

## Cisco MDS 9000 Family Highlights: Storage Virtualization Series

Highlighted Feature: Cisco Data Mobility Manager

### Purpose

The Cisco® MDS 9000 Family Highlights series provides both business and technical value propositions for the various Cisco MDS 9000 family features to address the business and technical challenges of the IT customer. This document begins by stating the relevant business and technical challenges that the Cisco Data Mobility Manager (DMM) feature addresses and then describes the solution benefits that help meet these challenges. To show how Cisco DMM can be used, a sample Cisco DMM configuration is provided.

### Audience

Readers of the Cisco MDS 9000 Family Highlights series should have some familiarity with the existing business and technical challenges in IT departments and perhaps some high-level understanding of storage and SAN technologies. For the business solution reader, the main challenges and solutions sections provide concise explanations of the problems faced by IT departments and how they can be solved with the Cisco MDS 9000 family products. For the more technical reader, sample configurations provide insight into the solution configuration.

### Background

In the day-to-day operation of a typical data center, one of the most complex tasks is migrating data from existing disk arrays to new disk arrays. Data migration may be needed because of storage consolidation, expiration of the existing disk array equipment lease, existing disk arrays that have reached capacity, or a shift in corporate standards toward a multivendor disk array strategy. Regardless of the reason for the storage data migration, this task usually involves coordination among different functional teams, complex logistical planning, and likely an long and unproductive outage for end users. Traditionally, the types of data migration can be divided into three major categories: host-based data migration, disk array-based data migration, and appliance-based data migration.

In host-based data migration, installed with a piece of volume management software, an application server is concurrently connected to both the existing and new disk arrays. Essentially, a mirror volume on the new disk array is created and synchronized with an active volume in the existing disk array. Since the host server operates on the volumes of disk arrays, the host becomes effectively indifferent to the make and uniqueness of the actual disk arrays. So long as the volume in the destination disk array is at least as large as the existing volume, the volume mirroring can be performed. However, this type of migration comes with the expense of performance degradation of the running application and directly affects the end user's experience during the migration period.

With disk array-based data migration, the data is moved directly between the two disk arrays and outside the data paths of the running application; thus, there is little noticeable performance degradation at the application host level. However, this method requires that additional ports on

the disk arrays be dedicated and configured specifically for the migration purpose. Perhaps the biggest limitation of disk array–based data migration is that the disk arrays must be from the same vendor and in a similar product family, thus making a multivendor disk array environment infeasible.

The third common data migration method relies on the replication appliance to perform the data migration. In this method, an appliance is inserted into the data path of the host application server and the disk arrays. The benefits of this approach are that it offloads the host application server and allows data migration across a heterogeneous disk array environment. However, the process of inserting the appliance itself requires additional outage time when the appliance is moved between the application servers. Even more challenging is the lack of scalability of the appliances for migrations across a large number of hosts and disk arrays. IT has to choose between moving the migration appliance to each set of application servers and disk arrays or incurring a large capital expenditure (CapEx) to procure an appliance for each set of application servers and disk arrays.

## **Executive Summary**

Data migration of disk arrays is a time consuming and complex activity faced by all data center operations teams. The Cisco MDS 9000 DMM tool enables highly scalable fabricwide disk array migration, simplifying the logistical coordination and reducing the potential outage time for end users during data migration,. By using a Cisco MDS 9000 DMM node as the main point of migration coordination, the application host server can continue to operate during the migration with no performance degradation. Since the solution is SAN fabric based, all host servers and all disk arrays continue to be operate without any configuration changes. Unlike with a storage-based migration solution, multivendor storage disk arrays are fully supported.

## **Challenges**

### **Business Challenges**

- Data migration requires a substantial amount of logistical planning and coordination between various IT operational groups and thus ties up valuable IT resources for the data center team.
- Traditional data migration solutions in general require an extensive amount of application end-user outage time.
- Storage-based migration solutions are proprietary, more expensive, and require a homogeneous disk array environment, therefore making a multivendor sourcing strategy infeasible.

