets must resolve to compatible devices. These devices may be within the same storage array, or they may be on different storage arrays. All extents for the source dataset(s) must reside on devices that are currently online and accessible.</p> <p>Note: The source dataset name must be different from the target dataset name. Neither the source mask nor any source dataset name should be a substring of the target mask. The source and target masks should have the same number of specific qualifiers. The specified qualifiers of the target mask replace the corresponding qualifiers of the source dataset name one for one. Therefore, it is not possible to insert a new qualifier.ules.</p> <p>“Restrictions on storage subsystems and dataset types” on page 93 provides a detailed listing of both supported and unsupported dataset types.</p>
Source Volume List Name	<p>Specify a 1- to 8-character name for the Source Volume List.</p> <p>If a valid name is entered, an additional panel will be displayed to allow entry of source volumes or volume masks.</p> <p>May not exceed 127 volumes or volume masks, with each entry delimited by a comma.</p> <p>To the right of the Source Volume List name are the names of the last four Source Volume Lists defined and saved:</p> <ul style="list-style-type: none"> • If you reuse one of the names, that definition's list will be presented on the panel that allows the volume to be entered. Any changes to the definition will be saved when the group definition is saved. • If you specify a new name and the group definition is saved, it will be saved in the first position and the others will move to the right with the original fourth entry deleted. <p>If a Source Volume List is specified, then the list of candidate datasets is built by scanning the VTOCs of the specified volumes. The catalog is then used to identify any other volumes containing the discovered datasets.</p> <p>Note: Source Volume List Name and Source Volumes are mutually exclusive. Either may be used with or without Storage Class as optional filters.</p>
Storage Class	<p>Define the source data volume storage class.</p> <p>This parameter acts as a filter that is applied after the catalog search. Only datasets returned by the catalog search and having a matching Storage Class will be selected for migration with the current group.</p> <p>Enter the storage class for the source dataset(s) represented by the dataset name mask. The Source Storage Class is a filter that further qualifies the datasets selected at <i>Activation</i>.</p> <p>Note: Source Volume List Name and Source Volumes are mutually exclusive. Either may be used with or without Storage Class as optional filters.</p>

Table 16 Source dataset options (4 of 4)

Parameter	Description
Source Volume(s)	<p>Specify one or more source data volume serial numbers. When specified, these volumes serve essentially as an unnamed Source Volume List, usable for just this dataset mask. Only datasets returned by the catalog search and having a matching volume serial number will be selected for migration with the current group.</p> <p>Be careful not to confuse this parameter with “Source Volume List Name.”</p> <p>Note: Source Volume List Name and Source Volumes are mutually exclusive. Either may be used with or without Storage Class as optional filters.</p>
Dataset Exclude Mask	<p>Define the dataset(s) to be excluded. This parameter can use dataset name masking.</p> <p>If the SOURCE parameter causes several datasets to be selected, the EXCLUDE parameter can be used to eliminate some of them from the selection list.</p> <p>A comma delimited list of up to 127 dataset names or masks can be specified.</p> <p>Use the IMPORT command to include an installation-defined set of datasets/masks. You will be allowed to append additional datasets to the end of the list after the import. If you append, make sure a comma (,) is the last character of the imported data.</p> <p>Examples:</p> <ul style="list-style-type: none"> • SYS1.PARMLIB • SYS1.PARMLIB,SYS2.PARMLIB,SYS3.PAR* • SYS1.PAR*,SYS1.PARMLIB.**
Rename Unconditional Prefix	<p>RENAMEU is needed when the naming conventions of components of a complex VSAM cluster (commonly the AIX components of a sphere) do not match the naming convention of the base cluster.</p> <p>Specifies the prefix (<i>ptx</i>) used to replace the first-level qualifier of the dataset name. It is optional, but if specified, must be the first parameter in the list of sub-fields. The prefix is used only if the (<i>oldnamemask,newnamemask</i>) parameters are not specified or the oldnamemask filters do not match the dataset name.</p>
Rename Unconditional Mask Pairs	<p>(<i>oldnamemask, newnamemask</i>) pairs.</p> <p>The <i>oldnamemask</i> is used to match existing PATH or AIX names, and then the corresponding <i>newnamemask</i> is used to transform the name into a new name.</p> <ul style="list-style-type: none"> • <i>oldnamemask</i> specifies a mask to be used as a filtering criterion to check if it matches the dataset name. • <i>newnamemask</i> specifies a mask used to derive the new dataset name when the existing dataset name matches the corresponding oldnamemask filtering criterion.

When using a Source Volume or Source Volume List, the Source DSN mask must be general enough to match a dataset's FMT1 DSCB entry (or entries) in the VTOC. For example:

Incorrect:

```
GROUP(LDSRCV) MODE(LMIGR)
SRCVLIST LIST5200(SMS0A0,SMS0A1,SMS0A2,SMS0A3,SMS0A7)
SET TRACE(N)
SOURCE(DSN(SRCE.LD5200.VSAM.CL1 ) - <-- The cluster name is a catalog
                                         entity only.
        SRCVLIST(LIST5200) ) -
TARGET(VOL(SMS0B0 ) -
        DSN(TGT.LD5200.VSAM.CL1 ) )
```

The error condition is reported (in the server log and group error log) as follows:

```
EFM2024I LDSRCV is now NULL (80->00) at ES_ACTIVATE_GROUP+08D4
EFM0555I Activation failed for group LD5200. Return code: 1029.
```

The Help message for EFM0555I is:

```
EFM0555I Activation failed for group <group>. Return code:
<return-code>.
Explanation: An activation command was unsuccessful. The return code
in the message can be one of the following:

1029 No datasets selected for migration.
```

A correct alternative definition would be:

Correct:

```
GROUP(LDSRCV) MODE(LMIGR)
SRCVLIST LIST5200(SMS0A0,SMS0A1,SMS0A2,SMS0A3,SMS0A7)
SET TRACE(N) -
SOURCE(DSN(SRCE.LD5200.VSAM.CL1.** ) -
        SRCVLIST(LIST5200) ) -
TARGET(VOL(SMS0B0 ) -
        DSN(TGT.LD5200.VSAM.CL1.** ) ) \
```

The corrected Source DSN mask causes the actual VSAM components to be found on the specified volumes. The catalog locate performed during Activation returns the cluster name and all of the truenamerecords for the VSAM cluster. The truenamerecords have matching FMT1 DSCB records in the VTOC of the volume(s) where the VSAM components are located.

5. Once you have specified the desired values, press **Enter**. Depending on the options selected, a panel like one of the following displays.

```

EFMPDEF5                                Define z/OS Migrator Group
Command ===>                               Scroll ===> DATA

Primary Commands :

Group Name . . . . . CAPTURE   Source Volume List Name . . . . . SVLCAP
Source Volume List Mask sym719 sym71a sym71b sym71c
_____
_____
_____
_____
_____
_____
_____
_____

```

Figure 22 Define Source Volume List panel

```

EFMPDEF4                                Define z/OS Migrator Group
Command ===>                               Scroll ===> DATA

Primary Commands. . . IM = Import
Group Name . . . . . CAPTURE
Data Set Exclude Mask XYZTEST.DONT.** ,XYZTEST.UDD.** ,XYZTEST.VSAM2.DB.** ,XYZ
TEST.DOPEY.DB.** ,XYZTEST.AAA.** ,XYZTEST.XYZ.** ,XYZTEST.A000*.** ,XYZTEST.LMIG
RX08.** _____
_____
_____
_____
_____
_____
_____
_____

```

Figure 23 Specify Exclude Masks panel

```
Define z/OS Migrator Group
Command ===>                                Scroll ===> DATA

Group Name . . . . . BLDG
Rename Unconditional Prefix. . . _____
Rename Unconditional Mask Pairs. _____
_____
_____
_____
_____
_____
_____
_____
```

Figure 24 Specify Rename Unconditional Masks panel

6. Press **Enter** and proceed to the Target Specifications panel.

```
EFMPDEF7                                Define z/OS Migrator Group
Command ===>                                Scroll ===> DATA

Group Name . . . . . CAPTURE
Target Data Set Name/Mask . . XYZTEST2.** _____
Target Volume Storage Class . _____
Target Volumes(s) . . . . .
SYM719 SYM71A SYM71B SYM71C _____
_____
_____
_____
_____
_____

PF3: Exit      Enter: Process
```

Figure 25 Target Specifications panel

7. Specify the Target Dataset Mask, Target Volume Storage Class, and Target Volume(s) parameters, as described in [Table 17](#).

Table 17 Target parameters (1 of 2)

Parameter	Description
Target Dataset Name/Mask	<p>Specify a complete target dataset name, or use a mask. The mask is applied against each source dataset to create a target dataset name. A corresponding Target dataset is allocated and cataloged for each selected source dataset. Target dataset names must be valid dataset names and must be authorized by site security rules.</p> <hr/> <p>Note: You cannot specify the same dataset name on both the Source DSN and Target DSN parameters. Neither the source mask nor any source dataset name should be a substring of the target mask. The source and target masks should have the same number of specific qualifiers. The specified qualifiers of the target mask replace the corresponding qualifiers of the source dataset name one for one. Therefore, it is not possible to insert a new qualifier.</p> <hr/> <p>It is important that the target dataset name (mask) is easily identifiable and unique within the storage environment. Once the migration is complete, the target and source datasets will have exchanged names. Thus, once success has been verified, the datasets remaining with the target names can be deleted to make the source volumes ready for reuse or removal.</p> <p>Ensuring that target dataset names are unique and easily identifiable is important to avoid confusion with production datasets. “Selecting datasets” on page 83, and “Small versus large migration groups” on page 98 provide more information about target dataset masks.</p> <p>DB2 linear datasets have a specific naming convention. The second level qualifier denotes whether it is the cluster or the data portion of the dataset. z/OS Migrator can only be used against the cluster; an individual component of the cluster cannot be specified. All cluster components will automatically be selected.</p> <p>The following two examples use z/OS Migrator to duplicate DB2 linear datasets. Both of the options require the use of wildcarding that is available in z/OS Migrator.</p> <p>For the following source dataset: SOURCE(EMCDB2.DSNDBC.STORE.TSSTORA.I00001.A001)</p> <ul style="list-style-type: none"> • Option 1: Use a wildcard to denote the single unique level in the target dataset name. If the single level wildcard specification is used, it must be the second level qualifier: TARGET (BCVDB2 . * . STORE . TSSTORA . I 0 0 0 1 . A 0 0 1) • Option 2: Use a wildcard to denote the entire target dataset name after the specification of a unique high-level qualifier (<i>hlq</i>): TARGET (BCVDB2 . * *)

Table 17 Target parameters (2 of 2)

Parameter	Description
Target Volume Storage Class	<p>Enter the storage class for the target volume(s). This parameter substitutes this SMS storage class for the existing Storage Class associated with the source dataset in the source/target dataset pair.</p> <hr/> <p>Note: Local SMS ACS routines may place the target dataset in a storage class other than that specified by this parameter. As with all SMS datasets, specifying storage class is only a suggestion to SMS, and may or may not be accepted by SMS. If an existing target dataset is reused, its relevant storage class information will not be changed.</p> <hr/>
Target Volume(s)	<p>Enter the volume serial number where the target dataset(s) resides. You may specify up to 59 volume names. "Example 2: Modifying a group definition to add more target volumes" on page 166 shows how to specify more target volumes.</p> <hr/> <p>Note: Esoteric name and Unit Name are not valid values for Target Volume. This parameter is an absolute volume name, no masking is supported.</p> <hr/>

8. Once you have entered all the relevant group/pair parameters, press **Enter**.

The generated group definition displays.

```

EFMPDEF8                               Define z/OS Migrator Group                               Row 1 of 31
Command ==> save                               Scroll ==> DATA

Primary Commands  : EXit  MOrE  EDit  SAve  VErify  PRomote

GROUP(CAPTURE) -
MODE(LMIGR()) -
  TOLERATE_ALLOCATION_FAILURE(YES) MAXRC(8) -
  REPLACE(YES)
SOURCE_VOLUME_LIST SVLCAP ( -
  SYM719,SYM71A,SYM71B,SYM71C -
)
SET -
ALLOCSAQ(NONE) -
TRACE(NO) -
SPHERE(YES) -
SOURCE ( -
  DSN (XYZTEST.***) -
  SOURCE_VOLUME_LIST (SVLCAP) -
  EXCLUDE ( -
    XYZTEST.DONT.***, -
    XYZTEST.UDD.***, -
    XYZTEST.VSAM2.DB.***, -
    XYZTEST.AAA.***, -
    XYZTEST.XYZ.***, -
    XYZTEST.A000**.*, -
    XYZTEST.LMIGRX08.** -
  ) -
) -
TARGET ( -
  DSN (XYZTEST2.***) -
  VOLUME ( -
    SYM830, SYM831 -
  ) -
)

```

Figure 26 Generated group definition

9. At this point, you may reject the defined group and start again, or edit, save, verify, or promote the newly created group.

EXit	Abandon the group creation (confirmation will be requested).
MOre	Takes you back to the Source Data Set Options panel to select more datasets.
EDit	Invoke ISPF Edit to manually edit the group definition.
SAve	Store the group definition in the configuration PDS.
VERify	Submit the group definition to the z/OS Migrator server for syntax scanning.
PRomote	Load the group definition into the server memory and the z/OS Migrator database.

Table 12, “Define z/OS Migrator Group panel - Primary commands,” on page 141 provides more information on these commands.

Defining a Volume Migrate Group

Type the Create command (or CR) on the Manage Group Members panel and press **Enter**.

The Build Type pop-up appears.

```

+----- Build Type -----+
|                               |
|   Select with an S a build type.   |
|   _ Build Migrate Member           |
|   _ Build Mirror Member           |
|   _ Build Logical Migrate Member   |
|                               |
|   F3=Cancel                       |
+-----+

```

Figure 27 Build Type panel

1. Type an **S** at Build Migrate Member and press **Enter**.

The Specify Migrate Group Parameters panel displays.

```

=====Specify the Migrate Group parameters =====
Command ==>                               Scroll ==> PAGE

Command=>                               Scroll=>
Group Name . . . . . _____
Migration options . . (S) _ FastCopy _ WarmStart _ DeactOnSus _ NoAutoVary
Migration Completion. (S) _ SplitOn _ SplitOff _ Swap _ ConstantCopy
Swap options. . . . (Y/N) _ BYPCC _ BYPSNAP _ BYPSYSC _ Consistent
      CFW _____ VOLPfx .. _____
Source CUU . . . . _____ and Count _____ or Volser . . . . _____
Target CUU . . . . _____ New Volser _____ or Volser Prefix . . _____
Verify Options (S) _ None or _ Empty Volser _____ '%' matches any char

***** Primary Commands: EDit Save VERify PROMote *****
***** Bottom of data *****

```

Figure 28 Specify Migrate Group parameters panel

2. Enter the desired parameters along with the source and target device specifications.

[Table 18, “Migrate Group parameters,”](#) on page 158 shows the parameters you need to supply to create a Migrate group.

Table 18 Migrate Group parameters

Parameter	Description
Group name	Eight-character name this group will be known by once it has been promoted.
Fastcopy	Migration should be accomplished by copying only the allocated tracks.
WarmStart	For future use – currently ignored.
SplitOn	At the conclusion of migration, vary the target online.
SplitOff	At the conclusion of migration, leave the target offline.
Swap	At the conclusion of migration, initiate a swap between source and target.
ConstantCopy	After initial migration has been completed, continue to monitor changed tracks and periodically copy them from source to target.
BYPCC	This and the next four are really options to SDAS for controlling the swap process. BYPCC means do not check for concurrent copy operations before swapping.
BYPSNAP	Bypass the test for SNAP operations.
BYPSYSC	Do not invalidate a swap because of the system count.
CFW option	Swap behavior relative to CacheFastWrite. (CacheFastWrite is described in the Glossary.)
VOLPfx	A prefix used by SDAS to alter the source volume serial after swap.
Source CUU	Device address of the first or only source device to be migrated.
And Count	Count of sequential addresses to be generated by this definition.
Or Volser	Alternatively, specify the volser of a single source volume.
Target CUU	Device address of the first or only target device.
New Volser	A mask that serves as the basis for generating the new volser of the target.
Or Volser Prefix	Alternatively, a prefix to be combined with the device number of the target to form a new volser.
None	Do not validate any of the mirror content before using.
Empty	Verify that the mirror volume is empty before using.
Volser	Verify the volume serial number of the mirrors before using.

3. After entering the first set of values, the generated group definition statements will be shown below the parameter entry fields. Continue entering additional source and target device specifications to extend the group.
4. When satisfied that the group definition is complete, proceed with one of the primary command actions shown on the line

***** Primary Commands: EDit SAve VERify PRomote *****

EDit	Invoke ISPF Edit to manually edit the group definition.
SAve	Store the group definition in the configuration PDS.
VERify	Submit the group definition to the z/OS Migrator server for syntax scanning.
PRomote	Load the group definition into the server memory and the z/OS Migrator database.

Table 12, “Define z/OS Migrator Group panel - Primary commands,” on page 141 provides more information on these commands.

Defining a Volume Mirror Group

Type the CREATE command (or CR) on the Manage Group Members panel and press Enter.

The Build Type pop-up appears.

```

+----- Build Type -----+
|
|   Select with an S a build type.
|   _ Build Migrate Member
|   _ Build Mirror Member
|   _ Build Logical Migrate Member
|
|   F3=Cancel
|
+-----+

```

Figure 29 Build Type panel

1. Type an S at Build Mirror Member.

The Specify Mirror Group Parameters panel is displayed.

```

=====Specify the z/OS Migrator Group parameters =====
Command ==>                                     Scroll ==> PAGE

Group Name . . . . . _____
Mirror options . . . (S) _ FastCopy _ WarmStart _ DeactOnSus _ NoAutoVary

Primary CUU . . . ___ and Count . . ___ or Volser . . . . . ___
Mirror CUU . . . ___ New Volser . . ___ or Volser Prefix . . ___
Verify Options . . _ None or ___ Emplty (S) Volser . . ___ '%' matches any char

*****Primary CommandsEDit Save VERify PROMote *****
*****Bottom of data *****

```

Figure 30 Specify Migrate Group parameters panel

2. Enter the desired parameters along with the source and target device specifications.

Table 19, “Mirror group parameters,” shows the parameters you need to supply to create a mirror group:

Table 19 Mirror group parameters

Parameter	Description
Group name	Eight-character name this group will be known by once promoted.
Fastcopy	Initial synchronization should be accomplished by copying only the allocated tracks.
WarmStart	For future use – currently ignored.
DeactOnSus	If a suspend occurs for any reason, automatically suspend mirroring for the affected units.
NoAutoVary	Do not automatically vary the mirror to match the online state of the primary. You may still manually vary the mirror.
Primary CUU	Device address of the first, or only, primary device to be mirrored.
And Count	Count of sequential addresses to be generated by this definition.
Or Volser	Alternatively, specify the volser of a single primary.
Mirror CUU	Device address of the first or only mirror device.
New Volser	A mask that serves as the basis for generating the “apparent” volser of the mirrors.
Or Volser Prefix	Alternatively, a prefix to be combined with the device number of the mirror to form an “apparent volser.”
None	Do not validate any of the content of the mirror before using.
Empty	Verify that the mirror volume is empty before using.
Volser	A mask used to verify the volume serial number of the mirrors before using.

3. After entering the first set of values, the generated group definition statements will be shown below the parameter entry fields. Continue entering additional source and target device specifications to extend the group.

4. When satisfied that the group definition is complete, proceed with one of the primary command actions shown on the line

***** Primary Commands: EDit SAve VERify PRomote *****

EDit	Invoke ISPF Edit to manually edit the group definition.
SAve	Store the group definition in the configuration PDS.
VERify	Submit the group definition to the z/OS Migrator server for syntax scanning.
PRomote	Load the group definition into the server memory and the z/OS Migrator database.

Table 12, “Define z/OS Migrator Group panel - Primary commands,” on page 141 provides more information on these commands.

Saving a group definition

When the group definition is complete, type **SAVE** (or **SA**) in the command field, and press **Enter**.

You are prompted to enter the 6-character configuration member name suffix:

```

EFMPDEF8                      Define z/OS Migrator Group                      Row 1 of 31
Command ==>>                      Scroll ==>> DATA

Primary Commands  : EXit  MOre  EDit  SAve  VERify  PRomote

-----
GROUP (CAPTURE) -
MODE(LMIGR()) -
TOLERATE_ALLOCATION_FAILURE(YES) MAXRC(8) -
REPLACE(YES) +----- Create Member Name -----+
SOURCE_VOLUME_L | EIBPSAV 6 characters to be used to create |
SYM719,SYM71A  | a member name to contain the data to be |
)               | saved ==>> captur |
SET -          |
ALLOCSEQ(NONE) | NOTE: Member Name is prefixed with LD |
TRACE(NO) -    | and then your 6 characters. |
SPHERE(YES) - +-----+
SOURCE ( -
DSN (XYZTEST.***) -
SOURCE_VOLUME_LIST (SVLCAP) -

```

Figure 31 Define z/OS Migrator Group panel

Specify a member name from 1 to 6 characters long and press **Enter**. The member name will be prefixed with “LD” and stored in the configuration PDS. Specify a member name that you can easily associate with the group that you have defined.

The new group definition is listed (in alphanumeric order) in the **Manage Group Members** panel.

Verifying a group definition

You must have UPDATE, ALTER, or CONTROL authority, or the security must be inactive (SAF Not Active, or No Security) for you to be able to perform the following functions: Promote, Verify, and Restore. If not, you will get the following message:

```
This user does not have the required authority for the
z/OS Migrator server database. You may proceed with
limited functionality. For full functionality, you must
change your security setting to be able to update this
dataset.
```

z/OS Migrator provides the VERIFY line command allowing you to validate a group definition’s configuration parameters and syntax to ensure that they match system requirements. The VERIFY line command is only available to group definitions that have not yet been promoted. You cannot verify a promoted group.

Note: Group definitions that have not been promoted are identified by “LD”, prefixed to the member name.

To verify an unpromoted group definition:

- ◆ On the **Manage Group Members** panel, place a **V** next to the group definition that you want to verify and press **Enter**.

The group definition is then parsed and checked for syntax and configuration errors as shown in [Figure 32](#). If an inconsistency is found, a syntax error message is displayed and the error is marked .

```

VIEW          SYS07155.T201707.RA000.JDOE1.R0138644
Command ==>
*****
***** Top of Data *****
000001
000002  z/OS MIGRATOR VERIFY MEMBER EXECUTED BY JDOE1 ON 06/04/07 AT 20:17:07
000003
000004 EFMM021I Processing config file parameters
000005 EFMP001I * GROUP NAME:CAPTURE          MEMBER GENERATED BY JDOE1
000006 EFMP001I * ON Sunday, 10 Jun 2007 AT 17:40:45 FROM: MIRROR CONTROL COMMANDS
000007 EFMP001I GROUP (CAPTURE) -
000008 EFMP001I  MODE(LMIGR()) -
000009 EFMP001I  TOLERATE_ALLOCATION_FAILURE(YES) MAXRC(8) -
000010 EFMP001I  REPLACE(YES)
000013 EFMP001I  SOURCE_VOLUME_LIST SVLCAP ( -
000014 EFMP001I  SYM719,SYM71A,SYM71B,SYM71C -
000015 EFMP001I  )
000016 EFMP001I SET -
000017 EFMP001I  ALLOCSEQ(NONE) -
000018 EFMP001I  TRACE(NO) -
000019 EFMP001I  SPHERE(YES) -
000020 EFMP001I  SOURCE ( -
000021 EFMP001I  DSN (XYZTEST.***) -
000022 EFMP001I  SOURCE_VOLUME_LIST (SVLCAP) -
000023 EFMP001I  EXCLUDE ( -
000024 EFMP001I  XYZTEST.DONT.***, -
000025 EFMP001I  XYZTEST.UDD.***, -
000026 EFMP001I  XYZTEST.VSAM2.DB.***, -
000028 EFMP001I  XYZTEST.AAA.***, -
000029 EFMP001I  XYZTEST.XYZ.***, -
000030 EFMP001I  XYZTEST.A000*.**, -
000031 EFMP001I  XYZTEST.LMIGRX08.*** -
000032 EFMP001I  ) -
000033 EFMP001I  ) -
000034 EFMP001I  TARGET ( -
000035 EFMP001I  DSN (XYZTEST2.***) -
000036 EFMP001I  VOLUME ( -
000037 EFMP001I  SYM830, SYM831 -
000038 EFMP001I  ) -
000039 EFMP001I  )
000040 *** END OF DISPLAY ***

```

Figure 32 Verify group definition

Promoting a group definition

You must have UPDATE, ALTER, or CONTROL authority, or the security must be inactive (SAF Not Active, or No Security) for you to be able to perform the following functions: Promote, Verify, and Restore. If not, you will get the following message:

```
This user does not have the required authority for the
z/OS Migrator server database. You may proceed with
limited functionality. For full functionality, you must
change your security setting to be able to update this
dataset.
```

Once you have created a group definition, you must promote the group to initiate the migration process. Promoting a group definition transfers the selected group to the z/OS Migrator database and communicates the group's configuration parameters to the service task. The promoted group is then identified by "P" to the right of its name in the Manage Group Members list, signifying its promoted status.

To promote a group definition:

- ◆ On the **Manage Group Members** panel, place a **P** next to the group definition that you want to promote and press **Enter**. You will be asked to confirm the action.

Before it is promoted, the group definition is initially parsed and checked for syntax and configuration errors. You will be notified once the group has been successfully promoted.

Modifying a group definition

z/OS Migrator provides the Modify line command, allowing you to edit an existing group definition. The Modify line command is only available to group definitions that have not yet been promoted. You cannot edit a promoted group.

Note: Group definitions that have not been promoted are identified by "LD" prefixed to the member name.

To modify a group definition:

1. On the **Manage Group Members** panel, place an **M** next to the group definition that you want to modify and press **Enter**. The group definition is then displayed in a standard ISPF *Edit* panel.

Note: [Chapter 8, "Defining z/OS Migrator Groups via Text Editor,"](#) describes how to use the text editor to create or edit group definitions.

2. Modify the group definition as desired.
3. Once you have finished editing the group definition, type the **SAVE** command, or press **PF3**.

Example 1: Modifying a group definition to exclude multiple datasets

The following group was created with only one Exclude mask:

```
GROUP (DOC6) MODE (LMIGR ())
SET TRACE (NO) -
SOURCE (DSN (**)) -
VOLUME (TD5441) -
EXCLUDE (SYS1.DATA.TD5441) -
TARGET (VOLUME (TD5270) -
DSN (KZP11.**))
```

Using ISPF Edit, invoked from the Manage Group list, simply add a second Exclude mask on the same line or a continuation line. The list is comma-separated and may contain up to 127 masks. Continuations are indicated with a '-' after the last value on a line, for example:

```
GROUP (DOC6) MODE (LMIGR ())
SET TRACE (NO) -
SOURCE (DSN (**)) -
VOLUME (TD5441) -
EXCLUDE (SYS1.DATA.TD5441, -
KZP10.EFMM.EXTMAP.LMTD5441.DB) -
TARGET (VOLUME (TD5270) -
DSN (KZP11.**))
```

Example 2: Modifying a group definition to add more target volumes

The following is an example of adding five additional target volumes to a group previously defined with five target volumes.

The original definition:

```
GROUP (TEST105) MODE (LMIGR ())
SET TRACE (NO) -
```

```
SOURCE (DSN (TEST. **)) -
TARGET (VOLUME (RWY352, RWY532, RWY253, RWY115, RWY235) -
        DSN (MIGR. **))
```

The extended definition. Note that up to 59 volumes may be specified in a comma-separated list. Continuation is indicated with a '-' at the end of any line.

```
GROUP (TEST105) MODE (LMIGR ())
SET TRACE (NO) -
SOURCE (DSN (TEST. **)) -
TARGET (VOLUME (RWY352, RWY532, RWY253, RWY115, RWY235 -
                SWT352, SWT532, SWT253, SWT215, SWT325) -
        DSN (MIGR. **))
```

Deleting a group definition

z/OS Migrator provides the Delete line command, which allows you to delete an existing group definition. The Delete line command in the Manage Group Members panel is only available to group definitions that have not yet been promoted. To delete a promoted group, you must use [“Option 2 – Monitor Promoted Groups” on page 168](#).

Note: Group definitions that have not been promoted are identified by “LD” prefixed to the member name.

To delete a group definition:

- ◆ On the **Manage Promoted Groups** panel, place a **D** next to the group definition that you want to delete and press **Enter**.
You will be asked to confirm the deletion.

Option 2 – Monitor Promoted Groups

Select the Monitor Promoted Groups option from the z/OS Migrator Functions panel to display the Monitor Promoted Groups panel. This panel displays the current status of promoted groups and their corresponding datasets or volumes.

Table 20 describes the kind of group as defined by mode.

Table 20 Modes of groups

MODE	Kind of Group
FASTMIR	Volume Mirror group without other options specified
MIGRATE	Volume Migrate group without other options
MIG/FC	Volume Migrate group with Fastcopy specified
LOGMIGR	Logical Migrate group, that is a dataset migration

You can use the list of groups to activate migration for one or more groups; simply use the **A** line command on the desired groups to activate them.

```

EFMPM1                               z/OS Migrator Monitor                               Row 1 of 24
Command===>                           17:33:40 06/03/07

Line Commands: A = Activate C = Config D = Deact  E = Errors J = Jobs
                L = Volumes  O = VolSum P = Suspend R = Resume S = ShowExt
                T = Reset    U = Report V = Divert X = Delete Z = Complete

  Group      Current      Dev      Num      Num
  Name      Status      Mode      Action  Owner Sets  I Dev  Comp
*****
_ ABC      NOT ACTIVE      LOGMIGR
_ XYZ      NOT ACTIVE      LOGMIGR  CONST
_ BUILD    NOT ACTIVE      LOGMIGR
_ CAPTURE  NOT ACTIVE      LOGMIGR
_ DUPE     NOT ACTIVE      LOGMIGR
_ LARGE    NOT ACTIVE      LOGMIGR
***** Bottom of data *****

```

Figure 33 Monitor Promoted Groups panel

The Monitor Promoted Groups panel provides several line commands allowing you to perform a variety of actions on both groups and their corresponding datasets, including the following:

- ◆ Deleting a group
- ◆ Diverting a group
- ◆ Activating and deactivating a group
- ◆ Suspending and resuming migration for a group
- ◆ Displaying detailed dataset configuration information

[Table 21, “Monitor Promoted Groups - line commands for groups,” on page 169](#) provides a complete list of actions available.

Line commands for groups

[Table 21](#) lists the group line commands that you can enter on the Monitor Promoted Groups panel. Note that the J, L, O, S, V, and Z line commands are specific to Logical Migration groups (that is dataset groups).

Table 21 Monitor Promoted Groups - line commands for groups (1 of 2)

Command	Description
A	Activate — Place an A next to a group to activate the selected group and begin the migration process. The Monitor Promoted Groups panel displays the progress of the migration process. “Activating a group” on page 175 provides more information.
C	Config — Place a C next to a group to display the dataset list for the group. Additional actions may be selected from the Dataset List panel.
D	Deactivate — Place a D next to a group to deactivate the selected group. To be deactivated, a group must first be suspended. “Suspending a group” on page 176 provides more information. Note: You cannot deactivate a group during certain phases of the migration.
E	Errors — Place an E next to a group to display the group error message log.
J	Jobs — Place a J next to a group in diversion to display any jobs blocking completion.
L	Volumes — Place an L next to a group to display information about the first extent on each volume for each dataset.
O	Volume Summary — Place an O next to a group to display the datasets, source and target, that are on the volume.

Table 21 Monitor Promoted Groups - line commands for groups (2 of 2)

Command	Description
P	Suspend — Place a P next to a group to pause or suspend migration for that group. “Diverting a Logical Migrator group” on page 176 provides more information. Note: If you suspend a group when it is in Mirror or Copy state, you break the mirroring. Updates are not copied to the target dataset until you resume mirroring.
R	Resume — Place an R next to a group to restart or resume the migration process for the group, after migration was paused or suspended.
S	Show Extents — Place an S next to a group to show all the extent information for the datasets in the group.
T	Reset — For volume migration only. Place a T to specify a new option for constant copy groups.
U	Report — Place a U next to a group to produce a printable report that formats all of the group’s parameters and status values.
V	Divert — Place a V next to a group to divert data from the source dataset(s) to the target dataset(s). This will cause the group to enter divert state until the target dataset(s) is recataloged and renamed. A group will remain in divert state if there are jobs allocated to the source dataset(s). Use the J and/or N option to display a list of jobs allocated to the source dataset(s). “Diverting a Logical Migrator group” on page 176 provides more information.
X	Delete — Place an X next to a group to delete the selected group. Note: You cannot delete a group for which a migration is currently in progress.
Z	Complete — If the NOAUTOTRM option has been selected (Config member or SET command), placing a Z next to a group manually triggers the completion processing.

Mirror/Migrate volume configuration display

Using the C line command for a volume oriented group in the Monitor Promoted Groups panel results in a display of the devices that comprise the group.

```

----- Group: PAL Configuration Information ----- Row 1 of 2
Command =>                                     Scroll =>

      Symm      M1
CUU  U Volser Sdev Serial CUU                Status          New    Num  Num
*****
_ 6D50 S JWA090 0090 06185 6D51                J@A+  0  0
_ 6D51 T JWA090 0091 06185 6D50                J@A+  0  0
***** Bottom of data *****

```

Figure 34 Device List panel

In this panel you may type an S on any line to show more detailed information about the device.

Logical Migration dataset configuration display

Using the C line command for a Logical Migrator (that is dataset oriented), group in the Monitor Promoted Groups panel results in a display of the datasets that comprise the group.

```

----- z/OS Migrator Dataset List----- Row 1 to 1 of

Line Commands: D = Detail      J = Jobs      L = Volumes
                M = Map Vol    O = VolSum    S = Extents
Filter Type:  _ G = Group      M = Dsn Mask V = Volser   S = Status
Value: _____
Source DS/Target DS                Group      Volser    Status
*****
_ APQATEST.CG2.NAA3.AD6721          MOVE1      AD6721    MIRROR
  APQATEST.CG2.AD6721.MIGR          TC71FF
_ XYZTEST.ABC                       CAPTURE    SYM71B    MIRROR
  XYZTEST2.ABC                      SYM719
_ XYZTEST.ABC.CROSS                 CAPTURE    SYM71A    MIRROR
  XYZTEST2.ABC.CROSS                SYM71C
_ XYZTEST.ABC4                     CAPTURE    SYM71C    MIRROR
  XYZTEST2.ABC4                     SYM71B
_ XYZTEST.ABC8                     CAPTURE    SYM71C    MIRROR
  XYZTEST2.ABC8                     SYM71B
_ XYZTEST.EMPTY                    CAPTURE    SYM71C    MIRROR
  XYZTEST2.EMPTY                    SYM71B
_ XYZTEST.GDG.DATASET              CAPTURE    SYM719    MIRROR
  XYZTEST2.GDG.DATASET              SYM71C
_ XYZTEST.GDG.DATASET.G0003V00     CAPTURE    SYM719    MIRROR
  XYZTEST2.GDG.DATASET.G0003V00    SYM71C
_ XYZTEST.GLARF2                   CAPTURE    SYM71B    MIRROR
  XYZTEST2.GLARF2                   SYM71C
***** Bottom of data *****

```

Figure 35 Dataset List panel

Dataset List line commands

[Table 22](#) lists the line commands you can enter from the Dataset List panel.

Table 22 Dataset List line commands

Command	Description
D	Detail — Displays a formatted list of all dataset attributes and status. This is the equivalent of the command Display DSNDDET <i>dsname</i> .
J	Jobs — If the group is in Diversion, display any jobs using this dataset and blocking completion.
L	Volumes — Show information on the first event of this dataset on each volume it occupies.
M	Map — Displays a map of z/OS Migrator-managed extents on the first dataset volume. This is the equivalent of the command D EXTMAP <i>ccuu</i> for the first device. Use the Operator Interface to request this type of display for other managed devices.
O	Volume Summary — Displays a volume summary for the dataset showing all the source and target volumes for the dataset, listed by volume.
S	Select Extents — Display the detail of all extents for the dataset.

Using Dataset List filters

Logical migration groups can contain a large number of datasets. z/OS Migrator provides a filtering mechanism to limit the datasets displayed in the dataset list. In the Filter Type field, specify one of the filter types listed in Table 23. Then in the Value field specify an appropriate group name, dataset mask, or volume serial number.

Table 23 Dataset List filters

Filter	Description																																						
G	Display the datasets for any other group without returning to the Monitor Group List.																																						
M	Select datasets for display using a dataset name mask. The Value field should contain a mask in the same format used for Source DSN selection.																																						
V	Select datasets based on whether any part of the dataset resides on the specified volume.																																						
S	<p>Select datasets based on the current status, for example, MIRROR. The value field should contain one of the status values that normally appear in the Dataset List Status field:</p> <table border="0"> <tr> <td>DROPPED</td> <td>Dataset has been dropped from the group</td> </tr> <tr> <td>DELETED</td> <td>Dataset deleted during migration</td> </tr> <tr> <td>SKIPPED</td> <td>Dataset skipped (dropped during selection)</td> </tr> <tr> <td>GDGBASE</td> <td>Dataset is GDG base</td> </tr> <tr> <td>CLUSTER</td> <td>Dataset is VSAM cluster</td> </tr> <tr> <td>SPHERE</td> <td>Dataset is base of sphere</td> </tr> <tr> <td>NORUN</td> <td>Dataset discovered during ACT NORUN processing</td> </tr> <tr> <td>SYNC</td> <td>Being synchronized</td> </tr> <tr> <td>PNDMIRR</td> <td>Synchronizing and pending mirror transition</td> </tr> <tr> <td>MIRROR</td> <td>Mirror mode</td> </tr> <tr> <td>PNDDVRT</td> <td>In process of Diversion</td> </tr> <tr> <td>DVRT</td> <td>I/O now being diverted to target</td> </tr> <tr> <td>DONE</td> <td>Migration complete</td> </tr> <tr> <td>DVRTFAIL</td> <td>Diversion failed</td> </tr> </table> <p>In addition, some of these status values can have a suffix indicating suspend or deactivation is in process, for example, MIRRSUSP = Mirroring suspended.</p> <table border="0"> <tr> <td>SUPR</td> <td>Suspend requested</td> </tr> <tr> <td>SUPN</td> <td>Suspend pending</td> </tr> <tr> <td>DEPN</td> <td>Deactivate pending</td> </tr> <tr> <td>SUSP</td> <td>Suspended</td> </tr> <tr> <td>TERM</td> <td>Termination - Free in process</td> </tr> </table>	DROPPED	Dataset has been dropped from the group	DELETED	Dataset deleted during migration	SKIPPED	Dataset skipped (dropped during selection)	GDGBASE	Dataset is GDG base	CLUSTER	Dataset is VSAM cluster	SPHERE	Dataset is base of sphere	NORUN	Dataset discovered during ACT NORUN processing	SYNC	Being synchronized	PNDMIRR	Synchronizing and pending mirror transition	MIRROR	Mirror mode	PNDDVRT	In process of Diversion	DVRT	I/O now being diverted to target	DONE	Migration complete	DVRTFAIL	Diversion failed	SUPR	Suspend requested	SUPN	Suspend pending	DEPN	Deactivate pending	SUSP	Suspended	TERM	Termination - Free in process
DROPPED	Dataset has been dropped from the group																																						
DELETED	Dataset deleted during migration																																						
SKIPPED	Dataset skipped (dropped during selection)																																						
GDGBASE	Dataset is GDG base																																						
CLUSTER	Dataset is VSAM cluster																																						
SPHERE	Dataset is base of sphere																																						
NORUN	Dataset discovered during ACT NORUN processing																																						
SYNC	Being synchronized																																						
PNDMIRR	Synchronizing and pending mirror transition																																						
MIRROR	Mirror mode																																						
PNDDVRT	In process of Diversion																																						
DVRT	I/O now being diverted to target																																						
DONE	Migration complete																																						
DVRTFAIL	Diversion failed																																						
SUPR	Suspend requested																																						
SUPN	Suspend pending																																						
DEPN	Deactivate pending																																						
SUSP	Suspended																																						
TERM	Termination - Free in process																																						

Activating a group

Once a group has been defined and promoted, you must activate the group in order to start the migration process for that group.

