

Achieving Platform-Independent Data Migration in a Heterogeneous SAN Environment

Providing cost-effective, SAN-wide data migration while
maximizing application availability



Introduction

The Internet and the networked information economy are placing increasing demands on many organization's storage infrastructure. These demands include the need to serve more applications with more data, to support more I/O's and to be a part of the organization's disaster recovery solution.

Since an organization's storage infrastructure is often changing on a continual basis, the issue of effectively migrating data from one storage device to another is of paramount importance. The process of upgrading and/or replacing storage devices usually involves the challenge of "how to do it with minimal or no disruption of business continuity." To address these issues, some storage vendors provide utilities to handle the data migration process, usually between devices from the same vendor. In a SAN environment, however, where devices from multiple vendors are commonly used, the question often becomes "how to do it in a heterogeneous environment with minimal or no disruption."

Within many existing solutions, data migration often results in substantial downtime and/or involves large amounts of time and money. There are three key reasons for this:

- The traditional method for migrating data is by using a 'backup & restore' mechanism (typically via tape drives or libraries). However, this approach can be very tedious due to limited throughput rates, and often results in taking production applications offline for extended periods of time.
- Vendors of high-end storage solutions usually offer an alternative to 'backup & restore' via their migration utilities. While these utilities serve their purpose, the downsides can include high licensing fees and incompatibility with other vendors' products, often forcing customers to buy a solution from each storage vendor.
- Vendors of mid-range storage solutions typically do not provide data migration utilities, often forcing customers to use the 'backup & restore' approach discussed above.

However, the Storage Virtualization Manager (SVM™) solution set from StoreAge includes a number of capabilities that can dramatically simplify the data migration between storage devices, while using the high speed Fibre Channel SAN infrastructure. This eliminates many of the traditional barriers encountered by IT professionals, providing:

- A data migration process that minimizes the disruption of production applications
- A single, uniform data migration solution compatible with multiple platforms and storage devices from multiple vendors
- A highly scalable, high performance method for migrating large amounts of data very quickly between devices
- A phased migration process, which allows data centers to migrate data in an orderly manner to minimize the disruption of production applications

This paper will explore the process of migrating data between different storage devices in a simple and virtually non-disruptive manner, using StoreAge SVM combined with the StoreAge **multiCopy** data replication module.

Achieving Platform-Independent Data Migration While Maximizing Application Availability

1. Plan the Migration

The first step in achieving a platform-independent data migration capability is to virtualize the SAN environment with a network-based virtualization solution, such as StoreAge SVM. This will create a vendor-neutral layer within the SAN that will allow uniform services to be performed across all storage within a SAN, such as on-the-fly provisioning, data replication and data migration.

In planning any migration from a non-virtualized to a virtualized SVM environment, a number of key processes should be discussed beforehand, in order to ensure a smooth transition. This normally includes:

- Planning for the installation of SVM appliances on the SAN
- Planning for installation of the SVM agent package on each host, at a scheduled down time
- What Quality of Service criteria should be used to assign storage resources (LUNs) to different pools
- Volume and Pool naming conventions
- Defining the timing and procedure for migrating each server and its existing LUNs into the virtualized environment.

A written project plan defining all the steps of the SVM implementation is strongly advised. Note that the virtualization process can take place over a period of time, since SVM allows the virtualization of the SAN in phases that best suit each customer's unique requirements.

2. Connect SVM(s) to the SAN

SVM can be connected to a SAN in a completely unobtrusive manner. Subject to SAN security measures already in place, such as switch zoning and LUN masking, the SVM will discover the SAN elements and display what it sees via its GUI. This GUI is accessible via a standard web browser on any computer that can connect to an SVM appliance via TCP/IP.

The information displayed by the SVM interface includes:

- LUNs presented by any Storage Devices within the SAN (including any disks exposed by JBOD devices)
- HBAs installed in Hosts
- SVM Appliances

Most production SVM environments will contain two SVM appliances, to ensure high availability. Both SVMs should be given the same view of the SAN, so that as they discover and display the objects in the SAN, they arrive at the same results. If there are any discrepancies in the results, these will be displayed, and can be very useful in troubleshooting any connectivity and other types of problems that should be addressed before continuing.

This is also the time to begin installing the SVM client agent package on the hosts that will later be using virtualized volumes. The agents enable the SVM to see each host, and will be used to store information about which virtualized volumes each host is permitted

SVM creates a vendor-neutral layer within the SAN that allows uniform services to be performed across all storage within a SAN, such as on-the-fly provisioning, data replication and data migration.


to access. Note that even with the agent package installed, the host will continue to access existing non-virtualized LUNs until those LUNs are assigned to virtual volumes and re-mounted by the host. If existing path failover and/or load balancing drivers are in place, the SVM agent's failover and load balancing features are deactivated until the host's LUNs are virtualized.

In addition, the SVM "auto-zoning" feature that prevents hosts from accessing storage unless validated by SVM will be deactivated during the data migration process, but will later be activated once the migration phase is completed

In a case where high data rate is required, one or more host servers in the SAN may need to be used "data movers" during the migration process, in order to meet throughput requirements. Each one of these data movers will need to support the FC-to-FC copy command, and StoreAge can recommend appropriate data movers and configurations upon request.

The figure below illustrates the components involved in the data migration process.