### **Technical Challenges**

- A host-based migration solution takes additional CPU processing away from application servers, resulting in performance degradation for all the application end users.
- An appliance-based migration solution limits scalability because hosts and storage must be directly connected to the appliance, and it requires additional outage time during the setup and reconfiguration of the migration tool.
- An appliance-based migration solution limits overall scalability because increasing numbers of appliances will need to be inserted and removed between the migrating hosts and storage disk arrays, a process that also requires more application outage time. Solution

Network-based data migration using the Cisco DMM running on a Cisco MDS 9000 DMM node provides all the benefits of traditional data migration approaches but mitigates their respective risks. Since the Cisco DMM runs on top of the Cisco MDS 9000 DMM node, the performance of the application host server is not affected. End users can continue to use the application during the data migration up to the actual point of cutoff between the disk arrays. Because it is a SAN-based solution, it does not require any particular proprietary vendor or model for the connected storage disk array. Unlike with appliance data migration solutions, host servers and disk arrays are not disrupted by the insertion and removal of appliances. Instead, all host server and disk array I/O traffic flows to the Cisco MDS 9000 node, which performs the actual data migration. In fact, even the insertion and configuration of the Cisco MDS 9000 DMM node is completely transparent to the host servers and disk arrays.

From a business perspective, the data migration solution investment is protected regardless of the choice of disk arrays and host servers and the potential changes in their connections. As long as any Cisco MDS 9000 family switch within the SAN fabric contains a Cisco MDS 9000 DMM node, any existing or new servers and disk arrays can take advantage of the network-based Cisco DMM feature. Also, with online server host migrations, end users can continue to access their applications even during the disk-to-disk copying process, thus greatly reducing the amount of application outage time.

In a traditional data migration, coordination and intermediary procedures need to be performed in a sequential manner by the application administrator, the database administrator, and the system and storage administrator, and end users are prevented from accessing their applications for the duration of the data migration. In contrast, with Cisco DMM online host server-based migration, end users can continue to work, and administrators can continue with migration preparations until the point that the cutoff occurs on the existing disks.

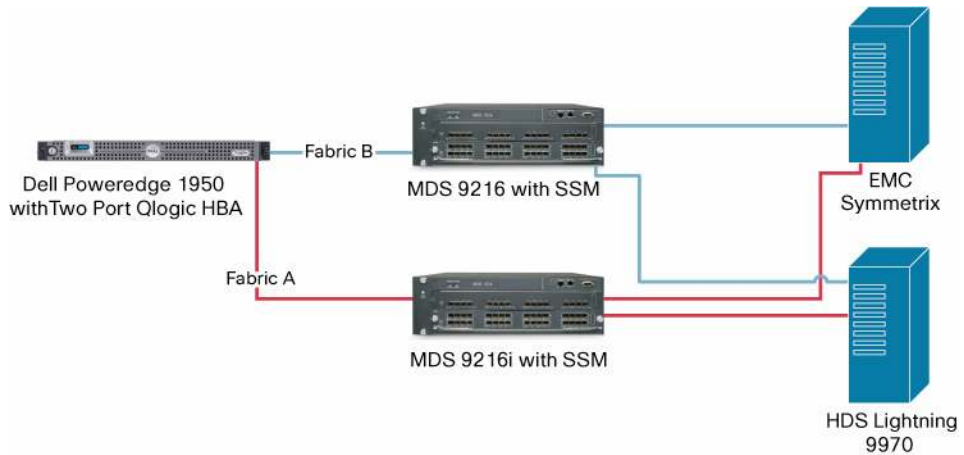
In addition, central management of the data migration from the Cisco Fabric Manager eliminates the complexity of interfacing directly with the individual host server, the appliance, or the disk array management server. The risk of an unplanned outage when changing the application data path during the insertion and removal of the data migration appliances is also mitigated. To accomplish this, the Cisco DMM uses the Fibre Channel Redirect (FC Redirect) capability of the Cisco MDS 9000 family to set up on-demand traffic redirection from the host server and storage disk arrays to the Cisco MDS 9000 DMM node. DMM node could either be a Storage Services Module (SSM) or a MSM-18/4 module or MDS 9222i switch.