To activate a group:

- ◆ Type **A** in the line command entry text field next to the group that you want to activate and press **Enter**.

For any of the major action commands, including Activate, you will see a pop-up message requesting that you confirm your intention to execute the command, for example:

```

+----- Confirm Activate Request -----+
|                                         |
| Are you sure you want to Activate   the Migrate of |
| TEST3                               ? Y (Y=Yes N=No) |
|                                         |
+-----+

```

Figure 36 Confirm Activate Request panel

Type **Y** in the input field and press **Enter** to proceed.

Active, Copy, and Mirror states

When the Activate request is accepted, you will see "Activate Issued" in the upper right corner of the dialog and the Activate processing will begin.

Once the group is activated, z/OS Migrator will begin the migration process for that group. During this time, the group will transition through *Active*, *Copy*, and then *Mirror* states. During the *Copy* phase, the status will display as Syncing(*nnn%*) showing the amount of data that has been copied. Pressing **Enter** in this screen will enable the monitor to update the view of migration status.

Volume Mirror groups proceed from the Copy state to the Mirror state and remain there. Volume Migrate groups, other than Constant Copy groups, proceed from the Copy state to a Swap that completes the migration.

A Logical Migrate group will remain in *Mirror* state until a *Divert* command is issued. Once mirroring has completed it will display "Moved=100% Mirror." When a *Divert* command is issued the migration process will enter the *Pending-Divert* state, followed automatically by the *Divert* state. Migration will remain in the *Divert*

state until all direct job allocations to the source datasets are released, after which the migration will enter *Pending-Complete* state (waiting for allocations locally), followed by *Pending All Complete* (waiting for allocations globally), and finally, *Complete* (all source allocations relinquished).

If the group status shows that an error has occurred, use the E line command to display detained information and messages. You can always use the E line command to show possible warning messages that may have been issued during processing.

Suspending a group

You may temporarily suspend a group while it is in the *Sync*, *Pending-Mirror*, or *Mirror* states by using the P line command in the group list. This pauses all migration activity for the group until you issue the Resume command. Changes to the migrating data continue to be monitored during the suspension so that any changed tracks can be copied once the group has been resumed. While in Suspended state, the group and dataset displays will show the appropriate suspend status:

SYNCSUSP: synchronization suspended

PNDMSUSP: pending mirror suspended

MIRRSUSP: mirroring suspended

Resuming a suspended group

To resume a suspended group:

- ◆ Type **R** in the line command field next to the group you wish to resume and press **Enter**.

The group will revert back to the *Sync* or *Pending-Mirror* state as remaining uncopied or changed tracks are copied. Once the copy is complete, the group will enter *Mirror* state normally.

Diverting a Logical Migrator group

Once a group has reached *Mirror* state, you may divert it. While in *Divert* state, all I/O from ongoing applications will automatically be diverted to the target datasets. Any newly started applications will directly allocate the target datasets.

Note: You may wish to wait to enter the *Divert* state until you are ready to restart any applications which are actively using the source datasets. While technically not necessary, it does simplify the process by shortening the amount of time that application I/O must be diverted by z/OS Migrator.

To divert a group:

- ◆ Type **V** next to the group that you want to divert and press **Enter**.

Once a group has been diverted, you can display a list of the allocated datasets in the group and the jobs that have them allocated by using the J line command at either the group level or the dataset level.



CAUTION

A Divert command will be rejected if a new extent allocation is in progress (**BUSY**). A Divert command will also be rejected if an LPAR on which the group is known is not currently active (**SERVER INACTIVE**). In this case, use the Operator Interface (Option 4 in the z/OS Migrator main menu) to show the status of the group on all LPARs, as shown below.

```

z/OS Migrator Display Group parameters                Row 1 of 54
Command =>                                           Scroll => DATA

Enter your group name or hit enter for all group status
Group .. _____
*****
D G TEST3
EFMP001I D G TEST3
NAME      1ST-PAIR  COUNT  STATUS          OWNER DEFINITION
TEST3     4D99/4D9C   16  MIRROR ACTIVE   X04  MODE(LMIG)

SMFID  SSID  OWN  DSNS          STATUS-FLAGS          STATUS
X06   DB01      0      VER              SERVER INACTIVE
X04   DB01 X04  16      ACT,MIR           MIRROR ACTIVE

```

Figure 37 z/OS Migrator Display Group parameters

If the LPAR is known to be "uninvolved", you may use the REMOVE command in the Operator Commands, Z dialog to remove it from consideration.

```
REMOVE smfid
```



CAUTION

Take caution when performing this procedure. Data integrity will be compromised if any application is accessing any of the group datasets from the removed system.

The actual Divert is preceded by a verification process to ensure that the catalog and VTOC activities of Diversion will be successful. If an error is found, the Divert will be rejected and the group will remain in Mirror state. Any error messages generated during the diversion processing can be displayed using the E line command in the group list:

- ◆ Type **E** in the line command field for the group and press **Enter**.

Completing a Logical Migrator group

Depending on selected options, Diverted groups will either proceed directly through the completion process or remain in *Divert* state until you enter the Complete command. In particular, if NOAUTOTRM has been specified, you will need to use the Z line command to move group completion along. (“[OPTIONS](#)” in “[Step 10: Allocate the z/OS Migrator system information dump dataset](#)” on page 58 provides details.)

- ◆ Type **Z** in the line command field next to the group and press **Enter**.

Once the command has been accepted the group will start progressing through the Pending Complete, Pending All Complete, and Complete phases.

Deactivating a group

You must *suspend* a group before you can deactivate it. Deactivation is only an option during the *Copy (Sync and Pending Mirror)* and *Mirror* states.

To deactivate a group:

- ◆ Type **P** next to the group that you want to deactivate and press **Enter**.
- ◆ Once the Suspend has been acknowledged, type **D** in the line command field and press **Enter**.

Note: You cannot deactivate a group once it has been diverted.

Option 3 - Display Configuration Information

This option allows you to produce formatted displays of the following configuration elements:

- | | |
|----------------------|--|
| 1 System | CPU serial, SMF id, Subsystem name |
| 2 Global Information | Certain z/OS Migrator server parameters |
| 3 Group | Similar to the interact with groups list |
| 4 Device | Device-related information for managed devices |
| 5 Dataset | Formatted display of dataset information |
| 6 Extent | Formatted display of extent information |

Select the desired item and enter the identifying information as requested.

Option 4 - Operator Interface

This option allows you to request z/OS Migrator Help information and interact with the server via the command interface. In addition, it presents a number of options for entering pre-formatted commands:

1 Display Device	A Display DEVIce command
2 Display Hosts	A Display HOSTS command
3 Display Group	A Display Group command
4 Display LM Paths	A Display PATHS LM command
5 Display Dataset Detail	A Display DSDET command
6 Display Extent	A Display EXTMAP command
7 Display Dataset	A Display DSN command
8 Debug Options	Display selected debugging information and set debug flags
9 Trace Options	A SET TRACE command
H Help	Request Help information
Z Console Interface	Enter free-form z/OS Migrator operator commands

If you are using the monitor log dataset, then the output of all commands submitted is logged to the log dataset and may be viewed using the **V** option in the main menu.

If you are using the option **Stack Console output . . Y** as specified in the Set User Session Options dialog, then while in the Help or Console Interface dialogs, all commands and responses will be viewable in push-down fashion. Once you switch to another part of the monitor, the push-down log will be reinitialized. However, the commands and responses will still be viewable in the background log.

Option 5 – Display Host Messages

This option allows you to display z/OS Migrator messages for a specific host system for which the current z/OS Migrator ISPF Monitor is defined. When this option is invoked, the initial screen that is displayed lists all valid host systems currently connected.

Command ==>>		Subsystem @\$ Host Messages			Row 1 to 2 of 2	
SMFID	SSID	CPUID		Last Server	Start Time	Scroll ==>> DATA
_ X06	DB01	0635DE2096	Idle	06/11/2007	17:23:14.258	
_ X04	DB01	0435DE2096		06/11/2007	02:17:16.693	

The display shows the following data.

SMFID	The SMF ID of the system
SSID	The subsystem ID of z/OS Migrator
CPUID	The significant digits of the hardware CPU serial number
HEARTBEAT	Time of the last update to the EFFM database by the system

The notation IDLE at the end of a line means that the observing system considers that system to be idle as far as z/OS Migrator is concerned. This means that the started task is not running.

Note: The integrity of z/OS Migrator operations depends on the z/OS Migrator server being active on all systems with access to any of the devices involved. If a required system appears to be idle with regard to z/OS Migrator, make sure that the system is actually down or the subject volumes are all offline to that system. Otherwise, affected groups should be deactivated and restarted when all systems are operational.

To display messages for a specific host system:

- ◆ Enter **S** next to the relevant host system line, and press **Enter**.

The messages for the selected host are displayed in a separate panel with a date and time stamp for each message. The date and time stamp can either be in local time or GMT depending on your setting.

(“Option S – Set User Session options” on page 137 provides more information on settings.)

When you first enter the Display Host Messages dialog for a host, you may see a pop-up similar to the following:

```
By default only messages after 06/11/2007 02:17:16.693 are
displayed, the time stamp of last z/OS Migrator server cold start.

Oldest message available to display is 06/01/2007 18:41:01.690

After the initial display, you may enter command all or depress
PF 5 to see all messages in the data base for host X04

                                Press ENTER to continue.
This message only shown once per z/OS Migrator monitor session.
```

Figure 38 Message pop-up panel

When you press **Enter**, you will see a list of messages from the selected host, similar to those shown in [Figure 39](#). Date changes in the message stream are indicated by the *****Current date***** statement.

```

                                Messages for Host X04                                Row 1 to 17 of 74
Command ==>>                                                                Scroll ==>> DATA

*** Current date is now 06/11/2007 ***
02:17:16 EFM1727I Installed software for this system is: ZOSM FM07040 Level.
02:17:16 EFM0001I z/OS Migrator initializing.
02:17:16 EFM0018I Processing execution parameters.
02:17:16 EFM0028I Entering the configuration file process.
02:17:16 EFM0029I The configuration file has been successfully opened.
02:17:16 EFM0030I All configuration files that were previously allocated are
                                being freed.
02:17:16 EFM0021I Processing all configuration file parameters.
02:17:16 EFM0056I Processing the command prefix (CPFX) parameters, the length
                                is: 02 and the CPFX is: @$$.
02:17:16 EFM0031I Entering the main initialization process.
02:17:16 EFM0151I Symmetrix Control Facility (SCF) is not active.
02:17:16 EFM0189I The return from the EMCKFI routine was: R15=0004005F.
02:17:16 EFM0191I The old SSCT Control Table was located at: 00C51040.
```

Figure 39 Host system messages

Note: Find and repeat find functions are available in the messages panel so that you can search for specific messages without the need to scroll through the entire messages listing.

Option 6 – System Change Summary

This option displays a history of the z/OS Migrator product release in use. The PTFs are identified by a fix number.

```
Command ===>                System Change Summary                Row 1 to 12 of 12
                                                             Scroll ===> DATA
                               Version 3.0.0
                               Report Date and Time
                               06/11/2007 14:16:12.34

Base System built April 5, 2007.

PTF FM070405 built April 5, 2007.
```

Figure 40 System Change Summary option

Option 7 - Message Help

This option contains most of the z/OS Migrator messages within the system. These messages provide a more detailed explanation than what is displayed in the z/OS Migrator ISPF Monitor or in the output listing.

See the *EMC z/OS Migrator Message and Code Guide*. For other types of messages, refer to the message guide for the corresponding product.

```

                                Message Display
Command =>                                Scroll => DATA

Refer to the z/OS Migrator Messages and Codes manual for the following message prefixes.

Chapter 2 - EMC SNAP Messages      Chapter 3 - EMC TimeFinder Utility Messages
AEXT  EQCA  EDSS   ESNP           BCVS  BCVU
EMCS  ESVP  Error Codes

z/OS Migrator messages have a prefix of EFM      Note
Enter a 4 digit                                Within the message text there may be
z/OS Migrator message number ____ replaced when the message is generated.
  
```

Figure 41 z/OS Migrator message display

To search for a specific message:

- ◆ Type the 4-digit value, contained in the message number, on the Message Number line and press **Enter**. For example, EFM0006I would be searched as 0006, as shown in [Figure 42 on page 185](#).

```

Command =>
Message Display
Row 1 to 10 of 10
Scroll => DATA

Refer to the z/OS Migrator Messages and Codes manual for the following message prefixes.

Chapter 2 - EMC SNAP Messages      Chapter 3 - EMC TimeFinder Utility Messages
AEXT  EQCA  EDSS   ESNP           BCVS  BCVU
EMCS  ESVP  Error Codes

z/OS Migrator messages have a prefix of EFM    Note
                                           Within the message text there may be
Enter a 4 digit                          special string variables that will be
z/OS Migrator message number 0006 replaced when the message is generated.
EFM0006I

Message Text:
z/OS Migrator now has a new environment active.

Explanation:
z/OS Migrator is now executing in a new environment.

System Action:
z/OS Migrator execution continues.

```

Figure 42 z/OS Migrator message display example

The last character in a z/OS Migrator message indicates the severity of the message.

Table 24 z/OS Migrator message severity table

Character	Return code	Description
A		Action message — WTO/WTOR option selected; requires a response from the z/OS console or the z/OS Migrator ISPF monitor.
I	0	Informational message — all migrations successful; no action required.
W	4	Warning message — all migrations successful; one or more warning messages was issued; an action may be required.
E	8	Error message — a migration did not complete successfully; a corrective action may be necessary.
S	12	Severe error — the session did not start or complete successfully; a corrective action must be taken.

For example, if a SET statement fails to find any matching datasets, return code 8 will be returned. Depending on the TolerateAllocationFailure parameter, this may lead to the group

activation failing or proceeding. In general, W or E messages indicate the situation must be reviewed to ensure that you get the result you expected. An S level severity generally indicates a server level failure.

Option 8 – Monitor Command Line Help

The z/OS Migrator ISPF Monitor does not support the standard ISPF commands on most screen displays. Select this option to obtain detailed usage information for the line commands that the z/OS Migrator ISPF Monitor does support and how they may be different from standard ISPF commands. [Figure 43 on page 188](#) shows the display.

```
Monitor Command Line Help

More:      +

Most screens presented by the monitor will not perform the normal ISPF
commands. Detailed below are the commands that the Monitor supports
on the command line and the differences in their action.

Monitor functions that invoke the ISPF "browse" support all ISPF
commands normally.

Monitor functions that build panel display support the following
commands as described.

FIND (F is the short command)
Enter FIND followed by your search argument(s). Unlike ISPF, the
FIND doesn't require a ' (hex '7d') if the user wished to specify
more than one word.
All of the following FINDs are valid

FIND copy
F   volume serial
FIND average seek time

Both the FIND argument and data are translated to uppercase prior
to the comparison. The line containing the match will be the first
line displayed.

FIND supports no parameters (ie: PREVIOUS).

RFIND (set to PF5)
Entering RFIND as a line command will give the ISPF message
"RFIND not active". The ISPF RFIND command is NOT passed to the
monitor. The monitor simulates this function by entering FIND with
no search argument. Depressing PF5 will present a FIND command
with no search argument. When used, the first line searched is the
2nd line on the display. Once at the bottom of the data, the next
PF5 will cause the search to wrap to the top.

PRINT
PRINT is an ISPF command and will only print the current physical screen.

ZOSMPRT (P is the short command)
Unlike ISPF, ZOSMPRT will print all the lines contained in that panel
to ISPF's print dataset. Each invocation of ZOSMPRT starts on a new
page.
```

Figure 43 Monitor command line help

Option 9 – Display Installation Options

Using this option, you may review the current default options set by the optional SYSOPTN1 batch job.

The initial z/OS Migrator display contains the following default dataset values:

- ◆ Company: *YOUR COMPANY*
- ◆ Site: *YOUR SITE*
- ◆ WTO/WTOR for AutoOps (messages for automated operations required): *NO*
- ◆ WTO AutoOps Route Codes (MVS routecodes): *2, 4*
- ◆ Local Time on Messages (display time): *LOCAL*

You can change these values by running the SYSOPTN1 batch job, with the new values you want displayed for your company.

Additionally, the display option includes the operating system environment in which the z/OS Migrator ISPF Monitor is executing.

```

                                Installation Options                                Row 1 to 13 of 13
Command =>                                                                Scroll => DATA

Company : EMC CORPORATION
Site . : HEADQUARTERS
Site ID : 00001 Date Initialized. . . : APRIL 02, 2007
WTO/WTOR for AutoOps . . . : NO
WTO AutoOps Route Codes . . : 2,4
Local Time on Messages . . : NO

                                Operating Environment

z/OS Migrator Version : 2.1.0
z/OS Migrator DB CPFY : @$
z/OS Migrator DB DSN  : JDOE1.SFMM110.BSAM.DB4096
z/OS Migrator DB VOL  : UWD012
z/OS Migrator DB SAF  : No Security
USER ID      : JDOE1
CPU ID      : 000435DE 2096
SCP NAME    : SP7.0.4
SCP FMID    : HBB7707
ETR ID      : 00
Local Time  : 06/11/2007 14:48:43.06
GMT Time    : 06/11/2007 18:48:43.06
Local Offset : -04:00:00
Leap Seconds : +000

```

Figure 44 Installation options

Option V- View or Browse Log Dataset

This option allows you to browse a member in the log PDS. The naming convention for dataset logs is as follows:

Myyyyddd

where *yyyy* is the year and *ddd* is the Julian date; for example, M2007138 is 18 May 2007.

<u>M</u> enu	<u>F</u> unctions	<u>U</u> tilities	<u>H</u> elp			
ISREP001 SXN80.EFM120.LOG			Row 00001 of 00005			
Command ==>			Scroll ==> DATA			
Name	Prompt	Size	Created	Charged	ID	
_ *M2007138						
_ *M2007165						
_ *M2007178						
_ *M2007179						
_ *M2007180						
END						

Figure 45 Dataset logs

To browse or view the contents of a dataset log:

- ◆ Type **S** to the left of the dataset log that you wish to view and press **Enter**.

Note: The find and repeat find functions are available in the browse panel so that you can search for specific dataset logs without the need to scroll through the entire listing.

Each command that is logged is time-stamped with a message of the following format:

```
z/OS Migrator QUEUE COMMAND EXECUTED BY JDOE1 ON 06/11/07 AT 13:17:06
```

This message shows the user-ID, date, and time.

This chapter lists and describes the commands available via the z/OS Migrator Command Line Interface (CLI). Topics are:

- ◆ Overview 192
- ◆ General syntax 193
- ◆ Help facility 194
- ◆ z/OS Migrator active commands 195
- ◆ z/OS Migrator environmental commands 200
- ◆ z/OS Migrator display commands 207
- ◆ z/OS Migrator diagnostic commands 216

Overview

This chapter lists and describes the commands available via the z/OS Migrator command line interface (CLI).

In addition to the ISPF Monitor dialog, z/OS Migrator has several other modes by which you can control migrations and display status and other information. Chief among these is the command line interface.

The commands of the command line interface can be submitted to the z/OS Migrator server from an operator console or any equivalent command environment, such as an SDSF session if it is authorized for submitting z/OS commands.

To submit commands to the server, prefix the command with the command prefix defined for that server via the using execution or config file CPMX parameter. For example, if CPMX=@\$, then display groups with the command:

```
@$DG
```

z/OS Migrator commands may also be submitted using the z/OS Migrator batch utility, EFMFBAT. When including commands in the batch utility input stream do not prefix them with the command prefix. Using the batch utility allows many common z/OS Migrator migration tasks to be automated. [Chapter 10 on page 295](#) provides more information on the batch utility.

General syntax

z/OS Migrator commands generally consist of a command verb followed by a resource type, such as GROUP or DSN (where appropriate) and a name or mask.

Dataset name masks include the following characters:

- ◆ A single asterisk (*) in the place of any number of consecutive characters, up to an entire single qualifier in any position.
- ◆ A double asterisk (**) to represent any number of consecutive qualifiers.
- ◆ One or more question marks (?) to represent a single character in any position.

Any dataset matching the mask provided will be displayed or acted upon.

Certain commands have optional "switches" that affect the execution of the command. These generally take the form of a single character preceded by a slash (/). The syntax is forgiving about the placement of such switches relative to other parameters.

In the descriptions in this chapter, the required characters of any keyword are indicated by capitalization. Commands are not case sensitive.

For convenience, the description of the z/OS Migrator commands is divided into four sections, immediately following a description of the user-extendable Help facility.

- ◆ [“z/OS Migrator active commands” on page 195](#)
Commands that initiate or modify the behavior of migration group processing.
- ◆ [“z/OS Migrator environmental commands” on page 200](#)
Commands that can be used to set environmental parameters affecting group management.
- ◆ [“z/OS Migrator display commands” on page 207](#)
Commands that can be used to display migration resources and their status.
- ◆ [“z/OS Migrator diagnostic commands” on page 216](#)
Commands that can be used to aid in problem solving, if necessary.

Help facility

z/OS Migrator includes a user-extendable help facility. SAMPLIB contains a number of help objects that can be loaded into the z/OS Migrator database to serve as basic command descriptions and supporting information. See the following members for samples:

HELP	A basic, high-level command summary, primarily of the Display commands.
#MODEL	A model for the control parameters for defining a user-specified help object.
#LOAD	Sample JCL for loading all the pre-formatted sample help objects or an individual user-defined help object.

Once the help objects are loaded to the z/OS Migrator database, you can enter the HELP command to display an individual object, as described next.

HELP

Syntax

HELP *object-name*

Parameters

object-name

The lowest level name qualifier of an object in the z/OS Migrator database under the name hierarchy: SYSTEM.HELP. For example, HELP DISDSN will display the object loaded from SAMPLIB member DISDSN, describing the Display DSN command.

Note: You may define and maintain your own object called SYSTEM.HELPINDEX to serve as a high level index of help objects. If you choose not to do this, the HELP INDEX command will simply display a sorted list of the names of all the help objects currently in the z/OS Migrator database.

z/OS Migrator active commands

The z/OS Migrator active commands are those that initiate or modify the behavior of z/OS Migrator migration group processing.

ACTivate

Activate a migration group.

Syntax `ACTivate groupname [NORUN]`

Parameters *groupname*

Specify the name of the group you want to activate. This is the name specified in the GROUP(name) group definition parameter. Activation causes the group definition to be interpreted, all source datasets to be discovered, and all target datasets to be allocated as defined by the group parameters.

NORUN

An optional parameter that causes the group to be parsed, source datasets to be identified, and target dataset names to be generated. However, group processing ends at that point without actual allocation and data movement. This can be used to evaluate a group definition prior to full activation.

COMPLete

If not using the AUTOTRM option, the COMP (or TERM) command is needed to signal the server to attempt to move the group from Diverted mode to Completed mode.

Syntax `COMPLete groupname`

Parameters *groupname*

Specify the name of the group you want to complete. This is the name specified in the GROUP(name) group definition parameter.

The progress of the group from this point depends on whether the ACTVCMP option has been selected and whether all tasks with source allocations have released their allocations.

DEACTivate

Deactivate a suspended migration group. Note that an active migration group must be suspended before using the DEACT command. Groups that have been deactivated cannot be resumed; they must be reactivated using the Activate command.

Syntax DEACTivate *groupname*

Parameters *groupname*

Specify the name of the group you want to deactivate. This is the name specified in the GROUP(name) group definition parameter.

DELeTe

Delete a group from active server memory. A group must be inactive to be deleted. Once a group has been deleted, it must be re-promoted to be known to the migration server.

Syntax DELeTe *groupname*

Parameters *groupname*

Specify the name of the group you want to delete. This is the name specified in the GROUP(name) group definition parameter.

DIVert

Request that a migration group be moved to the Diverted state. This involves the swapping of the identities of source and target datasets. If the AUTOTRM option is selected, then the group will automatically proceed to Completion as all tasks allocating the source datasets release their allocations. If the ACTVCMP option is selected, then source resources may be freed or reused prior to Completion.

Syntax DIVert *groupname*

Parameters *groupname*

Specify the name of the group you want to divert. This is the name specified in the GROUP(name) group definition parameter.

Note: The Divert command will be rejected unless all participating migration servers are active and acknowledge the command.

DROP

Dynamically exclude one or more datasets from an active group. This command is only accepted once a group has reached Mirror state and before the group is fully in the Diverted state.

Syntax `DROP groupname data-set-mask`

Parameters *groupname*

Specify the name of the group you want to affect. This is the name specified in the GROUP(name) group definition parameter.

data-set-mask

A dataset selection mask. All datasets and any related components will be dynamically excluded from further group processing. Such datasets will have to be remigrated in a subsequent group, if desired.

Note: The DROP command would commonly be used to exclude one or more datasets for which irregularities in metadata are blocking Diversion for the entire group. Once such datasets are dropped, Diversion will proceed normally for the rest of the group.

RESume

Resume processing for a suspended group. Groups may be suspended in the initial copy phase and while in Mirror state. Once the group is resumed, any remaining tracks as well as any changed tracks will be copied and the group will move to Mirror state.

Syntax `RESume groupname`

Parameters *groupname*

Specify the name of the group you want to resume. This is the name specified in the GROUP(name) group definition parameter.

SHUTDOWN

The same as [“STOP” on page 198](#).

SPIN

Spin off any accumulated log messages to the SYSPRINT dd and reallocate it.

Syntax SPIN COLD | ALL | PRINT <RESET>

Parameters COLD | ALL | PRINT

Specify the starting point of the messages to be written to the dynamically allocated EFMMLOG dd and spun off:

- COLD
All messages since last COLD start – the default.
- ALL
All messages currently in the log.
- PRINT
Only messages since the last time they were printed.

RESET

Reset the last-printed timestamp.

STOP

Causes the migration server to terminate. This command will only be accepted if there are no logical migration groups in either the diversion or mirror phase.

Syntax STOP <NOPRINT>
or
SHUTDOWN

Parameters NOPRINT

Normally, all log messages since the last print will be spun off at shutdown. NOPRINT overrides this and leaves the messages queued for later printing.

SUSpend

Temporarily stop processing for a group. Groups can only be suspended during the Copy and Mirror phases.

Syntax `SUSpend groupname`

Parameters `groupname`

Specify the name of the group you want to suspend. This is the name specified in the GROUP(name) group definition parameter.

TERM

The same as [“COMPLETE” on page 195](#).

z/OS Migrator environmental commands

Commands that can be used to set environmental parameters affecting group management.

IDLE

Reset the amount of time a system must be out of contact before it will be considered idle. Note that most group-oriented commands cannot be serviced unless all servers are known to be active.



CAUTION

If a system is taken out of service and the migration server appears idle, you must REMOVE that server in order to proceed. Please do so with extreme caution. If migrating datasets are, in fact, accessed from a system whose migration server has been removed, data integrity will be compromised.

Syntax

IDLE *seconds*

Parameters

seconds

Specify the number of seconds, integer only, after which an unresponsive server will be considered idle.

You should not specify a time shorter than the default, which is 45 seconds.

Note: Console messages identify the transition from active to idle and vice versa. The state of participating servers can be displayed using the Display HOSTS command.

Specifying a multi-minute time may cause certain action commands to take unacceptably long to time out if a system is actually IDLE.

MAXCIO

Reset the maximum number of concurrent I/O requests that will be scheduled on any one channel path.

[“MAX_CHANNEL_IO” on page 327](#) provides a full description.

Syntax

MAXCIO *count*

Parameters *count*

The number of concurrent requests.

MAXDIO

Reset the maximum number of concurrent I/O requests that will be scheduled for any one device pair. “MAX_DEVICE_IO” on page 328 has more information.

Syntax MAXDIO *count*

Parameters *count*

The number of concurrent requests.

REMOVE

Remove an idle migration server from active participation in migration activities.

Syntax REMOVE *smfid* <*ssid*>

Parameters *smfid*

The SMF ID of the server to be removed.

ssid

If more than one migration server is running on a single LPAR, and sharing the same z/OS Migrator database, then you must specify the subsystem ID of the server you want to remove.



CAUTION

Do not remove a server unless you know that there is no I/O to migrating datasets or volumes.

SET

Set various processing options and debugging controls.

Syntax

SET *option* | *DEBUG(features)* ON | OFF

Parameters

option

The available options are:

ACTVCMP or
NOACTVCMP

Active Completion allows the source volumes to potentially be varied offline without recycling the processes using source datasets. NOACTVCMP is the default setting.

AUTOTRM or
NOAUTOTRM

Automatic Termination means that groups will immediately attempt to go to completion as soon as divert status is achieved. Otherwise the COMPLETE command must be issued. AUTOTRM is the default setting.

TRMCENQ or
NOTRMCENQ

TRMCENQ means that z/OS Migrator will issue an ENQ request against the target datasets during the migration. TRMCENQ is the default setting.

AUTODEL or
NOAUTODEL

AUTODEL causes the group control blocks to automatically be freed once the group completes. AUTODEL is the default setting.

AUTOHK or
NOAUTOHK

AUTOHK causes a group activation to automatically re-activate the system interfaces. AUTOHK is the default setting.

AUTOCLN or
NOAUTOCLN

AUTOCLN causes the server to automatically remove interfaces and delete common resources at shutdown with no active groups. AUTOCLN is the default setting.

SDUMP or
NOSDUMP

SDUMP results in an SDUMP being taken at any server failure. SDUMP is the default setting.

ECHO ECHO ON indicates the joblog/syslog will contain results from various display commands issued from the ISPF interface or the batch interface. Results are automatically displayed for commands issued using the command line interface. ECHO OFF will display the command issued without the results of the command. The default setting is ON.

PREVENTVARY The PREVENTVARY | NOPREVENTVARY statement in the z/OS Migrator configuration file prevents duplicate volume labels by invalidating the VOL1 label on the target volume at the beginning of the migration. During the completion of the migration the VOL1 label is corrected.

The SET PREVENTVARY | NOPREVENTVARY command dynamically changes the setting. The default is NOPREVENTVARY.

Note: The PREVENTVARY option only applies to Volume Migrations groups and not to Mirror or Consistent groups.

DEBUG(feature)**CAUTION**

Debug features should only be set under the direction of EMC support personnel, since they produce voluminous output.

The basic DEBUG features are:

CONFIG	Trace parameter configuration activities.
CPLDS	Messages related to couple datasets.
TASKMGT	Trace task management, internal to migrator.
STGMGR	Debugging related to storage manager.
DIVERT	Debugging detail for group diversion.
XMAP	Produce a detailed dump of extent map for EXTMAP command.
DBASE	Trace I/O activity for the migrator database.
SYNC	Trace SYNC and SYLM task activity.
PCR1	Trace the migrator PC routine.
COMM	Debugging info for command processing.
SDDF	Additional debugging for SDDF related processing.
DTERM	Trace Diversion termination.
TRANS	Externalize state transition to Mirror and Divert.
DBHB	Messages relating to the heartbeat task.

Additional debug features related to I/O interfaces are:

<u>Feature</u>	<u>Tracing related to module</u>
IOTR	EFMMIOTR
IOTR_OK	
IOTR_ERR	
ENF	EFMMENFV
ENF_OK	
ENF_ERR	
ESDDF	EFMMSDDX
ESDDF_OK	
ESDDF_ERR	
HPST	EFMMHPST
HPST_OK	
HPST_ERR	

<u>Option</u>	<u>Tracing related to module</u>
HTIO	EFMMHTIO
HTIO_OK	
HTIO_ERR	
HPRV	EFMMHPRV
HPRV_OK	
HPRV_ERR	
HM00	EFMMHM00
HM00_OK	
HM00_ERR	
HPRE	EFMMHPRE
HPRE_OK	
HPRE_ERR	
HS00	EFMMHS00
HS00_OK	
HS00_ERR	
IOSM	EFMMIOSM
IOSM_OK	
IOSM_ERR	
SPMW	EFMMSPMW (suspend)
SPMW_OK	
SPMW_ERR	
IORQ	EFMMIORQ
IORQ_OK	
IORQ_ERR	
IOTM	EFMMIOTM
IOTM_OK	
IOTM_ERR	
IOVM	EFMMIOVM
IOVM_OK	
IOVM_ERR	

<u>Option</u>	<u>Tracing related to module</u>
IODR	EFMMIODR
IODR_OK	
IODR_ERR	
HSIO_FE	EFMMHSIO
HSIO_FE_OK	
HSIO_FE_ERR	
HSIO_BE	
HSIO_BE_OK	
HSIO_BE_ERR	
HMIH	EFMMHMIH
HMIH_OK	
HMIH_ERR	
HDIE	EFMMHDIE
HDIE_OK	
HDIE_ERR	
HEOS	EFMMHEOS
HEOS_OK	
HEOS_ERR	
CPTK	Trace track copy activity

Note: The options may also be set via the OPTIONS statement of the INI config file. The debugging features may be activated using the DEBUG execution parm also available within the INI config file.

z/OS Migrator display commands

The z/OS Migrator display commands can be used to display migration resources and their status.

The general format of the display commands is -

```
Display resource-type resource-mask options
```

For example:

```
Display DSN USER1.TEST.** /W
```

The above command will display any dataset under management by the migrator that fits the mask USER1.TEST.** and which has active allocations blocking movement from Diverted state to Complete.

Display DEBUG

Display internal configuration and status data.

Syntax

```
Display DEBUG <  
MOD | ANCH | BUILD | STOR | TASK | SERV | INTERFACE | SYS>
```

Parameters

MOD	Module IDs, entry points, and build dates
ANCH	Storage anchors
BUILD	Build level and date
STOR	Storage management parameters
TASK	Task control, debug settings, and so forth
SERV	Like MOD, but for internal services
INTERFACE	I/O interfaces and status
SYS	System environmental information

Display DEV

Display device summary and detail information.

Syntax `Display DEV <ccuu|ccuu-ccuu|ALL> <PRIM|MIRROR|BOTH>`

Parameters

ccuu

Specify a single unit address to get detailed information about the status and usage of the device.

ccuu-ccuu

Specify a range of devices to get summary information about the included devices.

ALL

Display summary information about all migration devices.

PRIM|MIRROR|BOTH

Limit the display to Primary (source) devices, Mirror (target) devices, or both.

Display DSDET

Display detailed information about one dataset.

Syntax `Display DSDET data-set-name`

Parameters

data-set-name

Specify a dataset to display.

Note: A mask may be used but only the first dataset encountered will be displayed and it may not be lowest in the collating sequence.

Display DSN

Display status information for all datasets matching the selection criteria.

Syntax `Display DSN <group-name | data-set-mask | ccuu | ccuu cchh> </M mode | /S | /W | /V | /X>`

Parameters

group-name

The name of the migration group. All datasets in the group will be displayed, subject to the optional switches.

data-set-mask

All datasets of any groups and matching the specified mask will display.

ccuu

All datasets whose source or target matches the specified device address will display.

ccuu cchh

The dataset having an extent on the specified unit and cylinder/head address (in hex) will display.

/M mode

Only display datasets whose current status matches the specified comparend. The comparend may contain mask characters. The valid comparends include:

SYNC	Synchronizing
PNDMIRR	Pending-mirror
MIRROR	Mirror
PNDDVRT	Pending divert
DVRT	Diverted
DONE	Done
DVRTFAIL	Divert failed

The rightmost characters of the status may be qualified with one of the following:

-SUPR	Suspend requested, for example MIRR-SUPR
-SUPN	Suspend pending
-DEPN	Deact pending
-SUSP	Suspended
-TERM	Active termination

/S

Only display "skipped" datasets.

/W

Only display datasets with allocation blocking completion.

/V

Display first extent on each volume of selected datasets.

/X

Display all extents of selected datasets (same as Display EXT).

Display EXT

Display extent level detail for datasets in migration.

Syntax

```
Display EXT <group-name | data-set-mask | ccuu | ccuu
cchh> </M mode | /S | /W | /V>
```

Parameters

group-name

The name of the migration group. All datasets in the group will be displayed, subject to the optional switches.

data-set-mask

All datasets of any groups and matching the specified mask will display.

ccuu

All datasets whose source or target matches the specified device address will display.

ccuu cchl

The dataset having an extent on the specified unit and cylinder/head address (in hex) will display.

/M mode

Only display datasets whose current status matches the specified comparend. The comparend may contain mask characters. The valid comparends include:

SYNC	Synchronizing
PNDMIRR	Pending-mirror
MIRROR	Mirror
PNDDVRT	Pending divert
DVRT	Diverted
DONE	Done
DVRTFAIL	Divert failed

The rightmost characters of the status may be qualified with one of the following:

-SUPR	Suspend requested, for example MIRR-SUPR
-SUPN	Suspend pending
-DEPN	Deact pending
-SUSP	Suspended
-TERM	Active termination

/S

Only display "skipped" datasets.

/W

Only display datasets with allocation blocking completion.

/V

Display first extent on each volume of selected datasets.

Display EXTMAP

Display the extent map for datasets on a particular device.

Syntax Display EXTMAP *ccuu* <*cchh*>

Parameters *ccuu*

The device address for the device of interest.

cchh

Optional - a *cchh* on the device. If specified, then only the next extent at or beyond the specified cylinder and head address is displayed along with a hex dump of its extent control block.

Display GROUP

Display the status and related information for one or all migration groups.

Syntax Display Group <*group-name*> </C | /E>

Parameters *group-name*

The name of the migration group. If this is omitted, then summary status of all groups will be displayed. Otherwise, detailed information will display for the specified group, including:

- High level error messages
- Group status on all LPARs as reflected in the z/OS Migrator database
- Group definition statements

/C

Display summary counts of datasets and tracks.

/E

Error message detail for activation and diversion.

Display HOSTS

Display the status and last heartbeat time-stamps for each participating migration server.

Syntax `Display HOSTS`

Parameters None

Note: The displayed status includes whether the server is considered IDLE or has been Removed.

Display JOBNames

Display the names of address spaces that still have source datasets allocated and that may be blocking group completion.

Syntax `Display JOBNames group-name | data-set-mask`

Parameters *group-name*

Specify the name of the group you want to activate. This is the name specified in the GROUP(name) group definition parameter.

data-set-mask

A dataset name mask. If specified, all datasets of any group will be scanned for comparison to the mask.

Display PATHS

Display path concurrency statistics for Logical Migration or Volume Migration. The values displayed should be evaluated relative to the MAXCIO and MAXDIO values.

Syntax `Display PATHS [RESET] [LM]`

Parameters RESET

Reset the statistics after the display.

LM

Display statistics for paths used by Logical Migration.

The default is statistics for Volume Migration.

Display PROC

Display migration status information for the unique device pairs of a group.

Note: This command displays diagnostic information not of general use without the assistance of EMC Customer Support.

Syntax `Display PROC group-name`

Parameters *group-name*

Specify the name of the group you want to activate. This is the name specified in the GROUP(name) group definition parameter.

Display TRACE

Display the control parameter settings and current status of z/OS Migrator user trace.

Note: The z/OS Migrator trace facility writes user trace records to GTF.

Syntax `Display TRACE`

Parameters None.

Display VOLumes

Display volume usage across groups or within a group.

Syntax `Display VOLumes [vmask] [/G group] [/D dsn] [/M mode] [/S
| /T | /B]`

Parameters

vmask

A volume mask for selecting the volume or volumes to be reported.

/G group

Limit the display to the volumes in use by the specified group.

/D dsn

Limit the display to the volumes in use by datasets matching the specified dataset name mask.

/M mode

Limit the display to extents that have a current status matching the specified mask. The */M* parameter of [“Display DSN” on page 209](#) provides details.

/S

Display only volumes and datasets currently serving as migration source.

/T

Display only volumes and datasets currently serving as migration target.

/B

Display volumes and datasets of both source and target usage. This is the default.

z/OS Migrator diagnostic commands

These commands can be used to aid in problem solving, if necessary.



CAUTION

These commands should only be used under the direction of EMC Customer Support.

COMPARE

Compare source and target data. This command is only of use during the Mirror phase.

