

Using the DMM CLI for Data Migration

This chapter describes how to use DMM CLI commands to configure and monitor data migration jobs. This chapter includes the following sections:

- About DMM CLI Commands, page 7-1
- Selecting Ports for Server-Based Jobs, page 7-2
- Configuring Data Migration Using the CLI, page 7-2
- Controlling DMM Jobs, page 7-14
- Monitoring DMM Jobs, page 7-16
- Completing DMM Jobs, page 7-17

About DMM CLI Commands

The DMM feature includes CLI commands to configure and perform data migration jobs. Job and session configuration commands are entered at the switch CLI prompt.

A DMM job can be active on more than one switch. For example, in a dual-fabric topology with multipath configurations, the DMM job runs on a switch in each fabric. To configure the job, you enter DMM CLI commands on both switches.

The DMM feature runs on an SSM in the switch. Each session runs on only one SSM. Enter the session configuration commands on the MDS switch that will perform the session migration.

The DMM **show** commands are accessed directly from the SSM. From the command prompt in the switch, you must attach to the SSM module before entering these commands.

When using the DMM CLI commands, note the following guidelines:

- In DMM job configuration mode, the job configuration is not saved until you enter the **commit** command. If you exit DMM configuration mode without issuing the **commit** command, all job configuration changes are discarded. You only need to enter the **commit** command when configuring a new job.
- For a storage-based migration, all servers that use the selected storage enclosure must use the same operating system (for example, all AIX or all Solaris).
- If the MDS switch (hosting the storage or the server) performs a restart after the migration but before the job is destroyed, you must restart the data migration from the beginning.

Selecting Ports for Server-Based Jobs

When creating a server-based migration job, you must include all possible paths from the server HBA ports to the LUNs being migrated because all writes to a migrated LUN need to be mirrored to the new storage until the cutover occurs, so that no data writes are lost.

For additional information about selecting ports for server-based jobs, see the "Ports in a Server-Based Job" section on page 3-11.

Configuring Data Migration Using the CLI

When you enter the command to create a data migration job, the CLI enters DMM job configuration submode. This submode provides commands to configure the server HBA ports, storage ports, and job attributes. The job is only created on the SSM when you enter the **commit** command.

In a dual-fabric topology with redundant paths, the data migration job runs on an SSM in each fabric. You need to configure the job on both SSMs.

In this chapter, the examples and command descriptions use the following terminology (see Figure 7-1):

- The dual fabric configuration includes Fabric 1 and Fabric 2.
- Switch A (on Fabric 1) contains the SSM for data migration jobs.
- Switch B (on Fabric 2) contains the SSM for data migration jobs.
- H1 and H2 are the server HBA ports to each fabric.
- ES1 and ES2 are the existing storage ports.
- NS1 and NS2 are the new storage ports.





The steps to configure a data migration job are described in the following sections:

- Configuring the Virtual Initiator (Storage-Based Migration), page 7-3
- Creating the Data Migration Job, page 7-4

- Configuring the Job, page 7-4
- Committing the Job, page 7-5
- Configuring the Peer SSM, page 7-6
- Configuring Sessions, page 7-6
- Server-Based Migration Example, page 7-8
- Storage-Based Migration Example, page 7-10

Configuring the Virtual Initiator (Storage-Based Migration)



This step is not required for a server-based data migration job.

Prior to creating a storage-based data migration job, you must retrieve the virtual initiator (VI) port world wide name (pWWN) and create a new zone containing the pWWNs of the VI and the storage ports. To use the new zone, add the new zone to a zone set and activate the zone set.