### **Sample Configuration**

Configuring Cisco DMM on Cisco MDS 9000 DMM node to Migrate Disk Volumes Between EMC and HDS Disk Arrays. The DMM node used in this configuration is the SSM module.

In Figure 1, a host server (DMM\_SVR) with a dual port Fibre Channel host bus adapter (HBA) is connected to two Cisco MDS 9216i Multilayer Fabric Switches that are physically separated the Fibre Channel fabric. The Fibre Channel ports on the existing and new disk arrays are also connected to each of the Cisco MDS 9216i switches. Note that communication between the Cisco MDS 9000 SSMs of the two switches is established as IP traffic through IP over Fibre Channel (IPFC) on VSAN 1.

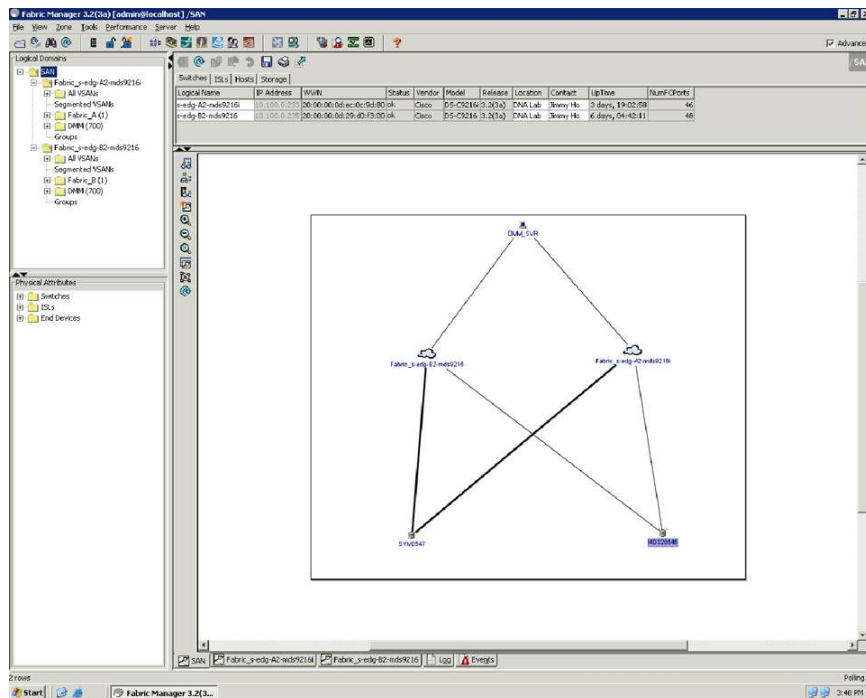
**Figure 1.** Cisco DMM Data Migration Setup



### Viewing the Cisco DMM Data Migration Topology

In Figure 2, Cisco Fabric Manager is shown managing two physical fabrics: Fabric\_s-edg-A2-mds9216i and Fabric\_s-edg-B2-mds9216. Notice that the enclosure name of the two ports on the HBA of the server is DMM\_SVR. Fibre Channel connectivity from the EMC and HDS arrays is established to both fabric switches: s-edg-A2-mds9216i and s-edg-B2-mds9216. Regardless of which path is used for the data migration, there will always be an independent path for the host to access the existing disk array.