Syntax

COMPARE *datasetname*

Parameters

datasetname

Specify the name of the dataset you want to compare. z/OS Migrator will do a track comparison of the dataset between source and target. The result will be reported through WTO. Additional information may be written to the EFMMCMR optional sysout file.

DEBUG

Reset the DEBUG switches.

Syntax DEBUG xxxxxxxxxx, xxxxxxxxxx

Parameters xxxxxxxx, xxxxxxxx

Specify the DEBUG bit mask. The bitmap has the following structure:

1. 1... .. CONFIG	2. 1... .. SYNC	3. 1... .. IOTR
.1... .. CPLDS	.1... .. COMM	.1... ..
.1. TASKMGT	.1. PCR1	.1.
...1 ... STGMGR	...1 ... SDDF	...1 ...
... 1... DIVERT	... 1... DTERM	... 1...
... .1..1.. TRANS1..
... .1. DB1.1. HERROR
... ...11 DBHB1 HOK
4. 1... .. HM00	5 & 6 ARE LIKE 4, LEFT TO RIGHT:	
.1... .. HPRV	BYTE 5 - IOSM, SPMW, IORQ, IOTM,	
.1. HTIO	IOVM, IODR, IERROR, IOK	
...1 ... HPST	BYTE 6 - INTERFACEMGT, HSIO_FE, HSIO_BE, HMIH	
... 1... HPRE	HDIE, HEOS, HERROR, HOK	
... .1..		
... .1. HERROR		
... ...1 HOK		
7. 1... 1... CPTK	8. UNUSED
.1...1..	
.1.1.	
...11	

DUMP

Schedule an SVC dump of the server address space.

Syntax DUMP

INTERFACE

Enable or disable system interfaces on demand.

Syntax INTERFACE <interface-name|ALL> ACTIVATE|DEACTIVATE

Parameters *interface-name* | ALL

Specify the name of the z/OS Migrator system interface or ALL as directed by EMC Customer Support.

ACTIVATE | DEACTIVATE

Specify ACTIVATE to enable the system interface and DEACTIVATE to disable it.

RELEASE

Issue DEQ (dequeue) and reset IOSLEVEL for a migration device.

Syntax RELEASE *ccuu*

Parameters *ccuu*

Specify the device to be released.

TRACE

Activate or deactivate z/OS Migrator internal trace to GTF.

Syntax TRACE ON|OFF, EID=xxxx, FID=xx

Parameters EID=xxxx

Specify the EVENT-ID for coordination with the IPCS GTF subcommand for extracting and printing the resulting user trace records, for example, EID=123.

FID=xx

Specify the FORMAT-ID for coordination with the IPCS GTF subcommand, for example, FID=00. (FID=00 is the default.)

VERify

Trigger a self-evaluation cycle for a migration group. This causes the status information for all resources within the group to be surveyed and the group status possibly being reset accordingly.

Syntax `VERify group-name`

Parameters *group-name*

Specify the name of the migration group to be verified.

Defining z/OS Migrator Groups via Text Editor

This chapter lists and describes the commands available via the z/OS Migrator Command Line Interface (CLI). Topics are:

- ◆ Overview 222
- ◆ General syntax rules 222
- ◆ Syntax for defining logical migration groups 223
- ◆ Defining groups for a volume migration 234
- ◆ Syntax for defining mirror groups 235
- ◆ Syntax for defining migrator groups 240

Overview

You may define z/OS Migrator migration groups using any text editor. Once you have defined several groups with the assistance of the ISPF Monitor, you may find this method easy to use when defining or extending groups that are like other groups you have previously defined.

Groups defined using a text editor may be stored in the configuration PDS and promoted using the ISPF Monitor. They may also be promoted using the EFMMBAT utility program, either from the configuration PDS or from an in-stream definition. [Chapter 10, "z/OS Migrator Batch Utility"](#) on [page 295](#) provides more information about the EFMMBAT utility.

This chapter describes:

- ◆ The group definition syntax for logical migrations beginning on [page 223](#).
- ◆ The group definition syntax for volume mirroring and migrations, beginning on [page 234](#).

General syntax rules

- ◆ Comments begin with an asterisk (*) in column one and may appear anywhere.
- ◆ Parameters are generally of form keyword(value) and are separated by white-space.
- ◆ Continuations are indicated with a dash (-) after the last keyword or value on a line.
- ◆ Lists can be comma or white-space separated, for example, VOL(TGT001 TGT002 TGT003).

Syntax for defining logical migration groups

A group definition is a collection of statements which can include comments. The first non-comment statement must be the GROUP statement. The GROUP statement defines certain global settings for the group.

Following the GROUP statements can be a variable number of SET statements defining the source and target dataset name masks and other parameters associated with each pair. An optional SRCVLIST statement is also available for defining Source Volume Lists. Any named SRCVLIST must be defined before the SET statement that refers to it.

The following is a general example:

```
* GROUP NAME:TEST3                MEMBER GENERATED BY  JDOE1
* ON THURSDAY, 25 OCT 2007 AT 18:20:04 FROM: LOGMIGR CONTROL COMMANDS
GROUP(TEST3) MODE(LMIGR()) REPLACE(YES) REUSE(NO) TOLALLOF(Y) -
  DEBUG(EXTRA)
  SRCVLIST ABC(SYM719)
  SRCVLIST ABC(SYM71A)
  SRCVLIST ABC(SYM71B)
  SRCVLIST ABC(SYM71C)
  SET SOURCE(DSN(XYZTEST.**)) -
    EXCLUDE(XYZTEST.DONT.** ,XYZTEST.UDD.** , -
            XYZTEST.VSAM2.DB.** , -
            XYZTEST.AAA.** ,XYZTEST.XYZ.** , -
            XYZTEST.LMIGRX09.** ) -
    SRCVLIST(ABC) -
    VOL(U6J1AD) -
    ) -
  TARGET(VOL( -
          SYM719 -
          SYM71A -
          SYM71B -
          SYM71C -
          ) DSN(XYZTEST2.**)) -
  TRACE(Y)
*
  SRCVLIST ABC2(SYM71A)
  SRCVLIST ABC2(SYM71B)
  SET SOURCE(DSN(XYZTEST.LMIGRX09.**)) -
    SRCVLIST(ABC2) ) -
  TARGET(DSN(XYZTEST2.**)) -
    VOL(SYM71C SYM719)) -
  CONSALL
*
  SET SOURCE(DSN(XYZTEST.ABC4)) -
  TARGET(DSN(XYZTEST2.**)) -
    VOL(SYM719, SYM71A, SYM71B))
*
  SET SOURCE(DSN(XYZTEST.GDG.DATASET.G0003V00)) -
  TARGET(DSN(XYZTEST2.**)) -
    VOL(SYM719, SYM71A, SYM71B))
```

GROUP

The GROUP statement delineates and defines the group. All subsequent SET statements define the datasets in the group.

Syntax

```
GROUP (groupname) MODE (LMIGR) [REPLACE (Y|N)] -  
      [TOLERATE_ALLOCATION_FAILURE (Y|N)] -  
      [SELECTMULTI (ALL|ANY|FIRST)] [CONSVOL|CONSALL] -  
      [DEBUG (ALL|TRACE|DUMP|EXTRA|ERROR)]
```

Note: Currently, CONSALL is not a supported parameter, but may be added in a future release.

Parameters

groupname

Specifies a one- to eight-character name by which this group will be known, once promoted. All commands that operate on the group reference it by this name.

MODE(LMIGR)

Denotes a Logical Migration group.

REPLACE

This parameter causes z/OS Migrator to delete and replace an existing dataset with the target name. Options are Y and N. The default is N.

Note: If the default N is accepted or specified, then an existing dataset with the target name for any source/target pair will abort the migration of that dataset. If TOLALLOF (N) is specified, the entire group activation will be aborted.

TOLERATE_ALLOCATION_FAILURE (or TOLALLOF)

Options are Y and N. The default is N. The default behavior of z/OS Migrator - TOLALLOF(N) - is to terminate group activation if any error is encountered in the discovery of source datasets or the allocation of target datasets.

If Tolerate_Allocation_Failure(Y) is specified in the group definition, then group activation will continue in spite of such errors and all successfully discovered and allocated datasets will be fully serviced. The failing datasets will be noted in messages to the various z/OS Migrator message outlets and display facilities. These failing datasets will be marked as skipped and will not take part in the processing for the group.

Examples of the errors that are tolerated include:

- Source dataset is not properly cataloged.
- Cataloged source dataset is not actually found on a referenced volume (not including current "candidate" volumes).
- Source dataset is of unsupported type (for example, page dataset), or format (for example, unmovable).
- One or more source volumes are offline.
- Any error is encountered in attempting to allocate the target.

If `Tolerate_Allocation_Failure(Y)` is specified, then:

- All components of a complex entity (that is, cluster or sphere) will be skipped if any component is found to have an error.
- If the group would have had only one dataset or complex entity and that is found to have an error, then the group activation will fail with reason code 1029 - "No datasets selected."

Methods of displaying dataset errors include the "E" line command in the Interact with Promoted Groups ISPF panel and the console command "Display Group groupname /E", both of which show the error messages associated with group activation.

In most cases, the "skipped" datasets will be displayed in the detailed dataset displays and noted as "Skipped". The console command "Display DSN groupname /S" can also be used to display just the skipped datasets.

Note: Certain generic failures do not result in a dataset block being created and, therefore, cannot be displayed with the Display DSN command. These should still be viewable in the message displays previously mentioned.

SELECTMULTI (ALL | ANY | FIRST)

This specification can be made at the GROUP level or at the SET level. A specification at the SET level overrides the GROUP specification for any datasets selected by that SET statement.

SELECTMulti allows selection of datasets to be controlled relative to a Source Volume List in the following ways:

- SELECTM(ALL) — The dataset will be selected if all volumes of the dataset are in the specified list.
- SELECTM(ANY) — The dataset will automatically be included (the default behavior) if any volume is in the list.

- **SELECTM(FIRST)** — All volumes of the dataset will be included if the first volume of the dataset is in the list.

CONSVOL or CONSALL

This optional specification can be made at the GROUP level or at the SET level. A specification at the SET level overrides the GROUP specification for any datasets selected by that SET statement.

Select one or the other. They are mutually exclusive. If this parameter is omitted, the default of no extent consolidation will take place.

Select CONSVOL to allocate the target space in consolidated extents on the corresponding target volumes for each source volume. The target dataset will still require as many volumes as the source.

Select CONSALL to consolidate the target space on as few volumes as possible for each source dataset. This results in a reduction of the number of volumes required for each dataset where possible.

Note: Currently, CONSALL is not a supported parameter, but may be added in a future release.

DEBUG(ALL | TRACE | DUMP | EXTRA | ERROR)

ALL	Produces the TRACE and DEBUG information needed for most situations.
TRACE	Produces normal TRACE output.
DUMP	Produces normal TRACE and DEBUG output.
EXTRA	Produces all possible TRACE and DEBUG information.
ERROR	Records some TRACE and DEBUG output in memory. This information is placed in the message log if an error occurs.

SRCVLIST

The SRCVLIST (or SOURCE_VOLUME_LIST) statement can be used to define lists of volumes to serve as a filter for source dataset selection. Source datasets are identified by applying the selection mask against the contents of the VTOC for all specified volumes. Note that once a dataset has been selected, the entire dataset will be migrated even if it extends to volumes not included in the source volume list.

Multiple source volume lists may be defined and referenced by different SET statements. The only requirement is that the source volume list be defined before the SET statement that references it.

Syntax

```
SRCVLIST name(volume1 volume2 ...)
```

Parameters

name

Specifies a one- to eight-character name for the source volume list. This name will be referenced by one or more subsequent SET statements. Multiple SRCVLIST statements may be specified having the same name and the definition of the list will be cumulative of all volumes.

volume list

Each statement may contain up to 127 six-character volume serial numbers or volume serial number masks. Use an '*' to match any number of characters; a '?' to match a single character.

SET

The SET statement defines the rule for selection of source datasets and naming of the corresponding target datasets during the migration. Once the migration is complete, each source dataset/target dataset pair will swap names.

Syntax

```
SET ALLOCSEQ(DATASET | NONE | SIZE) -
  TRACE(YES | NO) SPHERE(YES | NO) -
  [SELECTMULTI (ALL|ANY|FIRST)] [CONSVOL|CONSALL] -
  SOURCE(DSN(src-mask-list) -
    SRCVLIST(name) | VOLUME(src-vol) | STGCL(stgclass list) -
    EXCLUDE(exc-mask-list)) -
  TARGET(DSN(mask) -
    VOLUME(tgt-vol-list) | STGCL(sms stgclass))
SET ...
```

Note: Currently, CONSALL is not a supported parameter, but may be added in a future release.

Parameters**ALLOCSEQ**

Specifies the allocation sequence, with the following options:

DATASET	Specifies to process datasets in ascending name sequence.
NONE	Specifies to process datasets in the order in which they are selected for processing. This may appear random. This is the default setting.
SIZE	Specifies to process datasets in descending size sequence. The largest datasets are processed first and the smallest are processed last.

TRACE

If selected, z/OS Migrator will produce diagnostic output during the migration cycle; no output is the default setting. Each dataset defined for the group may use a different setting.

SPHERE

This parameter specifies whether, for any VSAM cluster copied, all associated AIX clusters and paths are to be copied. Individual names of sphere components do not need to be specified. Only the base cluster is eligible to be migrated if SPHERE(N) is specified, individual AIX clusters are ineligible. Each SET statement comprises a complete selection specification, and each may have a different SPHERE value. Valid input characters are YES and NO:

- YES specifies that all associated AIX clusters and PATHs are to be copied.
- NO specifies that only the selected base cluster will be copied.

Note: YES is the default setting; however, when migrating non-VSAM datasets, z/OS Migrator ignores the SPHERE parameter. If SPHERE(Y) is specified, then the selection of the base cluster is determined by the source selection mask, however, multiple related components (KSDS | ESDS + AIX + PATH) will be migrated in a single action, and they may not have a common naming structure. Unless there is a single AIX whose name is base-cluster-name.AIX, then the RENAMEUnconditional parameter must be used to specify the naming rule for the associated target PATH and AIX components.

SELECTMULTI

These specifications can be made at the GROUP level or at the SET level. A specification at the SET level overrides the GROUP specification for any datasets selected by that SET statement.

SELECTMulti allows selection of datasets to be controlled relative to a Source Volume List in the following ways:

- SELECTM(ALL) — The dataset will be selected if all volumes of the dataset are in the specified list.
- SELECTM(ANY) — The dataset will automatically be included (the default behavior) if any volume is in the list.
- SELECTM(FIRST) — All volumes of the dataset will be included if the first volume of the dataset is in the list.

CONSVOL or CONSALL

These optional specifications can be made at the GROUP level or at the SET level. A specification at the SET level overrides the GROUP specification for any datasets selected by that SET statement.

Select one or the other. They are mutually exclusive. If this parameter is omitted, the default of no extent consolidation will take place.

Select CONSVOL to allocate the target space in consolidated extents on the corresponding target volumes for each source volume. The target dataset will still require as many volumes as the source.

Select CONSALL to consolidate the target space on as few volumes as possible for each source dataset. This results in a reduction of the number of volumes required for each dataset where possible.

Note: Currently, CONSALL is not a supported parameter, but may be added in a future release.

SOURCE

The SOURCE parameter is used to define the selection criteria for the datasets to be migrated. Each source and target dataset must resolve to compatible devices. These devices may be within the

same storage array, or they may be on different storage arrays. All source and target dataset(s) must reside on devices that are currently online and accessible.

Note: The source dataset name must be different from the target dataset name. Neither the source mask nor any source dataset name should be a substring of the target mask. The source and target masks should have the same number of specific qualifiers. The specified qualifiers of the target mask replace the corresponding qualifiers of the source dataset name one for one. Therefore, it is not possible to insert a new qualifier.

[“Selecting datasets” on page 83](#) provides detailed information on dataset masking rules. [“Restrictions on storage subsystems and dataset types” on page 93](#) lists both supported and unsupported dataset types.

Subparameters for SOURCE are as follows:

- DSN(*mask list*)
A list of dataset name masks separated by commas or white space. All additional subparameters of the SOURCE parameter and all subparameters of the corresponding TARGET parameter will apply to each specified mask.
- SRCVLIST(*name*)
Optionally, specify the name of a previously defined source volume list. The datasets to be migrated will be discovered by reading the VTOCs of all specified volumes and applying the specified dsname masks.
- VOL(*volume list*)
As an alternative, you may specify up to 16 volume serial numbers to serve as a self-contained, unnamed source volume list.

Note: If neither SRCVLIST nor VOL is specified, then the datasets to be migrated will be discovered solely by searching the catalog system.

- STGCL(*sms stgclass list*)
You may specify up to 16 SMS storage class names to serve as an additional filter for selecting the source datasets. Only datasets of the specified storage class will be selected.

- EXCLUDE(*exclude mask list*)

This parameter may be used to prevent certain datasets that would otherwise have been selected from being migrated. Enter up to 127 dsname masks to be used as exclusion rules.

TARGET

The TARGET parameter is used to specify the rule for naming the target datasets that correspond to the source datasets selected by this particular SET statement. It also specifies the target volumes or the target storage class.

Once the migration is complete, the target and source dataset names will be "swapped" leaving the target with the name of the original source dataset and the source with the name generated by the Target DSN rule.

Note: You cannot specify the same dataset name on both the SOURCE and TARGET parameters. Neither the source mask nor any source dataset name should be a substring of the target mask. The source and target masks should have the same number of specific qualifiers. The specified qualifiers of the target mask replace the corresponding qualifiers of the source dataset name one for one. It is not possible to insert a new qualifier.

Where possible, you may want to consider having a common high level qualifier for the target masks and create a distinct user catalog for this alias.

Subparameters for TARGET are as follows:

- DSN(*mask*)

Specify a dsname mask to be applied against each source dataset selected by the SOURCE parameters to create the target dsname.

- VOL(*volume list*)

Specify from one to sixty volume serial numbers to serve as a pool for allocating the target datasets. Make sure that you avoid including volumes that may also be the source of any migration running concurrently. Overlapping source and target volume usage can result in deadlock conditions.

- STGCLS(*sms storage class*)

Enter the storage class for the target volume(s). This parameter requests that the target volumes be selected using the specified SMS Storage Class. If you specify a storage class,

you are not required to provide source volume serial number for the VOL parameter. Be aware that your SMS ACS routine may place the target dataset in a storage class other than that specified by this parameter. As with all SMS datasets, a specified storage class is only a suggestion to SMS, and may or may not be accepted by SMS.

RENAME Unconditional

RENAMEUnconditional specifies the rules for naming the target AIX and PATH components related to a selected base cluster when SPHERE(Y) is specified. Selection of spheres should be done using a SOURCE(DSN(*mask*)) suitable for selecting the base cluster. The additional components will automatically be identified using the catalog information.

Since the alternate index and path components may not match the general naming pattern of the base source and target cluster names, RENAMEU allows additional rules to be specified for building their corresponding target names. The basic mechanism is that when an *oldnamemask* matches a related source component, the *newnamemask* is used to build the corresponding target name. The *px* parameter provides a simplified solution where substitution of a replacement high-level qualifier is sufficient to generate unique names.

Note: RENAMEUnconditional is similar in usage and syntax to the DFDSS RENUNC parameter.

Syntax

```
RENAMEUnconditional(...)
  RENAMEUnconditional(px) |
  RENAMEUnconditional((px) (oldnamemask, newnamemask)...) |
  RENAMEUnconditional((oldnamemask, newnamemask)...)
```

Parameters *px*

Specifies a prefix used to replace the first-level qualifier of the source dataset name when building the target dataset name. This parameter is optional, but if specified, it must be the first parameter in the list of sub-fields. The prefix is used only if the (*oldnamemask*, *newnamemask*) parameters are not specified or no *oldnamemask* filter matches the source dataset name.

oldnamemask

Specifies a mask to be used to compare against source dataset names.

newnamemask

Specifies a mask used to derive the target dataset name when the existing dataset name matches the corresponding *oldnamemask*.

Example

```
GROUP (SPHERE) MODE (LMIGR ()) REPLACE (Y)
  SRCVLIST LIST1 (SRC001)
  SET TRACE (YES) SPHERE (YES) -
    SOURCE ( -
      DSN (TEST.XYZ.C01.VUDD0B3.@04K.KSDS0001.***) -
      SRCVLIST (LIST1) -
    ) -
    TARGET ( -
      VOLUME (TGT001) -
      DSN (TEST2.XY.***) -
    ) -
  RENAMEU (TEST2, (TEST.AAA.C01.VUDD0B3.@04K.KSDSAAAA.***, -
    TEST2.AA.C01.VUDD0B3.@04K.KSDSAAAA.***) )
```

Any sphere components matching the source mask TEST.AAA.C01.VUDD0B3.@04K.KSDSAAAA.** will have a corresponding target dataset name generated using the target mask TEST2.AA.C01.VUDD0B3.@04K.KSDSAAAA.**.

Any sphere components not matching the source mask (such as TEST.KSDSPATH) will have the first index level of the generated target dataset name changed to TEST2 (such as TEST2.KSDSPATH).

- ◆ RENAMEUnconditional is needed when the naming conventions of components of a complex VSAM cluster (commonly the AIX components) do not match the naming convention of the base cluster.
- ◆ RENUNC is a valid alias for the RENAMEUnconditional parameter.
- ◆ The RENUNC dataset name masks follow DFDSS rules for RENUNC.
- ◆ A maximum of 127 (*oldnamemask*, *newnamemask*) pairs may be specified.

Defining groups for a volume migration

When you establish your z/OS Migrator application, you must set configuration parameters to define your configuration and processing needs.

All Volume Mirror and Volume Migrator operations are initiated and controlled on a volume-group basis. Each volume group has a user defined name and consists of two basic kinds of statements:

- ◆ Statements that define group characteristics.
- ◆ Statements that define the device pairs that make up the group.

Definitions for Mirror and Migrator groups are similar but have different specific parameters.

[“Syntax for defining mirror groups” on page 235](#) defines those parameters that apply to mirror groups.

[“Syntax for defining migrator groups” on page 240](#) defines those parameters that apply to migrator groups.

Syntax for defining mirror groups

The following parameters define the volume groups for mirroring. This is done by way of the ISPF interface. The parameter sets are initially stored as members of the configuration PDS. After promotion, a canonical representation of the group is stored in the z/OS Migrator database.

Before promotion, you may view or edit the volume group members as needed. You may retrieve, revise, delete and re-promote promoted volume groups.

GROUP

The GROUP statement delineates and defines the group. Subsequent PAIR statements define the devices that make up the group.

Syntax

GROUP(groupname) MODE(FASTMIRROR(subparameters))

Parameters

groupname

Specifies a one- to eight-character name by which this group will be known, once promoted. All commands that operate on the group will reference it by this name.

MODE(FASTMIRROR | FMIR)

Denotes a Volume Mirror group.

subparameters

Note: When the parameter contains more than one subparameter, separate the subparameters by commas.

One or more of the following:

- FASTCOPY

Optionally, copy only the allocated tracks during the synchronization phase.

- DEACTONSUSP | DOS

Automatically deactivate mirroring for the group if mirroring is suspended. You can suspend mirroring in response to a command, or if any exceptional condition is detected in servicing the mirror I/O.

Remember that SUSPEND leaves the device pairs in a state in which mirroring may be resumed with the RESUME command. To reactivate mirroring on a *deactivated* group, an ACTIVATE command is required, and the group will have to be completely re-synchronized.

- NoAutoVary

The default behavior of Volume Mirror groups is to vary the mirror device automatically to match the online status of the primary device. After the device pair is synchronized, this service is ongoing.

To vary the mirror online, you must have specified the NEWVOLSER or VOLSER PREFIX value to allow z/OS Migrator volume Mirror to simulate a unique volume serial number for the mirror.

In this case, although physically the mirror device has the same serial number as the primary device, z/OS Migrator returns the simulated serial number in response to all attempts to read the volume label. The simulated serial number is known as the apparent volser. Select the NoAutoVary option to override z/OS Migrator vary management.

Use NoAutoVary with care. When the primary device is online and has PAV aliases associated (that is, bound) with it, you should also allow the mirror device to come online and have aliases bound. In this case, WLM can monitor and reassign PAV aliases as needed. If the mirror device is not varied online, aliases are not bound and longer IOS queue times to the mirror device may result. This, in turn, could affect application I/O response times.

Note: You can use the z/OS VARY operator command to change the online/offline state of the mirror device. If you do not specify the NEWVOLSER or NEWVOLP values, the mirror device does not vary online if the primary device is online.

PAIR

One or more PAIR statements define the source-target volume pairs. You can define one or more single pairs or ranges of consecutive pairs. Devices may be defined via ccuu device numbers or by volume serial numbers.

Syntax PAIR PRIMARY(ccuu[,n]) | PVOL(volser) MIRROR(ccuu) –
VERMIR | VERTARG ([NONE | EMPTY][,volser]) –
NEWVOLSER(volser) | NEWVOLP(pref)

Parameters PRIMARY(ccuu [,n])

The primary unit, that is, the unit that is to be mirrored. The first subparameter specifies the device address of the primary unit. If you are defining a range of devices, use the second subparameter to specify the number of consecutive device addresses whose definition should be generated by this PAIR statement. The addresses begin with the *ccuu* specified in the first subparameter and continue through the next *n* devices. You may not have gaps in the range.

PVOL(volser | volmask)

An alternate way to specify the primary volume.

Note: You can only define a single pair of devices when you use PVOLser (or PRIMARYVOLser). You cannot define a range of devices.

MIRROR(ccuu)

The unit that will be the mirror for this primary. If you specify a unit count (*n*) on the PRIMARY parameter, then the mirror units begin with the unit specified and continue through the next *n* devices. You may not have gaps in the range.

VERMIR | VERTARG

Optionally, specify the following mirror verification options:

- *NONE*

Do not do any verification on the target volume. This is the default behavior.

- EMPTY

Verify that the target volume is empty except for the VTOC and possible volume index.

- volser

Verify that the target volume has the specified volume serial number before beginning the copy.

- volmask

Verify that the target volume has a volume serial that matches the mask specified. A percent (%) character in any position of the mask matches any character in that position.

To check ranges of devices, you can substitute the mask character for volser characters. For example, VERTARG (EMPTY, UWD%%%) causes the target volumes to be checked to see that they are essentially empty, and have serial numbers with UWD in the first three positions.

Note: The z/OS Migrator wildcard character is the percent (%) character. A percent character in any position is a wildcard match for that position.

NEWVOLSER

The mirror volume may be varied online to take advantage of PAV access for performance. While the mirror volume will physically have the same volume serial number as the primary volume after synchronization, z/OS Migrator volume Mirror supplies the specified volume to the system on volume label reads in order to circumvent the duplicate volume condition. This is known as the apparent volser.

The mirror device may automatically be varied online and offline by the FMM Server, as long as you do not specify NoAutoVary, or may be varied online and offline by the z/OS VARY operator command.

You can use the % mask character in NEWVOLSER. The % mask causes the mirror volser to be changed to the specified value with the corresponding characters from the primary volser replacing any mask characters.

For example, mirroring from UWC012 to UWD012 with a parameter value of NEWVOLSER(%@%%) results in an apparent new volser value of U@C012.

NEWVOLP

Similar to NEWVOLSER but allows the new volume serial to be composed of the specified characters (up to 5) followed by the remaining characters of the source volume serial.

For example, a value of NEWVOLP(U@C) with mirror Symmetrix device number 0012 changes device UWC012 into U@C012.

Wildcards are not supported with NEWVOLPrefix.

Remember that, physically, the mirror volume will always have the same volume serial number as the primary.

Syntax for defining migrator groups

The following parameters define the volume groups for migration. This is done by way of the ISPF interface. The parameter sets are initially stored as members of the configuration PDS. After promotion, a canonical representation of the group is stored in the z/OS Migrator database.

Before promotion, you may view or edit the group members as needed. You may retrieve, revise, delete and re-promote promoted groups.

GROUP

The GROUP statement delineates and defines the group. Subsequent PAIR statements define the devices that make up the group.

Note: In defining groups, you can use the SWAPOPT statement, described on [page 243](#), to specify parameters that override the corresponding default values set in the global configuration parameters or in SCF(DAS) itself. The SWAPOPT statement, when used, should be placed between the GROUP and PAIR statements.

Syntax	GROUP(groupname) MODE(MIGRATE(action [,FASTCOPY]))
Parameters	<p>groupname</p> <p>Specifies a one- to eight-character name by which this group will be known, once promoted. All commands that operate on the group will reference it by this name.</p> <p>MODE(MIGRATE(</p> <p>Denotes a Volume Migrate group.</p> <p>action</p> <p>One of the following, to specify what action should be taken to conclude the migration:</p> <ul style="list-style-type: none"> • SWAP <ul style="list-style-type: none"> At completion, perform a device swap from the source to the target device. • SPLIT(ON)

Once migration is complete, terminate the operation and mount the target volume. NEWVOLSER must be supplied on each PAIR statement.

- SPLIT(OFF)
At completion, terminate the migration and leave the target offline.
- CONSTANTcopy
At completion of the initial migration, continue to periodically copy any modified tracks from source to target.

FASTCOPY

Optionally, copy only the allocated tracks during the synchronization phase.

PAIR

One or more PAIR statements define the source-target volume pairs. You can define one or more single pairs or ranges of consecutive pairs. Devices may be defined via ccuu device numbers or by volume serial numbers.

Syntax PAIR SOURCE(ccuu[,n]) | SOURCEVOL(volser) TARGET(ccuu) –
VERTARG([NONE | EMPTY][,volser]) –
NEWVOLSER(volser) | NEWVOLP(pref)

Parameters SOURCE(ccuu [,n])

The source unit, that is, the unit that is to be migrated. The first parameter specifies the device address of the source unit. If you are defining a range of devices, use the second subparameter to specify the number of consecutive device addresses whose definition should be generated by this PAIR statement. The addresses will begin with the *ccuu* specified in the first subparameter and continue through the next *n* devices. You may not have gaps in the range.

SOURCEVOL(volser | volmask)

An alternate way to specify a single source unit.

Note: You can only define a single pair of devices when you use SOURCEVOLser. You cannot define a range of devices.

TARGET(ccuu)

The unit that will be the target for this source. If you specify a unit count (*n*) on the SOURCE parameter, then the target units begin with the unit specified and continue through the next *n* devices. You may not have gaps in the range.

VERTARG

Optionally, specify the following target verification options:

- *NONE*

Do not do any verification on the target volume. This is the default behavior.

- EMPTY

Verify that the target volume is empty except for the VTOC and possible volume index.

- volser

Verify that the target volume has the specified volume serial number before beginning the copy.

- volmask

Verify that the target volume has a volume serial that matches the mask specified. A percent (%) character in any position of the mask matches any character in that position.

To check ranges of devices, you can substitute the mask character for volser characters. For example, VERTARG (EMPTY, UWD%%%) causes the target volumes to be checked to see that they are essentially empty, and have serial numbers with UWD in the first three positions.

Note: The z/OS Migrator wildcard character is the percent (%) character. A percent character in any position is a wildcard match for that position.

NEWVOLSER

Specify the new volume serial for the target. This is required for SPLIT(ON) or to terminate a CONSTANTCOPY operation with the SPLITMOUNT command (the equivalent of SPLIT(ON)).

You can also specify the volser in other circumstances as well, but it will be ignored for the SWAP case, which requires that the target volume have the same volser as the source volume.

You can use the % mask character in NEWVOLSER. The % mask causes the target volser to be changed to the specified value with the corresponding characters from the source volser, replacing any mask characters.

For example, migration from UWC012 to UWD012 with a parameter value of NEWVOLSER(%@% %%) results in a new volser value of U@C012.

NEWVOLP

Similar to NEWVOLSER but allows the new volume serial to be composed of the specified characters (up to 5) followed by the remaining characters of the source volume serial.

For example, a value of NEWVOLP(U@C) with target Symmetrix device number 0012 changes device UWC012 into U@C012.

Note: Wildcards are not supported with NEWVOLPrefix.

SWAPOPT

The SWAPOPT statement allows you to specify parameters that override the corresponding default values set in the global configuration parameters or in SCF(DAS) itself. The SWAPOPT statement, if provided, should be placed between the GROUP and PAIR statements.

Syntax SWAPOPT(subparameters)

Specify one or more of the following options to override default AutoSwap settings for the group:

Parameters [NO]BYPASSCONCURRENTCOPY | BYPCC

Specifies whether Concurrent Copy may be bypassed by swap processing using the BypassConcurrentCopy option. This option causes active concurrent copy jobs to fail at the time of swap.

[NO]BYPASSSYSTEMSCOUNT | BSC

Specifies whether to bypass the requirement that all other LPARs with connectivity to Symmetrix have EMCSCF running.

[NO]BYPASSSNAPSESSION | BSS

Indicates whether EMC Snap sessions may be bypassed by swap processing.

VOLUMEPrefix | VOLP=PP

Specifies a two-character prefix to be used to modify the source device's volser after a successful swap. This can be used to assure that if the source volume is mounted, it has a unique volser from the migration target. To bypass, use "--" for *pp*.

The volume prefix value is applied to the z/OS device number (CUU) to form the new volser. For example, if the z/OS device number of the source device is 055C and you select the VOLP(ZZ) value, the resulting volser is ZZ055C.

CFW=NO | OFF | RESume | BYPASS | OFFVAL

Controls cache fast write (CFW).

- **NO**
Indicates that CFW must be inactive for the SWAP to occur.
- **OFF**
z/OS Migrator will automatically turn off CFW on the source device controller, if it is active.
- **RESume**
CFW will automatically be turned off on the source controller but will be turned on a completion if it had been found active.
- **BYPASS**
Indicates that CFW processing will be ignored.
- **OFFVAL**
CFW will automatically be turned off during group validation.

Note: If you do not specify a value using the SWAPOPT parameter SCF(DAS) sets the values for these options.

This chapter provides sample usage scenarios for z/OS Migrator.
Topics are:

- ◆ Overview 246
- ◆ Setting session options 246
- ◆ Migrating a single dataset, non-SMS 249
- ◆ Migrating a single multi-volume dataset, non-SMS 252
- ◆ Migrating multiple datasets, non-SMS 255
- ◆ Migrating multiple datasets, SMS volumes 261
- ◆ Volume Mirror example 265
- ◆ Volume Migrator examples 275

Overview

Occasionally, changes in the usage of a dataset put it in contention with other datasets located on the same volume or volumes. z/OS Migrator gives you the capability to selectively move problematic datasets to other volumes where their activity will not come in conflict with other datasets. z/OS Migrator also gives you the capability to mirror volumes from one Symmetrix device to another and to migrate volumes from any device to a Symmetrix device.

This chapter gives you a selection of examples.

Setting session options

The first time you enter the z/OS Migrator dialog you are required to visit the Set User Session Options panel to specify various control parameters.

1. Type **S** on the z/OS Migrator main panel and press **Enter**.

```

                                EMC z/OS Migrator Migration

Function Number                    11/08/07 11:07:44
or Command ==> 2                  SMFID: X04
                                   CPMX: @$

                                z/OS Migrator 3.0.0 FUNCTIONS

1 Manage Groups                    5 Display Host Messages
2 Monitor Promoted Groups          6 System Change Summary
3 Display Configuration Information 7 Message Help
4 Operator Interface               8 Monitor Command Line Help
                                   9 Display Installation Options

                                   S Set User Session Options
                                   V View or Browse Log data set

                                   F1=Help F3=Exit

```

The Session parameters panel appears.

2. Set your session parameters.

```

                                Specify the Session parameters
Command ==>

Command Prefix . . . . . @$
Configuration PDS. . . . 'USER01.XYZMM.DW001'
Console Message Timeout. 1 (0-10)

Use Log PDS . . . . . Y (Y/N) Days to save . . 7 (1-30)
    Background Logging . . Y (Y/N)
    Use Browse or View . . V (B/V)
    Log PDS Name. . . . . 'USER01.EFMM300.TFMLOG'

ZOSMLoad Library . . . 'USER01.EFMM300.LINKLIB'
ZOSM Security Library .. 'USER01.EFMM300.SECCOM'
UNIT . . . SYSALLDA and VOLUME . . _____ (optional) for allocations

Stack Console output . . Y (Y/N)
Clear variables. . . . . N (Y/N)
Simulate Mode. . . . . N (Y/N)
Debug Mode . . . . . N (Y/N)

                                F1=Help    ENTER=Save Session    F3=Exit

```

The session parameters have the following meanings:

- **Command Prefix**
Used to communicate with the Volume Migrator service task.
This must match the prefix specified at run-time.
- **Configuration PDS**
Configuration members are stored in two hierarchical datasets. The Configuration PDS is used to store members that have not been promoted to use by the FMM server task.
- **Console Message Timeout**
This is the console message wait time in seconds from 0-10.
When using the Operator Interface and issuing a console command, this is the amount of time to wait for a response.
- **Use Log PDS**
Save messages to a PDS for subsequent review.
- **Background logging**
A value of "Y" enables logging that accumulates display data each time a panel is displayed. When this option is set, a logging member is created and display data is sent to both the panel and the log member. This option requires that "Use log PDS" be set to "Y".
- **Days to save**
Number of days to save the log message members.

- `Browse/View`
Use for viewing configuration members.
 - `Log PDS name`
Name of the optional log dataset.
 - `Product Load Library`
Name of the z/OS Migrator load library.
 - `Security Library`
Specify the library containing the output of the SYSOPTn job used to install the z/OS Migrator license code. Enter your PDS name in standard TSO format.
 - `UNIT`
The default unit for allocations.
 - `Volume`
Optional volume for allocations.
 - `Stack Console`
When using the operator Interface, set this option to accumulate output in the current dialog display. The newest output is displayed first. The command response stream is reset when you leave the operator Interface.
 - `Clear Variables`
Setting this value to Y causes the panel variables to be cleared when you move between panels. A value of N will retain any panel variables that you entered.
 - `Simulation Mode`
This mode allows you to create command simulations. Specify Y to enable simulation mode. When the commands are built, they are echoed back to the terminal, but NOT processed.
 - `Debug Mode`
Setting this value to Y enables the debug trace facility.
3. Press **ENTER** to submit your values.
 4. Press **PF3** to exit.

Note: Pressing **PF3** before you press **ENTER** discards your entries.

Migrating a single dataset, non-SMS

This section provides a scenario involving the migration of a single dataset that is not managed by SMS.

Storage environment

Non-SMS environment with the following source datasets to be migrated:

```
VOLSER=SRCVL1
  PROD.BANKING.INQ
  PROD.BANKING.INQ.CNTL
  PROD.BANKING.ATM
  PROD.BANKING.ATM.CNTL
  PROD.xxxxxxxx.Dyyddd (where xxxxxxxx can be
    anything, and Dyyddd is the date in Julian
    format (where ddd is the number of day,
    from 1 to 365))
VOLSER=TRGVL1
  (empty)
```

Note: If the GROUP has SPHERE set to Y, z/OS Migrator will capture the CONTROL and INDEX components; for example, PROD.BANKING.INQ.CNTL. [Table 16, “Source dataset options,” on page 146](#) provides more information.

Migration objective

To move the PROD.BANKING.INQ dataset to a different volume, because it is causing contention with other datasets on the volume.

Defining a Migration group

This is a simple migration. A single migration group can be used to achieve the movement.

1. Type **1** in the command field of the z/OS Migrator main panel and The Manage Group Members panel appears.
2. Type the **CREATE** command (or **CR**) on the Primary command line, and press **Enter**. The Build Type screen appears.

```

+----- Build Type -----+
|                               |
|   Select with an S a build type.   |
|   _ Build Migrate Member           |
|   _ Build Mirror Member            |
|   s Build Logical Migrate Member   |
|                               |
|   F3=Cancel                        |
+-----+

```

3. Type an **S** before the type of group you want to create and press **Enter**.

When you select Build Logical Migrate Member, the Define z/OS Migrator Group panel appears.