To configure the VI in Fabric 1, follow these steps:

Command	Purpose	
switchA# config t	Enters configuration mode.	
switchA(config)# dmm module module-id job job-id get-vi vsan 0-4093	Retrieves the VI information for the specified SSM module.	
	You must specify a unique job identifier.	
	The command output displays the assigned VI node WWN and port WWN.	
<pre>switchA(config)# zone name name vsan 0-4093</pre>	Creates a new zone.	
switchA(config-zone)# member pwwn <i>value</i>	Uses the member command multiple times to add the VI pWWN, the existing storage pWWNs, and the new storage pWWNs.	
<pre>switchA(config-zone)# exit</pre>	Exits zone configuration submode.	
switchA(config)# zoneset name name vsan 0-4093	Enters configuration mode for the active zone set. Specify the name of the active zone set.	
<pre>switchA(config-zoneset)# member name</pre>	Adds the named zone to the zone set.	
<pre>switchA(config-zoneset)# exit</pre>	Exits zone set configuration submode.	
switchA(config)# zoneset activate name name vsan 0-4093	Reactivates the zone set.	

Prior to creating the data migration job, you must complete the following configuration tasks on the storage devices:

- 1. Configure the existing storage to give the VI pWWN access to LUNs that need to be migrated.
- 2. Configure the new storage to give the VI pWWN access to LUNs that need to be migrated.



For a dual-fabric topology, you must repeat the same set of configuration steps on switch B. Retrieve the VI information and create a new zone on switch B and configure the storage to allow the VI to access the LUNs exposed in fabric B. For an example configuration, see the "Storage-Based Migration Example" section on page 7-10.

Creating the Data Migration Job

To configure a data migration job, first create the job on Switch A. After creating the job, the CLI enters DMM job configuration mode, where you enter the commands for configuring the job.

To create the data migration job, follow these steps:

	Command	Purpose
Step 1	switchA# config t	Enters configuration mode.
Step 2	<pre>switchA(config)# dmm module module-id job job-id create</pre>	Creates a migration job on the specified SSM module and enters DMM job configuration mode.
		Specify a unique job identifier. For a storage-based job, use the same job identifier that you specified when retrieving the VI information (in the previous task).

Configuring the Job

Use the commands in DMM job configuration mode to add the server and storage ports to the job.



To prevent data corruption, the job must contain all the server HBA ports that can access the set of LUNs being migrated, and all storage ports that expose these LUNs:

- Add all server HBA ports in this fabric that can access the LUNs being migrated.
- Add all storage ports in the fabric that expose the set of LUNs being migrated.

For additional information, see the "Checking the Storage ASL Status" section on page 4-3.

In a dual-fabric topology, configure the IP address of the peer SSM (the DMM peers communicate using the management IP network).

	Command	Purpose	
	<pre>switchA(config-dmm-job)# server vsan 0-4093 pwwn pWWN</pre>	Specifies the VSAN and pWWN of the server HBA porto include in the migration.	
		Note	All server HBA ports (in fabric 1) that can access the LUNs to be migrated need to be added to this job.
	<pre>switchA(config-dmm-job)# storage vsan 0-4093 pwwn pWWN existing</pre>	Specif port.	ies the VSAN and pWWN of the existing storage
		Note	All existing storage ports (in fabric 1) that expose the LUNs to be migrated need to be added to this job.
	switchA(config-dmm-job)# storage vsan 0-4093 pwwn pWWN new	Specif port.	ies the VSAN and pWWN of the new storage
		Note	All new storage ports (in fabric 1) that expose the new LUNs need to be added to this job.
<pre>switchA(config-dmm-job)# attributes</pre>		Specifies the job type, job mode, and job rate:	
job_type {1 2} job_mode job_rate {1 2 3 4} : {1 2 3}	job_type {1 2} job_mode {1 2} job_rate {1 2 3 4} job_method {1 2 3}	For jo for sto	b_type, enter 1 for server-based migration or 2 rage-based migration.
		For jo l migrat	b_mode, enter 1 for online or 2 for offline ion.
		For jo l mediu	b_rate, enter 1 for best effort, 2 for slow, 3 for m, and 4 for fast data migration.
		For jo l and 3 f	b_method, enter 1 for Method 1, 2 for Method 2, for Method 3.
		For ad see the page 2	ditional information about data migration rate, e "Configuring Migration Rate" section on -6.
	<pre>switchA(config-dmm-job)# peer</pre>	Config	gures the IP address of the SSM on switch B.
	<i>IP_address</i>	For inf see the	formation about configuring SSM IP addresses, e "Configuring IP Connectivity" section on

To configure the data migration job, use the following steps:

Committing the Job

The next step is to commit the data migration job on switch A. To commit the job, use the **commit** command.