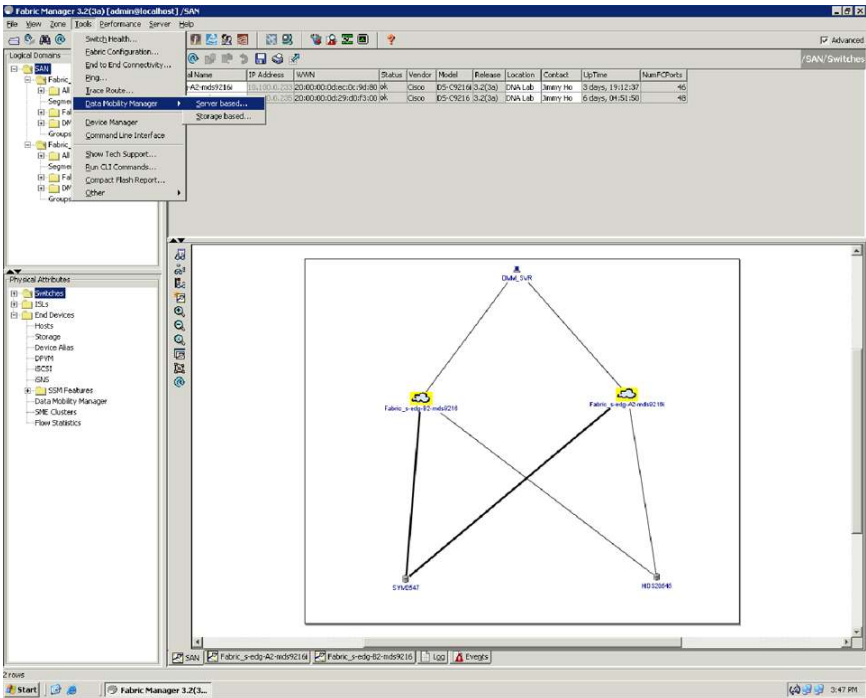
**Figure 2.** Cisco Fabric Manager View of Cisco DMM Setup



### Configuring Cisco DMM Server-Based Online Data Migration in Cisco Fabric Manager

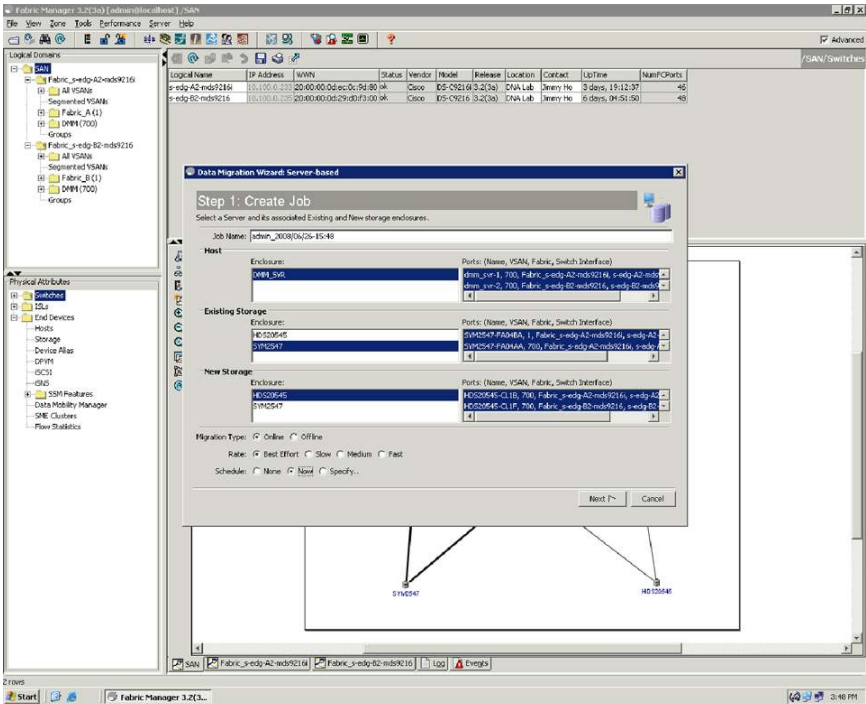
1. Select a migration. In Figure 3, a server-based Cisco DMM migration is selected.

**Figure 3.** Initiating Server-Based Cisco DMM Migration from Cisco Fabric Manager




2. The Cisco DMM GUI displays the selections of the selectable host server and its respective enclosure, the ports of the existing storage, and the ports of the new storage. Specify the migration. In Figure 4, an immediate online migration with the best effort rate is selected.