4. Type **MOVE1** as the Group Name.

```

                                Define z/OS Migrator Group
Command ==> NE                                Scroll ==> CSR

Primary Commands : EXit  NExt

Group Name . . . . MOVE1

Group Options . . . N Replace Existing Data Sets (Y/N)
                   Y Tolerate Allocation Failure (Y/N)

```

5. Type the **NEXT** command (or NE) on the Primary command line, and press **Enter**.
6. Type the Source Dataset Name and Source Volume as follows: **PROD.BANKING.INQ** and **SRCVL1**.

```

                                Define z/OS Migrator Group
Command ==> NE                                Scroll ==> CSR

Primary Commands : EXit  NExt

Group Name . . . . . MOVE1
Source Data Set Options . . N Trace (Y/N)      N AllocSeq (D/S/N)
                           . Y Sphere (Y/N)    N Rename UnConditional (Y/N)
                           . N Build Data Set Exclude list (Y/N)

Source Data Set Name/Mask . PROD.BANKING.INQ_____
Source Volume List Name . . _____
Storage Class . . . . . _____
Source Volume(s) . . SRCVL1 _____
                           _____
                           _____

```

Migrating a single multi-volume dataset, non-SMS

This section provides a scenario involving the migration of a single dataset that is not managed by SMS and that resides on multiple volumes.

Storage environment

Non-SMS environment with the following source datasets to be migrated. PROD.BANKING.ATM, PROD.BANKING.INQ and PROD.BANKING.HISTORY datasets are VSAM Key Sequence Datasets (KSDS). A KSDS dataset consists of multiple components that have catalog entries, but the dataset is referenced by its cluster name.

```
VOLSER=SRCVL1
  PROD.BANKING.HISTORY
  PROD.BANKING.INQ.CNTL
  PROD.BANKING.INQ
  PROD.BANKING.ATM.CNTL
  PROD.BANKING.ATM
VOLSER=SRCVL2
  PROD.BANKING.HISTORY
VOLSER=SRCVL3
  PROD.BANKING.HISTORY
VOLSER=SRCXYZ
  PROD.BANKING.HISTORY
VOLSER=TRGVL1
  (empty)
VOLSER=TRGVL2
  (empty)
VOLSER=TRGVL3
  Various other datasets
VOLSER=TRGVL4
  Various other datasets
```

Migration objective

To move the multi-volume dataset PROD.BANKING.HISTORY from the four volumes (SRCVL1, SRCVL2, SRCVL3, and SRCXYZ) upon which it currently resides, to the target volumes (TRGVL1, TRGVL2, TRGVL3, and TRGVL4).

Defining a Migration group

This is a simple migration operation. A single migration group can be used to achieve the movement, illustrated as follows:

1. Type the **CREATE** command (or CR) on the Primary command line of the Manage Group Members panel, and press **Enter**.
2. Type **MOVMULTI** as the Group Name.

```

Define z/OS Migrator Group
Command ==> NE                               Scroll ==> CSR
Primary Commands : EXit  NExt
Group Name . . . . MOVMULTI
Group Options . . . N Replace Existing Data Sets (Y/N)
                   Y Tolerate Allocation Failure (Y/N)

```

3. Type the **NEXT** command (or NE) on the Primary command line, and press **Enter**.
4. Type the Source Dataset Name as follows:
PROD.BANKING.HISTORY

```

Define z/OS Migrator Group
Command ==> NE                               Scroll ==> CSR
Primary Commands : EXit  NExt
Group Name . . . . . MOVMULTI
Source Data Set Options . . N Trace (Y/N)      N AllocSeq (D/S/N)
                           . Y Sphere (Y/N)    N Rename UnConditional (Y/N)
                           . N Build Data Set Exclude list (Y/N)
Source Data Set Name/Mask . PROD.BANKING.HISTORY_____
Source Volume List Name . . _____
Storage Class . . . . . _____
Source Volume(s) . . _____
                    _____

```

5. Type the **NEXT** command (or NE) on the Primary command line, and press **Enter**.

Migrating multiple datasets, non-SMS

This section provides a scenario involving the migration of several datasets that are not managed by SMS.

Storage environment

Non-SMS environment with the following source datasets to be migrated:

```
VOLSER=SRCVL1
  PROD.BANKING.ATM.CNTL
  PROD.BANKING.INQ
  PROD.BANKING.DCTR
VOLSER=SRCVL2
  PROD.BANKING.INQ.CNTL
  PROD.BANKING.ATM
  PROD.BANKING.DCTR
VOLSER=SRCVL3
  PROD.BANKING.ATM
  PROD.BANKING.INQ
VOLSER=TRGVL1
  (empty)
```

Migration objectives

1. To move the production banking datasets starting with PROD.BANKING.INQ and PROD.BANKING.ATM, to a single volume to off-load the volumes upon which the datasets currently reside. The SRCVOL1 and SRCVOL2 volumes will be used for the new application with datasets named PROD.BANKING.DCTR.**.
2. To move only the datasets on SRCVL1 and SRCVL2.

Defining a Migration group

To achieve this migration, we use two migration groups.

1. Migrate the INQ datasets.

The first migration group migrates the INQ datasets.

- a. Type the **CREATE** command (or CR) on the Primary command line of the Manage Group Members panel, and press **Enter**.
- b. Type **MOVEVOL1** as the Group Name.

```

z/OS Migrator Group                                     Define
Command ==>
NE
                Scroll ==> CSR

Primary Commands : EXit  NExt

Group Name . . . . MOVEVOL1

Group Options . . . N Replace Existing Data Sets (Y/N)
                  Y Tolerate Allocation Failure (Y/N)

```

- c. Type the **NEXT** command (or NE) on the Primary command line, and press **Enter**.
- d. Type **Y** next to Build Dataset Exclude list (Y/N). You will specify the Dataset Exclude Mask in a subsequent panel.

- e. Type the Source Dataset Name and Source Volume as follows:
PROD.BANKING.** and **SRCVL1**.

```

Define z/OS Migrator Group
Command ==> NE                               Scroll ==> CSR

Primary Commands : EXit  NExt

Group Name . . . . . MOVEVOL1
Source Data Set Options . . N Trace (Y/N)      N AllocSeq (D/S/N)
                          . Y Sphere (Y/N)      N Rename UnConditional (Y/N)
                          . Y Build Data Set Exclude list (Y/N)

Source Data Set Name/Mask . PROD.BANKING.** _____
Source Volume List Name . . _____
Storage Class . . . . . _____
Source Volume(s) . . SRCVL1 _____
                          _____
                          _____
                          _____

```

- f. Type the **NEXT** command (or **NE**) on the Primary command line, and press **Enter**.
- g. Type the Dataset Exclude Mask as follows:
PROD.BANKING.DCTR.**.

```

Define z/OS Migrator Group
Command ==> NE                               Scroll ==> CSR

Primary Commands : EXit  IMport  NExt

Group Name . . . . . MOVEVOL1
Data Set Exclude Mask PROD.BANKING.DCTR.** _____
_____
_____
_____
_____
_____
_____
_____

```

- h. Type the **NEXT** command (or **NE**) on the Primary command line, and press **Enter**.

- i. Type the Target Dataset Name and Target Volume as follows: **MIGR.BANKING.INQ.**** and **TRGVL1**.

```

Define z/OS Migrator Group
Command ==>  BU                               Scroll ==> CSR

Primary Commands : EXit  BUild

Group Name . . . . . MOVEVOL1
Target Data Set Name/Mask . . MIGR.BANKING.INQ.**
Target Volume Storage Class .
Target Volumes(s) . . . . .
TRGVL1 _____
_____
_____
_____
_____
_____

```

- j. Type the **BUILD** command (or BU), and press **Enter**.
 - k. Type the **SAVE** command (or SA), and press **Enter**.
 - l. Specify a member name by following the displayed instructions.
2. Migrate the ATM datasets.

The second migration group migrates the ATM datasets.

- a. Type the **CREATE** command (or CR) on the Primary command line of the Manage Group Members panel, and press **Enter**.
- b. Type **MOVEVOL2** as the Group Name.

```

Define z/OS Migrator Group
Command ==>  NE                               Scroll ==> CSR

Primary Commands : EXit  NExt

Group Name . . . . . MOVEVOL2

Group Options . . . N Replace Existing Data Sets (Y/N)
                   Y Tolerate Allocation Failure (Y/N)

```

- c. Type the **NEXT** command (or NE) on the Primary command line, and press **Enter**.

- d. Type **Y** next to Build Dataset Exclude list (Y/N). You will specify the Dataset Exclude Mask in a subsequent panel.
- e. Type the Source Dataset Name and Source Volume as follows:
PROD.BANKING.** and **SRCVL2**.

```

Define z/OS Migrator Group
Command ==>  NE                               Scroll ==> CSR
Primary Commands : EXit  NExt
Group Name . . . . . MOVEVOL2
Source Data Set Options . . N Trace (Y/N)      N AllocSeq (D/S/N)
                          . Y Sphere (Y/N)      N Rename UnConditional (Y/N)
                          . Y Build Data Set Exclude list (Y/N)
Source Data Set Name/Mask . PROD.BANKING.** _____
Source Volume List Name . . _____
Storage Class . . . . . _____
Source Volume(s) . . SRCVL2 _____

```

- f. Type the **NEXT** command (or **NE**) on the Primary command line, and press **Enter**.
- g. Type the Dataset Exclude Mask as follows:
PROD.BANKING.DCTR.**.

```

Define z/OS Migrator Group
Command ==>  NE                               Scroll ==> CSR
Primary Commands : EXit  IMport  NExt
Group Name . . . . . MOVEVOL2
Data Set Exclude Mask PROD.BANKING.DCTR.** _____
_____
_____
_____
_____
_____

```

- h. Type the **NEXT** command (or **NE**) on the Primary command line, and press **Enter**.

Migrating multiple datasets, SMS volumes

This section provides a scenario involving the migration of several datasets that are managed by SMS.

Note: In cases where the target dataset takes extents that cannot be accommodated on volumes selected at *Activation*, the selection of additional target volumes comes from the storage group containing current target volumes.

Storage environment

SMS environment with the following source datasets to be migrated:

```
VOLSER=SRCVL1, STORGRP1
  PROD.BANKING.ATM.CNTL
  PROD.BANKING.INQ
  PROD.BANKING.DCTR.D05281
```

```
VOLSER=SRCVL2, STORGRP1
  PROD.BANKING.INQ.CNTL
  PROD.BANKING.ATM
  PROD.BANKING.DCTR.D05283
```

```
VOLSER=SRCVL3, STORGRP1
  PROD.BANKING.ATM
  PROD.BANKING.INQ
```

```
VOLSER=TRGVL1
  (empty)
```

Migration objectives

For this exercise, migration objectives include:

1. To move the production banking datasets from the SRCVL1, SRCVL2 and SRCVL3 volumes that reside in the SMS Storage Group STORGRP1.
2. The data is being moved to a new Storage Array. A high capacity volume has been provided.
3. The SMS ACS routines at this installation are coded in such a way that all datasets starting with PROD.BANKING are assigned to the SCLASS1 Storage Class, which then maps to the STORGRP1 Storage Group.

Defining a Migration group

To achieve this migration, we use the following process:

To consolidate volumes in a SMS environment:

1. The volume TRGVL1 is added to the STORGRP1 SMS Storage Group.
2. The SRCVL1, SRCVL2, and SRCVL3 volumes, which are in the STORGRP1 SMS Storage Group, are set to volume status DISNEW. This SMS volume status does not allow new allocations to be made to the volumes.

Note: SMS provides another option that does not allow datasets on the volumes to extent to new extents; this is volume status QUINEW. That option is not used here, as z/OS Migrator automatically moves parts of datasets as they grow during a migration.

3. The first migration group definition migrates all the datasets from SRCVL1.
 - a. Type the **CREATE** command (or CR) on the Primary command line of the Manage Group Members panel, and press **Enter**.
 - b. Type **MOVBANK1** as the Group Name.

```

Define z/OS Migrator Group
Command ==> NE                               Scroll ==> CSR
Primary Commands : EXit  NExt
Group Name . . . MOVBANK1
Group Options . . . N Replace Existing Data Sets (Y/N)
                  Y Tolerate Allocation Failure (Y/N)

```

- c. Type the **NEXT** command (or NE) on the Primary command line, and press **Enter**.

- d. Type the Source Dataset Name and Storage Class as follows:
PROD.BANKING.** and **SCLASS1**.

```

Define z/OS Migrator Group
Command ==> NE                               Scroll ==> CSR

Primary Commands : EXit  NExt

Group Name . . . . . MOVBank1
Source Data Set Options . . N Trace (Y/N)      N AllocSeq (D/S/N)
                          . Y Sphere (Y/N)      N Rename UnConditional (Y/N)
                          . N Build Data Set Exclude list (Y/N)

Source Data Set Name/Mask . PROD.BANKING.** _____
Source Volume List Name . . _____
Storage Class . . . . . SCLASS1
Source Volume(s) . . _____
                    _____

```

- e. Type the **NEXT** command (or **NE**) on the Primary command line, and press **Enter**.
- f. Type the Target Dataset Name, Target Volume Storage Class, and Target Volume as follows: **MIGR.BANKING.****, **SCLASS1**, and **TRGV11**.

```

Define z/OS Migrator Group
Command ==> BU                               Scroll ==> CSR

Primary Commands : EXit  BUild

Group Name . . . . . MOVBank1
Target Data Set Name/Mask . . MIGR.BANKING.** _____
Target Volume Storage Class . SCLASS1
Target Volumes(s) . . . . .
TRGV11 _____
_____
_____
_____
_____
_____

```

- g. Type the **BUILD** command (or **BU**), and press **Enter**.
- h. Type the **SAVE** command (or **SA**), and press **Enter**.
- i. Specify a member name by following the displayed instructions.

The definition for the migration group MOV BANK1 can be explained as follows:

- ◆ A dataset mask is specified that selects all data on the volume.
- ◆ The Storage Class, SCLASS1, is specified to pick up all datasets on the SRCVL1, SRCVL2 and SRCVL3 volumes.
- ◆ A target dataset name of MIGR.BANKING.** is specified.
 - The DF/SMS ACS routines route dataset allocations with the high-level qualifiers of PROD.BANKING to the SMS Storage Class, SCLASS1.
 - The old source datasets will be easy to identify once the migration group reaches the *Completion* phase on both migrations, as they will all start with the unique MIGR.BANKING high-level qualifier.
 - During the allocation process, SMS target volumes in the storage group(s) are associated with SMS storage class 'SCLASS1'. You may influence the SMS volume selection process by assigning a DISNEW volume status to volumes containing other source dataset candidates, that is, volumes that are planned to be source volumes in subsequent migration operations.

Following SMS volume selection rules, volumes within the same storage group that are enabled for new allocations will be available for z/OS Migrator target dataset allocations.

Volume Mirror example

The following panels illustrate how to create, verify, promote, and activate a Mirror volume group. The Mirror volume group consists of five primary and five mirror devices.

Initially, the primary devices will be copied to the mirror devices so their contents can be synchronized. After synchronization, Mirror will continue to monitor write I/O to the primary and replicate it on the mirror.

Step 1: Start the process

1. Type **1** on the z/OS Migrator main panel and press **Enter**.

```
EMC z/OS Migrator Migration

Function Number                11/08/07 11:07:44
or Command ==> 2                SMFID: X04
                                CPMX: @$

                                z/OS Migrator 3.0.0 FUNCTIONS

1 Manage Groups                 5 Display Host Messages
2 Monitor Promoted Groups       6 System Change Summary
3 Display Configuration Information 7 Message Help
4 Operator Interface            8 Monitor Command Line Help
                                9 Display Installation Options

S Set User Session Options
V View or Browse Log data set

F1=Help F3=Exit
```

The Manage Group Members panel appears.

Step 2: Create the volume group

1. To create the new member type **cr** and press **Enter**.

```

                                Manage Group Members                                Row 1 of 40
Command=> cr                                                                Scroll=> CSR
Primary Cmds: CR = Create EX = Exit    F = Find  FIL = Filter SW = Switch
Line   Cmds:  C = Copy   D = Delete E = Restore M = Modify      SMFID X04
                P = Promote R = Rename S = Display V = Verify      CPIX @ $
*****
- ABC      P
  * GROUP NAME:ABC          MEMBER GENERATED BY USER01
  * ON THURSDAY, 25 OCT 2007 AT 18:20:04 FROM: LOGMIGR CONTROL COMMANDS
- BIPPY    P
  * GROUP NAME:BIPPY       MEMBER GENERATED BY USER01
  * ON MONDAY, 17 DEC 2007 AT 17:30:32 FROM: MIRROR CONTROL COMMANDS
- LARGE    P
  * GROUP NAME:LARGE      MEMBER GENERATED BY PTN
  * ON FRIDAY, 16 NOV 2007 AT 15:26:25 FROM: z/OS Migrator CONTROL COMMANDS
- LDBACKG1
  * MIGRATION DIRECTION:   FROM MOD 9'S TO MOD 3'S
  * ALL APPLICATION DATA (1520-1527) TO (1500-151F)
- LDBLAH
  * GROUP NAME:BLAH       MEMBER GENERATED BY USER01
  * ON THURSDAY, 25 NOV 2007 AT 16:27:31 FROM: MIGRATE CONTROL COMMANDS
- LDCAPTUR
  * GROUP NAME:CAPTURE    MEMBER GENERATED BY USER01

```

The Build Type pop-up appears.

2. Enter **s** in the Build Mirror Member selection and press **Enter**.

```

+----- Build Type -----+
| Select with an S a build type. |
| _ Build Migrate Member         |
| s Build Mirror Member         |
| _ Build Logical Migrate Member |
| F3=Cancel                      |
+-----+

```

The Mirror Group parameters panel appears.

Type the following values in the parameters panel.

Group Name	mirgrp1
Mirror Options	s Fastcopy
Primary CCUU	A012
and Count	5
Mirror CCUU	9812
Verify Options	NONE

```

----- Specify the Migrate Group parameters -----
Command ==>                                     Scroll ==>

Group Name . . . . . MIRGRP1
Mirror options . . (S) S FastCopy _ WarmStart _ DeactOnSys _ NoAutoVary
Primary CUU . . . A012 and Count . . 5__ or Volser . . . . . _____
Mirror CUU . . . 9812 New Volser . . M%%%% or Volser Prefix . . _____
Verify Options . S None or _ Empty (S) Volser . . _____ '%' matches any char
***** Primary Commands:EDit SAvE VERify PROMote *****
***** Bottom of data *****

```

Specify the information on the Mirror Group parameters panel, and then press **Enter**

The information just entered displays at the end of the panel.

```

----- Specify the Migrate Group parameters -----
Command ==>                                     Scroll ==>

Group Name . . . . . MIRGRP1
Mirror options . . (S) S FastCopy _ WarmStart _ DeactOnSys _ NoAutoVary
Primary CUU . . . _____and Count . . _____ or Volser . . . . . _____
Mirror CUU . . . _____ New Volser . . _____ or Volser Prefix . . _____
Verify Options . S None or _ Empty (S) Volser . . _____ '%' matches any char
***** Primary Commands:EDit SAvE VERify PROMote *****
GROUP(MIRGRP1) MODE(FASTMIRROR(FASTCOPY))
PAIR PRIMARY(A012,5) MIRROR(9812) VERTARG( *NONE*) NEWVOL (M%%%%%)
***** Bottom of data *****

```

At this point the group definition is complete.

- To save the definition, enter **sa** on the command line and press **Enter**.

```

----- Specify the Migrate Group parameters -----
Command ==>SA                                     Scroll ==>

Group Name . . . . . MIRGRP1
Mirror options . . (S) S FastCopy _ WarmStart _ DeactOnSys _ NoAutoVary
Primary CUU . . . _____and Count . . _____ or Volser . . . . . _____
Mirror CUU . . . _____ New Volser . . _____ or Volser Prefix . . _____
Verify Options . S None or _ Empty (S) Volser . . _____ '%' matches any char
***** Primary Commands:EDit SAvE VERify PROMote *****
GROUP(MIRGRP1) MODE(FASTMIRROR(FASTCOPY))
PAIR PRIMARY(A012,5) MIRROR(9812) VERTARG( *NONE*) NEWVOL (M%%%%%)
***** Bottom of data *****

```

The Create Member Name pop-up appears.

- In the Create Member Name pop-up, enter a 6-digit member name. In this example, the member name is MIRGP1.

Press **Enter** and the characters MD will be added as a prefix to create the actual member name.

```

+----- Create Member Name -----+
| Specify 6 characters to be used to create |
| a member name to contain the data to be |
| saved ==> mirgp1 |
| NOTE: Member Name is prefixed with MD |
| and then your 6 characters. |
+-----+

```

A confirmation message appears with the new member name, MDMIRGP1.

```

----- Specify the Migrate Group parameters -----
Command ==>                               Scroll ==>

Group Name . . . . . MIRGRP1
Mirror options . . (S) S FastCopy _ WarmStart _ DeactOnSys _ NoAutoVary
Primary CUU . . . ___and Count . . ___ or Volser . . . . . ___
Mirror CUU . . . ___ New Volser . . ___ or Volser Prefix . . ___
Verify Options . S None or _ Empty (S) Volser . . ___ '%' matches any char
***** Primary Commands:EDit SAve VERify PROMote *****
MIR COMMAND(S) SAVED TO: LFU0C01.EFMM300.CNTL '(LDMIRGP1)
***** Bottom of data *****

```

5. You can optionally verify MDMIRGP1 by entering v next to the member name and pressing **Enter**.

```

                                Manage Group Members                                Row 1 of 40
Command=>                                                                Scroll=> CSR
Primary Cnds: CR = Create EX = Exit   F = Find  FIL = Filter SW = Switch
Line   Cnds: C = Copy   D = Delete   E = Restore M = Modify   SMFID X04
                P = Promote R = Rename   S = Display V = Verify   CPFIX @$
*****
-  ABC          P
  * GROUP NAME:ABC          MEMBER GENERATED BY USER01
  * ON THURSDAY, 25 OCT 2007 AT 18:20:04 FROM: LOGMIGR CONTROL COMMANDS
-  BIPPY        P
  * GROUP NAME:BIPPY       MEMBER GENERATED BY USER01
  * ON MONDAY, 17 DEC 2007 AT 17:30:32 FROM: MIRROR CONTROL COMMANDS
-  LARGE        P
  * GROUP NAME:LARGE       MEMBER GENERATED BY PTN
  * ON FRIDAY, 16 NOV 2007 AT 15:26:25 FROM: z/OS Migrator CONTROL COMMANDS
-  LDBACKG1
  * MIGRATION DIRECTION:   FROM MOD 9'S TO MOD 3'S
  * ALL APPLICATION DATA (1520-1527) TO (1500-151F)
  V MDMIRGP1
  * GROUP NAME:MIRGRP1     MEMBER GENERATED BY USER01
  * ON THURSDAY, 25 NOV 2007 AT 16:27:31 FROM: MIGRATE CONTROL COMMANDS
-  LDCAPTUR
  * GROUP NAME:CAPTURE     MEMBER GENERATED BY USER01
  * ON FRIDAY, 29 JUN 2007 AT 23:03:47 FROM: MIRROR CONTROL COMMANDS

```

The VERIFY results panel appears.

```

Menu Utilities Compilers Help
-----
BROWSE   SYS04106.T162332.RA000.USER1 .R0132643   Line 00000000 Col 001 080
Command ===>                                                                Scroll ==> CSR
*****
***** Top of Data *****
MIGRATE VERIFY MEMBER EXECUTED BY USER1 ON 04/15/07 AT 16:23:32...

EMCF021I PROCESSING CONFIG FILE PARAMETERS
EMCP001I * GROUP NAME:MIRGRP1          MEMBER GENERATED BY USER1
EMCP001I * ON Monday, 17 Dec 2007 AT 16:54:37 FROM: MIRROR CONTROL COMMANDS
EMCP001I GROUP(MIRGRP1) MODE(FASTMIRROR(FASTCOPY))
EMCP001I PAIR SOURCE(A012,5) TARGET(9812) VERTARG(*NONE*) NEWVOL(M%$$$)
*** END OF DISPLAY ***
***** Bottom of Data *****

```

Any errors or events encountered display on the panel.

6. Press F3 to proceed.

Step 3: Promote the volume group

1. To promote the volume group, in the Manage Group Members panel, enter **p** next to the member name and press **Enter**.

```

                                Manage Group Members                                Row 1 of 40
Command=>                                                                Scroll=> CSR
Primary Cnds: CR = Create EX = Exit   F = Find  FIL = Filter SW = Switch
Line   Cnds:  C = Copy   D = Delete E = Restore M = Modify SMFID X04
                P = Promote R = Rename S = Display V = Verify CPFY @$

*****
-  ABC          P
   * GROUP NAME:ABC          MEMBER GENERATED BY  USER01
   * ON THURSDAY, 25 OCT 2007 AT 18:20:04 FROM:  LOGMIGR CONTROL COMMANDS
-  BIPPY        P
   * GROUP NAME:BIPPY        MEMBER GENERATED BY  USER01
   * ON MONDAY, 17 DEC 2007 AT 17:30:32 FROM:   MIRROR CONTROL COMMANDS
-  LARGE        P
   * GROUP NAME:LARGE        MEMBER GENERATED BY  PTN
   * ON FRIDAY, 16 NOV 2007 AT 15:26:25 FROM:   z/OS Migrator CONTROL COMMANDS
-  LDBACKG1
   * MIGRATION DIRECTION:    FROM MOD 9'S TO MOD 3'S
   * ALL APPLICATION DATA   (1520-1527)   TO (1500-151F)
p MDMIRGP1
   * GROUP NAME:MIRGRP1      MEMBER GENERATED BY  USER01
   * ON THURSDAY, 25 OCT 2007 AT 10:47:52 FROM: MIRROR CONTROL COMMANDS

```

The Confirm Promote pop-up appears.

2. Type **y** in the PROMOTE confirmation pop-up and press **Enter**.

```

+----- Confirm Promote -----+
| Are you sure you want to PROMOTE the member MDMIRGP1 |
|           to                                           |
|           The Database                                 |
|                                     ? y (Y=Yes N=No)   |
+-----+

```

The confirmation pop-up panel appears.

```

Menu  Utilities  Compilers  Help
-----
BROWSE   SYS04106.T165027.RA000.USER1 .R0132768      Line 00000000 Col 001 080
Command ===>                                     Scroll ===> CSR
***** Top of Data *****

MIGRATE PROMOTE MEMBER EXECUTED BY USER1  ON 04/15/07 AT 16:50:27...

EMCF021I PROCESSING CONFIG FILE PARAMETERS
EMCP001I * GROUP NAME:MIRGRP1             MEMBER GENERATED BY  USER1
EMCP001I * ON Monday, 17 Dec 2007 AT 16:54:37 FROM:  MIRROR CONTROL COMMANDS
EMCP001I GROUP(MIRGRP1) MODE(FASTMIRROR(FASTCOPY))
EMCP001I  PAIR SOURCE(A012,5) TARGET(9812) VERTARG(*NONE*) NEWVOL(M%)
* UPDATED ON 2007.106 AT 16:50:25, SYSTEM=Z04
GROUP(MIRGRP1) -
    MODE(FASTMIRROR(FASTCOPY))
*   DEFINING SYSTEM ADDRESSES: A012 TO 9812
    PAIR SSYMM(000184600045,A00012,5) TSYMM(000184600058,980012) -
        VERTARG(*NONE*) NEWVOL(M% % % %)
*GROUP PROMOTED, STOW LENGTH IS 251
*** END OF DISPLAY ***
***** Bottom of Data *****

```

3. Press F3 to proceed.

Step 4: Activate the volume group

To perform actions on a promoted device group or to follow the progress and status of promoted groups, select Monitor Promoted Groups

1. In the z/OS Migrator main panel select option 2. All promoted groups are managed through this option.

```

EMC z/OS Migrator Migration

Function Number                               11/08/07 11:07:44
or Command ==> 2                             SMFID: X04
                                              CPMX: @$

z/OS Migrator 3.0.0 FUNCTIONS

1 Manage Groups                               5 Display Host Messages
2 Monitor Promoted Groups                     6 System Change Summary
3 Display Configuration Information           7 Message Help
4 Operator Interface                          8 Monitor Command Line Help
                                           9 Display Installation Options

S Set User Session Options
V View or Browse Log data set

F1=Help F3=Exit

```

The monitor panel appears.

2. Type an **A** next to the group name **MIRGRP1** and press **Enter** to activate the group

```

z/OS Migrator Monitor                               Row 1 of 5
Command===>                                       11:27:24 11/08/07

Line Commands: A = Activate C = Config D = Deact  E = Errors J = Jobs
                L = Volumes O = VolSum P = Suspend R = Resume S = ShowExt
                T = Reset   U = Report V = Divert X = Delete Z = Complete

Group   Current   Mode   Action   Owner Sets I Dev  Comp
Name    Status      Mode   Action   Owner Sets I Dev  Comp
*****
- ABC    NOT ACTIVE  LOGMIGR              0      0      0
- BIPPY  NOT ACTIVE  LOGMIGR  CONST      0      0      0
A MIRGRP1 NOT ACTIVE  FASTMIR              0      0      0
- TEST3  DEACTIVATED LOGMIGR             16     14      0
- SPHERE MIRROR ACTIVE  FASTMIR              0      0      0
- UDD    NOT ACTIVE  LOGMIGR              0      0      0
- UDDMIGR NOT ACTIVE LOGMIGR              0      0      0
***** Bottom of data *****

```

A confirmation pop-up appears.

- Type **y** and press **Enter** to confirm the activation of MIRGRP1.

```

+----- Confirm Activate Request -----+
|
| Are you sure you want to Activate   the Migrate of
| MIRGRP1                             ? y (Y=Yes N=No)
|
+-----+

```

Once activated, the mirror group initially displays a Current Status of SYNCING to show that the volumes in the group are synchronizing. The percentage complete also displays.

Note: While a group is active the Owner field displays the system ID of the system from which the group was activated.

```

-----z/OS Migrator Monitor ----- Row 1 of 5
Command===>                               17:24:18 04/15/04

Line Commands: A = Activate  C = Config  D = Deactivate  E = Exception
                P = Suspend  R = Resume   T = Reset      U = Report      X = Delete
Group          Current      Mode   Action   Owner Sets  I Dev  Num  Num
Name           Status                               Comp
*****
- KCHMIG       NOT ACTIVE      MIGRATE  CONST
- KCHMIR       NOT ACTIVE      FASTMIR
- MIRGRP1     ACTIVE(8%)       FASTMIR           Z04     5     10   0
- MIRRX       DEACTIVATED     FASTMIR           2       0     0
- MIRR2       MIRROR ACTIVE   FASTMIR           Z06     2     4    2
***** Bottom of data *****

```

- To refresh the percentage, press **Enter**.

When synchronization is complete, the group displays a current status of Mirror Active.

```

                               z/OS Migrator Monitor           Row 1 of 27
Command===>                               11:27:24 11/08/07

Line Commands: A = Activate  C = Config  D = Deact  E = Errors  J = Jobs
                L = Volumes  O = VolSum P = Suspend R = Resume  S = ShowExt
                T = Reset    U = Report V = Divert  X = Delete  Z = Complete

Group          Current      Mode   Action   Owner Sets  Dev  Num  Num
Name           Status                               Comp
*****
- ABC          NOT ACTIVE      LOGMIGR           2       4     0
- BIPPY       NOT ACTIVE      LOGMIGR  CONST
- MIRGRP1     MIRROR ACTIVE   FASTMIR           Z04     5     10   5
- TEST3       DEACTIVATED     LOGMIGR           2       0     0
- SPHERE      MIRROR ACTIVE   FASTMIR           Z06     2     4    2
***** Bottom of Data *****

```

While a Mirror volume group is active you may use the following line commands on the owning system:

C — Display component devices.

D — Deactivate mirroring for a suspended volume group.

E — Display devices pairs with errors.

P — Suspend mirroring for the volume group.

R — Resume a suspended volume group.

Note: The volume group will be re-synchronized first.

U — Produce a volume group report.

You can only use the 'X' (Delete) command to remove a deactivated (or NOT ACTIVE) group from the list of managed volume groups and the z/OS Migrator database.

Volume Migrator examples

The following sections describe:

- ◆ A split offline example
- ◆ A constant copy example

Split offline example

The following panels illustrate how to create, verify, promote and activate a Migrator volume group. The group consists of four source and four target devices.

The source devices will be copied to the target devices. After the copy the target devices will be split offline; that is, migration will stop and the target volumes will be left offline.

Step 1: Start the process

1. To start the process, enter **1** on the z/OS Migrator main panel.

```

EMC z/OS Migrator Migration

Function Number                11/08/07 11:07:44
or Command ==> 2                SMFID: X04
                                   CPIX: @S

z/OS Migrator 3.0.0 FUNCTIONS

1 Manage Groups                 5 Display Host Messages
2 Monitor Promoted Groups       6 System Change Summary
3 Display Configuration Information 7 Message Help
4 Operator Interface            8 Monitor Command Line Help
                                   9 Display Installation Options

S Set User Session Options
V View or Browse Log data set

F1=Help F3=Exit

```

The Manage Group Members panel appears.

Step 2: Create the volume group

1. To create a new group, type **cr** on the command line and press **Enter**.

```

----- Manage Group Members ----- Row 1 of 25
Command====> cr                               Scroll====> PAGE
PrimaryCmds: CR=Create EX=Exit F=Find FIL=Filter SW=Switch
Line   Cmds:  C = Copy      D = Delete E = Restore M = Modify      SMFID X04
          P = Promote R = Rename S = Display V = Verify      CPF# @$
*****
_  KCHMIG      P
    * UPDATED ON 2007.085 AT 08:32:57, SYSTEM=Z04
    *
_  KCHMIR      P
    * UPDATED ON 2007.084 AT 11:25:02, SYSTEM=Z04
    *
_  MDCAMAC1
    * GROUP NAME:CAM1FM           MEMBER GENERATED BY  TSMT11
    * ON FRIDAY, 16 NOV 2007 AT 14:45:32 FROM:  MIRROR CONTROL COMMANDS
_  MDCLONE
    * GROUP NAME:CLONE           MEMBER GENERATED BY  JDOE1
    * ON THURSDAY, 25 OCT 2007 AT 19:57:04 FROM:  MIGRATE CONTROL COMMANDS
_  MDDXB1
    * GROUP NAME:JED1            MEMBER GENERATED BY  JDOE1
    * ON FRIDAY, 16 NOV 2007 AT 13:49:41 FROM:  MIRROR CONTROL COMMANDS
_  MDFMGRP2
    * GROUP NAME:FMGRP2          MEMBER GENERATED BY  JDOE1

```

The Build Type panel appears.

2. Type **s** in the pop-up to select Build Migrate Member and press **Enter** :

```

+----- Build Type -----+
|
| Select with an S a build type.
|   s Build Migrate Member
|   _ Build Mirror Member
|   _ Build Logical Migrate Member
|
| F3=Cancel
|
+-----+

```

The Specify the Migrate Group parameters panel appears.

3. Specify the information for the new group, and then press **Enter**.

The following values are used in this example:

Group Name	fmgrp1
Migration Completion	S Split Offline
Source CCUU	5970
Target CCUU	5980

```

----- Specify the Migrate Group parameters -----
Command ==>                               Scroll ==> PAGE

Group Name . . . . . fmgrp1__
Migration options . . (S) _ FastCopy _ WarmStart _ DeactOnSus _ NoAutoVary
Migration Completion. (S) _ SplitOn S SplitOff _ Swap _ ConstantCopy
Swap options. . . . (Y/N) _ BYPCC _ BYPSNAP _ BYPSYSC _ Consistent
CFW option . . . . . VOLPfx . _
Source CUU . . . 5970 and Count . . . _ or Volser . . . . . _
Target CUU . . . 5980 New Volser . . . _ or Volser Prefix . . . _
Verify Options . _ None or _ Empty (S) Volser . . . _ '%' matches any char
***** Primary Commands:EDit SAVE VERify PROMote *****
***** Bottom of data *****

```

The new information displays at the bottom of the panel.

4. Enter **ed** at the command line and add three additional pairs to the group.

```

----- Specify the Migrate Group parameters ----- Row 1 of 2
Command ==> ed                               Scroll ==> PAGE

Group Name . . . . . FMGRP1__
Migration options . . (S) _ FastCopy _ WarmStart _ DeactOnSus _ NoAutoVary
Migration Completion. (S) _ SplitOn S SplitOff _ Swap _ ConstantCopy
Swap options. . . . (Y/N) _ BYPCC _ BYPSNAP _ BYPSYSC _ Consistent
CFW option . . . . . VOLPfx . _
Source CUU . . . _ and Count . . . _ or Volser . . . . . _
Target CUU . . . _ New Volser . . . _ or Volser Prefix . . . _
Verify Options . _ None or _ Empty (S) Volser . . . _ '%' matches any char
***** Primary Commands:EDit SAVe VERify PROMote *****
GROUP(FMGRP1) MODE(MIGRATE(SPLIT(OFFLINE)))
PAIR SOURCE(5970) TARGET(5980)
***** Bottom of data *****

```

Note: There are several ways to define the additional pairs in the group including: use of the panels for each pair and, if the devices are sequentially addressed, the **and Count** field. Use of the **ed** option is purely illustrative.

- Exit this screen by pressing PF3.

The newly added devices are shown on the panel. At this point, the group definition is complete.

```
File Edit Confirm Menu Utilities Compilers Test Help
-----
EDIT          SYS04106.T180719.RA000.JDOE1.R0133085          Columns 00001 00072
Command ==>                                     Scroll ==> PAGE
***** ***** Top of Data *****
000001 GROUP (FMGRP1) MODE (MIGRATE (SPLIT (OFFLINE)))
000002 PAIR SOURCE (5970) TARGET (5980)
000003 PAIR SOURCE (5971) TARGET (5981)
000004 PAIR SOURCE (5972) TARGET (5982)
000005 PAIR SOURCE (5973) TARGET (5983)
***** ***** Bottom of Data *****
```

- Save this definition by typing **sa** on the command line in the Specify the z/OS Migrator Group parameters panel and pressing **Enter**.

```
----- Specify the Migrate Group parameters ----- Row 1 of 2
Command ==> sa                                     Scroll ==> PAGE

Group Name . . . . . FMGRP1__
Migration options . . (S) _ FastCopy _ WarmStart _ DeactOnSus _ NoAutoVary
Migration Completion. (S) _ SplitOn S SplitOff _ Swap _ ConstantCopy
Swap options. . . . (Y/N) _ BYPCC _ BYPSNAP _ BYPSYSC _ Consistent
                          CFW option . . _____ VOLPfx . . _____
Source CUU . . . _____ and Count . . _____ or Volser . . . . _____
Target CUU . . . _____ New Volser . . _____ or Volser Prefix . . _____
Verify Options . _ None or _ Empty (S) Volser . . _____ '%' matches any char
***** Primary Commands:EDit SAVE VERify PROMote *****
GROUP (FMGRP1) MODE (MIGRATE (SPLIT (OFFLINE)))
  PAIR SOURCE (5970) TARGET (5980)
  PAIR SOURCE (5971) TARGET (5981)
  PAIR SOURCE (5972) TARGET (5982)
  PAIR SOURCE (5973) TARGET (5983)
***** Bottom of data *****
```

The Create Member Name pop-up appears.