To achieve platform-independent data migration, the SVM appliance and client host agents are added to the SAN environment. This allows data to be moved between any SAN storage device in a uniform manner, and eliminates incompatibilities between different vendors' devices. Depending on throughput requirements, one or more servers may act as "data movers" during the migration process.



3. Connect New Storage Devices

SVM enables data to be migrated between any storage devices within the SAN. If a new device is added to the SAN, SVM will automatically discover the new device and display its status as "Present" in the SVM GUI. The administrator then uses the GUI to assign the device's LUNs to a virtualized Storage Pool(s). At this point, the new device's SVM status is updated to "Configured", as is now available for production use.

For maximum protection, the SVM configuration data (Virtualization Metadata) is stored on the new storage device, rather than on the SVM itself. This protects the configuration information in the event an SVM should ever be powered off or replaced.

4. Actual Migration Process

Most migrations will be phased on a host-by-host basis or on a LUN-by-LUN basis. During the process, LUN masking and zoning may need to be adjusted, as SVM must always have access to the LUNs being migrated between the source and destination


storage devices. In addition, for any volume presented to a host to be accessible, the host must also have access to all the LUNs contained within that volume.

Part 1: Existing, non-virtualized data on the old storage device is virtualized. During this process, the affected production volumes will need to be dismounted for a brief period of time (1-2 minutes), while the LUNs are converted to SVM virtualized volumes, as follows:

- Within the SVM GUI, any non-virtualized LUNs accessible by SVM will be listed with a status of “Present”. The appropriate LUNs are then selected from the list, and are assigned to a virtualized Storage Pool
- SVM reserves approximately 6 MB of space at the beginning and end of each LUN that is virtualized. Any data currently residing in that space is automatically moved to any free space within the storage pool the LUN is joining, and SVM then writes its configuration data into that reserved space.
- Once a LUN is virtualized and is added to the production Storage Pools through the migration feature in Part 2 below, its status in the SVM GUI will be automatically changed to “Configured”.

During this process the user should not set any permissions to these new volumes.


In Part 1 of the process, the production server on the left continues in operation with the old storage system, while the non-virtualized LUNs 3+4 owned by the other server are converted to SVM virtualized volumes.



Part 2: Create an instant copy of the new virtual volumes, return the applications to production status, and migrate the data between the old and new storage devices.

- Using StoreAge **multiCopy**, instant “copy jobs” of the source virtual volumes are created on the new storage system.
- The instant copies on the new device have read/write permissions, may be immediately mounted by the production host, and the applications restarted while the migration process occurs in the background.
- As the data is migrated from the old device to the new, full data integrity is ensured. When write requests are received from the application, the new data will be written to the new storage device and upon migration completion will be merged with the old data.
- If data movers have been deployed, SVM will use them to increase migration performance, shrinking the migration timeframe.


In Part 2 of the migration process, StoreAge multiCopy creates instant “copy jobs” that begin mirroring the old data to the new storage system. The instant copies on the new system have read/write permissions, may be immediately mounted by the production host(s), and the applications restarted while the migration process occurs in the background. As the data is migrated from the old storage to the new, full data integrity is ensured. When write requests are received from the application, the new data will be written to the new storage device, and upon migration completion, will be merged with the old data. The red line indicates the control path between SVM and any data movers being used in the migration process.



Part 3: As the migration process for each application is completed, SVM features such as path failover and load balancing can be activated, providing high availability and any-to-any connectivity between host servers and storage devices.

Part 4: SVM allows the old storage devices to remain connected to the SAN, and be used as either primary or secondary storage (for online backup and restore, application testing, data warehouse updates, etc.). Or, if they are no longer needed, the old devices may be disconnected from the SAN entirely.

After all the data on the old storage device has been migrated, the old storage may be re-purposed and remain connected to the SAN, or be removed from the SAN entirely.



Summary

When combined, StoreAge SVM and **multiCopy** work together to deliver cost-effective, platform-independent data migration in heterogeneous SAN environments.

Together, these products enable comprehensive, end-to-end management of the entire data migration process, virtually eliminate application downtime, and allow continuous business operation and data availability while data is migrated between storage systems.

In addition, by supporting storage devices from virtually all vendors, StoreAge provides one unified method for migrating data between different devices, while avoiding the traditional reliance on 'backup and restore' mechanisms. For scalable, high performance data migration, the StoreAge solution enables the next generation of "best practices" to be implemented today.



U.S. Headquarters

One Technology Drive, Bldg. C-515
Irvine, CA 92618
T: (949) 754-0640
F: (949) 754-0645

International Headquarters

63 Bar Yehuda Road
Nesher 36651 Israel
T: +972 4 8203454
F: +972 4 8203464

For more information:

www.StoreAge.com
info@StoreAge.com

© 2004 StoreAge Networking Technologies Ltd. All rights reserved.

No part of this publication may be reproduced in any material form (including photocopying, or storing it in any medium by electronic means whether or not transiently or incidentally to some other use of this publication) without the prior written permission of the copyright owner, or under the terms of a license issued by the copyright owner.

This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any product, product feature, or service offered or to be offered by StoreAge. StoreAge reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use.

StoreAge, SVM, *multiView*, *multiCopy*, *multiMirror*, SVM PM for Exchange and SVM App-Pack are trademarks of StoreAge Networking Technologies Ltd. All other marks mentioned in this material are the property of their respective owners.