**Figure 4.** Selection for Host Enclosure, Existing Storage, and New Storage



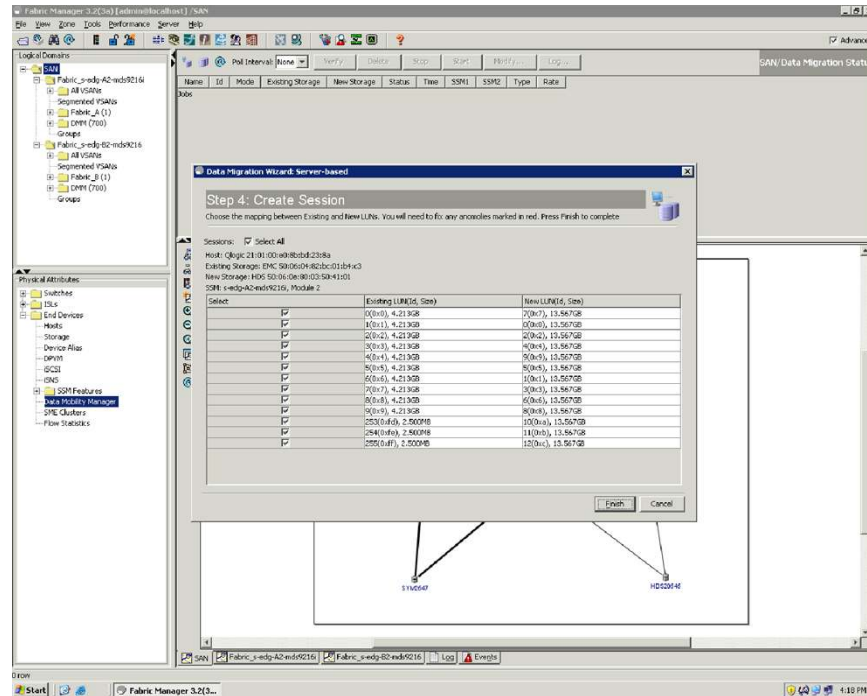
- Figure 5.** Selection of the Cisco MDS 9000 SSM for the Migration Job



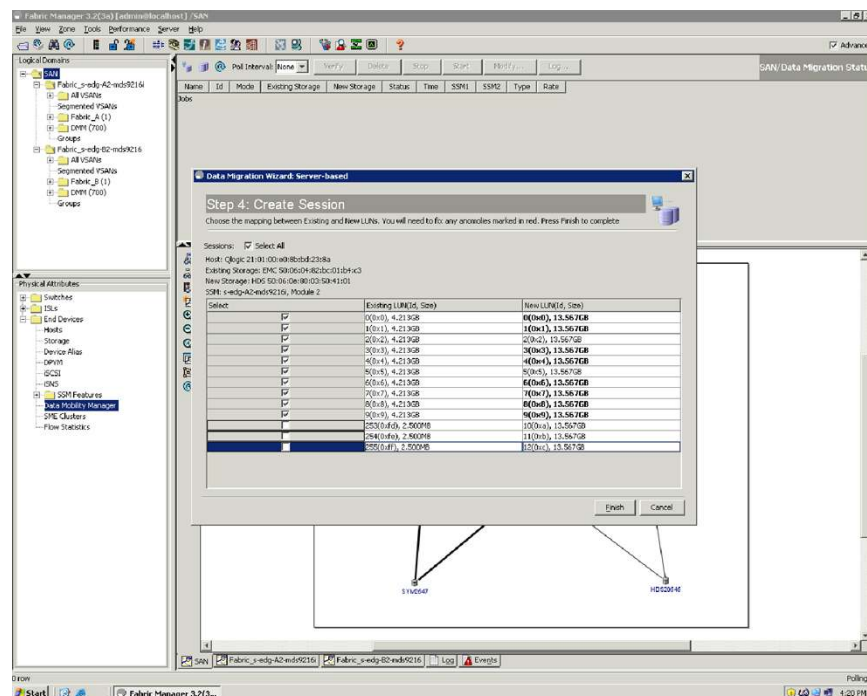
The screenshot shows the Solaris Manager 3.2.1 GUI. The main window displays the 'Data Migration Wizard: Server-based' dialog box. The wizard is at Step 3: Choose Existing and New Path. The dialog box has a title bar 'Data Migration Wizard: Server-based' and a close button. The main content area is titled 'Step 3: Choose Existing and New Path' and contains the text 'Default path is pre-selected. Click on a new row to change selection for Existing and New Storage Path.' Below this text are two tables: 'Existing Storage Path:' and 'New Storage Path:'. The 'Existing Storage Path:' table has columns '# of LUNs', 'Host', 'Storage', 'Switch', 'SMP', and 'Path'. It contains two rows: '12 Cgbe 21:01:00:00:88:94:238A, EMC 50:05:04:52:bc:01:04:02, 10.100.0.233, Module 2' and '13 Cgbe 21:00:00:00:88:94:238A, EMC 50:05:04:52:bc:01:04:02, 10.100.0.236, Module 2'. The 'New Storage Path:' table has the same columns and contains two rows: '12 Cgbe 21:01:00:00:88:94:238A, HDS 50:06:0e:80:63:50:41:05, 10.100.0.233, Module 2' and '13 Cgbe 21:00:00:00:88:94:238A, HDS 50:06:0e:80:63:50:41:05, 10.100.0.236, Module 2'. The 'Next' button is highlighted. The background shows the Solaris Manager interface with a tree view on the left and a table on the right.