When you enter the **commit** command, the switch sends the job configuration to the SSM.

The DMM feature sends configuration information to other switches in the fabric as required, so that all traffic between the server HBA port and the existing storage is redirected to the SSM.

The SSM performs discovery of all existing and new storage LUNs visible to the server HBA ports/VIs in this job.



The **commit** command may require a noticeable amount of time to complete, depending on the number of LUNs to be discovered.

Configuring the Peer SSM

The next step is to configure the data migration job on the peer SSM.



You must use the same job number that you created on switch A.

To configure the data migration job on the peer SSM, follow these steps:

	Command	Purpose	
Step 1	switchB# configuration terminal	Enters configuration mode.	
Step 2	<pre>switchB(config)# dmm module module-id job job-id create</pre>	Creates a migration job on the specified SSM module and enters DMM job configuration mode.	
		Note Enter the same job ID that you created on switch A.	
Step 3	switchB(config-dmm-job)# server vsan 0-4093 pwwn pWWN	Specifies the VSAN and pWWN of the server HBA port to include in the migration.	
		Note All server HBA ports (in Fabric 2) that can access the LUNs to be migrated need to be added to this job.	
Step 4	<pre>switchB(config-dmm-job)# storage vsan 0-4093 pwwn pWWN existing</pre>	Specifies the VSAN and pWWN of the existing storage port.	
		Note All existing storage ports (in Fabric 2) that expose the LUNs to be migrated need to be added to this job.	
Step 5	switchB(config-dmm-job)# storage vsan 0-4093 pwwn pWWN new	Specifies the VSAN and pWWN of the new storage port.	
		Note All new storage ports (in Fabric 2) that expose the new LUNs need to be added to this job.	
Step 6	<pre>switchB(config-dmm-job)# attributes job_type {1 2} job_mode {1 2} job_rate {1 2 3 4} job_method</pre>	Specifies the job type, job mode, job rate, and job method.	
	{1 2 3}	Note The configuration values for the attributes and the schedule must match on both switches.	
Step 7	<pre>switchB(config-dmm-job)# peer IP_address</pre>	Configures the IP address of the SSM on switch A.	
Step 8	switchB(config-dmm-job)# commit	Commits the data migration job on switch B.	

Configuring Sessions

The next step is to configure sessions in the data migration job. For a server-based migration, configure all of the sessions on one SSM.

For a storage-based migration, you can manually balance the load on the SSMs by configuring sessions on both SSMs.



For a storage-based migration, use the pWWN of the VI as the server in the session configuration.

To verify that the SSM has discovered the LUNs correctly, enter the **show dmm job job-id** *job id* **storage** command from the SSM CLI.

To configure sessions, follow these steps:

	Command	Purpose	
Step 1	switchA(config)# dmm module module-id job job-id session	Enters session configuration mode for the specified job on the specified SSM.	
Step 2	<pre>switchA(config-session)# server pWWN src_tgt pWWN src_lun num dst_tgt pWWN dst_lun num</pre>	server <i>pWWN</i> Configures a session. The server HBA port, existing storage port, and new storage port must all belong to th same VSAN.	
		• server is the server pWWN (server-based job) or VI pWWN (storage-based job).	
		• src_tgt is the existing storage pWWN.	
		• src_lun is the LUN number in the existing storage. Enter this value in hexadecimal notation.	
		• dst_tgt num is the new storage pWWN.	
		• dst_lun is the LUN number in the new storage. Enter this value in hexadecimal notation.	

Server-Based Migration Example

The topology for the example shown in Figure 7-2, is dual fabric with multipath ports defined in the server and redundant paths to the storage devices.