- Type the member name in the pop-up window and press **Enter**. In this example, the member name is `fmgrp1`. The characters `MD` will prefix this name.

```

+----- Create Member Name -----+
| Specify 6 characters to be used to create |
| a member name to contain the data to be |
| saved ==> fmggrp1 |
| NOTE: Member Name is prefixed with MD |
| and then your 6 characters. |
+-----+

```

A confirmation message is shown with the new member name, MDFMGRP1.

```

----- Specify the Migrate Group parameters -----
Command ==>Scroll ==> PAGE

Group Name . . . . . _____
Migration options . . (S) _ FastCopy _ WarmStart _ DeactOnSus _ NoAutoVary
Migration Completion. (S) _ SplitOn S SplitOff _ Swap _ ConstantCopy
Swap options. . . . (Y/N) _ BYPCC _ BYPSNAP _ BYPSYSC _ Consistent
CFW option . . _____ VOLPfx . . _____
Source CUU . . . _____and Count . . _____orVolser . . . . . _____
Target CUU . . . _____New Volser . . _____ or Volser Prefix . . _____
Verify Options . _ None or _ Empty (S) Volser . . _____ '%' matches any char
***** Primary Commands:EDit SAVE VERify PROMote *****
MIR COMMAND(S) SAVED TO: 'EMC.SFMM.CONFIG'(MDFMGRP1)
***** Bottom of data *****

```

8. You can optionally verify MDFMGRP1 by entering v next to the member name in the Manage Group Members panel.

```

----- Manage Group Members -----Row 1 of 27
Command===>Scroll===> PAGE
Primary CMds: CR = Create  EX = Exit  F = Find  FIL = Filter  SW = Switch
Line  CMds:  C = Copy  D = Delete  E = Restore  M = Modify  SMFID X04
          P = Promote  R = Rename  S = Display  V = Verify CPIX @$
*****
_  KCHMIGP
*  UPDATED ON 2007.085 AT 08:32:57, SYSTEM=Z04
*
_  KCHMIRP
*  UPDATED ON 2007.084 AT 11:25:02, SYSTEM=Z04
*
_  MDCAMAC1
*  GROUP NAME:CAM1FMMEMBER GENERATED BY  TSMIT1
*  ON FRIDAY, 16 NOV 2007 AT 14:45:32 FROM:  MIRROR CONTROL COMMANDS
_  MDCLONE
*  GROUP NAME:CLONEMEMBER GENERATED BY  JDOE1
*  ON THURSDAY, 25 OCT 2007 AT 19:57:04 FROM:  MIGRATE CONTROL COMMANDS
_  MDDXB1
*  GROUP NAME:JED1MEMBER GENERATED BY  JDOE1
*  ON FRIDAY, 16 NOV 2007 AT 13:49:41 FROM:  MIRROR CONTROL COMMANDS
v  MDFMGRP1

```

The VERIFY results panel appears.

```

Menu  Utilities  Compilers  Help
-----
BROWSESYS04106.T182351.RA000.JDOE1.R0133229Line 00000000 Col 001 080
Command ===>Scroll ===> CSR
***** Top of Data *****
MIGRATE VERIFY MEMBER EXECUTED BY JDOE1 ON 10/26/07 AT 18:23:51...

EMCF021I PROCESSING CONFIG FILE PARAMETERS
EMCP001I * GROUP NAME:FMGRP1MEMBER GENERATED BY  JDOE1
EMCP001I * ON Thursday, 25 Oct 2007 AT 16:52:20 FROM:  MIRROR CONTROL COMMANDS
EMCP001I GROUP (FMGRP1) MODE(MIG(SPLIT(OFFLINE)))
EMCP001IPAIR SOURCE(5970) TARGET(5980)
EMCP001IPAIR SOURCE(5971) TARGET(5981)
EMCP001IPAIR SOURCE(5972) TARGET(5982)
EMCP001IPAIR SOURCE(5973) TARGET(5983)
*** END OF DISPLAY ***
***** Bottom of Data *****

```

Any errors or events encountered display on the panel.

9. Press F3 to continue.

Step 3: Promote the volume group

1. To migrate the group, you must first promote it. Type **p** next to the member name and press **Enter**.

```

----- Manage Group Members -----Row 1 of 28
Command====>Scroll====> PAGE
Primary CMds: CR = Create  EX = Exit    F = Find    FIL = Filter  SW = Switch
Line   CMds: C = Copy    D = Delete  E = Restore  M = Modify   SMFID X04
          P = Promote  R = Rename  S = Display  V = Verify   CPFX @$
*****
_  KCHMIGP
*  UPDATED ON 2007.085 AT 08:32:57, SYSTEM=Z04
*
_  KCHMIRP
*  UPDATED ON 2007.084 AT 11:25:02, SYSTEM=Z04
*
_  MDCAMAC1
*  GROUP NAME:CAM1FMMEMBER GENERATED BY  TSMIT1
*  ON FRIDAY, 16 NOV 2007 AT 14:45:32 FROM:  MIRROR CONTROL COMMANDS
_  MDCLONE
*  GROUP NAME:CLONEMEMBER GENERATED BY  JDOE1
*  ON THURSDAY, 25 OCT 2007 AT 19:57:04 FROM:  MIGRATE CONTROL COMMANDS
_  MDDXB1
*  GROUP NAME:JED1MEMBER GENERATED BY  JDOE1
*  ON FRIDAY, 16 NOV 2007 AT 13:49:41 FROM:  MIRROR CONTROL COMMANDS
p  MDFMGRP1
*  GROUP NAME:FMGRP1MEMBER GENERATED BY  JDOE1

```

The PROMOTE confirmation pop-up appears.

2. Type **y** and press **Enter**.

```

+----- Confirm Promote -----+
| Are you sure you want to PROMOTE the member MDMIRGP1 |
|           to                                           |
| The Database                                           |
|                                     ?  y  (Y=Yes N=No)   |
+-----+

```

The PROMOTE informational panel appears.

```

Menu  Utilities  Compilers  Help
-----
BROWSE   SYS04106.T165027.RA000.JDOE1.R0132768      Line 00000000 Col 001 080
Command ===>                                         Scroll ===> CSR
***** Top of Data *****

MIGRATE PROMOTE MEMBER EXECUTED BY DBURT1 11/15/07 AT 16:50:27...

EMCF021I PROCESSING CONFIG FILE PARAMETERS
EMCP001I * GROUP NAME:FMGRP1             MEMBER GENERATED BY  JDOE1
EMCP001I * ON Thursday, 25 Oct 2007 AT 16:45:20 FROM:  MIGRATE CONTROL COMMANDS
EMCP001I GROUP(FMGRP1) MODE(MIG(SPLIT(OFFLINE)))
EMCP001I  PAIR SOURCE(5970) TARGET(5980)
EMCP001I  PAIR SOURCE(5971) TARGET(5981)
EMCP001I  PAIR SOURCE(5972) TARGET(5982)
EMCP001I  PAIR SOURCE(5973) TARGET(5983)
* UPDATED ON 2007.106 AT 16:50:25, SYSTEM=Z04
GROUP(FMGRP1) -
  MODE(MIGRATE(SPLIT(OFFLINE))
*   DEFINING SYSTEM ADDRESSES: 5970 TO 5980
  PAIR SSYM(000184600045,590070,4) TSYM(000184600058,590080) -
  VERTARG(*NONE*)
*GROUP PROMOTED, STOW LENGTH IS 238

```

A confirmation message also displays, indicating that MDFMGRP1 has been promoted.

Once promoted, the member name prefix changes from MD to MW.

3. Press **Enter** to proceed.

Step 4: Activate the volume group

To perform actions on a promoted group or to follow the progress and status of the split, select the **Monitor Promoted Groups** function.

1. Return to the z/OS Migrator main panel and select option 2. All promoted groups are managed through this option.

```

EMC z/OS Migrator Migration

Function Number              11/08/07 11:07:44
or Command ==> 2             SMFID: X04
                              CPFX: @$

z/OS Migrator 3.0.0 FUNCTIONS

1 Manage Groups              5 Display Host Messages
2 Monitor Promoted Groups    6 System Change Summary
3 Display Configuration Information 7 Message Help
4 Operator Interface         8 Monitor Command Line Help
                              9 Display Installation Options

S Set User Session Options
V View or Browse Log data set

F1=Help F3=Exit

```

The Monitor panel appears.

2. To activate group FMGRP1, type A on the command line and press **Enter**.

```

Command===>                               16:55:14 04/15/07

Line Commands: A = Activate  C = Config  D = Deact   E = Errors  J = Jobs
                L = Volumes   O = VolSum  P = Suspend R = Resume  S = ShowExt
                T = Reset     U = Report  V = Divert  X = Delete  Z = Complete

  Group   Current      Mode   Action   Owner Sets  I Dev  Num  Comp
  Name    Status
*****
_ KCHMIG  NOT ACTIVE      MIGRATE  CONST          2      4    0
_ KCHMIR  NOT ACTIVE      FASTMIR          2      4    0
a FMGRP1  NOT ACTIVE      MIGRATE  SPLIT(OFF)    5     10    0
_ MIRRX   DEACTIVATED      FASTMIR          2      0    0
_ MIRR2   MIRROR ACTIVE    FASTMIR          Z06    2      4    0
***** Bottom of data *****

```

A confirmation pop-up appears.

3. To confirm activation, type Y and press **Enter**.

```
+----- Confirm Activate Request -----+
| Are you sure you want to Activate the Migrate of |
| FMGRP1                                     ? y (Y=Yes n=No) |
+-----+-----+

```

Once activated, the group displays a current status of ACTIVE in the Monitor panel.

4. Press **Enter** to refresh the percentage.

```

Group   Current                               Dev Num Num
Name    Status   Mode      Action   OwnerSwtsDevComp
_ FMGRP1 ACTIVE(8%)MIGRATE  SPLIT(OFF)Z04  5  10  5

```

When the migration is complete, the group displays a current status of COMPLETED:

```
_ FMGRP1 COMPLETED MIGRATE  SPLIT(OFF)Z04  5  10  5
```

Constant Copy example

The following screens illustrate how to create, promote, activate and reset a Volume Migrator Constant Copy group. The group consists of nine primary and nine target devices. The source devices are constant copied to the target devices.

This means that after the initial migration is complete, additional changed tracks are periodically copied from the source to the target. In this example, the source and target devices are then swapped on demand.

Step 1: Start the process

1. To start the process, enter **1** on the command line of the z/OS Migrator main panel and press **Enter**.

```

EMC z/OS Migrator Migration

Function Number                11/08/07 11:07:44
or Command ==> 2                SMFID: X04
                                CPFY: @$

                                z/OS Migrator 3.0.0 FUNCTIONS

1 Manage Groups                 5 Display Host Messages
2 Monitor Promoted Groups       6 System Change Summary
3 Display Configuration Information 7 Message Help
4 Operator Interface            8 Monitor Command Line Help
                                9 Display Installation Options

                                S Set User Session Options
                                V View or Browse Log data set

                                F1=Help F3=Exit

```

The Manage Group Members panel appears.

Step 2: Create the volume group

1. To create the group, type **cr** on the command line and press **Enter** to create the new member.

```

----- Manage Group Members -----      Row 1 of 1
Command==> cr                               Scroll==> PAGE
Primary Cnds: CR = Create  EX = Exit    F = Find  FIL = Filter  SW = Switch
Line   Cnds: C = Copy    D = Delete  E = Restore M = Modify  SMFID X04
                P = Promote  R = Rename  S = Display  V = Verify   CPFY @ $
*****
- MDCAMAC1
  * GROUP NAME:CAM1FM      MEMBER GENERATED BY TSMIT1
  * ON FRIDAY, 16 NOV 2007 AT 14:45:32 FROM: MIRROR CONTROL COMMANDS
- MDCLONE
  * GROUP NAME:CLONE      MEMBER GENERATED BY JDOE1
  * ON THURSDAY, 25 OCT 2007 AT 19:57:04 FROM: MIGRATE CONTROL COMMANDS

- MDDXB1
  * GROUP NAME:DJB1MEMBER GENERATED BY JDOE1
  * ON FRIDAY, 16 NOV 2007 AT 13:49:41 FROM: MIRROR CONTROL COMMANDS
- MDFMGRP1
  * GROUP NAME:FMGRP1     MEMBER GENERATED BY JDOE1
  * ON THURSDAY, 25 OCT 2007 AT 11:21:17 FROM: MIGRATE CONTROL COMMANDS

```

A Build Type pop-up appears.

2. Enter **s** in the pop-up to select **Build Migrate Member**:

```

+----- Build Type -----+
|                               |
| Select with an S a build type. |
|   s Build Migrate Member      |
|   _ Build Mirror Member      |
|   _ Build Logical Migrate Member |
|                               |
| F3=Cancel                    |
+-----+

```

The Specify the Migrate Group parameters panel appears.

- Specify the information on the Group Parameters panel, and then press **Enter**.

The following values are used in this example:

Group Name	FMGRP2
Migration Completion	S Constant Copy
Source CCUU	5978
Count	9
Target CCUU	5988
Verify Volser	MIG%%%

```

----- Specify the Migrate Group parameters -----
Command ==>                               Scroll ==> PAGE

Group Name . . . . . fngrp2__
Migration options . . (S) _ FastCopy _ WarmStart _ DeactOnSus _ NoAutoVary
Migration Completion. (S) _ SplitOn S SplitOff _ Swap          s ConstantCopy
Swap options. . . . (Y/N) _ BYPCC _ BYPSNAP _ BYPSYSC _ Consistent
CFW option . . . . . _____ VOLPfx . . _____
Source CUU . . . . 5978 and Count . . 9_____ or Volser . . . . . _____
Target CUU . . . . 5988 New Volser . . _____ or Volser Prefix . . _____
Verify Options . . _ None or _ Empty (S) Volser . . MIG%%% '%' matches any char
***** Primary Commands:EDit SAve VERify PROMote *****

```

At this point the group definition is complete.

- Save this definition by entering **sa** on the command line.
- Enter the member name in the pop-up window. In this example it is FMGRP2. The characters MD will prefix this name.

```

+----- Create Member Name -----+
| Specify 6 characters to be used to create |
| a member name to contain the data to be  |
| saved ==> fmggrp2                |
|                                           |
| NOTE: Member Name is prefixed with MD    |
| and then your 6 characters.              |
+-----+

```

A confirmation message is then shown with the new member name, MDFMGRP2.

```

----- Specify the Migrate Group parameters -----
Command ==>                                     Scroll ==> PAGE

Group Name . . . . . _____
Migration options . . (S) _ FastCopy _ WarmStart _ DeactOnSus _ NoAutoVary
Migration Completion. (S) _ SplitOn S SplitOff _ Swap _ ConstantCopy
Swap options. . . . (Y/N) _ BYPCC _ BYPSNAP _ BYPSYSC _ Consistent
                          CFW option . . _____ VOLPfx . _____
Source CUU . . . _____ and Count . . _____ or Volser . . . . . _____
Target CUU . . . _____ New Volser . . _____ or Volser Prefix . . _____
Verify Options . _ None or _ Empty (S) Volser . . _____ '%' matches any char
***** Primary Commands:EDit Save VERify PROMote *****
MIR COMMAND(S) SAVED TO: 'EMC.SFMM.CONFIG'(MDFMGRP2)
***** Bottom of data *****

```

Step 3: Promote the volume group

To migrate the group, you must first promote it.

1. To promote the member, type **P** next to the member name MDFMGRP2 and press **Enter**.

```

----- Manage Group Members ----- Row 1 of 1
Command===>                               Scroll===> PAGE
PrimaryCmds: CR=Create EX=Exit F=Find FIL=Filter SW=Switch
Line  Cmds: C = Copy   D = Delete E = Restore  M = Modify   SMFID X04
          P = Promote R = Rename S = Display  V = Verify   CPF# @$
*****
_  MDCAMAC1
    * GROUP NAME:CAM1FM           MEMBER GENERATED BY TSMIT1
    * ON FRIDAY, 16 NOV 2007 AT 14:45:32 FROM: MIRROR CONTROL COMMANDS
_  MDCLONE
    * GROUP NAME:CLONE           MEMBER GENERATED BY JDOE1
    * ON THURSDAY, 25 OCT 2007 AT 19:57:04 FROM: MIGRATE CONTROL COMMANDS
_  MDDXB1
    * GROUP NAME:DJB1           MEMBER GENERATED BY JDOE1
    * ON FRIDAY, 16 NOV 2007 AT 13:49:41 FROM: MIRROR CONTROL COMMANDS
_  MDFMGRP1
    * GROUP NAME:FMGRP1         MEMBER GENERATED BY JDOE1
    * ON THURSDAY, 25 OCT 2007 AT 11:21:17 FROM: MIGRATE CONTROL COMMANDS
p  MDFMGRP2
    * GROUP NAME:FMGRP2         MEMBER GENERATED BY JDOE1
    * ON THURSDAY, 25 OCT 2007 AT 13:11:19 FROM: MIGRATE CONTROL COMMANDS

```

The confirmation pop-up appears.

2. Enter **y** to confirm new member.

```

+----- Confirm Promote -----+
| Are you sure you want to PROMOTE the member MDMIRGP1 |
|           to                                           |
|           The Database                                 |
|                                     ?  y  (Y=Yes N=No)  |
+-----+

```

The Promote Member informational panel appears.

```

Menu Utilities Compilers Help
-----
BROWSE   SYS04097.T131127.RA000.JDOE1.R0132768   Line 00000000 Col 001 080
Command ==>                                     Scroll ==> CSR
***** Top of Data *****

MIGRATE PROMOTE MEMBER EXECUTED BY JDOE1 ON 10/25/07 AT 13:12:27...

EMCF021I PROCESSING CONFIG FILE PARAMETERS
EMCP001I * GROUP NAME:FMGRP2             MEMBER GENERATED BY JDOE1
EMCP001I * ON THURSDAY, 25 OCT 2007 AT 13:11:19 FROM:  MIGRATE CONTROL COMMANDS
EMCP001I GROUP (FMGRP2) MODE(MIGRATE(CONSTANT))
* UPDATED ON 2007.097 AT 13:11:19, SYSTEM=Z04
GROUP (FMGRP2) -
  MODE(MIGRATE(CONSTANT))
*   DEFINING SYSTEM ADDRESSES: 5978 TO 5988
  PAIR SSYMM(000184600045,590070,9) TSYMM(000184600058,590080) -
  VERTARG(,MIG%%%)

*GROUP PROMOTED, STOW LENGTH IS 192

```

A confirmation message displays, indicating that MDFMGRP2 has been promoted. Once promoted, the group appears in the Manage Group Members list under its group name and with a **p** to the right of the group name. The configuration member remains in the configuration PDS for future use.

```

----- Manage Group Members ----- Row 1 of 1
Command==>                                     Scroll==> PAGE
Primary Cmds: CR = Create  EX = Exit    F = Find  FIL = Filter  SW = Switch
Line   Cmds:  C = Copy    D = Delete  E = Restore  M = Modify    SMFID X04
                P = Promote  R = Rename  S = Display  V = Verify    CPFY @$
*****
-  MDTSMIT1
   * GROUP NAME:TSMIT1             MEMBER GENERATED BY TSMIT1
   * ON FRIDAY, 16 NOV 2007 AT 14:45:32 FROM:  MIRROR CONTROL COMMANDS
-  MDCLONE
   * GROUP NAME:CLONE             MEMBER GENERATED BY JDOE1
   * ON THURSDAY, 25 OCT 2007 AT 19:57:04 FROM:  MIGRATE CONTROL COMMANDS
-  MDJADXB1
   * GROUP NAME:JADBX1           MEMBER GENERATED BY JDOE1
   * ON FRIDAY, 16 NOV 2007 AT 13:49:41 FROM:  MIRROR CONTROL COMMANDS
-  MDFMGRP1
   * GROUP NAME:FMGRP1           MEMBER GENERATED BY JDOE1
   * ON THURSDAY, 25 OCT 2007 AT 11:21:17 FROM:  MIGRATE CONTROL COMMANDS
-  MDFMGRP2  P
   * GROUP NAME:FMGRP2           MEMBER GENERATED BY JDOE1
   * ON THURSDAY, 25 OCT 2007 AT 13:11:19 FROM:  MIGRATE CONTROL COMMANDS
                                     +-----+
                                     | Member MDFMGRP2 Promoted |
                                     +-----+

```

Step 4: Activate the volume group

To perform actions on a promoted group or to follow the progress and status of the swap, select the Monitor Promoted Groups function.

1. Return to the z/OS Migrator main panel and select option 2. All promoted groups are managed through this option.

```

EMC z/OS Migrator Migration

Function Number                               11/08/07 11:07:44
or Command ==> 2                               SMFID: X04
                                                CPFY: @$

z/OS Migrator 3.0.0 FUNCTIONS

1 Manage Groups                               5 Display Host Messages
2 Monitor Promoted Groups                     6 System Change Summary
3 Display Configuration Information           7 Message Help
4 Operator Interface                          8 Monitor Command Line Help
9 Display Installation Options

S Set User Session Options
V View or Browse Log data set

F1=Help F3=Exit

```

The Monitor panel appears.

2. To activate group FMGRP2, type **a** next to the group name and press **Enter**.

```

Command====>                               16:55:14 04/15/07

Line Commands: A = Activate  C = Config  D = Deact  E = Errors  J = Jobs
                L = Volumes  O = VolSum  P = Suspend R = Resume  S = ShowExt
                T = Reset    U = Report  V = Divert  X = Delete  Z = Complete

  Group      Current      Mode      Action      Owner Sets  Dev  Num  Num
  Name       Status       Mode      Action      Owner Sets  I Dev  Comp
*****
_ FMGRP1     COMPLETED      MIGRATE  SPLIT(OFF)      4      8   4
a FMGRP2     NOT ACTIVE       MIGRATE  CONST           2      4   0
*****
***** Bottom of data *****

```

A confirmation pop-up appears.

3. To confirm activation, type **y** and press **Enter**.

```

+----- Confirm Activate Request -----+
|
| Are you sure you want to Activate the Migrate of
| MIRGRP2                                     ? y (Y=Yes N=No)
|
+-----+

```

Once activated, the group has a current status of `ACTIVE` in the Monitor panel. The percentage complete is also shown.

```

Group Name      Current Status      Mode      Action      Dev Owner Sets  Num I Dev  Num Comp
*****
_ FMGRP1        COMPLETED    MIGRATE    SPLIT(OFF)    4      8      4
_ FMGRP2        ACTIVE(6%)    MIGRATE    CONST         2      4      0
***** Bottom of data *****

```

4. Press **Enter** to refresh the percentage

In constant copy mode the process remains active even after migration is 100% complete.

```

Group Name      Current Status      Mode      Action      Dev Owner Sets  Num I Dev  Num Comp
*****
_ FMGRP1        COMPLETED    MIGRATE    SPLIT(OFF)    4      8      4
_ FMGRP2        ACTIVE(100%)  MIGRATE    CONST         2      4      0
***** Bottom of data *****

```

5. You can reset the process, or change the constant copy action, by typing **t** (for reset) next to the group name and pressing **Enter**.

```

Group Name      Current Status      Mode      Action      Dev Owner Sets  Num I Dev  Num Comp
*****
_ FMGRP1        COMPLETED    MIGRATE    SPLIT(OFF)    4      8      4
t FMGRP2        ACTIVE(100%)  MIGRATE    CONST         2      4      0
***** Bottom of data *****

```

A pop-up panel with new action choices for resetting migration action appears.

```

+- Reset Migration Action +-
|
| New Action . . . . w
|
|      S Split
|      M SplitMount
|      W Swap
|
| F1=Help   F3=Cancel
|
+-----+

```

6. To select Swap, type **w** after New Action and press **Enter**.

A Confirm Swap Request pop-up appears.

```

+----- Confirm Swap Request -----+
|                                     |
| Are you sure you want to Swap      the Migrate of      |
| FMGRP2                             ? y (Y=Yes N=No)      |
|                                     |
+-----+

```

- To confirm the swap request, type **y** and press **Enter**.

The action changes to **SWAP** and the percentage continues to increment if necessary.

The **Reset** command (**T**) may be used even before the group is fully migrated. The requested action takes place when the migration is complete.

```

-----z/OS Migrator Monitor ----- Row 1 of 1
Command==>                                     17:05:45 04/15/07

Line Commands: A = Activate  C = Config  D = Deact   E = Errors  J = Jobs
                L = Volumes  O = VolSum  P = Suspend R = Resume  S = ShowExt
                T = Reset    U = Report  V = Divert  X = Delete  Z = Complete

  Group      Current          Dev      Num      Num
  Name      Status           Owner Sets I Dev  Comp
*****
_ FMGRP1    COMPLETED        MIGRATE  SPLIT(OFF)      4      8      4
_ FMGRP2    ACTIVE(100%)          MIGRATE  SWAP      *      2      4      0

```

The asterisk next to the **SWAP** status indicates that the action was temporarily changed from Constant Copy. If you subsequently reactivate this group, the action reverts to **CONSTANT COPY**.

After the **SWAP** is complete, the group displays a current status of **COMPLETED**.

```

-----z/OS Migrator Monitor ----- Row 1 of 1
Command==>                                     16:59:37 04/15/07

Line Commands: A = Activate  C = Config  D = Deact   E = Errors  J = Jobs
                L = Volumes  O = VolSum  P = Suspend R = Resume  S = ShowExt
                T = Reset    U = Report  V = Divert  X = Delete  Z = Complete

  Group      Current          Dev      Num      Num
  Name      Status           Owner Sets I Dev  Comp
*****
_ FMGRP1    COMPLETED        MIGRATE  SPLIT(OFF)      4      8      4
_ FMGRP2    COMPLETED        MIGRATE  SWAP      *      2      4      0
***** Bottom of data *****

```

This chapter explains the execution parameters, SYSIN commands, and provides sample batch jobs. Topics are:

- ◆ Overview 296
- ◆ Execution parameters 297
- ◆ Group definition dataset 298
- ◆ SYSIN commands 299
- ◆ Logical sample batch jobs 303

Overview

A batch utility, EFMMBAT, is provided with z/OS Migrator. This batch utility allows operations to be activated using a batch job. The utility submits commands to the z/OS Migrator service task and retrieves the responses.

Note: Although possible, EMC recommends that you do not run the z/OS Migrator as a batch job since this negates certain safeguards.

Any command that can be submitted using the operator console can be submitted using the batch utility. In addition, there are a number of specialized commands for the batch utility environment, including the following:

- ◆ VERIFY
- ◆ PROMOTE
- ◆ WAIT
- ◆ MAXWAIT
- ◆ IF
- ◆ QUIT

The batch utility is driven by three input sources:

- ◆ Execution parameters
- ◆ Group definition dataset
- ◆ SYSIN commands

Text group definitions may also be entered from the batch utility program, although they won't be saved to a partitioned dataset if entered this way.

Execution parameters

The following is a z/OS Migrator execution parameter:

```
PARM= ' CPMX=cccccccc, SSID=cccc, DEBUG=YES | NO '
```

Where:

- ◆ CPMX is a one- to eight-character command prefix for the z/OS Migrator subsystem.
- ◆ SSID identifies the z/OS Migrator subsystem ID.

Note: You need only specify either CPMX or SSID; however, if both are specified, they must match the settings for a z/OS Migrator subsystem in order to proceed.

- ◆ DEBUG: If set to YES, it will cause diagnostic messages to be written to an optional EFMMBAT dd statement. The default is NO. Include a CONFIG DD statement to identify the library from which members will be promoted.

Group definition dataset

All configuration members to be acted upon must reside in the Configuration PDS. You can create these members using either:

- ◆ A normal editor, or
- ◆ The z/OS Migrator ISPF Monitor

Note: Configuration members coded at the command line are promoted directly to the z/OS Migrator service task and stored in the z/OS Migrator database. These members are not stored in the Configuration PDS.

SYSIN commands

The SYSIN DD statement allows you to submit any z/OS Migrator command to the z/OS Migrator server. In addition, there are a number of specialized commands for controlling the Batch Utility environment.

Descriptions of z/OS Migrator SYSIN commands are as follows:

VERIFY *group-member-name*

Causes the specified group member to be submitted to the z/OS Migrator service task to verify syntax and compatibility with previously promoted members.

PROMOTE *group-member-name*

Submits the specified member to the z/OS Migrator service task for addition to the active configuration.

PROMOTE *DLM=xx*

The statements immediately following the PROMOTE statement are to be interpreted as a group definition and promoted accordingly. The instream group definition ends when a statement beginning with the two characters specified via the DLM=xx parameter is encountered. This group will be added to the z/OS Migrator configuration database but not to the configuration PDS.

IF RC *operator value*

This command tests the most recent return code against a specified value. The operators are '>', '<', and '='. The value is a positive number. If the condition is met, then the next statement is executed. Otherwise, the next command is skipped and processing continues with the subsequent command.

IF GROUP *groupname* **(NOT) EXIST**

This command functions like IF RC; however, the condition being tested is the existence of the specified group.

IF GROUP *groupname* STATUS (NOT) *groupstatus*

This command functions like IF RC, however, the condition being tested is the status of the specified group.

The following is a list of *groupstatus* values:

- ◆ SUSPEND — The group is suspended.
- ◆ ACTIVE — The group is active; this is generally a transient state when the group is first started.
- ◆ PNDACT — Waiting for other systems to acknowledge the activation.
- ◆ PNDMIRR — Synchronizing.
- ◆ MIRROR — All extents are mirroring.
- ◆ PNDIVERT — Waiting for other systems to acknowledge DIVERT request.
- ◆ DIVERT — Fully diverted.
- ◆ PNDTERM — Waiting for other systems to acknowledge TERMINATION request.
- ◆ TERM — The group is fully terminated.
- ◆ DEACT — Group has been deactivated (by command or error).
- ◆ COMPLETE — Volume migration is complete.

WAIT *nnnnn*

Waits the specified number of seconds before submitting the next command.

You can enter a number from 0 to 99999. A value of zero (0) means that there will be no wait. Any other value will cause z/OS Migrator to wait the specified number of seconds before executing the next command.

WAIT GROUP *groupname* STATUS (NOT) *groupstatus*

The specified group will be queried and batch command processing will be suspended until the specified status has been observed. A limit for this wait may be specified using the MAXWAIT command. If the MAXWAIT limit is reached before the requested condition has

been met, then the return code (RC) is set to 4. Refer to the IF GROUP command for a description of the status values.

MAXWAIT *nnnnn*

Specifies the maximum number of seconds to wait for group status. If the time is exceeded, then the command return code (RC) is set to 4 and processing continues with the next statement.

QUIT

Causes command interpretation to immediately end.

Usage examples

```
IF GROUP MYGROUP EXISTS
  DELETE MYGROUP
  DISPLAY GROUPS

PROMOTE DLM=$$
  GROUP(LD6255) MODE(LMIGR)
    SET SOURCE(DSN( SG2.LD6255.VSAM.* )) -
      TARGET(STGCLS(SCBASE          ) -
        DSN( SG2.LD6255.TEMP.* ))
  $$

ACTIVATE MYGROUP
IF RC > 0
  QUIT
MAXWAIT 300
WAIT GROUP MYGROUP STATUS COMPLETED
IF RC = 4
  QUIT
```

Note: You can use the continuation character—a dash (-) at the end of the line as shown in the above example—wherever blanks are permitted in the syntax for batch interfaces. Up to 16 dataset filters can be specified for the source-dsn list.

Logical sample batch jobs

The two sample jobs included for batch interface are the minimum number of separate jobs that are required to complete a z/OS Migrator migration.

Note: Each step could be a separate batch job.

The sample jobs listed here are used to migrate three different Partitioned Datasets (PDS). The space allocation of the datasets will be in records, tracks and cylinders. The three PDS will have more members added after the *Activate* operation has been completed.

- ◆ The first job is ZOSM5070.
- ◆ The second job is ZOSM5071.

ZOSM5070

The first job, ZOSM5070, will perform the following tasks:

1. Remove the datasets if they exist, since this is a sample job.
2. Allocate and populate the three PDSs with the same small PDS.
3. *Delete* the LD5070 group if it exists.
4. *Define* and *Promote* the LD5070 group.
5. Wait for a few seconds to allow the group to be recognized in the other active z/OS Migrator servers in other LPARS.
6. *Activate* the LD5070 group.
7. Copy more members into each source PDS after the *Activate* operation is complete. Add enough new members to cause each PDS to go into secondary extents.
8. Wait for the LD5070 group to *Mirror* after the additional members are added to the source PDS.
9. *Divert* the LD5070 group.

ZOSM5070 sample batch job¹

```

//ZOSM5070 JOB 'YOUR JOB STATEMENT'
//*****
//*          ZOSM5070 - ZOSM5071                      *
//*          *                                         *
//*          PDS IN RECORDS, TRACKS AND CYLINDERS    *
//*          z/OS Migrator TO COPY                    *
//*          EXTEND ALL PDS TO MULTIPLE EXTENTS AFTER ACTIVATE *
//*          *                                         *
//*          1. DELETE GROUP IF IT EXISTS              *
//*          2. CREATE GROUP GROUP NAME LD5070        *
//*          3. PROMOTE GROUP                          *
//*          4. ACTIVATE GROUP                        *
//*          5. DIVERT GROUP                          *
//*          *                                         *
//*****
//* SOURCE VOL : EF24                                 *
//* TARGET VOL : EF14                                 *
//* VOLSER : MVEF24 MVEF14                           *
//*****
//DELETE EXEC PGM=IDCAMS,REGION=1024K
//SYSPRINT DD SYSOUT=*
//SYSIN      DD *
        DEL LDSRCE.MVEF24.PDS.BLOCKS
        DEL LDSRCE.MVEF24.PDS.CYL
        DEL LDSRCE.MVEF24.PDS.TRK
        DEL LDTGT.MVEF14.PDS.BLOCKS
        DEL LDTGT.MVEF14.PDS.CYL
        DEL LDTGT.MVEF14.PDS.TRK
        SET MAXCC=0
/*
//STEP2B EXEC PGM=IEBCOPY
//*      BLOCKS
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN=TEST.SMALL.PDS,DISP=SHR
//SYSUT2 DD DSN=LDSRCE.MVEF24.PDS.BLOCKS,
//      DISP=(NEW,CATLG,DELETE),
//      UNIT=3390,VOL=SER=(MVEF24),
//      SPACE=(24000,(0004,0010,200))
//SYSIN DD DUMMY
/*
//STEP2C EXEC PGM=IEBCOPY
//*      CYLINDERS
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN= EMC.SMALL.PDS,DISP=SHR
//SYSUT2 DD DSN=LDSRCE.MVEF24.PDS.CYL,
//      DISP=(NEW,CATLG,DELETE),
//      UNIT=3390,VOL=SER=(MVEF24),

```

1. ZOSMvrm = version, release, and modification of the current z/OS Migrator tool.

```

//          SPACE=(CYL,(01,01,200))
//SYSIN    DD DUMMY
/*
//STEP2T   EXEC PGM=IEBCOPY
//*        TRACKS
//SYSPRINT DD SYSOUT=*
//SYSUT1   DD DSN= TEST.SMALL.PDS,DISP=SHR
//SYSUT2   DD DSN=LDSRCE.MVEF24.PDS.TRK,
//          DISP=(NEW,CATLG,DELETE),
//          UNIT=3390,VOL=SER=(MVEF24),
//          SPACE=(TRK,(07,12,200))
//SYSIN    DD DUMMY
/*
//PROMOTE   EXEC PGM=EFMMBAT,PARM='CPFX=WJ,DEBUG=NO',COND=(7,LE)
//BB        DD DISP=SHR,DSN=EMC.ZOSM.TEST.LINKLIB
//EFMMDB    DD DISP=SHR,DSN=EMC.ZOSMvrm.DB
//CONFIG    DD DISP=SHR,DSN=EMC.ZOSM.CONFIG.PDS
//SYSPRINT  DD SYSOUT=*
//EFMMBAT   DD SYSOUT=*
//SYSUDUMP  DD SYSOUT=*
//SYSIN     DD *
IF GROUP LD5070 EXISTS
DELETE LD5070
*****
* CREATE GROUP AND PROMOTE GROUP *
*****
PROMOTE DLM=ZZ
GROUP(LD5070) MODE(LMIGR)
SET SOURCE(DSN(LDSRCE.MVEF24.PDS.*)) -
TARGET(VOL(MVEF14)) -
DSN(LDTGT.MVEF14.PDS.*)
ZZ
D GROUPS
/*
//WAIT20    EXEC PGM=WAIT,PARM=20
//ACTIVATE  EXEC PGM=EFMMBAT,PARM='CPFX=WJ,DEBUG=NO',COND=(7,LE)
//STEPLIB   DD DISP=SHR,DSN=EMC.ZOSMvrm.TEST.LINKLIB
//EFMMDB    DD DISP=SHR,DSN=EMC.ZOSM.DB
//CONFIG    DD DISP=SHR,DSN=EMC.ZOSM.CONFIG.PDS
//SYSPRINT  DD SYSOUT=*
//EFMMBAT   DD SYSOUT=*
//SYSUDUMP  DD SYSOUT=*
//SYSIN     DD *
*****
* ACTIVATE GROUP - WAIT FOR ACTIVATION *
*****
ACTIVATE LD5070
WAIT GROUP LD5070 STATUS MIRROR
D GROUPS
/*
//*****
//* AFTER ACTIVATE CAUSE THE 3 PDS TO GO INTO EXTENTS. *

```

```

//*****
//STEP3B EXEC PGM=IEBCOPY
//*      BLOCKS
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN=TEST.LARGER.PDS,DISP=SHR
//SYSUT2 DD DSN=LDSRCE.MVEF24.PDS.BLOCKS,
//      DISP=SHR,
//      UNIT=3390,VOL=SER=(MVEF24),
//      SPACE=(24000,(0004,0010,200))
//SYSIN DD DUMMY
/*
//STEP3C EXEC PGM=IEBCOPY
//*      CYLINDERS
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN=TEST.LARGER.PDS,DISP=SHR
//SYSUT2 DD DSN=LDSRCE.MVEF24.PDS.CYL,
//      DISP=SHR,
//      UNIT=3390,VOL=SER=(MVEF24),
//      SPACE=(CYL,(01,01,200))
//SYSIN DD DUMMY
/*
//STEP3T EXEC PGM=IEBCOPY
//*      TRACKS
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DSN=TEST.LARGER.PDS,DISP=SHR
//SYSUT2 DD DSN=LDSRCE.MVEF24.PDS.TRK,
//      DISP=SHR,
//      UNIT=3390,VOL=SER=(MVEF24),
//      SPACE=(TRK,(07,12,200))
//SYSIN DD DUMMY
/*
//WAIT21 EXEC PGM=WAIT,PARM=20
//MIRROR EXEC PGM=EFMMBAT,PARM='CPFX=WJ,DEBUG=NO',COND=(7,LE)
//STEPLIB DD DISP=SHR,DSN=EMC.ZOSMvrm.TEST.LINKLIB
//EFMMDB DD DISP=SHR,DSN=EMC.ZOSM.DB
//CONFIG DD DISP=SHR,DSN=EMC.ZOSM.CONFIG.PDS
//SYSPRINT DD SYSOUT=*
//EFMMBAT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
*****
* WAIT FOR MIRROR STATUS AFTER UPDATE OF SOURCE *
*****
D GROUPS
  WAIT GROUP LD5070 STATUS MIRROR
D GROUPS
/*
//*
//DIVERT EXEC PGM=EFMMBAT,PARM='CPFX=WJ,DEBUG=NO',COND=(7,LE)
//STEPLIB DD DISP=SHR,DSN=EMC.ZOSMvrm.TEST.LINKLIB
//EFMMDB DD DISP=SHR,DSN=EMC.ZOSM.DB
//CONFIG DD DISP=SHR,DSN=EMC.ZOSM.CONFIG.PDS

```

```
//SYSPRINT DD SYSOUT=*
//EFMMBAT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
*****
* DIVERT GROUP - *
*****
DIVERT LD5070
WAIT GROUP LD5070 STATUS DIVERT
D GROUPS
/*
//
```

ZOSM5071

The second sample job, ZOSM5071, is built with the assumption that ZOSM5070 was run first and has completed with a return code of 4 or less.