- The discovered logical unit numbers (LUNs) for the existing and new disk array are listed. Notice in Figure 7 that the numbering of the LUNs is not necessarily sequential. In Figure 8, the first 11 LUNs of the new disk array have been manually numbered to match the existing disk array. After the Finish button is clicked, the job is immediately initiated.

**Figure 7.** Default Mapping of Existing LUNs to New LUNs



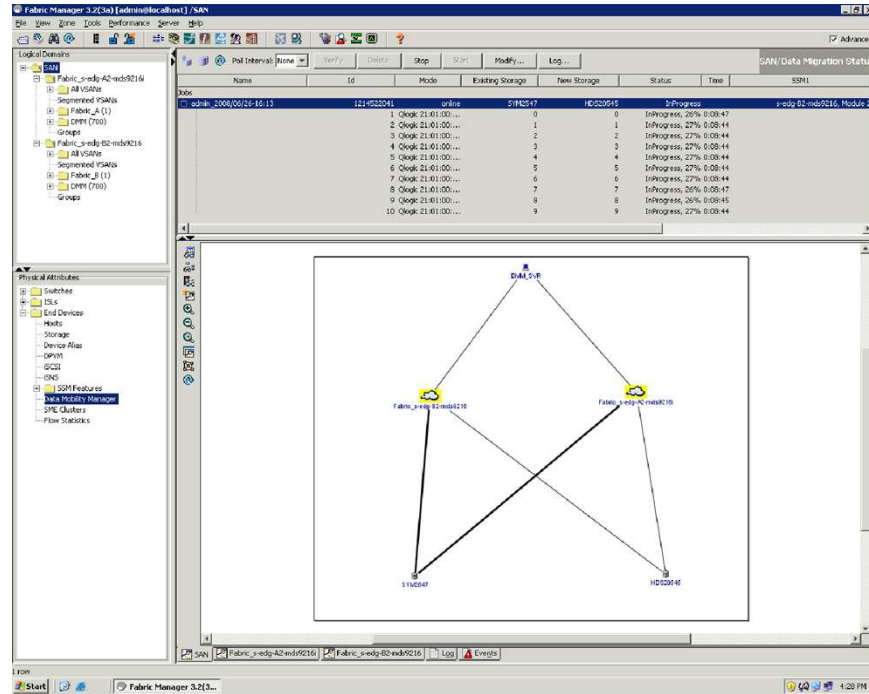
**Figure 8.** Modified Mapping of Existing LUNs to New LUNs





6. Figures 9 and 10 show the progress of the migration monitored when the Cisco DMM is selected in the physical attribute pane. Figure 11 shows the status after the data migration is completed.