Figure 7-2 Topology for the Example



On both switches, the SSM module is located in slot 8. The pWWNs for the ports are listed here:

Port	pWWN
H1	21:00:00:e0:8b:0a:5d:e7
ES1	50:06:04:82:bf:cf:e0:43
NS1	50:06:0e:80:03:4e:95:13
H2	21:01:00:e0:8b:0a:5d:e7
ES2	50:06:04:82:bf:cf:e0:5d
NS2	50:06:0e:80:03:4e:95:03

The following example shows how to configure a data migration job on switch A:

```
switchA# configure terminal
switchA(config)# dmm module 8 job 2345 create
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
switchA(config-dmm-job)# server vsan 100 pwwn 21:0d:00:0d:ec:02:2d:82
switchA(config-dmm-job)# storage vsan 100 pwwn 50:06:04:82:bf:cf:e0:43 existing
switchA(config-dmm-job)# storage vsan 100 pwwn 50:06:0e:80:03:4e:95:13 new
switchA(config-dmm-job)# peer 10.10.2.4
switchA(config-dmm-job)# attributes job_type 1 job_mode 1 job-rate 1 job-method 1
switchA(config-dmm-job)# commit
switchA(config-dmm-job)# end
Ending DMM Job Configuration.
```

If the Job was not committed, it will be required to reconfigure the job. The following example shows how to configure a data migration job on switch B:

```
switchB# configure terminal
switchB(config)# dmm module 8 job 2345 create
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
switchB(config-dmm-job)# server vsan 100 pwwn 21:0d:00:0d:0a:01:2b:82
switchB(config-dmm-job)# storage vsan 100 pwwn 50:06:04:82:bf:cf:e0:5d existing
switchB(config-dmm-job)# storage vsan 100 pwwn 50:06:0e:80:03:4e:95:03 new
switchB(config-dmm-job)# peer 10.10.1.8
switchB(config-dmm-job)# attributes job_type 1 job_mode 1 job-rate 1 job-method 1
switchB(config-dmm-job)# end
Ending DMM Job Configuration.
If the Job was not committed, it will be required to reconfigure the job.
switchB#
```

The following example shows how to configure data migration sessions on switch A:

```
switchA(config)# dmm module 4 job 2345 session
switchA(config-session)# server 21:00:00:e0:8b:0a:5d:e7 src_tgt 50:06:04:82:bf:cf:e0:43
src_lun 0x5 dst_tgt 50:06:0e:80:03:4e:95:13 dst_lun 0x0
switchA(config-session)# server 21:00:00:e0:8b:0a:5d:e7 src_tgt 50:06:04:82:bf:cf:e0:43
src_lun 0x6 dst_tgt 50:06:0e:80:03:4e:95:13 dst_lun 0x1
switchA(config-session)# exit
```

The following example shows how to start a data migration job on switch A:

```
switchA(config)# dmm module 8 job 2345 start
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
switchA(config)# exit
```

Storage-Based Migration Example

The topology for the example shown in Figure 7-3, is dual fabric with multipath ports defined in the server and redundant paths to the storage devices.

Figure 7-3 Storage-Based Migration Example



On both switches, the SSM module is located in slot 8. The pWWNs for the ports are listed here:

Port	pWWN	
Host 1	21:00:00:e0:8b:0a:5d:e7	
VI 1	21:0d:00:0d:ec:02:2d:82	
ES1	50:06:04:82:bf:cf:e0:43	
NS1	50:06:0e:80:03:4e:95:13	
Host 2	21:01:00:e0:8b:0a:5d:e7	
VI 2	21:0d:00:0d:0a:01:2b:82	
ES2	50:06:04:82:bf:cf:e0:5d	
NS2	50:06:0e:80:03:4e:95:03	

The following example shows how to configure the VI on switch A:

```
switchA# configure terminal
switchA(config)# dmm module 8 job 2345 get-vi vsan 100
DMM Storage Job:0x929 assigned following VI -
VI NodeWWN: 21:0c:00:0d:ec:02:2d:82
VI PortWWN: 21:0d:00:0d:ec:02:2d:82
sjc7-9509-6(config)#
```