ZOSM5071 will perform the following tasks:

1. Compare the PDS to make sure they are complete.
2. *Delete* the LD5070 group after completion.

ZOSM5071 sample batch job

```
//ZOSM5071 JOB 'YOUR JOB STATEMENT'
//*****
//* ZOSM5070 - ZOSM5071 *
//* * *
//* 6. DELETE GROUP *
//* * *
//*****
//TERMGRP EXEC PGM=EFMMBAT, PARM='CPFX=WJ, DEBUG=NO', COND=(7, LE)
//STEPLIB DD DISP=SHR, DSN=EMC.ZOSMvrm.TEST.LINKLIB
//EFMMDB DD DISP=SHR, DSN=EMC.ZOSM.DB
//CONFIG DD DISP=SHR, DSN=EMC.ZOSM.CONFIG.PDS
//SYSPRINT DD SYSOUT=*
//EFMMBAT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
*****
* COMPARE THE PDS *
*****
//COMPARE1 EXEC PGM=IEBCOMPR
//* COMPARE DATASETS AND DELETE
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD DISP=(SHR), DSN=LDSRCE.MVEF24.PDS.TRK
//SYSUT2 DD DISP=(SHR), DSN=LDSRCE.MVEF24.PDS.CYL
//SYSIN DD *
```

```

                                COMPARE TYPORG=PO
/*
//COMPARE2 EXEC PGM=IEBCOMPR
/* COMPARE DATASETS AND DELETE
//SYSPRINT DD SYSOUT=*
//SYSUT1   DD DISP=(SHR),DSN=LDSRCE.MVEF24.PDS.TRK
//SYSUT2   DD DISP=(SHR),DSN=LDSRCE.MVEF24.PDS.BLOCKS
//SYSIN    DD *
                                COMPARE TYPORG=PO
/*
//DELGROUP EXEC PGM=EFMMBAT, PARM='CPFX=WJ,DEBUG=NO',COND=(7,LE)
//STEPLIB  DD DISP=SHR,DSN=EMC.ZOSMvrm.TEST.LINKLIB
//         DD DISP=SHR,DSN=EMC.SNAP.ZOSM.LINKLIB
//EFMMDB   DD DISP=SHR,DSN=EMC.ZOSM.DB
//CONFIG   DD DISP=SHR,DSN=EMC.ZOSM.CONFIG.PDS
//SYSPRINT DD SYSOUT=*
//EFMMBAT  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSIN    DD *
*****
* DELETE GROUP ON WAY OUT *
*****
IF GROUP LD5070 EXISTS
DELETE LD507
/*
//

```

Volume sample batch jobs

This sample JCL displays the volume groups and their status. If group CF8 has already been promoted, then it is deleted. CF8 is then promoted and activated. After 10 seconds the summary group status displays.

The job waits up to 3600 seconds for the volume group (a migration group) to complete. If the volume group completes within the time limit then one last summary display is issued. Otherwise, command execution immediately ends.

```
//MYJOB1 JOB (EMC) , CLASS=A, MSGCLASS=X, LINES=200,
//          NOTIFY=&SYSUID
//*
//TESTPCR1 EXEC PGM=EFMMBAT, PARM=' CPFX=@$ '
//STEPLIB DD DSN=EMC.ZOSMvrm.TEST.LINKLIB, DISP=SHR
//SYSPRINT DD SYSOUT=*
//CONFIG DD DSN=EMC.ZOSM.CONFIG.PDS, DISP=SHR
//EFMMDB DD DSN=EMC.ZOSM.DB, DISP=SHR
//SYSUDUMP DD SYSOUT=*
//SYSIN DD *
*
* SAMPLE COMMAND STREAM
*
D GROUPS
IF GROUP CF8 EXISTS
    DELETE CF8
PROMOTE MWCF8
ACTIVATE CF8
MAXWAIT 3600
WAIT 10
D GROUPS
WAIT GROUP CF8 STATUS COMPLETE
IF RC = 4
    QUIT
D GROUPS
/*
```

Sample SYSPRINT output

This example shows a volume migration group.

```

==> D GROUPS
NAME 1ST-PAIR PAIRS STATUS DEFINITION
CF7 9813/9812 1 NEVER ACTIVE MODE(MIG(SWAP))
SYSRES 1541/2A93 1 COMPLETED MODE(MIG(SPLIT(OFF)),FASTCOPY)
** RC=0

==> WAIT 10
STARTING 10 SECOND WAIT AT 11:34:20

==> DELETE CF7
EMCF532I GROUP CF7 DELETED
** RC=0

==> VERIFY SANTA
MEMBER SANTA NOT FOUND
** RC=8

==> VERIFY LDCF7
EMCF021I PROCESSING CONFIG FILE PARAMETERS
EMCP001I * GROUP NAME:CF7 MEMBER GENERATED BY SMITH
EMCP001I * ON Monday, 17 Dec 2007 AT 07:17:00 FROM: MIGRATE CONTROL
COMMANDS
EMCP001I GROUP(CF7) MODE(MIGRATE(SWAP))
EMCP001I SWAPOPT(CFW(BYPASS))
EMCP001I PAIR SOURCE(9813) TARGET(9812) VERTARG(*NONE*)
** RC=0

==> PROMOTE SANTA
MEMBER SANTA NOT FOUND
** RC=8

==> PROMOTE LDCF7
EMCF021I PROCESSING CONFIG FILE PARAMETERS
EMCP001I * GROUP NAME:CF7 MEMBER GENERATED BY SMITH
EMCP001I * ON Monday, 17 Dec 2007 AT 07:17:00 FROM: MIGRATE CONTROL
COMMANDS
EMCP001I GROUP(CF7) MODE(MIGRATE(SWAP))
EMCP001I SWAPOPT(CFW(BYPASS))
EMCP001I PAIR SOURCE(9813) TARGET(9812) VERTARG(*NONE*)
** RC=0

==> D GROUPS
NAME 1ST-PAIR PAIRS STATUS DEFINITION
CF7 9813/9812 1 NEVER ACTIVE MODE(MIG(SWAP))
SYSRES 1541/2A93 1 COMPLETED MODE(MIG(SPLIT(OFF)),FASTCOPY))
** RC=0
*** MAXIMUM RETURN CODE WAS 8

```

Determining the CPU Serial Number

This appendix explains how to determine the serial number of a specific mainframe server.

- ◆ [Determining the CPU serial number.....](#) 312

Determining the CPU serial number

Authorization to issue z/OS commands is required in order to determine the serial number of a specific mainframe. The command can be issued from the operator console or via a product such as System Display and Search Facility (SDSF).

- ◆ Issue the z/OS command Display Matrix for CPU. The command syntax is:

```
D M=CPU
```

The operating system will issue a Store CPU ID (STIDP) instruction, which returns a double word (8 bytes; 16 digits) of information. The format of this information is as follows:

Table 25 Deciphering D M=CPU output

00	nn	xxxx	mmmm	0000
0-1	2-3	4-7	8-11	12-15
Version code	Logical partition identifier (LPID)	CPU serial number	CPU model number	zeros

The operating system will display the following information:

```
RESPONSE=MVS1
IEE174I 13.48.57 DISPLAY M 803
PROCESSOR STATUS
ID CPU SERIAL
0 + 405525995
2 + 605525995

CPC ND = 05995A.140.AMH.05.000000050686
+ ONLINE - OFFLINE . DOES NOT EXIST
CPC ND CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
```

The serial number displayed is 405525995. Using the previous chart, the CPU serial number maps out as follows:

- ◆ Version Code is not displayed
- ◆ 40 is the LPID
- ◆ 5552 is the CPU serial number
- ◆ 5995 is the CPU model type
- ◆ Zeros are not displayed

For more information on the Store CPU ID instruction, refer to the appropriate level of the IBM *z/Architecture Principles of Operations* manual.

Detailed Parameter Reference

Some parameters are for logical migrations only; some are for volume mirrors and migrations only; and some are for both. The following parameters are for both volume and logical migrations unless otherwise noted.

- ◆ [Execution parameters](#) 314
- ◆ [Global configuration parameters](#) 322

Execution parameters

As shown in “Volume sample batch jobs” on page 309, the sample startup JCL contains a series of z/OS Migrator execution parameters for a volume migration job.

The following sections list the execution parameters.

Note: You can also specify several of these execution parameters as configuration parameters. The following sections indicate when an execution parameter can also be a configuration parameter.

CPFX= (FM1 , SYSPLEX)

CFG

Description

This is an execution parameter for volume mirror/migrations.

You specify certain global configuration parameters in a parameter file. The parameter file must be a partitioned dataset. This file may be specified with the //CONFIG JCL statement or, optionally, with the CFG execution parameter.

CFG is an alias of PARMLIB. You can use either parameter to specify global configuration values.

Syntax

CFG=*dsname* (*member*)

Where:

dsname

The name of a partitioned dataset that is the parameter file.

member

The member of the partitioned dataset that holds the parameter library.

Example

None.

CPFX

Description

CPFX is a parameter that can be used as either an execution or a global configuration parameter. It identifies the command prefix for

entering z/OS Migrator server operator commands. The command prefix you enter is registered with z/OS. You can subsequently display the command prefix using the z/OS DISPLAY OPDATA.

Syntax `CPFX=command prefix`

or

`CPFX=(command prefix[,SYSPLEX] [,FORCE])`

Where:

command prefix

A unique one to eight character string formatted and named according to the standard rules for defining such prefixes.

SYSPLEX

Indicates the command prefix should have *SYSPLEX* scope. The default is System scope.

FORCE

Indicates that use of the command prefix by this server should override pre-existing usage. **USE WITH EXTREME CAUTION.**

Example

`CPFX=FM`

Choosing a command prefix

z/OS Migrator supports a command prefix of up to eight characters, any or all of which can be alphabetic, however, a one-or-two-character, non-alphabetic prefix is usually the best choice. The DISPLAY OPDATA,PREFIX operator command (which can be abbreviated as "D O") can be used to determine which command prefixes are currently in use. For example:

D O

IEE603I *hh.mm.ss* OPDATA DISPLAY *nnn*

PREFIX	OWNER	SYSTEM	SCOPE	REMOVE	FAILDSP
(OPSSATM	LDM2	SYSTEM	NO	PURGE
(OPSSATM	LDM1	SYSTEM	NO	PURGE
!	OPSSOSF	LDM2	SYSTEM	NO	PURGE
!	OPSSOSF	LDM1	SYSTEM	NO	PURGE
\$	JES2	LDM2	SYSTEM	NO	SYSPURGE
\$	JES2	LDM1	SYSTEM	NO	SYSPURGE

Default None.

Note: If you receive message EFM0047S at server startup, the usual response is to choose another command prefix. It is possible, however, for the server to terminate in a manner that does not deactivate the command prefix. In this case, code “CPFX=(xx,FORCE)” in the startup configuration file.

CONGROUP

Description CONGROUP is a parameter used to specify the started task ID of the EMC Consistency Group address space on the LPAR for a z/OS Migrator server. The started task ID may be different on different LPARs, so change the value accordingly if different LPARs are used as owning servers. A modify SETSWAP command will be issued against this task ID to DISABLE/ENABLE congroups to allow the swap to occur, completing the migration.

Syntax CONGROUP=xxxxxxxx

Where:

xxxxxxxx

Is the EMC Consistency Group started task ID to be used in the z/OS Modify (F) command.

Example CONGROUP=EMCCNGRP

DB

Description The DB parameter can be either an execution or a global configuration parameter. It identifies the fully qualified, cataloged dataset that is the z/OS Migrator database. [“Step 11: Allocate the z/OS Migrator database” on page 60](#) provides more details on the z/OS Migrator database.

Syntax DB=*dsname*

Where:

dsname is the name of a dataset that is the z/OS Migrator database.

OR

DB= *DB_DSN*

Where:

DB_DSN

A one to 44 character z/OS Migrator database DSN.

Example

`DB=hlg.ZOSM300.DB`

Default

None.

DEBUG

Description

DEBUG is a parameter used as either an execution or a global configuration parameter. It is used to set various debugging switches, if needed.



CAUTION

DEBUG should only be used if requested by EMC Customer Support.

DEBUG specifies the debug options that can be utilized for problem isolation purposes.

Syntax

`xxDEBUG=mask,mask`

Where:

- ◆ *mask* is an 8-bit pattern used internally to turn on/off certain tracing information and messages, and is provided by EMC Customer Support.
- ◆ *xx* is the user's CPFY value.

Example

`xxDEBUG=A8A10000,00000000`

Default

None.

DIVSUB

Description

The Diversion Subset (DIVSUB) option allows the renaming process to be divided into smaller units, serializing each catalog pair for shorter periods. The number of data sets to be served within one enqueue bracket is now defined using a new parameter specified on the OPTION statement in the z/OS Migrator configuration member. The syntax for this option is:

`DIVSUB = n`

Where n is an integer between 0 and 32,767. The default for n is 0, indicating that the renaming process will continue to work as it always has, where all the data sets for a given catalog pair will be serviced within one continuous catalog enqueue.

If n is set to any *non-zero* value, the renaming process will be broken up into units of n datasets at a time. All components of a dataset, for example, VSAM, sphere or GDG, will be included within a diversion subset even if that number exceeds the subset value.

Example. DIVSUB = 50

If a generation data group includes 50 generations, all of them will be included within the same subset even if the diversion subset value is less than 50. However, if the DIVSUB number is reached or exceeded after the inclusion of one or a group of data sets, then selection will be broken off and processing for a catalog pair bracket will proceed.

PARMLIB

Description PARMLIB is an alias of CFG and is only used as an execution parameter.

Syntax PARMLIB=*dsname* (*member*)

Where:

dsname

The name of a partitioned dataset that is the parameter file.

member

The member of the partitioned dataset that holds the parameter library.

Example None.

REFRESH

Description REFRESH is a parameter only used as an execution parameter.

It is designed to be included in the EXEC statement of a new z/OS Migrator server execution that is intended to take over from an existing z/OS Migrator server execution. The primary purpose is to allow maintenance to be applied or a newer version of z/OS Migrator

to be started non-disruptively; that is, started while any active mirroring and migration continues to operate in the new z/OS Migrator server.

Syntax REFRESH

Example None.

SSID

Description SSID can be used as either an execution parameter or as a configuration parameter. The volume migrator and the logical migrator operate within a subsystem. This parameter identifies the subsystem you are using.

SSID is an alias of SUBSYS. You can use either parameter to specify the subsystem.

Syntax *SSID=subsysID*

Where:

subsysID

The name of the subsystem.

Example None.

START

Description The START parameter is an execution parameter. z/OS Migrator stores its operating parameters in common storage. This parameter determines whether these operating parameters are retained between executions.

Format START=WARM | COLD

Where:

WARM

(Default) The operating parameters persist between executions of z/OS Migrator. (HOT start can be used interchangeably with a WARM start.)

COLD

The operating parameters are discarded between executions of z/OS Migrator and new I/O interfaces are installed. If you perform a COLD start after a prior startup in the same IPL, all devices in all migrator volume groups are deactivated. However, deactivation does not occur if this start is the initial start for the IPL.

You can perform a non-disruptive FMM Server restart with the REFRESH operator command. EMC recommends using REFRESH rather than using the START=COLD parameter.

Note: [“SSID” on page 319](#) provides more information.

Example None.

SUBSYS

Description SUBSYS can be used as either an execution parameter or a configuration parameter. The z/OS Migrator server operates within a subsystem. The SUBSYS parameter identifies the subsystem you are using.

SUBSYS is an alias of SSID. You can use either parameter to specify the subsystem.

Format `SUBSYS=`*subsysID*

Where:

subsysID

The name of the subsystem.

Example `SUBSYS=EMCF`

TRACE

TRACE activates a diagnostic trace in volume migrations.



CAUTION

TRACE should only be used under the direction of EMC Customer Support.

Syntax TRACE

Example None.

Global configuration parameters

You can enter global configuration parameters in an input file, generally a member of a partitioned dataset. You can specify the file name with the `PARMLIB` execution parameter or with the `//CONFIG` dd statement.

The following parameters are used for both volume and logical migrations unless otherwise specified.

Note: A sample member, `CONFIG`, can be found in 'hlq.prodid.SAMPLIB' located in the Mainframe Enablers `SAMPLIB`.

Global configuration parameters are specified according to the following format:

- ◆ Each record in the configuration member may contain only one operation.
- ◆ Comment records are allowed and are indicated with an asterisk (*) in column 1.

CONSTANT_COPY_INT

Description `CONSTANT_COPY_INT` is a global configuration parameter for volume mirror/migrations. It specifies the update interval in seconds for constant copy mode. The default value is 60 seconds.

Syntax `CONSTANT_COPY_INT=seconds`

Note: You can abbreviate `CONSTANT_COPY_INT` as `CCINT`.

Example `CONSTANT_COPY_INT=55`

DWA_ADD_PAGE

Description

DWA_ADD_PAGE is a global keyword used only on disruptive start (START=COLD) or non disruptive start (REFRESH). It specifies the number of disabled 4K work area pages to be added to the default number.

Note: “START” on page 319 and “SSID” on page 319 provide more information.

By default, the number of 4K disabled work area pages allocated in E/SQA is set to 21*4K for each processor. Of the 21 pages, 20 are for use by the channel scan logic. A total of 20 pages allows for 648 channel command words (CCWs) for each channel program.

When a channel program exceeds this CCW limit, the EFMM310I message is written to the system log, for example:

```
EFMM310I CP SCAN CONDITION: SIO,NZD00CPY,543C,BUFREJ,SHORTAGE,9E,CCW,648
```

Note: Message EFMM310I in the *z/OS Migrator Message and Code Guide* provides more information.

The action taken by z/OS Migrator for a volume migration depends on whether the primary or mirror device is being accessed. If the primary device is being accessed, z/OS Migrator suspends, and optionally deactivates the mirroring with the primary device. If the mirror device is being accessed, z/OS Migrator rejects the job's read request to the mirror with a EFMM311I message. For example:

```
EFMM311I REQUEST REJECTED:PALFDR ,9956,EXCEED
```

Note: Message EFMM311I in the *z/OS Migrator Message and Code Guide* provides more information.

You can set DWA_ADD_PAGE to increase this allowance. The value you specify for DWA_ADD_PAGE is added to the default 21 pages to allow for larger channel programs. For each additional page added, z/OS Migrator can process an additional minimum of 35 CCWs. The total CCWs capable of being processed is:

$$\text{INT}((4096/116) * \text{DWA_ADD_PAGE}) + 648$$

As the pages are allocated to E/SQA, you must carefully consider the value you assign to `DWA_ADD_PAGE`, along with the `SQA` parameter in the `z/OS IEASYSxx PARMLIB` member.

Syntax `DWA_ADD_PAGE=additional 4K page count`

Where:

additional 4K page count is a value between 0 and 128.

Examples

- ◆ If you specify:

`DWA_ADD_PAGE=16`

`z/OS Migrator` can process an additional 564 CCWs. With the base 648 CCWs, this allows `z/OS Migrator` to process channel programs up to 1212 CCWs in size.

- ◆ If you specify:

`DWA_ADD_PAGE=64`

`z/OS Migrator` can process an additional 2259 CCWs. With the base 648 CCWs, this allows `z/OS Migrator` to process channel programs up to 2907 CCWs.

HEARTBEAT_INT

Description

`HEARTBEAT_INT` controls how quickly cross-system group status changes are recognized by `z/OS Migrator`. You should specify a fairly short interval to enhance product responsiveness; for instance, in recognizing configuration and volume group state charges. `HEARTBEAT_INT` also controls how long the host must be dormant before other FMM Servers stop waiting for it to respond.

Format

`HEARTBEAT_INT=interval`

Where:

interval

Two meanings:

- The number of seconds between timer-driven reads of the `z/OS Migrator` database.

- A third of the time in seconds that pass before other systems consider the current system is dormant. In other words, after 3 * *intervals*, other z/OS Migrator systems assume that the current system is no longer functioning and stop waiting for it to respond.

EMC advises that you set HEARTBEAT_INT to the same value on all systems.

The minimum value you can use is zero (0). An interval value of zero gives HEARTBEAT_INT a 15 second default. There is no set maximum value, however EMC advises that you do not use a number much higher than the default of 15 seconds. Status changes would take too long to take effect.

Example

```
HEARTBEAT_INT=10
```

IDLE**Description**

IDLE=nnn is used for logical migrations. It indicates the minimum number of seconds another LPAR must be mute before it is considered IDLE. This value may be altered using the IDLE command and displayed using the D DEBUG MISC command.

If system contention is suspected to cause z/OS Migrator servers to be incorrectly considered idle, the threshold for such judgment can be increased using either the INI parameter IDLE=nnnn, or the operator command IDLE nnnn. If this is an ongoing condition, then the INI file presents the most effective solution. If transient or when first noticed, the command provides an immediate remedy.

The default is 45 seconds; increasing the value by multiples of 15 seconds is the most sensible approach. This value affects the timeout of certain command processes, therefore you should not casually set the value higher than necessary, otherwise the command processor could be adversely affected.

Syntax

```
IDLE=nnnn
```

Where:

nnnn is the minimum number of seconds another LPAR must be mute before it is considered IDLE.

Example

```
IDLE=45
```

Default

```
45 seconds
```

LMPFX

Description The LMPFX parameter is used for logical migrations.



CAUTION

The value of LMPFX must be identical for all z/OS Migrator systems.

The z/OS Migrator server places a dataset requiring 10 cylinders of space on each volume that is the target of a migration. z/OS Migrator uses this dataset to communicate the status of mirroring operations for managed dataset extents on the volume.

Syntax `LMPFX=dsn_prefix`

Where:

dsn_prefix is concatenated with ".LMvolser.DB" to form the unique name that the z/OS Migrator server can allocate, as required, on each target volume to communicate track mirroring status.

Each generated dataset name must be unique within your environment and if multiple configuration datasets are being used, the *dsn_prefix* value must be specified in exactly the same way for every participating z/OS Migrator server.

Once a prefix has been decided upon it is very difficult to change it. If this was required, it would be necessary to complete or deactivate all migration groups, shut down all the active z/OS Migrator servers, delete and reallocate the main z/OS Migrator database, and then perform a z/OS Migrator cold start.

Example `LMPFX=hlq.EXTMAP`

This creates a *hlq*.EXTMAP.LMxxxxxx.DB extent map dataset on every target volume, where *xxxxxx* is the target volume serial number.

Default None.

MAX_CHANNEL_IO

Description

Specifies the maximum number of concurrent I/O requests that should be issued to a channel during the *Copy* phase or *Mirror Synchronization* phase.

The MAX_CHANNEL_IO limit applies to any active I/O against a channel group, whether read or write. If the source and target devices are on the same channel group, z/OS Migrator allows total concurrent I/O requests on the channel to equal the value specified.

Together, MAXIO, MAX_CHANNEL_IO, and MAX_DEVICE_IO control the scheduling and distribution of concurrent I/O requests.

Syntax

`MAX_CHANNEL_IO=Requests`

Note: You can also write this parameter as MAXCIO.

Where:

Requests is the maximum number of concurrent I/O requests. The smallest value is **0**. There is no theoretical maximum value; however, the largest practical value is the current MAXIO value.

Example

`MAX_CHANNEL_IO=5`

Default

25

Note: The MAXCIO, MAXDIO, IDLE, and DEBUG parameters can be changed on each LPAR using a system console command.

The syntax for MAXCIO, MAXDIO, and IDLE is as follows:

`xxMAXCIO nn`

`xxMAXDIO nn`

`xxIDLE=nnnn`

Where *xx* is the user's CPFY value, *nn* is an integer, and *nnnn* is the number of seconds.

The DEBUG bit mask uses specific bit positions set using Technical Support's direction, and the syntax is:

`xxDEBUG ***** ,*****`

Note that the SET command is not used for these parameters.

MAX_DEVICE_IO

Description MAX_DEVICE_IO specifies the number of concurrent I/O requests that can be issued to a device during migration or mirror synchronization. If you do not specify a value, z/OS Migrator uses a default of three (3). The maximum value is five (5). The minimum value is zero (0). If you enter zero or do not specify MAX_DEVICE_IO, z/OS Migrator uses the default of 3.

Together, MAXIO, MAX_CHANNEL_IO, and MAX_DEVICE_IO control the scheduling and distribution of concurrent I/O requests.

Syntax MAX_DEVICE_IO=*Requests*

Where:

Requests is the maximum number of concurrent I/O requests.

Note: You can also write this parameter as MAXDIO.

Example MAX_DEVICE_IO=2

Default 3

MAXIO

Description MAXIO determines the maximum overall number of I/O requests that can be active at one time on the server. Since I/O buffers and control areas are allocated based on the MAXIO value, choose a number appropriate to the resources available.



CAUTION

This parameter has a direct impact on the amount of page fixed memory used by z/OS Migrator. Too large a value can negatively affect overall system performance or exhaust available page space.

Approximately 1 MB of memory is allocated for each integer you add to the MAXIO specification; for example, 25 MB for a MAXIO value of 25. The memory allocation is fixed during active I/O.

Together, MAXIO, MAX_CHANNEL_IO, and MAX_DEVICE_IO control the scheduling and distribution of concurrent I/O requests.

Syntax MAXIO=*number*

Where: *number* is the maximum number of overall requests that can be active at one time on the server. The minimum value is 0. If you choose 0, or do not specify a MAXIO value, the number of overall requests defaults to 25.

Example MAXIO=15

Default 25

MAXTRK

Description This optional parameter specifies the size of I/O operation in tracks, that is, z/OS Migrator's I/O copy operations are to transfer less than a full cylinder (1 extent) of data. The MAXTRK value is used to reduce the application response time impact of z/OS Migrator Copy operations, immediately following *Activation*.

For example: MAXTRK=5 causes z/OS Migrator to move one extent in three I/O operations, dividing the time the device is unavailable to application I/O operations into three short windows, as opposed to a single, longer window.

Syntax MAXTRK=*n*

Where: *n* is a value from 1 to 15

Example MAXTRK=5

OPTIONS

Description Sets various processing options in logical migrations.



CAUTION

The default settings should be used unless otherwise directed by EMC Customer Support. You can reset option values using the z/OS Migrator operator SET command.

Syntax `OPTIONS=value`

Where *value* is one of the following:

<code>ACTVCMP</code> or <code>NOACTVCMP</code> (dataset migration only)	Active completion allows the source volumes to potentially be varied offline without recycling the processes using source datasets. <code>NOACTVCMP</code> is the default setting.
<code>AUTOTRM</code> or <code>NOAUTOTRM</code> (dataset migration only)	Automatic Termination means that groups will immediately attempt to go to completion as soon as divert status is achieved. Otherwise the <code>COMPLETE</code> command must be issued. <code>AUTOTRM</code> is the default setting.
<code>TRMCENQ</code> or <code>NOTRMCENQ</code> (dataset migration)	<code>TRMCENQ</code> means that z/OS Migrator will issue an ENQ request against the target datasets during the migration. <code>TRMCENQ</code> is the default setting.
<code>AUTODEL</code> or <code>NOAUTODEL</code> (dataset migration)	<code>AUTODEL</code> causes the group control blocks to automatically be freed once the group completes. <code>AUTODEL</code> is the default setting.
<code>AUTOHK</code> or <code>NOAUTOHK</code>	<code>AUTOHK</code> causes a group activation to automatically re-activate the system interfaces. <code>AUTOHK</code> is the default setting.
<code>AUTOCLN</code> or <code>NOAUTOCLN</code>	<code>AUTOCLN</code> causes the server to automatically remove interfaces and delete common resources at shutdown with no active groups. <code>AUTOCLN</code> is the default setting.
<code>SDUMP</code> or <code>NOSDUMP</code>	<code>SDUMP</code> results in an <code>SDUMP</code> being taken at any server failure. <code>SDUMP</code> is the default setting.

Example `OPTIONS=ACTVCMP`

Default All defaults are the positive values, except for `NOACTVCMP`.

SSUB#

Description SSUB# specifies the number of concurrent service subtasks within the z/OS Migrator server that can simultaneously perform work.



CAUTION

This option should only be used if requested by EMC Customer Support.

Syntax `SSUB#=number`

Where:

number is the number of subtasks that will be initiated.

Note: The SSUB # must be greater than 4. If the value is less than 4, a limitation error is issued and the process will skip to the end of the command.

Example: `SSUB#=8`

Default: 8

SUBNAME

Description z/OS Migrator server operates within a z/OS subsystem. The SUBNAME parameter identifies the subsystem identifier you choose.

Note: This does not correspond with the hardware SSID.

Syntax `SUBNAME=subsysID`

Where:

subsysID is the name you specify for the subsystem.

Example `SUBNAME=EMCF`

Choosing a subsystem identifier

You should decide on a three- or four-character subsystem identifier that is unique across the installation. The DISPLAY SSI,LIST operator command (can be abbreviated to “D SSI,L”) can be used to determine which subsystems are currently active and/or defined on each system.

For example:

```
D SSI,L
IEFJ100I  hh.mm.ss  SSI DISPLAY nnn
SUBSYS=JES2 (PRIMARY)
      DYNAMIC=YES      STATUS=ACTIVE      COMMANDS=REJECT
SUBSYS=MSTR
      DYNAMIC=NO       STATUS=ACTIVE      COMMANDS=N/A
SUBSYS=LOGR
      DYNAMIC=YES      STATUS=ACTIVE      COMMANDS=REJECT
SUBSYS=SMS
      DYNAMIC=YES      STATUS=ACTIVE      COMMANDS=REJECT
SUBSYS=RACF
      DYNAMIC=YES      STATUS=ACTIVE      COMMANDS=REJECT
SUBSYS=EMC
      DYNAMIC=YES      STATUS=ACTIVE      COMMANDS=ACCEPT
SUBSYS=GOAL
      DYNAMIC=NO       STATUS=INACTIVE    COMMANDS=N/A
```

It is important that the chosen subsystem name does not match the z/OS Migrator started task cataloged procedure name, unless you intend to start z/OS Migrator under control of the MSTR subsystem, with no SYSOUT data definitions, and with the cataloged procedure saved in the IEFPSI library specified in the Master JCL.

If, for any reason, you decide to use the same name (for example, z/OS Migrator) for both the subsystem and the cataloged procedure, either the SYSOUT DD statements must be removed and the cataloged procedure moved to the IEFPSI library, or the start command must specify the primary job entry subsystem. For example:

```
S ZOSM, SUB=JES2
```

Default None.

SWAPOPT

Description SWAPOPT specifies swap options at the global level in volume migrations. You can override these at the group level with the Volume Group Configuration parameter SWAPOPT.

Note: If you omit setting parameters through SWAPOPT, SCF(DAS) sets values for these parameters.

Syntax `SWAPOPT=(opt1[,opt2[,... ,optn]])`

Where:

opt can be as follows:

`[NO]BYPassConCurrentcopy`

Specifies whether Concurrent Copy may be bypassed by swap processing using the BypassConcurrentCopy option. The use of this option causes active concurrent copy jobs to fail at the time of swap.

`[NO]BYPASSSYSTEMSCOUNT | [NO]BYPSYSCNT | [NO]BSC`

Specifies whether to bypass the requirement that all other LPARs with connectivity to Symmetrix have EMCSCF running.

`[NO]BYPASSSNAPSESSION`

Specifies whether EMC TimeFinder® SNAP™ sessions may be bypassed by swap processing.

`CFW=NO | OFF | RESume | BYPASS |OFFVALidation`

Controls cache fast write (CFW).

`NO`

Indicates that CFW must be inactive for the SWAP to occur.

`OFF`

Indicates that CFW will be turned off on the source device controller, if it is active.

`RESume`

Indicates that CFW will be turned off on the source device controller if it is active. If CFW is active, then it is enabled for the target device controller.

BYPASS

Indicates that CFW processing will be ignored.

OFFVALidation

Indicates that CFW will be turned off during group validation.

VOLumePrefix(*pp*)

Specifies a two-character prefix (*pp*) to be used to modify the source device's volser after a successful swap. The full volser will be *ppccuu*, where the prefix will be followed by the device MVS *cuu*.

You can use this parameter to assure that, if mounted, the source volume has a unique volser from the migration target. If you use dashes for the two characters (VOLumePrefix(--)), there will be no volser relabel.

The volume prefix value is applied to the z/OS device number (CUU) to form the new volser. For example, if the z/OS device number of the source device is 054C and you select the VOLP(ZZ) value, the resulting volser is ZZ054C.

Example None.

Authorization Return Codes

This appendix provides all possible return codes and their meaning from authority checking. The topics are:

- ◆ Authority checking 336
- ◆ Reading/writing the z/OS Migrator security record..... 337

Authority checking

Table 1 documents all possible return codes and their meanings from authority checking.

Table 1 z/OS Migrator authority checking return codes

Return code decimal (hex)	Description	Reason
00 (00)	Passed authority checking	Normal return code. Authorization is good.
28 (1C)	z/OS Migrator license has expired	z/OS Migrator temporary license has expired. Contact the EMC Customer Support Center for technical assistance.
64 (40)	Program check	Verify input statements. Contact the EMC Customer Support Center for technical assistance.
516 (204)	Error loading EFFMUKEY	Verify that the correct library is pointed to.
520 (208)	Error locating CPU PCCA	Rerun job. Contact the EMC Customer Support Center for technical assistance.
524 (20C)	Error in caller identification	Verify input statements. Contact the EMC Customer Support Center for technical assistance.
552 (228)	Error with z/OS Migrator version definition	Verify that the proper z/OS Migrator version is specified.
556 (22C)	EFMMUKEY I/O error	Contact the EMC Customer Support Center for technical assistance.
560 (230)	Incorrect EFMMUKEY version	Contact the EMC Customer Support Center for technical assistance.
564 (234)	EFMMUKEY checksum error	Contact the EMC Customer Support Center for technical assistance.

Reading/writing the z/OS Migrator security record

Table 2 documents all possible return codes and their meaning from the program for reading/writing the z/OS Migrator security record.

Table 2 z/OS Migrator security record return codes

Return code decimal (hex)	Description	Reason/response
00 (00)	Successful	Normal return code. Authorization is good.
200 (C8)	BLDL error	Missing SECCOM DD statement or dataset does not exist or is misspelled.
204 (CC)	BLDL list in error	EFMMKEYS load module has multiple text records or the format is invalid. Contact the EMC Customer Support Center for technical assistance.
208 (D0)	TTR conversion error	Verify that JCL is correct. Contact the EMC Customer Support Center for technical assistance.
212 (D4)	Unsuccessful I/O operation (read)	The first eight bits of RC are copied from the ECB used for the EXCP request. Refer to description of IOBECBCC field in the DFSMS/z/OS documentation. Contact the EMC Customer Support Center for technical assistance.
216 (D8)	Unsuccessful I/O operation (write)	The first eight bits of RC are copied from the ECB used for the EXCP request. Refer to description of IOBECBCC field in the DFSMS/z/OS documentation. Contact the EMC Customer Support Center for technical assistance.
220 (DC)	GETMAIN error	Unable to allocate required storage. Rerun job.
224 (E0)	RDJFCB error	Dataset does not exist or is misspelled.
228 (E4)	Bad buffer address	Verify that JCL is correct. Contact the EMC Customer Support Center for technical assistance.
232 (E8)	Invalid function code	Verify that JCL and input statements are correct. Contact the EMC Customer Support Center for technical assistance.

Catalog Structure

This appendix describes the structure of the ICF catalog. Topics include:

- ◆ Overview 340
- ◆ Catalog management considerations 342
- ◆ Catalog diagnostic recommendations 343
- ◆ BCS/VVDS/VTOC synchronization 343
- ◆ Dataset alias synchronization 346
- ◆ BCS/VVDS structure check 348
- ◆ BCS integrity check 349

Overview

As z/OS Migrator is a dataset-level migration tool that includes updating the information in the ICF catalog, it is very important that you have a high-level outline of ICF catalog structure to help you perform diagnostics for the pre-migration process.

The ICF catalog records are stored in two components: the basic catalog structure (BCS) and the VSAM volume dataset (VVDS). The BCS can be considered the main catalog, whereas VVDS can be considered an extension of the VTOC. A BCS can point to multiple VVDSs and a VVDS can point to multiple BCSs, which signify a one-to-many relationship.

The BCS is a VSAM KSDS and its primary function is to point to the volumes on which a dataset is located. The BCS is created when either a user or master catalog is defined using Access Method Services (AMS). A BCS does not have to be on the same volume as the dataset it references, and there can be more than one BCS on a volume.

The VVDS is a VSAM ESDS that contains information required to process VSAM datasets. In a Storage Management Subsystem (SMS) environment, it also contains the volume-related information for non-VSAM SMS-managed datasets. There is one VVDS on each DASD volume that contains a VSAM or SMS-managed dataset cataloged in an ICF catalog. The VVDS is always on the same volume as the datasets it references. For this reason, any volume containing a BCS will also contain a VVDS because the BCS is itself a VSAM dataset.

The records in both the VVDS and BCS consists of variable length cells and subcells. The two cell types that are often referred to are the VSAM volume record (VVR) and the non-VSAM volume record (NVR), which are both held in the VVDS. The VVR contains information relating to VSAM datasets and the NVR contains information relating to non-VSAM SMS-managed datasets. Most datasets will have entries in only one VVDS except multi-volume datasets, which can have entries in the VVDS of each volume to which they are allocated (type Q for a secondary record).

Table 1 Dataset entries

Information	VSAM dataset	SMS-managed non-VSAM dataset	Non-SMS non-VSAM dataset	Uncataloged datasets
Volume	BCS	BCS	BCS	n/a
Dataset Type	BCS	BCS	BCS	n/a
Association	BCS	BCS	BCS	n/a
Ownership	BCS	BCS	BCS	n/a
SMS Class info	BCS & VVDS	BCS & VVDS	n/a	n/a
Dataset Attributes	VVDS	VVDS & VTOC	VTOC	VTOC and/or VVDS
Extent Description	VVDS & VTOC	VTOC	VTOC	VTOC and/or VVDS
Catalog Name	VVDS	VVDS	n/a	n/a

Catalog management considerations

Separate catalog(s) and unique target dataset high-level qualifiers should be established as part of the migration planning process. The use of dedicated catalogs and unique target high-level qualifiers will mitigate catalog growth impact of adding large numbers of target dataset entries and will simplify the management of post-migration dataset deletion tasks.

BCS record types

Each cell in each catalog record has a cell type field. The first cell type field in each record is also the record type field (or ID) identified by the DIAGNOSE command. The following are the possible BCS ID record types and their one-character identifiers:

- ◆ A–non-VSAM dataset
- ◆ B–generation data group
- ◆ C–cluster
- ◆ D–data component
- ◆ E–VSAM extension record
- ◆ G–Alternate index
- ◆ H–GDS
- ◆ I–Index component of a cluster
- ◆ J–GDG extension cell
- ◆ L–library
- ◆ R–path
- ◆ T–true name
- ◆ U–user catalog connector
- ◆ W–volume
- ◆ X–alias

VSAM volume record (VVR)

- ◆ Type **Z** for a primary record.
- ◆ Type **Q** for a secondary record and for datasets allocated with IMBED **
- ◆ Type **N** Non-VSAM Record (NVR) for non-VSAM datasets.

Catalog diagnostic recommendations

As the z/OS Migrator data movement process includes catalog alteration, it is essential that you start with a healthy catalog, as well as verifying that the candidate datasets are correctly cataloged on the system(s) to which z/OS Migrator will require access. It is recommended that EMC's Catalog Solution or IDCAMS be used to perform the required diagnostics.

Note: Under a licensing agreement with EMC, Rocket Software has taken over sales and support of EMC's Catalog Solution. Please contact Rocket to report an issue. You can call the Rocket Software Support Hotline at (617) 614-4323, or visit the Rocket Mainframe Productivity Suite website: <http://www.rocketsoftware.com/mainframe-productivity/rcp-access.html>

BCS/VVDS/VTOC synchronization

Catalog Solution's diagnostic routines check data structure, index integrity and BCS/VVDS/VTOC synchronization, and can be directed to generate IDCAMS control cards to fix detected errors. Specifically, the IDCAMS control cards for DEFINE RECATALOG can be generated to recreate missing VSAM and non-VSAM catalog entries.

Additionally, there are diagnostic routines to detect unnecessary catalog entries and generate IDCAMS control cards for DELETE NOSCRATCH to remove those entries. The diagnostic routines may be directed to the entire catalog environment or specifically assigned to a catalog, volume, or object (catalog, VVDS or volume) through the use of keywords that allow specific or generic parameters.

Running Catalog Solution's diagnostic programs on a regular basis will provide a consistent catalog synchronization.

Sample VVDS to BCS synchronization JCL

Use Catalog Solution's GENERATE-RECATALOG-CARD (GRC) Diagnostic utility to read the contents of the VVDS, and make sure a corresponding BCS entry exists and points to the correct volume. Both VVRs (VSAM) and NVRs (non-VSAM) are processed. If the BCS entries do not exist, they are re-created by generating IDCAMS DEFINE...RECATALOG control cards and writing them to a PS dataset for subsequent use by IDCAMS.