**Figure 9.** Data Migration in Progress at 26 Percent Complete



**Figure 10.** Data Migration in Progress at 90 Percent Complete

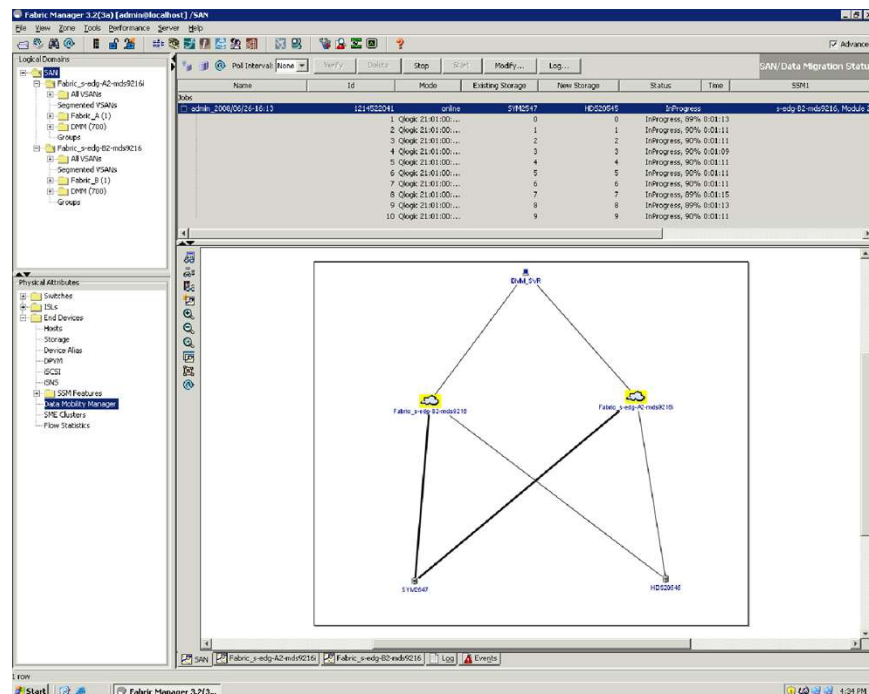
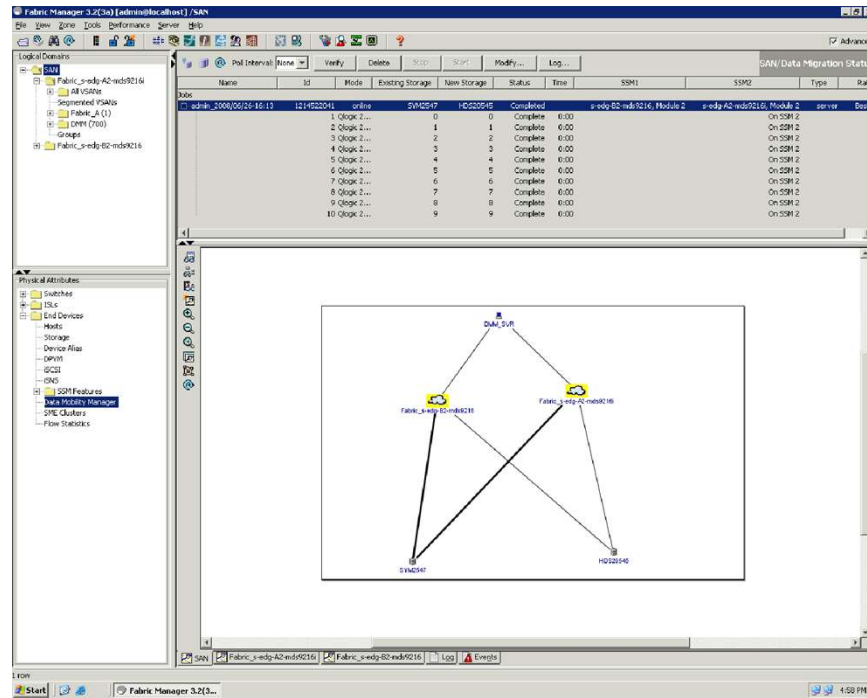




Figure 11. Completed Data Migration



## For More Information

<http://www.cisco.com/en/US/products/hw/ps4159/index.html>



Americas Headquarters  
Cisco Systems, Inc.  
San Jose, CA

Asia Pacific Headquarters  
Cisco Systems (USA) Pte. Ltd.  
Singapore

Europe Headquarters  
Cisco Systems International BV  
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumina, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco ICS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet QuikStart, IOS, iPhone, iQuikStart, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGC, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanel, ProConnect, ScriptShare, SenderBase, SMMRTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (081215)