The following example shows how to configure the zone and zone set on switch A:

```
switchA(config)# zone name DMM1 vsan 100
switchA(config-zone)# member pwwn 21:0d:00:0d:ec:02:2d:82 vi
switchA(config-zone)# member pwwn 50:06:04:82:bf:cf:e0:43 es
switchA(config-zone)# member pwwn 50:06:0e:80:03:4e:95:13 ns
switchA(config-zone)# exit
switchA(config)# zoneset name DMM1 vsan 100
switchA(config-zoneset)# member DMM1
switchA(config-zoneset)# exit
switchA(config-zoneset)# exit
switchA(config)#
```

The following example shows how to configure the data migration job on switch A:

```
switchA(config)# dmm module 8 job 2345 create
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
switchA(config-dmm-job)# server vsan 100 pwwn 21:00:00:e0:8b:0a:5d:e7
switchA(config-dmm-job)# storage vsan 100 pwwn 50:06:04:82:bf:cf:e0:43 existing
switchA(config-dmm-job)# storage vsan 100 pwwn 50:06:0e:80:03:4e:95:13 new
switchA(config-dmm-job)# peer 10.10.2.4
switchA(config-dmm-job)# attributes job_type 2 job_mode 1 job-rate 1 job-method 1
switchA(config-dmm-job)# end
Ending DMM Job Configuration.
If the Job was not committed, it will be required to reconfigure the job.
switchB#
```

The following example shows how to configure the VI on switch B:

```
switchB# configure terminal
switchB(config)# dmm module 8 job 2345 get-vi vsan 100
DMM Storage Job:0x929 assigned following VI -
VI NodeWWN: 21:0c:01:0e:ec:02:2d:82
VI PortWWN: 21:0d:00:0d:0a:01:2b:82
switchB(config)#
```

The following example shows how to configure the zone and zone set on switch B:

```
switchB(config)# zone name DMM1 vsan 100
switchB(config-zone)# member pwwn 21:0d:00:0d:0a:01 :2b:82 vi
switchB(config-zone)# member pwwn 50:06:04:82:bf:cf:e0:5d es
switchB(config-zone)# member pwwn 50:06:0e:80:03:4e:95:03 ns
switchB(config-zone)# exit
switchB(config)# zoneset name DMM1 vsan 100
switchB(config-zoneset)# member DMM1
switchB(config-zoneset)# exit
switchB(config)#
```

The following example shows how to configure the data migration job on switch B:

```
switchB# configure terminal
switchB(config)# dmm module 8 job 2345 create
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
switchB(config-dmm-job)# server vsan 100 pwvm 21:01:00:e0:8b:0a:5d:e7
switchB(config-dmm-job)# storage vsan 100 pwvm 50:06:04:82:bf:cf:e0:5d existing
switchB(config-dmm-job)# storage vsan 100 pwvm 50:06:0e:80:03:4e:95:03 new
switchB(config-dmm-job)# peer 10.10.1.8
switchB(config-dmm-job)# attributes job_type 2 job_mode 1 job-rate 1 job-method 1
switchB(config-dmm-job)# end
Ending DMM Job Configuration.
If the Job was not committed, it will be required to reconfigure the job.
switchB#
```

The following example shows how to configure the data migration sessions on switch A:

```
switchA(config)# dmm module 4 job 2345 session
switchA(config-session)# server 21:0d:00:0d:ec:02:2d:82 src_tgt 50:06:04:82:bf:cf:e0:43
src_lun 0x5 dst_tgt 50:06:0e:80:03:4e:95:13 dst_lun 0x0
```

switchA(config-session)# exit

The following example shows how to start the data migration job on switch A:

switchA(config)# dmm module 8 job 2345 start
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
switchA(config)# exit

The following example shows how to configure the data migration sessions on switch B:

```
switchB(config)# dmm module 4 job 2345 session
switchB(config-session)# server 21:0d:00:0d:0a:01:2b:82 src_tgt 50:06:04:82:bf:cf:e0:5d
src_lun 0x5 dst_tgt 50:06:0e:80:03:4e:95:03 dst_lun 0x0
```

switchB(config-session)# exit

The following example shows how to start the data migration job on switch B:

switchB(config)# dmm module 8 job 2345 start
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
switchB(config)# exit