The following Catalog Solution JCL checks all VVDS records and processes those that contain a reference to ICFCAT.PROD01. IDCAMS DEFINE ... RECATALOG control cards are then created for each dataset which has one or more VVDS records referring to that BCS but no corresponding entry in that BCS:

```
//STEP01 EXEC PGM=CSLMAIN
//STEPLIB DD DSN=authorized.loadlib,DISP=SHR
//CCFILE DD DSN=grc.cc,DISP=(NEW,CATLG),UNIT=SYSDA,
// SPACE=(TRK,(1,1),RLSE)
//SYSIN DD *
DIAGNOSE TEST=GRC,INCCAT=ICFCAT.PROD01,INCVOL=*
/*
```

- ◆ Supports both VVR and NVR VVDS entries.
- ◆ Cannot recreate exact BCS entries for multi-volume datasets.
- ◆ Cannot recreate exact BCS entries for path statements.

The following IDCAMS JCL corrects the discrepancies:

```
//STEP01 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD DSN=grc.cc,DISP=OLD
```

Note: Any multi-volume dataset control card needs to be modified to include ALL of the volumes it resides on in the correct order (use a prior LISTCAT to determine this).

Any PATH recatalog card needs to be modified to include the correct PATH name.

Sample BCS to VVDS synchronization JCL

Use the Catalog Solution GENERATE-DELETE-CARD (GDC) Diagnostic utility to read the contents of all online BCSs and ensure that required VVR records are in the VVDS for clusters on volume VSAM01. If datasets are not on the volume, the corresponding BCS entries should be removed with the IDCAMS DELETE ... NOSCRATCH control cards that are generated and written to CCFILE.

The following JCL compares all online BCSs to the VVDS on volume VSAM01:

```
//STEP01 EXEC PGM=CSLMAIN
//STEPLIB DD DSN=authorized.loadlib,DISP=SHR
//CCFILE DD DSN=gdc.cc,DISP=(NEW,CATLG),UNIT=SYSDA,
// VOL=SER=xxxxxx,SPACE=(TRK,(1,1),RLSE)
//SYSIN DD *
DIAGNOSE TEST=GDC,INCCAT=**,INCVOL=VSAM01
/*
```

The following JCL corrects the discrepancies:

```
//STEP01 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD DSN=gdc.cc,DISP=OLD
```

Sample VTOC to BCS synchronization JCL



CAUTION

The following diagnostic is not intended for use on SMS-managed volumes. For SMS-managed volumes, use DIAGNOSE TEST = GENERATE-RECATALOG- CARD (GRC).

Use the Catalog Solution GENERATE-DEFINE-NONVSAM (GDN) Diagnostic utility to read the contents of the VTOC on volume VSAM01 and ensure that a corresponding catalog entry points to the correct volume. If catalog entries do not exist, they are re-created by generating IDCAMS DEFINE ... NONVSAM control cards that are written to CCFILE.

The following JCL compares the VTOC on volume VSAM01 to all online BCSs:

```
//STEP01 EXEC PGM=CSLMAIN
//STEPLIB DD DSN=authorized.loadlib,DISP=SHR
//CCFILE DD DSN=gdn.cc,DISP=(NEW,CATLG),UNIT=SYSDA,
// VOL=SER=xxxxxx,SPACE=(TRK,(1,1),RLSE)
//SYSIN DD *
DIAGNOSE TEST=GDN,INCCAT=**,INCVOL=VSAM01
/*
```

Sample VTOC to VVDS synchronization JCL

Use the Catalog Solution VTOC to VVDS Comparison Diagnostic utility to read a VVDS and the VTOC on a volume. It indicates any missing or extra entries and compares the extent information.

The following JCL compares the VTOC to the VVDS:

```
//STEP01 EXEC PGM=CSLMAIN
//STEPLIB DD DSN=authorized.loadlib,DISP=SHR
//SYSIN DD *
DIAGNOSE TEST=VV,INCVOL=volser
/*
```

Dataset alias synchronization

If a dataset is created prior to the creation of an alias, then the full dataset reference is placed in the master catalog. Additionally, if these datasets were to have shared access and the aliases were only defined on one system, the dataset will not be accessible as required. In these situations, the Catalog Solution ALIASCHECK command allows you to verify that alias entries relating to specified user catalogs are the same in all specified master catalogs. ALIASCHECK identifies alias entries not defined for selected user catalogs. If alias entries are not found in all specified master catalogs, ALIASCHECK can automatically define these entries in those master catalogs.

If an alias relates to an incorrect user catalog, ALIASCHECK will identify it, but will not perform a resynchronization. The COMPAREMAST or MASTERCAT1 and MASTERCAT2 keywords specify the master catalogs being checked. The related user catalogs are specified by the INCCAT or CAT1 and CAT2 keywords. Use RESYNCH to force a resynchronization of the two master catalogs. Running this procedure weekly will help ensure alias synchronization.

Sample ALIASCHECK JCL

- ◆ To verify that the alias entries pointing to user catalog ICF.USERCAT1 are identical in the master catalogs ICF.MASTERA and ICF.MASTERB, type the following JCL:

```
//jobname JOB
//stepname EXEC PGM=CSLMAIN
//STEPLIB DD DSN=authorized.loadlib,DISP=SHR
//SYSIN DD *
ALIASCHECK MASTERCAT1=ICF.MASTERA,
MASTERCAT2=ICF.MASTERB,
CAT1=ICF.USERCAT1,CAT2=ICF.USERCAT1
/*
```

- ◆ To verify that the alias entries pointing to user catalog ICF.USERCAT3 in master catalog ICF.MASTX are identical to alias entries pointing to user catalog ICF.USERCAT4 in master catalog ICF.MASTY (and to resynchronize the two master catalogs if a discrepancy exists), type the following JCL:

```
//jobname JOB
//stepname EXEC PGM=CSLMAIN
//STEPLIB DD DSN=authorized.loadlib,DISP=SHR
//SYSIN DD *
ALIASCHECK MCAT1=ICF.MASTX,
MCAT2=ICF.MASTY,
CAT1=ICF.USERCAT3,CAT2=ICF.USERCAT4,
RESYNCH
/*
```

This allows the alias entries to be defined in either master catalog if it is required.

- ◆ To verify that user catalogs residing in four separate master catalogs have matching alias entries, type the following JCL:

```
//jobname JOB
//stepname EXEC PGM=CSLMAIN
//STEPLIB DD DSN=authorized.loadlib,DISP=SHR
//SYSIN DD *
ALIASCHECK CMAST=MASTERA,CMAST=MASTERB,
CMAST=MASTERC,CMAST=MASTERD,
INCCAT=**
/*
```

BCS/VVDS structure check

The Catalog Solution SYSTEMCHECK command invokes IDCAMS to DIAGNOSE ICFCATALOG (BCS to VVDS) the catalog environment specified by the OBJECT or INCCAT keywords.

As Catalog Solution processes each BCS, it automatically determines which VVDSs to process. If the VVDS's volume is online to the current system, the volume will be dynamically allocated.

Catalog Solution will issue messages informing you if a required volume was unavailable. When SYSTEMCHECK is executed for one or more BCSs (either by default or if BCS is explicitly specified), the INCVOL and EXCVOL keywords may be used to limit the volumes used for compare during DIAGNOSE ICFCATALOG processing.

Sample SYSTEMCHECK JCL

- ◆ To check the relationships in BCS ICFCAT.PROD01, use the following JCL:

```
//jobname    JOB
//stepname   EXEC PGM=CSLMAIN
//STEPLIB   DD DSN=authorized.loadlib,DISP=SHR
//SYSIN     DD *
SYSTEMCHECK OBJECT=ICFCAT.PROD01
/*
```

- ◆ To process all online catalogs on the system except BCS ICFCAT.TEST01, use the following JCL:

```
//jobname    JOB
//stepname   EXEC PGM=CSLMAIN
//STEPLIB   DD DSN=authorized.loadlib,DISP=SHR
//SYSIN     DD *
SYSTEMCHECK BCS,INCCAT=**,EXCCAT=ICFCAT.TEST01
/*
```

- ◆ To check all catalogs associated with the VVDS on DASD volume VSAM01, use the following JCL:

```
//jobname    JOB
//stepname   EXEC PGM=CSLMAIN
//STEPLIB   DD DSN=authorized.loadlib,DISP=SHR
//SYSIN     DD *
SYSTEMCHECK VVDS,INCVOL=VSAM01
/*
```

BCS integrity check

To ensure that a BCS and its associated backups are structured correctly, the EXAMINE parameter should be included on the Catalog Solution DUMP command for each scheduled BCS backup. The EXAMINE parameter causes IDCAMS EXAMINE to be invoked as part of CSL DUMP processing, and so allows both components of each BCS being dumped to be checked "on the fly" for KSDS-related structural problems.

Sample Catalog Solution DUMP JCL

This example shows how to back up catalog ICFCAT.PROD01 and check it for structural errors at the same time:

```
//jobname JOB . . .
//stepname EXEC PGM=CSLMAIN
//STEPLIB DD DSN=authorized.loadlib,DISP=SHR
//BACKUPDD DD DSN=backup.dataset.name,DISP=(,KEEP),...
//ERRFILE DD DSN=errfile.dataset.name,DISP=(,KEEP),...
//SYSIN DD *
    DUMP OBJECT=ICFCAT.PROD01,EXAMINE
/*
```

In this example, the data and index component of an ICF catalog is checked prior to the backup.

```
001      DUMP BCS,INCCAT=ICFCAT.PROD01,EXAMINE
        CSL03933I  CURRENT OBJECT BEING PROCESSED
                DSN=ICFCAT.PROD01

NONVSAM          367
GDG              0
CLUSTER         203
DATA             15
AIX              0
INDEX           11
ATL LIBENT      0
PATH            0
TRUENAME        0
UCAT            0
ALIAS           0
VOLUME          0
EXTENSION       0
ATL VOLENT      0
GDG EXT         0

          9  MASTCAT ALIASES BACKED UP
         596 DATA RECORDS BACKED UP

DUMP PROCESSING COMPLETED
```

RETURN CODE WAS 0

NUMBER OF CATALOGS PROCESSED: 1
 NUMBER OF VOLUMES PROCESSED: 0

HIGHEST RETURN CODE WAS 0

VERIFY

DATASET(ICFCAT.PROD01) -
 0IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0
 0

EXAMINE

NAME(ICFCAT.PROD01) -
 INDEXTTEST DATATEST) -

-IDC01700I INDEXTTEST BEGINS
 IDC11773I 37 KEYS PROCESSED ON INDEX LEVEL 1, AVERAGE KEY LENGTH: 6.4
 IDC11774I CURRENT INDEX CISIZE IS 4096, RECOMMENDED MINIMUM INDEX CISIZE IS 3584
 0IDC01724I INDEXTTEST COMPLETE - NO ERRORS DETECTED
 -IDC01701I DATATEST BEGINS
 IDC01709I DATATEST COMPLETE - NO ERRORS DETECTED
 IDC01708I 37 CONTROL INTERVALS ENCOUNTERED
 IDC01710I DATA COMPONENT CONTAINS 596 RECORDS
 IDC01711I DATA COMPONENT CONTAINS 0 DELETED CONTROL INTERVALS
 IDC01712I MAXIMUM LENGTH DATA RECORD CONTAINS 835 BYTES
 IDC01722I 99 PERCENT FREE SPACE
 0IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0
 0
 0IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0

Diagnostic Information and Utilities

This appendix describes how to run tests and report issues for diagnostic purposes. The topics are:

- ◆ Reporting results 352
- ◆ z/OS Migrator Group syntax (for Logical Migrations only)..... 352
- ◆ z/OS Migrator unresponsive, no apparent activity 357
- ◆ z/OS Migrator diagnostic facilities 360

Reporting results

It is expected that all executions of z/OS Migrator will be successful, however, unexpected results may occur. To ensure that any issues or problems are addressed as soon as possible, it is requested that all problems be documented with the appropriate supporting documentation to facilitate timely problem-resolution.

z/OS Migrator Group syntax (for Logical Migrations only)

z/OS Migrator Group syntax allows user selection of various levels of diagnostic messaging.

If recreating a problem situation, please include the DEBUG(EXTRA) setting on all test scenarios so that relevant data is produced for every test, successful or not.

```
SET ALLOCSEQ(DATASET | NONE | SIZE) -
TRACE(YES) SPHERE(Yes | No) DEBUG(EXTRA) -
SOURCE(DSN(MASK-LIST) VOLume(VOLUME) STGCL(SMSCLASS) -
EXCLUDE(MASK-LIST)) -
TARGET(DSN(mask) VOLume(VOL-LIST) STGCL(SMSCLASS))
```

Requested information for diagnostics

As part of the diagnostic process, the following information is requested:

1. All the SYSOUT from the z/OS Migrator server(s) output, as well as all JESMSG LG, JESJCL, and JESYSMSG. This should be saved in a dataset using the SDSF command, PRINT ODSN, or via the use of SDSF action "SE", and creating an ISPF member for each SYSOUT DD output.

The progress through the test phases can be annotated in the server log using a "COMMENT" command entered at the SDSF console command field. Enter the command prefix followed by an asterisk and text up to 70 characters, for example:

```
@$*I just started the test job
```

In Logical Migrations only, there are operator commands that will put the documentation in the server job log, for example:

```
Display Groups
Display Group <groupname>
Display DSN
Display DSN <groupname>
Display DSN <dsn-mask>
Display MEM EXT <dsn-mask>
Display PROC <groupname>
```

Users should familiarize themselves with these commands and think of them as situations present themselves. For example, if a group or test job seems to be hung, having a few of these displays in the log at the time will help.

2. In Logical Migration only, execute the z/OS Migrator COMPARE utility when the group state is "Mirror Active" for all datasets in the group.

Note: If the COMPARE utility is used in a multi-LPAR environment, it is possible to have mismatches as the other LPARs are not prevented from updating the data while the COMPARE is running. If a mismatch occurs, it is recommended that the comparison be rerun. If the same mismatch occurs, this may indicate a problem, but it is not conclusive.

3. If requested by EMC Customer Support, turn on the z/OS Migrator event tracing. Use the TRACE ON EID=123 command. The EID must match the value specified in the USR=GTF option, referenced in [Step 4](#).

Select the appropriate diagnostic information. The DEBUG flags may be specified to diagnose various functions in z/OS Migrator. Specify /@@DEBUG 0000FFFF,FFFF0000 to trace all I/O function processing. Other debug flags may be requested depending on the nature of the problem. Note that tracing will have a performance impact on z/OS Migrator processing.

Recommended DEBUG flags for typical testing:

```
DEBUG A8A10000,00000000
```

4. If a problem occurs that appears to be related to the I/O processing during the *Mirroring* or *Diversion* phases, then a GTF trace is recommended for the related volume:

```
//GTFMGS8   PROC TRACE='HLQ.GTF.TRACE28', MBR=GTFMGS
//*
//IEFPROC   EXEC PGM=AHLGTF, REGION=6M,
//          PARM='MODE=EXT, BLOK=500K, AB=40K, TIME=YES'
//IEFRDER   DD DISP=SHR, DSN=&TRACE
//SYSLIB    DD DSN=HLQ.JCL.LIB(&MBR), DISP=SHR
```

Where the SYSLIB input is as follows, change the device range in the IO=SSCH list to match the devices to be traced.

```
TRACE=IOP, SSCHP, CCWP, CSCH, HSCH, MSCH, PCI, USRP
IO=SSCH=(E9E3-E9EF, E9F3-E9FF)
IO=SSCH=(sourcecuu-sourcecuu, targetcuu-targetcuu)
* for a range of source and target cuu's
USR=(123)
CCW=(SI, CCWN=256, DATA=256, IOSB) END
```

To turn tracing on the server:

```
TRACE ON, EID=123
```

Note: Only start the trace once you are mirroring, otherwise it will get too many logs.

Set debug flags:

```
@@DEBUG A8A10000,00000000 where @@ is the CFPX
currently in use.
```

5. Provide application-related sysout or log showing errors.
6. Write an English language narrative of what the job was expected to do and how the result differed from the expected results. Make clear references to phases (*ACTIVATE*, *MIRROR*, *P-DIVERT*, and so forth), where reported error messages, anomalies or unexpected results occurred.
7. Provide LOGREC data. Many z/OS Migrator errors are captured as software records in LOGREC. This includes ABEND information in addition to symptom records, which will be created to describe particular error situations.

Ensure that the LOGREC on the current system is recording.

Following any tests, execute IFCEREP1 to generate LOGREC software records:

```
//XPALSFT PROC
//STEP1 EXEC PGM=IFCEREP1 , PARM= ' CARD '
//EREPPPT DD SYSOUT=*
//TOURIST DD SYSOUT=*
//SERLOG DD DSN=SYS1 . LOGREC , DISP=SHR
//ACCDEV DD DUMMY
//SYSIN DD DSN=SYS3 . DEVELOP . PROCLIB ( PALSFTI ) , DISP=SHR
```

Send the raw data to EMC Customer Support personnel.

8. In Logical Migration only, a dump of the VTOC, VTOCIX and VVDS for all volumes in the z/OS Migrator group before z/OS Migrator processing begins and after the anomaly. The following utilities are recommended; formatted LISTCAT output is useful as well:

- EMCSNAP:

```
DEBUG DATASET ( [ INDDNM(ddname) ] | [ SOURCE(dsname) ] )
```

Where either ddname or dsname can be used.

Run once with the source dataset and a second time with the target dataset.

- Dump a VTOC:

```
AMASPZAP
//PALPRINT JOB (EMC) , CLASS=A , MSGCLASS=X
//*
//PRNTVTOC EXEC PGM=AMASPZAP , PARM= 'IGNIDRFULL'
//SYSPRINT DD SYSOUT=*
//SYSLIB DD DISP=SHR , DSN=FORMAT4 . DSCB , KEYLEN=44 ,
// VOL=SER=PALSRC , UNIT=SYSALLDA
//SYSIN DD *
ABSDUMPT ALL
//*
```

- Catalog Solution — Dump VTOC entries:

```
DISPLAY INCVOL=&VOL1 , KEY=C' &SOURCEA.KSDS1 ' , RPTFRMT=DUMP , LIST=ALL
DISPLAY INCVOL=&VOL2 , KEY=C' &TARGETA.KSDS1 ' , RPTFRMT=DUMP , LIST=ALL
DISPLAY INCVOL=&VOL1 , KEY=C' &SOURCEB.ESDS2 ' , RPTFRMT=DUMP , LIST=ALL
DISPLAY INCVOL=&VOL2 , KEY=C' &TARGETB.ESDS2 ' , RPTFRMT=DUMP , LIST=ALL
DISPLAY INCVOL=&VOL1 , KEY=C' &SOURCEC.NONVSAM1 ' , RPTFRMT=DUMP , LIST=ALL
DISPLAY INCVOL=&VOL2 , KEY=C' &TARGETC.NONVSAM1 ' , RPTFRMT=DUMP , LIST=ALL
```

- AMASPZAP

Dump the VVDS and VTOCIX:

```
//PALIDCAM JOB , 'IDCAMS', CLASS=A, MSGCLASS=X
//STEP1 EXEC PGM=IDCAMS
//VVDSDD DD DSN=SYS1.VVDS.VPALSRC, DISP=SHR,
// VOL=SER=PALSRC, UNIT=SYSALLDA
//VTOCIXDD DD DSN=SYS1.VTOCIX.PALSRC, DISP=SHR,
// VOL=SER=PALSRC, UNIT=SYSALLDA
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
        PRINT INFILE(VVDSDD)
        PRINT INFILE(VTOCIXDD)
//
```

- Catalog Solution — Dump VVDS entries:

```
LIST OBJ=SYS1.VVDS.V&VOL1, KEY=C'&SOURCEA.KSDS1', GENERIC, LIST=ALL, OFFSETS
LIST OBJ=SYS1.VVDS.V&VOL2, KEY=C'&TARGETA.KSDS1', GENERIC, LIST=ALL, OFFSETS
LIST OBJ=SYS1.VVDS.V&VOL1, KEY=C'&SOURCEB.ESDS2', GENERIC, LIST=ALL, OFFSETS
LIST OBJ=SYS1.VVDS.V&VOL2, KEY=C'&TARGETB.ESDS2', GENERIC, LIST=ALL, OFFSETS
LIST OBJ=SYS1.VVDS.V&VOL1, KEY=C'&SOURCEC.NONVSAM1', GENERIC, LIST=ALL, OFFSETS
LIST OBJ=SYS1.VVDS.V&VOL2, KEY=C'&TARGETC.NONVSAM1', GENERIC, LIST=ALL, OFFSETS
```

- Catalog Solution — Dump catalog entries:

```
LIST OBJ=&CATHLQ.&SOURCE.USERCAT, KEY=C'&SOURCEA.KSDS1' GENERIC, LIST=ALL, OFFSETS
LIST OBJ=&CATHLQ.&SOURCE.USERCAT, KEY=C'&SOURCEB.ESDS2' GENERIC, LIST=ALL, OFFSETS
LIST OBJ=&CATHLQ.&SOURCE.USERCAT, KEY=C'&SOURCEC.NONVSAM1' GENERIC, LIST=ALL, OFFSETS
LIST OBJ=&CATHLQ.&TARGET.USERCAT, KEY=C'&TARGETA.KSDS1' GENERIC, LIST=ALL, OFFSETS
LIST OBJ=&CATHLQ.&TARGET.USERCAT, KEY=C'&TARGETB.ESDS2' GENERIC, LIST=ALL, OFFSETS
LIST OBJ=&CATHLQ.&TARGET.USERCAT, KEY=C'&TARGETC.NONVSAM1' GENERIC, LIST=ALL, OFFSETS
```

z/OS Migrator unresponsive, no apparent activity

If processing hangs, look at the job log and the messages and determine if time-stamped messages are being produced. If no activity is evident from the job log message stream, please issue the following commands at the console or via SDSF, using the defined command prefix.

```
Display Group groupname
Display MEM Group groupname
Display PROC groupname
```

Display any GRS conflicts - D GRS,C - and display the unit status of the z/OS Migrator control file volume and any volumes involved in z/OS Migrator testing if practical.

Interrogating the dataset allocation status

To interrogate the allocation status of a dataset:

- ◆ Use the following GRS command on all the LPARs participating in the migration. The dataset name can be specified in full, or as one or more high level qualifiers followed by an asterisk.

```
D GRS,RES=(SYSDSN,dataset name)
```

Requesting an address space dump

The DUMP COMM command takes an SVC dump that will be written to the dump datasets as defined by your site's system standards.

To request an SVC dump:

1. Take an SVC dump of the server address space:

```
DUMP COMM=(Description of the z/OS Migrator Symptom)
```

- At the prompt, enter the following:

Note: If you are not using the default subsystem name, z/OS Migrator, substitute the appropriate name for the DSPNAME= parameter in the command.

```
R YY JOBNAME=ZOSMjob, SDATA=(ALLNUC, PSA, CSA, SQA, TRT, RGN, SWA, GRSQ), CONT
R YY DSPNAME=(0001.*zos*), END
```

Where **ZOSMjob** is the server job name, and **zos** is the first three characters of your subsystem name ('ZOSM' is the default subsystem name).

Note: Technical Support personnel may require you to provide dumps of servers from all LPARs.

- A message similar to the following indicates that all requested data has been successfully dumped to the dump dataset:

```
IEA611I COMPLETE DUMP ON SYS1.X14.DMP00009 223
```

- A message similar to the following indicates that the dump is incomplete and not useful:

```
IEA611I PARTIAL DUMP ON SYS1.X14.DMP00009 223
```

You must resolve the problem that caused the partial dump to occur and take the dump again. Insufficient space on the dump device(s) is often the problem that causes partial dumps.

SYSDUMP dump

A dataset containing the output of a SYSDUMP contains a summary dump for the failing program, plus some system data for the failing task. SYSDUMP dumps are the only ABEND dumps that you can format with the Interactive Problem Control System (IPCS).

A SYSDUMP DD statement, utilizing a Generation Data Group (GDG) dataset is provided in the z/OS Migrator server started task JCL member, ZOSMPROC. Figure 12 (continued), displays a sample of ZOSMPROC.

Changing SYSDUMP defaults

The IBM-supplied IEADMR00 PARMLIB member for SYSDUMP defaults specifies "SDATA=(NUC,SQA,LSQA,SWA,TRT,RGN,SUM)", so dumped storage will only include CSA-resident control blocks that match the SYSDUMP criteria (pointed to by registers at the time of

the ABEND). Since a lot of important z/OS Migrator data are resident in ECSA, a short-term co-requisite is the Changedump command.

The following Changedump command may be used on each LPAR before starting z/OS Migrator for the first time. An accompanying change to PARMLIB member IEADMR00 will make the added CSA value effective at each IPL.

To change a dump:

```
CD SET, SYSMDUMP= (CSA) ,ADD
```

z/OS Migrator diagnostic facilities

z/OS Migrator uses the following diagnostic facilities available through z/OS Migrator and the mainframe operating system:

- ◆ Diagnostic messages
- ◆ z/OS dump services
- ◆ z/OS GTF tracing service
- ◆ z/OS logrec (EREP) data

Where a condition occurs that prevents continued mirroring, z/OS Migrator takes the following steps:

- ◆ Suspends processing
- ◆ Generates a symptom record in the logrec dataset

If an ABEND is associated with the error, z/OS Migrator formats VRA data and creates a logrec record.

To facilitate problem diagnosis, please retain the logrec software EREP records for EMC Customer Support. You can format the software records using the following JCL. The example assumes that SYS1.LOGREC is being used.

```
//STEP1 EXEC PGM=IFCEREPI, PARM='CARD'
//EREPT DD SYSOUT=*
//TOURIST DD SYSOUT=*
//SERLOG DD DSN=SYS1.LOGREC, DISP=SHR
//ACCDEV DD DUMMY
//SYSIN DD *
PRINT=PS
HIST=N
TYPE=S
ENDPARM
```

Note: The IBM publication, *EREP Reference* provides more information.

Interactions with EMC Software Products

This appendix explains how z/OS Migrator interacts with other EMC software products. The topics are:

◆ Introduction	362
◆ AutoSwap.....	362
◆ Catalog Solution.....	362
◆ Compatible Flash	362
◆ Concurrent Copy.....	362
◆ Consistency Groups.....	363
◆ Controller-based data movers	365
◆ ResourcePak Base.....	365
◆ TimeFinder/Clone Mainframe Snap and Concurrent Sessions	365
◆ TimeFinder Clone and z/OS Migrator operations.....	366
◆ TimeFinder and z/OS Migrator operations	369
◆ SRDF and z/OS Migrator operations	372

Introduction

The following sections explain how z/OS Migrator's logical and volume migrations rely on or interact with other EMC software products.

AutoSwap

Active mirror pairs and migration volumes prior to the migration SWAP are excluded from S/DAS SWAP actions. Because an SDDF session tracks changes to the primary, an S/DAS SWAP would lose this capability.

AutoSwap runs as a service under SCF, a service of ResourcePak Base.

Catalog Solution

The Catalog Solution SCAN exit CS#MULTI exit can be used to generate a Multi Volume VSAM Cluster Report.

[“z/OS Migrator dataset selection planning” on page 79](#) recommends the use of a listing utility such as SCAN in EMC Catalog Solution.

Migrations do not interact with Catalog Solution.

Note: Under a licensing agreement with EMC, Rocket Software has taken over sales and support of EMC's Catalog Solution. Please contact Rocket to report an issue. You can call the Rocket Software Support Hotline at (617) 614-4323, or visit the Rocket Mainframe Productivity Suite website: <http://www.rocketsoftware.com/mainframe-productivity/rcp-access.html>

Compatible Flash

The product interaction for EMC Compatible Flash® and FlashCopy (an IBM product) is the same as for TimeFinder Mirror and Snap.

FlashCopy and Compatible Flash operations are blocked for active volumes.

Concurrent Copy

See [“TimeFinder/Clone Mainframe Snap and Concurrent Sessions” on page 365](#).

Consistency Groups

In interactions between z/OS Migrator and EMC Consistency Groups, Volume Migrator does not support migrating volumes in a consistency group definition, unless the target is also in the same consistency group.

z/OS Migrator version 3.0 provides SETSWAP functionality to third parties for temporarily disabling AutoSwap to avoid conflicts between possible ConGroup/AutoSwap swaps and swaps done as part of a volume migration. The bits that are visible in z/OS Migrator are: New config member parameter, CONGROUP=cccccccc - specify the task name of the ConGroup server.

If the CONGROUP parameter is specified in the config member, then z/OS Migrator automatically issues the SETSWAP Disable command to ConGroup using the Modify (F) operator command just before initiating its swap to conclude the migration of a volume pair.

The user sees messages similar to the following in the z/OS Migrator JESMSGLG and the system log:

```
IEA630I OPERATOR EFMCMD NOW ACTIVE, SYSTEM=X04 , LU=EFMMMCS
F PALCGRP,DAS,SETSWAP GRP TSTGRP1 DISABLE
EFM9999I COMMAND COMPLETED 009
F PALCGRP,DAS,SETSWAP GRP TSTGRP1 DISABLE
CGRS598I (00007) SETSWAP DISABLE completed:
Group TSTGRP1 now DISABLED
Total groups processed : 1
Successful : 1
Failed : 0
IEA631I OPERATOR EFMCMD NOW INACTIVE, SYSTEM=X04 , LU=EFMMMCS
```

Once Autoswap is disabled for the associated group, it will remain disabled as long as there are no active z/OS Migrator swaps for devices in the group and at least 30 seconds has elapsed (to accommodate imminent additional swaps). Once there are no outstanding swaps, you will see messages similar to the following:

```
IEA630I OPERATOR EFMCMD NOW ACTIVE, SYSTEM=X04 , LU=EFMMMCS
F PALCGRP,DAS,SETSWAP GRP TSTGRP1 ENABLE
EFM9999I COMMAND COMPLETED 091
F PALCGRP,DAS,SETSWAP GRP TSTGRP1 ENABLE
CGRS598I (00008) SETSWAP ENABLE completed:
Group TSTGRP1 now ENABLED
Total groups processed : 1
Successful : 1
Failed : 0
IEA631I OPERATOR EFMCMD NOW INACTIVE, SYSTEM=X04 , LU=EFMMMCS
```

If the user specifies CONGROUP=, the devices involved in the migration are part of an AutoSwap group, and the SETSWAP Disable fails, then the user will see message EFMM261E:

```
EFMM261E SETSWAP failed for vvvvvv(ccuu)
```

Where 'vvvvvv' is the volume serial number of the source volume and 'ccuu' is the device address. If this occurs, then the z/OS Migrator migration group will be deactivated. Messages describing the failure will be found in the z/OS Migrator server log and the system log. After correcting the indicated problem, the migration may be restarted.

The following table lists the potential congroup migration arrangements and whether or not each configuration is allowed or denied (No).

Table 1 Consistency Group migration requirements

Consistency Group Migration Requirements			
Source	Target	Allowed	Notes
Not in CG/STAR/CAX	In CG/STAR/CAX	Yes	SET SWAP DISABLE/ENABLE commands issued to CG address space prior to and following SWAP.
In CG/STAR/CAX	Not in CG/STAR/CAX	No	Migrations are allowed within the same congroup. All other options are denied.
In same CG/STAR/CAX as Target	In same CG/STAR/CAX as Source	Yes	SET SWAP DISABLE/ENABLE commands issued to CG address space prior to and following SWAP.
In Different CG/STAR/CAX as Target	In different CG/STAR/CAX as Source	No	Migrations are allowed within the same congroup. All other options are denied.

Controller-based data movers

Because z/OS Migrator is host-based, it does not support the concurrent use of products that operate by moving data strictly within the controller.

ResourcePak Base

If you are using the volume migration feature of z/OS Migrator, you need to have SCF active and an implied or explicit association between z/OS Migrator and SCF, either by default or with the `//SCF$xxx DD` statement. EMCSCF is a service of ResourcePak Base.

[“Installing/upgrading EMCSCF \(Mainframe Enablers\)”](#) on page 43 explains the use of EMCSCF.

TimeFinder/Clone Mainframe Snap and Concurrent Sessions

Migrations with *Action on Complete* specifying swaps that cross Symmetrix system boundaries may terminate or produce unpredictable results on concurrent sessions.

This includes:

- ◆ Traditional snap
- ◆ Virtual device snap
- ◆ Concurrent copy operations

These technologies use Symmetrix features that cannot be dynamically moved from one Symmetrix to another.

TimeFinder Clone and z/OS Migrator operations

Tables 2, 3, and 4 describe which operations are permissible using TF Clone source volumes belonging to a z/OS Migrator Group.




Table 2 TF Clone and Migrator operations before migration promotion

When TF Clone volume pairs are created BEFORE MIGRATION PROMOTION, and the TF/Clone target volumes are non-z/OS Migrator volumes...

TF Clone Source	TF Clone Target	Promotion	During Activation	During AUTOSWAP Validation
z/OS Migrator Source: S1	Non-z/OS Migrator Volume	Clone and Promotion allowed	Activation allowed*	Validation allowed
z/OS Migrator Target: T1	Non-z/OS Migrator Volume	Clone and Promotion allowed	Activation allowed*	Validation allowed
*If BYPSNAP is specified, AUTOSWAP issues the following message: FMMS475W (00003)(PID 00063) "FROM" DEVICE xxxx HAS AN ACTIVE SNAP SOURCE DEVICE SESSION: ALLOWED.				
z/OS Migrator Source: S1	Non-z/OS Migrator Volume	Clone and Promotion allowed	Activation not allowed*	Validation not allowed
z/OS Migrator Target: T1	Non-z/OS Migrator Volume	Clone and Promotion allowed	Activation not allowed*	Validation not allowed
*If NOBPSNAP is specified, AUTOSWAP issues the following message: FMMS476E (00001)(PID 00002) "FROM" DEVICE xxxx CANNOT HAVE AN ACTIVE SNAP SOURCE DEVICE SESSION.				

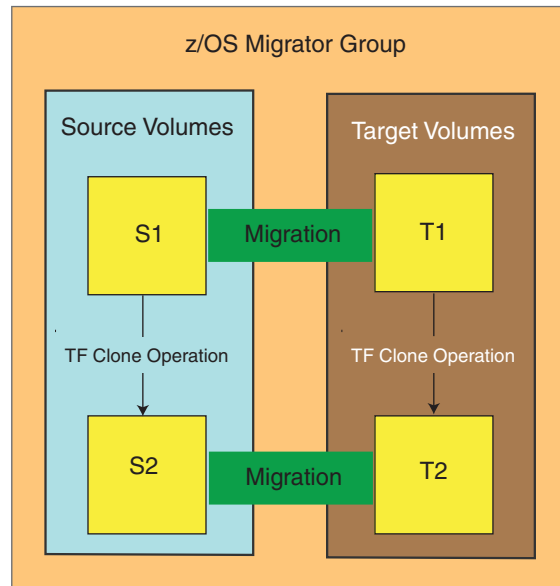


Table 3 TF Clone and z/OS Migrator operations between migrator source and target volumes after migration activation

When TF Clone volume pairs are created AFTER MIGRATION ACTIVATION, and the TF/Clone source and target volumes are located within the z/OS Migration Group...

TF Clone Source	TF Clone Target	Promotion	During Activation	During AUTOSWAP Validation
z/OS Migrator Source: S1	z/OS Migrator Source: S2	Not Applicable (N/A)	Clone not allowed*, Activation allowed	Clone not allowed*, Validation allowed
z/OS Migrator Target: T1	z/OS Migrator Target: T2	N/A	Clone not allowed*, Activation allowed	Clone not allowed*, Validation allowed

* Error message for SNAP not allowed condition:
ESNP404E TARGET VOLUME ("xxxx" SN xxxxxxx-xxxx/xxxxxxxx) IS IN USE BY ANOTHER PROCESS.

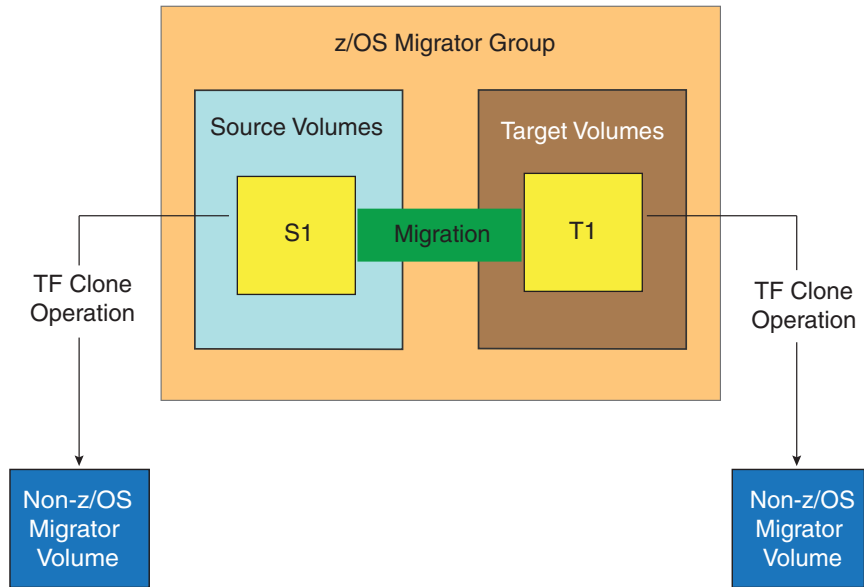


Table 4 TF Clone and z/OS Migrator operations between migrator and non-migrator volumes after migration activation

When TF Clone volume pairs are created AFTER MIGRATION ACTIVATION, and the TF Clone target volumes are non-z/OS Migrator volumes...				
TF Clone Source	TF Clone Target	Promotion	During Activation	During AUTOSWAP Validation
z/OS Migrator Source: S1	Non-z/OS Migrator Volume	Not Applicable (N/A)	Clone and Activation allowed	Validation allowed
z/OS Migrator Target: T1	Non-z/OS Migrator Volume	N/A	Clone and Activation allowed	Validation allowed

TimeFinder and z/OS Migrator operations

Note: *TimeFinder* refers to both TimeFinder Mirror and Clone Emulation operations. The permissible commands between Clone and Mirror actions and z/OS Migrator operations are identical.

Table 5 and Table 6 describe which z/OS Migrator operations are permissible between TimeFinder source and target volumes belonging to a z/OS Migrator Group.




Table 5 TimeFinder and z/OS Migrator operations inside the Migrator Group before promotion

When BCV pairs are established using z/OS Migrator source and target volumes BEFORE PROMOTION...					
TimeFinder Source	TimeFinder Target	Before Promotion	Promotion	During Activation	During AUTOSWAP Validation
z/OS Migrator Source: S1 (STD)	z/OS Migrator Source: S2 (BCV)	TF-ESTABLISH allowed	Promotion not allowed *	Not Applicable (N/A)	N/A
z/OS Migrator Target: T1 (STD)	z/OS Migrator Target: T2 (BCV)	TF-ESTABLISH allowed	Promotion not allowed *	N/A	N/A
* Errors during promotion: EFMM117E FC01 call failed (xxxxxxxx-xxxx-xxxx-xxxxxxxx-xxxx) EFMM116E CUU xxxx CAN NOT BE REACHED					

Table 6 TimeFinder and z/OS Migrator operations inside the Migrator Group during activation

When BCV pairs are established using z/OS Migrator source and target volumes DURING ACTIVATION...					
TimeFinder Source	TimeFinder Target	Before Promotion	Promotion	During Activation	During AUTOSWAP Validation
z/OS Migrator Source: S1 (STD)	z/OS Migrator Source: S2 (BCV)	Not Applicable (N/A)	N/A	TF-ESTABLISH is not allowed *, Activation allowed	Validation allowed
z/OS Migrator Source: T1 (STD)	z/OS Migrator Source: T2 (BCV)	N/A	N/A	TF-ESTABLISH is not allowed *, Activation allowed	Validation allowed
* Error during activation: BCVM120E BCV xxxx IN USE BY ANOTHER OPERATION					

Table 7 describes which z/OS Migrator operations are permissible between TimeFinder source and target volumes belonging to a z/OS Migrator Group and non-z/OS Migrator volumes.

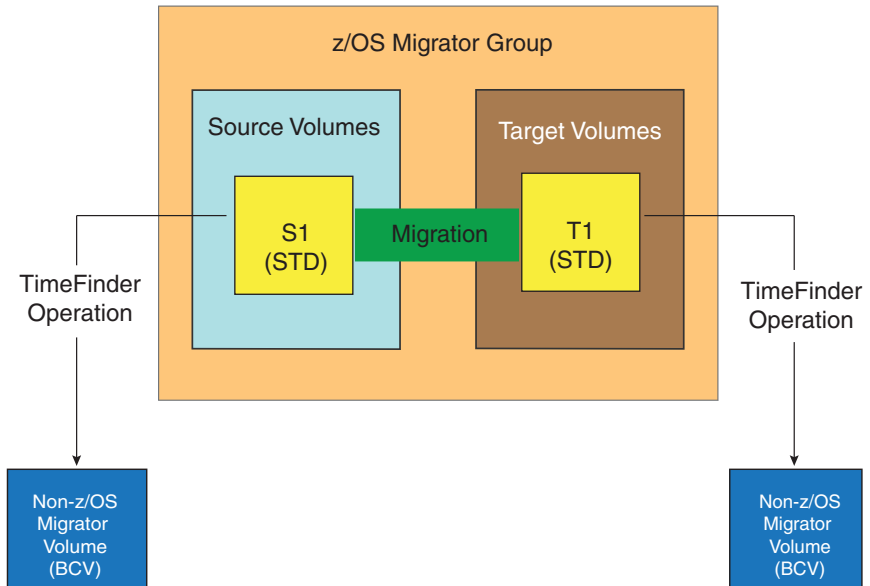


Table 7 TimeFinder and z/O Migrator operations outside of the Migrator Group

When BCV pairs are established, and the BCV target volumes are outside of the z/OS Migrator Group...				
TimeFinder Source	TimeFinder Target	Promotion	During Activation	During AUTOSWAP Validation
z/OS Migrator Source S1 (STD)	Non z/OS Migrator Volume (BCV)	TF-ESTABLISH and Promotion allowed	Activation allowed	Validation allowed
z/OS Migrator Target T1 (STD)	Non z/OS Migrator Volume (BCV)	TF-ESTABLISH and Promotion allowed	Activation allowed	Validation allowed

SRDF and z/OS Migrator operations

Table 8 and Table 9 describe which operations are permissible using z/OS Migrator source and target volumes and SRDF volumes outside of the Migrator Group.

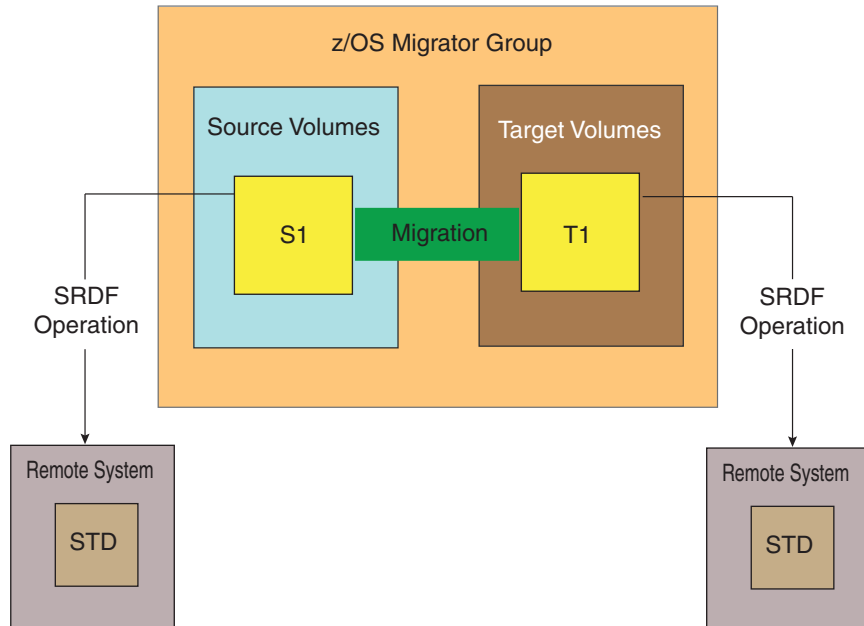


Table 8 SRDF and z/OS Migrator operations before promotion

When RDF pairs are created BEFORE PROMOTION...					
SRDF Source	SRDF Target	Before Promotion	Promotion	During Activation	During AUTOSWAP Validation
z/OS Migrator Source S1	Remote System (STD)	SRDF-HC CREATEPAIR allowed	PROMOTION allowed	ACTIVATION allowed	VALIDATION allowed
z/OS Migrator Target T1	Remote System (STD)	SRDF-HC CREATEPAIR allowed	PROMOTION allowed	ACTIVATION allowed	ACTIVATION allowed

Table 9 SRDF and z/OS Migrator operations after activation

When RDF pairs are created AFTER ACTIVATION...					
SRDF Source	SRDF Target	Before Promotion	During Promotion	During Activation	During AUTOSWAP Validation
z/OS Migrator Source S1	Remote System (STD)	Not Applicable (N/A)	N/A	SRDF-HC CREATEPAIR not allowed*, ACTIVATION allowed	SRDF-HC CREATEPAIR not allowed*, Validation allowed
z/OS Migrator Source T1	Remote System (STD)	N/A	N/A	SRDF-HC CREATEPAIR not allowed*, ACTIVATION allowed	SRDF-HC CREATEPAIR not allowed*, Validation allowed
*Error during activation: EMCCW3EE FMLM LOCK NOT AVAILABLE FOR LOCAL DEVICE					
AFTER ACTIVATION...many SC VOL commands are not allowed due to a z/OS Migrator device lock					

EMCSAFI Security Interface

This appendix describes the EMCSAFI Security Interface and discusses how to enable, disable, and customize that interface. Topics include:

- ◆ Overview 376
- ◆ Class names and resource names used 377
- ◆ ACF2 Requirements 378
- ◆ Disabling the security interface 378
- ◆ Customizing the security interface 380
- ◆ EMCSAFRB structure 382
- ◆ Enhanced SAF device number checking 387

Overview

The EMCSAFI Security Interface feature uses SAF calls (RACROUTE) to request authorization to use resources. The input to EMCSAFI is the EMCSAFRB request block that describes the authorization.

EMCSAFI requires that RACF version 1.9 or later, or an equivalent SAF-compliant security product, be installed and activated. The security interface is active by default. If you choose not to use it, you can find instructions to disable it in [“Disabling the security interface” on page 378](#).

Note: When running with the security interface active, check with your security administrator to ensure that the proper classes are active and the proper resources are defined. The class and resource names used by the z/OS Migrator Security Interface are described in [“Class names and resource names used” on page 377](#).

Class names and resource names used

[Table 1](#) summarizes the resource validation requests for each z/OS Migrator feature and function.

Note: DATASET is the default SAF Class and EMC.VALIDATE.ACCESS is the default SAF Profile. Both may be localized in the global initialization parameters. Sample jobs on how to do this in a RACF environment are in member EMCCLASS in the ResourcePak Base SAMPLIB.

Table 1 Resource validation requests for z/OS Migrator

Program	Function	Class	Resource	Attr	Volser	
EMCTFU ¹	Relabel volume	DASDVOL	<i>old-volser</i>	ALTER		
		DASDVOL	<i>new-volser</i>	ALTER		
		GDASDVOL	<i>old-volser-mask</i>	ALTER		
		GDASDVOL	<i>new-volser-mask</i>	ALTER		
	Rename	Non-VSAM dataset	DATASET	<i>old-dsname</i>	ALTER	<i>volser</i>
			DATASET	<i>new-dsname</i>	ALTER	<i>volser</i>
		VSAM cluster	DATASET	<i>old-dsname</i>	ALTER	<i>old-cat-volser</i>
			DATASET	<i>new-dsname</i>	ALTER	<i>new-cat-volser</i>
		VSAM path	DATASET	<i>old-pathname</i>	ALTER	<i>old-cat-volser</i>
			DATASET	<i>new-pathname</i>	ALTER	<i>new-cat-volser</i>
	SMS Indexed VTOC Management	FACILITY	STGADMIN.IGG.DIRCAT	ALTER		
	EMCSNAP ²	SNAP volume	DASDVOL	<i>old-volser</i>	READ	
DASDVOL			<i>new-volser</i>	ALTER		
GDASDVOL			<i>old-volser-mask</i>	READ		
GDASDVOL			<i>new-volser-mask</i>	ALTER		
SNAP dataset		DATASET	<i>old-dsname</i>	ALTER	<i>volser</i>	
		Normal z/OS security processing is performed for the output dataset				
SMS Indexed VTOC management		FACILITY	STGADMIN.IGG.DIRCAT	ALTER		
ICKDSF Indexed VTOC management		FACILITY	STGADMIN.ICK.BUILDIX	ALTER		

1. A component of EMC TimeFinder Utility
2. A component of EMC TimeFinder/Clone Mainframe SNAP Facility

The RACF resources and classes used for the QUERY operator command and the DISPLAY operator command are the same.

ACF2 Requirements

In order to maintain system integrity, CA-ACF2 requires that a user be authorized to access security definitions. To accommodate this situation, ACF2 will now allow the security administrator to define the specific calls for which the authorization check for STATUS=ACCESS will be bypassed.

The bypassing of the authorization check in ACF2 is done through use of the new NOAPFCHK keyword added to the GSO SAFDEF record. NOAPFCHK will only be honored for STATUS=ACCESS calls. The following is an example of the SAFDEF entry used for z/OS Migrator:

```
INSERT SAFDEF.zosm PROGRAM(EFMM**) RB(EFMM**) NOAPFCHK
RACROUTE (REQUEST=AUTH, CLASS=DATASET, STATUS=ACCESS)
```

Disabling the security interface

The following sections describe how to disable the security interface.

The EMCSAFD routine

Member EMCSAFD in the ResourcePak Base SAMPLIB contains assembler source code that you can employ to disable the EMCSAFI security interface. Use member #90SAFJB in the z/OS Migrator RIMLIB to assemble and link the modules. Make sure to change all occurrences of literal *SAFMBR* in #90SAFJB to EMCSAFD.

This source replaces CSECT EMCSAFI with a routine that returns a return code of 0 or 4, depending on the content in RNAME_TBL. Entries in RNAME_TBL return four (4). No attempt to call RACF is made and the following message is placed in ESRBMSG:

```
EMC SAF INTERFACE IS DISABLED
```

Note: The ESRBMSG message field is part of the EMCSAFRB structure. The ESRBMSG message field contains the messages related to the EMCSAFI routine. [Figure 2, “EMCSAFRB macro,” on page 383](#) shows a sample EMCSAFRB structure.

Disabling the EMCSAF interface

To disable the EMCSAF interface, take the following steps:

1. Change the JCL to conform to your installation standards.
2. Run the job.

Result: Program EFMMMAIN is relinked.

The sample EMCSAFD (Figure 1 on page 379) includes MNOTE statements to identify RNAME_TBL table. After reviewing this code, comment the MNOTE statements.

```

**
* IF THE SUPPLIED RESOURCE NAME MATCHES AN ENTRY IN THE RNAME_TBL,
* RETURN WITH RC 4.
**

MNOTE 12,'*-----* '
MNOTE 12,'* Please review the new functionality added with * '
MNOTE 12,'* the RNAME_TBL. A rc 4 will be returned for * '
MNOTE 12,'* each resource name in the table, allowing for * '
MNOTE 12,'* protection of these functions when a full SAF * '
MNOTE 12,'* interface is not desired. * '
MNOTE 12,'* * '
MNOTE 12,'* To run without this enhancement, replace the * '
MNOTE 12,'* first character of the first entry in the table * '
MNOTE 12,'* with a hex FF. * '
MNOTE 12,'* * '
MNOTE 12,'* Before submitting this module for assembly, * '
MNOTE 12,'* delete or comment these MNOTE statements. * '
MNOTE 12,'*-----* '
LA R2,RNAME_TBL PRIME POINTER TO RNAME TBL

```

Figure 1 Sample EMCSAFD

Customizing the security interface

The following sections describe how to customize the security interface.

The EMCSAFI routine

Member EMCSAFI in the ResourcePak Base SAMPLIB dataset contains assembler source code for the security interface. Use member #90SAFJB in the z/OS Migrator RIMLIB to assemble and link the modules, making sure to change all occurrences of literal *SAFMBR* to EMCSAFI.

The source code included in ResourcePak Base SAMPLIB is the code linked with EFMMMAIN. This source is provided to allow you to customize the behavior of the security code to match your installation requirements.

To customize the interface, take the following steps:

1. Change the JCL to conform to your installation standards.
2. Change the source for EMCSAFI to suit your needs.
3. Run the job.

Result: Program EFMMMAIN is relinked.

Note: Customization of the EMCSAFI routine should only be attempted by experienced system programmers who have extensive knowledge of assembler language and standard linkage conventions, and who understand the RACF RACROUTE interface. You need to take normal precautions to test changes in an isolated environment, and to protect working production code.

On entry to EMCSAFI, R1 points to a full word containing the address of the EMCSAFRB. The information in this request block is used to build a RACROUTE request.

On return, R15 contains one of the following return codes:

- | | |
|-----------|---|
| 0 | For ESRBATTR = T: Security subsystem is active
For ESRBATTR ≠ T: Access is allowed |
| 8 | Access is denied |
| 12 | EMCSAFRB failed validation |
| 16 | RACF is not active |

Note: Take care when customizing EMCSAFI to maintain re-entrancy.

Return code zero

If the return code is zero, then field ESRBMSG contains one of the messages shown in [Table 2](#).

Table 2 EMCSAFI routine: Return code zero

Message	Description
ACCESS ALLOWED	Access to the requested resource is allowed.
ACCESS ALLOWED - (WARN MODE)	Access to the requested resource would have been denied; however, warn mode is in effect, so access is allowed.
ACCESS ALLOWED - CLASS NOT ACTIVE	The requested class is not defined, and PROTECT ALL is not in effect.
ACCESS ALLOWED - RESOURCE NOT PROTECTED	The requested resource is not defined, and PROTECT ALL is not in effect.
SECURITY SUBSYSTEM IS ACTIVE	The request was to determine if the security subsystem is active and it is.

Return code non zero

If the return code is non zero, then field ESRBMSG contains one of the messages shown in [Table 3](#).

Table 3 EMCSAFI routine: Return code non zero (page 1 of 2)

Message	Description
ACCESS DENIED	The security subsystem has denied access to the resource. Contact your security administrator for proper access.
SECURITY SUBSYSTEM IS NOT ACTIVE	The security subsystem is not running. Start the security subsystem or run job #09SAFJB in the RIMLIB using member EMCSAFD in the ResourcePak Base SAMPLIB to disable the security feature.
EMCSAFRB ERROR - CLASS NOT SPECIFIED	The EMCSAFRB control structure passed to the security interface is in error. The Field ESRBCLAS is not filled in. If you have customized the SAF interface, review your changes for errors. If you have not customized the SAF interface, contact the EMC Customer Support Center for technical assistance.
EMCSAFRB ERROR - INVALID AUTHORITY LEVEL REQUESTED	The EMCSAFRB control structure passed to the security interface is in error. Field ESRBATTR has an invalid value. If you have customized the SAF interface, review your changes for errors. If you have not customized the SAF interface, contact the EMC Customer Support Center for technical assistance.

Table 3 EMCSAFI routine: Return code non zero (page 2 of 2)

Message	Description
EMCSAFRB ERROR - RESOURCE NAME NOT SPECIFIED	The EMCSAFRB control structure passed to the security interface is in error. Field ESRBRNAM is not filled in. If you have customized the SAF interface, review your changes for errors. If you have not customized the SAF interface, contact the EMC Customer Support Center for technical assistance.
EMCSAFRB ERROR - INVALID DSTYPE VALUE SPECIFIED	The EMCSAFRB control structure passed to the security interface is in error. Field ESRBDSTY has an invalid value. If you have customized the SAF interface, review your changes for errors. If you have not customized the SAF interface, contact the EMC Customer Support Center for technical assistance.
EMCSAFRB ERROR - DSTYPE IS NOT M AND VOLSER NOT SPECIFIED	The EMCSAFRB control structure passed to the security interface is in error. Field ESRBDSTY has an invalid value. The value is not M, and field ESRBVSER is not filled in. If you have customized the SAF interface, review your changes for errors. If you have not customized the SAF interface, contact the EMC Customer Support Center for technical assistance.

EMCSAFRB structure

EMCSAFRB describes the resource access request and is built by the caller and passed to the EMCSAFI routine. The macro is included in the ResourcePak Base SAMPLIB and is also shown in [Figure 2 on page 383](#).

When the class name is DATASET (or the class name specified in the initialization of z/OS Migrator) the following are true:

- ◆ ESRBCLAS — Set the class name (QNAME)
- ◆ ESRBRNAM — Set to the resource name (RNAME), as documented in [Table 6 on page 385](#)
- ◆ ESRBATTR — U (update) or R (read), as documented in [Table 4 on page 385](#)
- ◆ ESRBUTOK — Security token supplied in the CIB representing the operator command
- ◆ ESRBDSTY — M (model profile)
- ◆ ESRBUID — Set to spaces
- ◆ ESRBGID — Set to spaces

Depending on the action specified, some fields may not be filled in.
You can use these fields to customize the security exit.

```

MACRO
*****
.*
.* EMCSAFRB
.*
.* STRUCTURE PASSED AS INPUT TO THE EMCSAFI SECURITY INTERFACE
.* ROUTINE.
.*
.*-----*
.* CHANGE LOG:
.*
.* 09/09/04 BASE
.* 09/28/04 ADD OPTIONAL PARMS
.*
.*-----*
EMCSAFRB
EMCSAFRB DSECT
*
*****
*
* AUTHORITY LEVEL
*
*****
ESRBATTR DS CL1 REQUEST AUTHORITY LEVEL
* T - TEST RACF ACTIVE
* R - READ
* U - UPDATE
* A - ALTER
* C - CONTROL
* DS CL3 .. OPEN ..
*****
*
* REQUESTOR INFORMATION
*
*****
* SET ESRBUID AND ESRBGID TO BLANK UNLESS 3RD PARTY AUTHORIZATION
* CHECKING IS TO BE USED
*
ESRBUID DS CL8 USERID OR BLANK
ESRBGID DS CL8 GROUPID OR BLANK
ESRBACEE DS A | 0 OR ADDR OF ACEE OR 4X'FF'
* | IF ESRBUTOK POINTS TO UTOKEN
*

```

Figure 2 EMCSAFRB macro

```

*****
*
*                               RESOURCE INFORMATION                               *
*****
ESRBCLAS DS      CL8              RESOURCE CLASS
ESBRNAM  DS      CL44            RESOURCE NAME
ESRBVSER DS      CL6              VOLSER (IF CLASS=DATASET)
ESRBDSTY DS      CL1              DSTYPE: (IF CLASS=DATASET)
*
*                               N - NONVSAM
*                               V - VSAM
*                               M - MODEL PROFILE
*                               T - TAPE
*
*                               DS      CL5              | .. OPEN ..
*****
*
*                               RETURN CODES                                   *
*****
ESRBR15  DS      F                RETURN CODE FROM RACROUTE
ESBRRET  DS      F                RACF RETURN CODE
ESBRREA  DS      F                RACF REASON CODE
ESRSRET  DS      F                SAF RETURN CODE
ESRSREA  DS      F                SAF REASON CODE
*****
*
*                               ERROR MESSAGE                                 *
*****
ESRBMSG  DS      CL100            AREA FOR MESSAGE RETURN
*****
* REUSE ERROR MESSAGE AREA FOR ADDITIONAL PARAMETERS.  NOTE THAT          *
* THESE PARAMETERS WILL BE OVERLAYED ON RETURN FROM EMCSAFI.              *
*****
*                               ORG  ESRBMSG                                |
ESRBP1   DS      XL4              | OPTIONAL PARAMTERS PASSED
ESRBP2   DS      XL4              | . BY APPLICATION
ESRBP3   DS      XL4              | . . .
ESRBP4   DS      XL4              | . . .
ESRBUTOK DS      XL4              | A (USER TOKEN)
*                               ORG  ESRBMSG+L'ESRBMSG                       |
EMCSAFRL EQU    *-EMCSAFRB        | AREA LENGTH
*                               MEND

```

Figure 2 (continued)EMCSAFRB macro

Tables 4 through 7 provide descriptions of each field in the EMCSAFRB structure.

Table 4 Authority level

ESRBATTR	<p>1 byte field describing the level of access required. Valid values are:</p> <ul style="list-style-type: none"> • T — Test if the security subsystem is active <hr/> <p>Note: A return code of zero indicates that it is active, a non-zero return code indicates that it is not active.</p> <hr/> <ul style="list-style-type: none"> • R — requests READ access to the resource • U — requests UPDATE access to the resource • A — requests ALTER access to the resource • C — requests CONTROL access to the resource
----------	---

Table 5 Requestor information

ESRBUID	8 byte field containing the userid of the user requesting access. If ESRBACEE is specified, this field is ignored.
ESRBGID	8 byte field containing the groupid of the user requesting access. If this field is blank and the ESRBUID field is specified, the default group for the userid is used. If ESRBUID is blank or ignored, this field is ignored.
ESRBACEE	8 byte field containing the address of a valid ACEE for the user requesting access. If this field is zero, and ESRBUID is blank, the ACEE associated with the current address space is used.

Table 6 Resource information

ESRBCLAS	8 byte field containing the class name
ESRBNAM	44 byte field containing the resource name
ESRBVSER	Volume serial -- Used only when ESRBCLAS = DATASET and ESRBDSTY is not equal to 'M'
ESRBSTY	<p>Dataset type when ESRBCLAS = DATASET -- Valid values are:</p> <ul style="list-style-type: none"> • N — NonVSAM • V — VSAM • M — Model Profile • T — Tape dataset

Table 7 **Return codes**

ESRBR15	On return, contains the return code from the RACROUTE macro
ESRBRRET	On return, contains the RACF return code
ESRBRREA	On return, contains the RACF reason code
ESRBSRET	On return, contains the SAF return code
ESRBSREA	On return, contains the SAF reason code

Note: The RACROUTE, RACF, and SAF return codes and reason codes are documented in the IBM publication, *External Security Interface (RACROUTE) Macro Reference for MVS and VM*, or in later releases, *SecureWay Security Server RACROUTE Macro Reference*.

Enhanced SAF device number checking

The EMCSAFI Security Interface feature provides additional security checks for environments where multiple groups of users are using different devices in a single controller. EMC mainframe applications that support specifying the internal or Symmetrix device now check to ensure that devices are logically available for use only by an authorized user.

Specifying the SAF resource name

To implement this check at the Symmetrix device number level, the SAF checking provided by ResourcePak Base supports a new resource name. It must belong to the XFACILIT general resource class. The format of the SAF request is as follows:

```
EMC.DEVC.12digitserialnumber.ssid.dev#
```

Where:

12-digitserialnumber

Specifies the 12-digit Symmetrix controller ID.

ssid

Specifies the subsystem ID.

dev#

Specifies the internal Symmetrix device number. This must be a 4-digit value.

Example

```
EMC.DEVC.000000006185.0C02.0230
```

This identifies a Symmetrix with serial number 000000006185. The SSID is 0C02, and the specific internal device number is 0230.

You can use an asterisk (*) wildcard character to specify more than a specific internal device number. For example, the following statement protects an entire SSID:

```
EMC.DEVC.000000006185.0C02.*
```

The following wildcard protects the entire controller:

```
EMC.DEVC.000000006185.*
```

This glossary contains terms related to disk storage subsystems. Many of these terms are used in this manual.

A

- actuator** A set of access arms and their attached read/write heads, which move as an independent component within a head and disk assembly (HDA).
- adapter** Card that provides the physical interface between the director and disk devices (SCSI adapter), director and parallel channels (Bus & Tag adapter), director and serial channels (Serial adapter).
- alternate track** A track designated to contain data in place of a defective primary track. [See also "primary track."](#)
- APAR** Authorized Program Analysis Report. An official report of a software error to IBM. Also used to refer to the patch supplied by IBM to fix the error (PTF). [See also "SPE."](#)
- Asynchronous** Pertaining to two or more processes that do not depend upon the occurrence of specific events such as common timing signals.
- Asynchronous data transfer** A physical transfer of data to or from a device that occurs without a regular or predictable time relationship following execution of an I/O request.

C

cache	Random access electronic storage used to retain frequently used data for faster access by the channel.
cache fast write (CFW)	A form of fast write to cache where the data is written directly to cache without using nonvolatile storage and is available for later de-staging. This function is used for data of a temporary nature, or data that is readily recreated, such as the sort work files created by sort programs.
cache slot	Unit of cache equivalent to one track.
catalog, or basic catalog structure	<p>A dataset that contains information about other datasets, for example, primarily their location (volume) and SMS attributes.</p> <p>User Catalogs contain entries for any data object that can be accessed by z/OS. z/OS master catalogs contain entries for system datasets and catalog connector records that associate a dataset's alias, also referred to as high-level qualifier (hlq), with the user catalog that contains the dataset's entries. See also "ICF catalog," "VTOC," and "VVDS."</p>
channel command word (CCW)	A mechanism in which a channel command for I/O can be issued. One or more CCWs make up the channel program that directs data channel operations.
channel director	The component in the Symmetrix subsystem that interfaces between the host channels and data storage. It transfers data between the channel and cache.
CICS	Customer Information Control System. General purpose TP monitor for terminal-oriented and inter-system transaction processing in z/OS and VSE/ESA environments. Sits between user application programs, teleprocessing access method (for example, VTAM), and database managers; that is, CICS invokes user-written application programs in response to transactions entered at TP terminals. Originally developed for z/OS and VSE/ESA, but has evolved into a client/server product, in a number of ways.
concurrent copy	An extended function that produces a backup copy and allows concurrent access to data during the copy.

controller ID Controller identification number of the director the disks are channeled to for EREP usage. There is only one controller ID for Symmetrix.

count-key-data (CKD) A DASD data recording format employing self-defining record formats in which each record is represented by a count area that identifies the record and specifies its format, an option key area that may be used to identify the data area contents, and a data area that contains the user data for the record. CKD is also used to refer to a set of channel commands that are accepted by a device that employs the CKD recording format.

count-key-data (CKD) device A disk storage device that stores data in a format consisting of a count field, usually followed by a key field, followed by the actual data of a record.

D

DASD Direct access storage device, a device that provides nonvolatile storage of computer data and random access to that data. IBM wording equivalent to “disk”.

DASD fast write (DFW) A form of fast write to cache where the data is written concurrently to cache and nonvolatile storage and automatically scheduled for de-staging to the DASD. Both copies are retained in the storage control unit until the data is completely written to the DASD, providing data integrity equivalent to writing directly to the DASD.

data availability Access to any and all user data by the application.

dataset A unit of data storage and retrieval consisting of one or more data records. Outside of the IBM mainframe environment, they are called files.

data sharing The ability of concurrent subsystems (such as DB2 or IMS DB) or application programs to directly access and change the same data while maintaining data integrity.

DB2 Database/2. Relational database management system first announced for z/OS environments in 1983. Originally promoted as an end-user tool, it is now IBM's preferred DBMS for almost everything apart from the legacy applications serviced by IMS1. DB2 became a family in the early 1990s and there are now DB2 versions,

most under the name DB2 Universal Database, for a broad range of platforms.

delayed fast write

There is no room in cache for the data presented by the write operation.

dependent write

Most applications, and in particular database management systems (DBMS) have a dependent write logic embedded in them that ensures data integrity if a failure occurs in:

- ◆ The host processor
- ◆ The software
- ◆ The storage subsystem

A dependent write is a write that will not be issued by an application until a prior, related write I/O operation has completed.

An example of dependent write is a database update. When updating a database, a database management system (DBMS) takes the following steps:

1. The DBMS writes to the disk containing the transaction log.
2. The DBMS writes the data to the actual database dataset.
3. The DBMS writes again to the log volume to indicate that the database update was made.

These three write I/O operations (log, database, and log again) are related and each write I/O is not issued until the prior, related write I/O operation has successfully completed.

destage

The asynchronous write of new or updated data from cache to disk device.

device

A uniquely addressable part of the Symmetrix subsystem that consists of a set of access arms, the associated disk surfaces, and the electronic circuitry required to locate, read, and write data. [See also "volume."](#)

device address

The hexadecimal value that uniquely defines a physical I/O device on a channel path in an MVS environment. [See also "unit address."](#)

device number

The value that logically identifies a disk device in a string.

device support facilities program (ICKDSF)

A program used to initialize a Symmetrix storage subsystem at installation and provide media maintenance.

DFSMS

The Data Facility Storage Management Subsystem (SMS) automates storage management by providing constructs that describe logical and physical data management criteria. The logical criteria are described and implemented by Data Class, Storage Class and Management Class. The physical criteria is the location of the data, supported by Storage Group. These constructs are used to implement installation-defined storage management policies.

One of the mechanisms Data Facility Storage Management Subsystem Multiple Virtual Storage (DFSMS/MVS(tm)) provides to help automate storage management is the SMS constructs. They are defined using ISMF and are entities that are associated with datasets and objects based on an installation-defined policy. The intent behind these constructs is to separate the logical view of data from the physical view.

The logical view of data is concerned with what the data looks like and what services it requires. This view is supported by data class, storage class, and management class.

The physical view is concerned with where the data actually resides. It is supported by storage group. There is also a construct called the base configuration, which is not associated with individual datasets or objects. It contains both logical and physical information and applies to the system as a whole.

diagnostics

System level tests or firmware designed to inspect, detect, and correct failing components. These tests are comprehensive and self-invoking.

director

The component in the Symmetrix subsystem that allows Symmetrix to transfer data between the host channels and disk devices. [See also "channel director."](#)

disk director

The component in the Symmetrix subsystem that interfaces between cache and the disk devices.

dual copy

A high availability function made possible by the nonvolatile storage (NVS) in cached 3990 models. Dual copy maintains two functionally identical copies of designated DASD volumes in the logical 3990 subsystem, and automatically updates both copies every time a write operation is issued to the dual copy logical volume.

dual-initiator	A Symmetrix feature that automatically creates a backup data path to the disk devices serviced directly by a disk director, if that disk director or the disk management hardware for those devices fails.
dynamic path reconnect (DPR)	A function that allows disconnected I/O operations with Symmetrix storage subsystem to reconnect over any available channel path rather than be limited to the one on which the I/O operation was started. This function is available only on System 370/XA, System 370/ESA, and System 390/ESA systems.
dynamic sparing	A Symmetrix feature that automatically transfers data from a failing disk device to an available spare disk device without affecting data availability. This feature supports all non-mirrored devices in the Symmetrix subsystem.
E	
EREP program	The program that formats and prepares reports from the data contained in the Error Recording Dataset (ERDS).
ESCON	Enterprise Systems Connection, a set of IBM and vendor products that connect mainframe computers with each other and with attached storage, locally attached workstations, and other devices using optical fiber technology and dynamically modifiable switches called ESCON Directors. See also "ESCON director."
ESCON director	Device that provides a dynamic switching function and extended link path lengths (with XDF capability) when attaching an ESCON channel to a Symmetrix serial channel interface.
extended count-key-data	A set of channel commands that use the CKD track format. Extended count-key-data uses the Define Extent and Locate Record commands to describe the nature and scope of a data-transfer operation to the storage control to optimize the data-transfer operation.
extended count-key-data (CKD) device	A disk storage device that has a data transfer rate faster than some processors can utilize and that is connected to the processor through use of a speed matching buffer. A specialized channel program is needed to communicate with such a device.
extended remote copy (XRC)	A hardware- and software-based remote copy option that provides an asynchronous volume copy across storage subsystems for disaster recovery, device migration, and workload migration.

extended specify task abnormal exit (ESTAE) A macro-instruction that specifies a routine to receive control in the event of abnormal termination of the issuing task.

F

fast migrator A familiar term for the Volume Migration function in EMC's z/OS Migrator.

fast mirror A familiar term for the Volume Mirror function in EMC's z/OS Migrator.

fast write In Symmetrix, a write operation at cache speed that does not require immediate transfer of data to disk. The data is written directly to cache and is available for later destaging.

FBA Fixed Block Architecture, disk device data storage format using fixed-size data blocks.

field replaceable unit (FRU) Field Replaceable Unit, a component that is replaced or added by service personnel as a single entity.

frame Data packet format in an ESCON environment. [See also "ESCON."](#)

G

gatekeeper A small logical volume on a Symmetrix storage subsystem used to pass commands from a host to the Symmetrix storage subsystem. Gatekeeper devices are configured on standard Symmetrix disks.

GB Gigabyte, 10^9 bytes.

global resource serialization (GRS) A function in which resources can be shared across multiple operating systems and still maintain data integrity.

H

head and disk assembly A field replaceable unit in the Symmetrix subsystem containing the disk and actuator.

home address The first field on a CKD track that identifies the track and defines its operational status. The home address is written after the index point on each track.

hyper-volume extension The ability to define more than one logical volume on a single physical disk device making use of its full formatted capacity. These logical volumes are user-selectable in size. The minimum volume size is one cylinder and the maximum size depends on the disk device capacity and the emulation mode selected.

I

ICF catalog Integrated Catalog Facility. An ICF consists of one BCS and one or more VSAM Volume Datasets (VVDS). The BCS can be considered the catalog while the VVDS contains dataset attributes specific to the volume, complementing the Volume Table of Contents (see "[VTOC](#).")

ICKDSF See device support facilities program.

ID Identifier, a sequence of bits or characters that identifies a program, device, controller, or system.

IML Initial microcode program loading.

IMS DB Information Management System Database Manager. IBM's large system hierarchical DBMS.

index marker Indicates the physical beginning and end of a track.

index point The reference point on a disk surface that determines the start of a track.

INLINES An EMC-provided host-based Cache Reporter utility for viewing short and long term cache statistics at the system console.

I/O device An addressable input/output unit, such as a disk device.

ISPF Interactive System Productivity Facility. Menu and screen management system that is used on z/OS TSO. Provides facilities for developing and running menu-driven dialog systems on 3270 terminals and, more recently, workstations. ISPF is an element of z/OS.

J

JCL Job Control Language. The language used to describe the steps of a batch job (such as files to be used, programs to be run).

K

K Kilobyte, 1024 bytes.

L

least recently used algorithm (LRU) The algorithm used to identify and make available the cache space by removing the least recently used data.

logical migrate member The user's definition of any dataset oriented migration group in EMC's z/OS Migrator.

logical migrator The familiar term for the Dataset Migration function in EMC's z/OS Migrator.

logical volume A user-defined storage device. In the Model 5200, the user can define a physical disk device as one or two logical volumes.

long miss Requested data is not in cache and is not in the process of being fetched.

longitude redundancy code (LRC) Exclusive OR (XOR) of the accumulated bytes in the data record.

LPAR Logical partitioning. A method of dividing up a processor capacity under PR/SM. Available on eserver zSeries 900 and iSeries 400. Ability of products and programs running in an LPAR to query a unique identifier and the capacity for the LPAR, giving the LPAR the properties of a virtual server. This gives ISVs the option of licensing their software to an LPAR, and thus an additional method of implementing sub-Central Electronic Complex pricing.

M

MB Megabyte, 10⁶ bytes.

media The disk surface on which data is stored.

metadata Data about data. Metadata describes how and when and by whom a particular set of data was collected, and how the data is formatted.

migrate member The user's group definition for any Volume Migrate group in EMC z/OS Migrator.

mirror member The user's group definition for any Volume Mirror group in EMC z/OS Migrator.

mirroring The Symmetrix system maintains two identical copies of a designated volume on separate disks. Each volume automatically updates during a write operation. If one disk device fails, the Symmetrix system automatically uses the other disk device.

mirrored pair A logical volume with all data recorded twice, once on each of two different physical devices.

P

parallel processing The simultaneous processing of units of work by many servers. The units of work can be either transactions or subdivisions of large units or work (batch).

parallel sysplex A sysplex that uses one or more coupling facilities.

partitioned dataset (PDS) assist An IBM feature for 3990 Model 6 and 3990 Model 3 with Extended Platform units. PDS Assist improves performance on large, heavily-used partitioned datasets by modifying the directory search process.

physical ID Physical identification number of the Symmetrix director for EREP usage. This value automatically increments by one for each director installed in a Symmetrix system. This number must be unique in the mainframe system. It should be an even number. This number is referred to as the SCU_ID.

primary device In Volume Mirror, the "source" device in a mirroring pair.

primary track The original track on which data is stored. [See also "alternate track."](#)

promotion The process of moving data from a track on the disk device to cache slot.

R

read hit Data requested by the read operation is in cache.

read miss Data requested by the read operation is not in cache.

record zero The first record after the home address.

S

scrubbing	The process of reading, checking the error correction bits, and writing corrected data back to the source.
SCSI adapter	Card in the Symmetrix subsystem that provides the physical interface between the disk director and the disk devices.
SCU_ID	For 3880 storage control emulations, this value uniquely identifies the storage director without respect to its selection address. It identifies to the host system, using the EREP, the director detecting the failing subsystem component. This value automatically increments by one for each director installed. The SCU_ID must be a unique number in the host system. It should be an even number and start on a zero boundary.
sequential data striping	An extended function where the system writes consecutive tracks from datasets on different volumes and reads or writes them in parallel.
shared storage complex	Catalog environment of the systems involved in z/OS Migrator migration.
short miss	Requested data is not in cache, but is in the process of being fetched.
SMP	System Modification Program. See also "SMP/E."
SMP/E	System Modification Program - Extended. An element of z/OS that is used to install most software products. Applies the relevant parts of Change Management to the installation and maintenance of software delivered in executable (OCO) format, with no source code. Most notably, the ability to test a new version of software (or maintenance to an old one) on a production system while the old version is simultaneously being used by production applications.
source dataset name	A name which is given to the dataset in the source location. The note on page 121 offers a discussion of source and target names.
source device	When applied to Volume Migrator, the source device is the source in a migration pair.

source volume	One device of a migration pair. All channel commands to copy the volume are directed to the source volume. The data on the source volume is duplicated on the target volume.
SPE	Small Programming Enhancement. Typically an APAR, which adds new functionality rather than just corrects a bug.
SSID	For 3990 storage control emulations, this value identifies the physical components of a logical DASD subsystem. The SSID must be a unique number in the host system. It should be an even number and start on a zero boundary.
stage	The process of writing data from a disk device to cache.
storage control unit	The component in the Symmetrix subsystem that connects Symmetrix to the host channels. It performs channel commands and communicates with the disk directors and cache. See also "channel director."
string	A series of connected disk devices sharing the same disk director.
synchronization	The action of forcing certain points in the execution sequences of two or more asynchronous procedures to coincide in time.
synchronous data transfer	A physical transfer of data to or from a device that has a predictable time relationship with the execution of an I/O request.
system authorization facility (SAF)	Provides an interface between a product, subsystem, or component requesting access to a resource within the system and a security product (for example, RACF, ACF2, and Top Secret).
T	
target dataset name	A name which is given to the dataset in the target location. The note on page 121 offers a discussion of source and target names.
target device	When applied to Volume Migrator, the target device is the target in a migration pair.
target volume	One of the devices of a migration pair that will contain a duplicate of the data on the source volume.

U

unit address The hexadecimal value that uniquely defines a physical I/O device on a channel path in an MVS environment. [See also "device address."](#)

V

volume A general term referring to a storage device. In the Symmetrix subsystem, a volume corresponds to single disk device.

volume migrator A synonym for fast migrator.

volume mirror A synonym for fast mirror.

VSAM Virtual Storage Access Method. IBM mainframe proprietary software for direct (by key or by record number) or sequential processing of fixed and variable length records on DASD. VSAM gives device independence of data storage and eases migration of data to new devices.

VTOC The Volume Table Of Contents and the VTOC index (VTOCIX) are datasets that maintain space allocation information for the volume. The VTOC is used to find empty space for new dataset allocations and to locate existing non-VSAM datasets on the volume. [See also "ICF catalog."](#)

VVDS VSAM Volume Dataset. Resides on every volume that contains a [VSAM](#) or SMS-managed dataset cataloged in an [ICF catalog](#). It contains the dataset extent information and attributes that describe VSAM datasets, and SMS-managed non-VSAM datasets on the volume.

W

write hit There is room in cache for the data presented by the write operation.

write miss There is no room in cache for the data presented by the write operation.