DMM Method 3 Migration Example

Figure 7-4

Server SAN A SAN B Cisco MDS 9216 Cisco MDS Cisco MDS 5509 5509 Existing storage Cisco MDS Local data 5509 Center Replication/Migration _ _ _ _ _ _ _ _ _ SAN Remote data Cisco MDS Center 5509 192356 New storage

The topology for the example is shown in Figure 7-4.

DMM Method 3 Migration Example

This section describes how to configure a data migration job using DMM method 3. The job needs to be configured on the SSM/MSM in the migration fabric as well as the SSM(s)/MSM(s) in the production fabrics.

To configure the migration fabric, follow these steps:

```
Step 1 Select a VI from the VI list for the module.
```

migr-fab# show dmm module 2 vi-list				
DPP-Id	VI-pWWN	VI-nWWN	Outstanding jobs	
1	22:93:00:0d:ec:4a:63:83	22:92:00:0d:ec:4a:63:83	1	
2	22:8b:00:0d:ec:4a:63:83	22:8a:00:0d:ec:4a:63:83	0	
3	22:8d:00:0d:ec:4a:63:83	22:8c:00:0d:ec:4a:63:83	0	
4	22:95:00:0d:ec:4a:63:83	22:94:00:0d:ec:4a:63:83	0	
5	22:97:00:0d:ec:4a:63:83	22:96:00:0d:ec:4a:63:83	0	
6	22:8f:00:0d:ec:4a:63:83	22:8e:00:0d:ec:4a:63:83	0	
7	22:91:00:0d:ec:4a:63:83	22:90:00:0d:ec:4a:63:83	0	
8	22:99:00:0d:ec:4a:63:83	22:98:00:0d:ec:4a:63:83	0	

Step 2 Create a zone of the selected VI with the ES and NS ports.

Step 3 To create the DMM job.

```
migr-fab# config terminal
migr-fab(config)# dmm module 2 job 4 set-vi 22:93:00:0d:ec:4a:63:83
22:92:00:0d:ec:4a:63:83 vsan 3 (Use the VI selected above and Zoned with the existing
storage and new storage ports)
migr-fab(config)# dmm module 2 job 4 create
migr-fab(config-dmm-job)# attributes job_type 1 job_mode 1 job_rate 1 job_method 3
migr-fab(config-dmm-job)# peer 10.1.2.3 (SSM/MSM in production fabric 1)
migr-fab(config-dmm-job)# peer 10.1.1.4 (SSM/MSM in production fabric 2)
migr-fab(config-dmm-job) # storage vsan 3 pwwn 44:51:00:06:2b:02:00:00 existing
migr-fab(config-dmm-job)# storage vsan 3 pwwn 44:f1:00:06:2b:04:00:00 new
migr-fab(config-dmm-job)# commit
migr-fab(config-dmm-job)# exit
migr-fab(config)# dmm module 2 job 4 session
migr-fab(config-dmm-session)# server 22:93:00:0d:ec:4a:63:83 src_tgt
44:51:00:06:2b:02:00:00 src_lun 0 dst_tgt 44:f1:00:06:2b:04:00:00 dst_lun 0
migr-fab(config-dmm-session)# end
```

The following example shows the configuration of production fabric 1:

```
prod-fab1(config) # dmm module 3 job 4 create
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
prod-fab1(config-dmm-job) # attributes job_type 1 job_mode 1 job_rate 1 job_method 3
prod-fab1(config-dmm-job) # peer 10.1.3.2 (only the migration fabric is configured as peer)
prod-fab1(config-dmm-job) # server vsan 100 pwwn 21:01:00:e0:8b:28:5e:3e
prod-fab1(config-dmm-job) # storage vsan 100 pwwn 44:f0:00:06:2b:03:00:00 existing
(only the existing storage is required for production fabric, no new storage is required)
prod-fab1(config-dmm-job) # commit
prod-fab1(config-dmm-job) # end
```

The following example shows the configuration of production fabric 2:

```
prod-fab2(config) # dmm module 4 job 4 create
Started New DMM Job Configuration.
Do not exit sub-mode until configuration is complete and committed
prod-fab2(config-dmm-job) # attributes job_type 1 job_mode 1 job_rate 1 job_method 3
prod-fab2(config-dmm-job) # peer 10.1.3.2
prod-fab2(config-dmm-job) # server vsan 9 pwwn 21:00:00:e0:8b:08:5e:3e
prod-fab2(config-dmm-job) # storage vsan 9 pwwn 44:50:00:06:2b:01:00:00 existing
prod-fab2(config-dmm-job) # job_name name_job
prod-fab2(config-dmm-job) # commit
prod-fab2(config-dmm-job) # commit
```

Controlling DMM Jobs

The DMM CLI provides a set of commands to control jobs that have been configured. The job state determines which commands are valid to run. Table 7-1 shows job state values.

Job Status Value	Description
Created	The job has been created but has not been scheduled.
Scheduled	The job has been configured with a scheduled start time. It will automatically start at that time.
Complete	The job has been completed successfully.

Table 7-1 Job Status Values

Job Status Value	Description	
Verify	The completed job is being verified.	
Stopped	The job has been stopped manually by the user.	
Failed	The job has been stopped because of failures. See Table 6-5 for details.	
In_Progress	The job is currently running.	
Reset	The job has been reinitialized because of failures. See Table 6-6 for details.	
Finishing	The Method 2 job is in the final copy iteration.	
Verify_Stopped	The job verification has been stopped.	
Verify_Complete	The job verification has been completed.	
Verify_Failure	The job verification is unsuccessful.	

 Table 7-1
 Job Status Values (continued)

Table 7-2 describes the data migration commands.

|--|

Command	Description	Valid Job Status Values
Verify	Performs offline verification of the selected job, and you are prompted to confirm the verification command.	Completed, InProgress, VerifyStopped, Verify_Failure
Destroy	Deletes the selected job (or jobs) and retrieves the job execution log from the SSM, and you are prompted to confirm the delete command.	Stopped, Failed, Completed, Reset, VerifyStopped, Verify_Failure, Created, Scheduled
Stop	Stops the selected job.	InProgress
Start	Starts the selected job.	Created, Reset
Modify	Allows you to modify the job attributes or configure a start time for the selected job.	Created, Scheduled, Reset, Stopped
Schedule	Allows you to set up schedules.	Created, Scheduled, Stopped
Validate	Validates the stored configuration for a job in a Reset state.	Reset
Finish	Completes the selected job only in case of Method 2.	InProgress
Log	Opens the DMM log for the selected job.	All job status values



You must enter these commands on the switch with sessions configured. If both SSMs have sessions configured, enter the commands on both switches.

To control the data migration job, follow these steps:

	Command	Purpose
Step 1	switchA(config) # dmm module module-id job job-id start	Starts a data migration job or restarts a job that was stopped.
		Note For a job in the reset state, enter the validate command on both switches before restarting the job.
		The start command is ignored if the job is scheduled for a future time. Use the schedule now command to start a scheduled job.
Step 2	switchA(config) # dmm module module-id job job-id stop	Stops execution of the job.
		Note Job progress is not preserved. If you start the job again, the job restarts from the beginning.
Step 3	<pre>switchA(config)# dmm module module-id job job-id validate</pre>	If the job is in the reset state, enter the validate command. After validation, start the job using the start command.
		Note Always run the validate command on both SSMs (even if only one SSM has sessions).
Step 4	<pre>switchA(config)# dmm module module-id job job-id schedule [now hour br min min day day month month</pre>	Configures a scheduled start time for the data migration job.
	[year] reset]	Enter the schedule now to start the job immediately.
		Enter the reset to remove the scheduled start time from the job. The job remains in the created state until you manually start it.
		Note Enter the schedule command on each SSM with sessions.

Monitoring DMM Jobs

Use the **show dmm job** command in the SSM CLI to monitor the status of data migration jobs, and the current progress of jobs and sessions that are in progress.

To monitor data migration jobs, follow these steps:

	Command	Purpose
Step 1	switch# attach module module-id	Enters CLI mode on the SSM module.

	Command	Purpose
Step 2	module# show dmm job	Displays summary information about the data migration jobs configured on this SSM.
Step 3	3 module# show dmm job job-id job-id {detail session storage}	Displays information about the specified job. The detail command displays the job attributes, schedule, server HBA and storage ports, the job log, and job error log.
		The session command displays the sessions included in the job.
		The storage command displays the storage ports included in the job.

For additional information about monitoring and troubleshooting data migration jobs, see Chapter 6, "Troubleshooting Cisco MDS DMM."

Completing DMM Jobs

When all of the sessions in a job have completed successfully, you can delete the job in coordination with other post-migration tasks, which are described in the following sections:

- (Optional) Verifying the Completed Job, page 7-17
- Post-Migration Activities, page 7-18
- Deleting the Job, page 7-19

(Optional) Verifying the Completed Job

When all of the sessions in a job have completed successfully, you can optionally perform verification of the data in the new storage location. The SSM compares the data in the new storage with the data in the existing storage by reading each migration region from the existing and new storage, and then performing a comparison of the data.

To perform migration verification, follow these steps:

	Command	Purpose
Step 1	<pre>switchA(config)# dmm module module-id job job-id verify</pre>	Verifies the data migration by comparing the data in the new storage with the data in the existing storage. The verify command operates in offline mode.
Step 2	<pre>switch# show dmm job job-id session [session-id sess-id]</pre>	Displays the verification progress while verification is performed on a job.



Verification is performed in offline mode. Any service using the existing storage needs to be quiesced before you start verification.



Verification is recommended only for the test environment and is NOT recommended for the production environment because this process brings down all the existing applications.

Post-Migration Activities

After the data migration job has completed successfully, you need to reconfigure the server to use the new storage. The exact post-migration configuration steps vary depending on the operating system of the server.

To reconfigure the server, you might need to take the following steps:

- Perform a graceful shutdown on all server applications that use the migrated LUNs to ensure that there are no pending I/O operations on the existing storage.
- Unmount any file systems, existing LUNs, and the associated storage ports.
- Deport the volume groups or disk groups for the migrated storage for some of the volume managers.
- Use the DMM CLI to perform these tasks:
 - For Method 1— Delete the data migration job. DMM removes the FC-Redirect entries to the SMM. Server writes are no longer mirrored to the existing and new storage.
 - For Method 2— Finish the data migration job. When the job moves to Completed state, delete the data migration job. See the "Finishing Jobs" section on page 5-48 for more details.
 - For Method 3— Finish the data migration job. When the job moves to Completed state, delete the data migration job. See the "Finishing Jobs" section on page 5-48 for more details.
- Use either of the following options to remove server access:
 - Configure zoning to remove server access to the existing LUNs.
 - Use an appropriate array tool to remove the masking or mapping access. Choose this option if an application that is being migrated requires access to the existing storage after the first migration is completed.
- Before you configure a DMM job, ensure that the zoning is completed for any devices that require NS.
- Configure zoning to add server access to the new LUNs.
- From the server, scan for the new storage.
- Import the volume group or disk groups from the new storage.
- Mount the file system for the new storage.
- From the server, restart the server applications to access data from the new storage.

Finishing the Job

To finish the data migration job, follow this step:

	Command	Purpose
Step 1	<pre>switchA(config)# dmm module module-id job job-id finish</pre>	Finishes the specified data migration job. This is valid only for Method 2. See "Finishing Jobs" section on page 5-48 for more details.

Deleting the Job

To delete the data migration job, follow this step:

	Command	Purpose
Step 1	<pre>switchA(config)# dmm module module-id job job-id destroy switchA#</pre>	Deletes the specified data migration job.