WHITE PAPER

SEAGATE BEST PRACTICES AND INTEGRATION GUIDE FOR VMWARE

CONTENTS

Introduction

Today's data centers have embraced server virtualization to optimize hardware resources, energy resources, and real estate. Data centers large and small have moved most non-critical applications to virtualized environments. The next step in this evolution toward virtualization is the migration of mission-critical, Tier 1 applications to the virtual machine. However, in order to gain the trust of IT departments, the virtualization environment must offer the features of a five nines (99.999%), high-availability infrastructure and provide guarantees of scalability, performance, and reliability. In order to accomplish this, all components of the infrastructure must meet these new demands.

VMware's flagship vSphere virtualization product is ready to manage mission-critical applications. It provides world-class high-availability features and resource management tools that allow customers to give priority to Tier 1 applications. To complement these distributed, virtualized computing environments, Seagate Systems offers networked storage solutions that provide equivalent high-availability and reliability features for the underlying storage subsystem.

Seagate Systems is an active partner with VMware in bringing comprehensive virtualized solutions to customers. The Seagate® Exos® X series storage arrays are designed to be integral components in demanding high-availability computing environments, delivering data to servers from a robust, reliable, and high-performance platform.

This technology paper will outline the advantages of using Seagate Exos X series storage products in VMware environments. To begin with, this guide will highlight elements of the strategic partnership between Seagate and VMware, including the advantages gained through product integration and certification. Implementation and tuning of disparate components in a system can sometimes prove daunting.

This guide includes a targeted list of best practices and configuration tips specific to VMware configurations. Finally, this guide will demonstrate the performance characteristics of Exos X storage in VMware environments, based on actual test results from application-specific test tools.

Seagate and VMware Strategic Partnership

Seagate Systems maintains a strategic partnership with VMware through the VMware Technology Alliance Partnership (TAP) program. As an active VMware partner, Seagate actively integrates VMware-specific features into its Exos X storage products and ensures that Exos X systems integrate seamlessly into VMware environments.

The primary means used to ensure seamless integration is the participation in the Hardware Certification Program at VMware. All Exos X products are subjected to compatibility tests defined by VMware for each VMware vSphere version, and are added to VMware's hardware compatibility list when the storage system passes this battery of tests. These processes and programs help ensure the integrity and quality of the overall solution. Seagate ensures that all Exos X products are certified with the latest releases of VMware vSphere.

The Seagate Exos X family of products is designed for high-availability computing environments, with redundant, active/active processing components, RAID protected storage, multiple data paths, and built-in snapshot and replication features. In addition, Exos X systems offer an assortment of drive and interface options, providing custom solutions to a wide variety of applications. Finally, the Exos X storage system is architected for speed. Every system is optimized for high-performance delivery of data regardless of the workload. Figure 1 shows a few examples of base configurations of Exos X series storage systems.

Figure 1. Seagate Exos X Series Storage Products

In addition to the quality features built into the Exos X storage array, the VMware vSphere software suite includes many features to enhance the overall storage management solution. These features dovetail and integrate with the Exos X, providing a complete portfolio for customers requiring a robust, scalable, high-performance computing and storage environment. Table I outlines many of the common storage-related features that complement or are integrated with the Exos X storage systems.

Table 1. Exos X Storage Integration With VMware vSphere

Connectivity	All Seagate® Exos® X series products are certified with the latest revisions of VMware vSphere. Seagate offers connectivity options with FC, iSCSI, and SAS interfaces.
High Availability	VMware supports a clustered server environment with automatic failover. The Exos X series storage systems offer fully redundant components and storage, providing an ideal storage complement to a VMware cluster. In addition, storage volumes from the Exos X can be presented to multiple hosts, enabling VMware vSphere HA features, such as vSphere vMotion.
Multipathing	The Exos X series is ALUA-compliant and compatible with standard VMware multipathing drivers. This allows all data to be transferred along optimum paths for maximum speed and efficiency.
Thin Provisioning	The Exos X series is fully compatible with the thin provisioning features of VMware. This feature can be leveraged to help simplify the management of storage resources.
vSphere vMotion	The Exos X series is an ideal storage solution for VMware vSphere vMotion, as it allows virtual machines from one ESX host to another, while underlying storage remains online.
vSphere Storage vMotion	The Exos X series is completely compatible with VMware vSphere Storage vMotion as a storage source or a storage target. Storage is efficiently and quickly migrated to or from the Exos X as the need arises.
Storage I/O Control (SIOC)	The Exos X series is a high-performance solution for many VMware workloads. In cases where I/O needs to be managed to assure bandwidth for particular applications, Exos X is fully compatible with VMware vSphere Storage I/O Control.

Finally, a key aspect of the relationship between Seagate and VMware is the service and support available to customers to ensure seamless integration of the software and the hardware. Seagate is committed to providing world-class support and services to resolve issues, assist in implementing best practices, and provide proactive maintenance and upgrades to their infrastructure. For more information, visit **www.seagate.com**.

SAN Storage Best Practices

The Seagate Systems Exos X series storage arrays are ideal data repositories for any VMware deployment. They are especially suited for clustered deployments, where the customer can take advantage of the high-availability features of Exos X, such as redundant components. In addition, distributed, SAN-based storage is a key enabler for clustered servers, and in the case of VMware, it enables important features, such as vMotion. Configuring a SAN environment for use with VMware is not much different from configuring a SAN for any other type of server. This section covers some general purpose SAN configuration best practices, along with some VMware specific guidelines.

In brief, best practice guidelines aim to accomplish two specific goals:

- **Redundancy.** Having no single point of failure is critical in the design of high-availability systems.
- **Performance.** In many cases, this translates to distributing I/O loads among available resources to realize the highest overall system performance.

To get the most from Exos X storage in a VMware environment, the following best practice guidelines have been established. These guidelines should help in deploying a solution that provides optimum performance, reliability, and availability.

Ensure proper cabling for performance and redundancy

One of the most important, yet challenging, aspects of storage network design and implementation is in the cabling. Proper cabling provides redundancy, so that the failure of any one component does not compromise access to any of the storage. At the same time, this redundancy should also provide a mechanism for aggregating throughput, delivering higher performance to the host systems. Following are a few rules and guidelines to follow when connecting components:

- Deploy redundant switches. A switch in a network represents a single point of failure. A best practice in high-availability systems is to utilize two switches between a storage system and a host.
- On each storage controller in the storage system, connect at least one of the host ports to each switch. If a controller has two host port connections, then connect one port to switch A and the other port to Switch B.
- Each host server should have at least two ports dedicated for connectivity into the SAN. Once again, this eliminates single points of failure in the network. Connect each port to a separate switch.

Ensure proper cabling for performance and redundancy

These principles are demonstrated in the Fibre Channel configuration for Exos X storage connectivity to VMware servers (Figure 2).

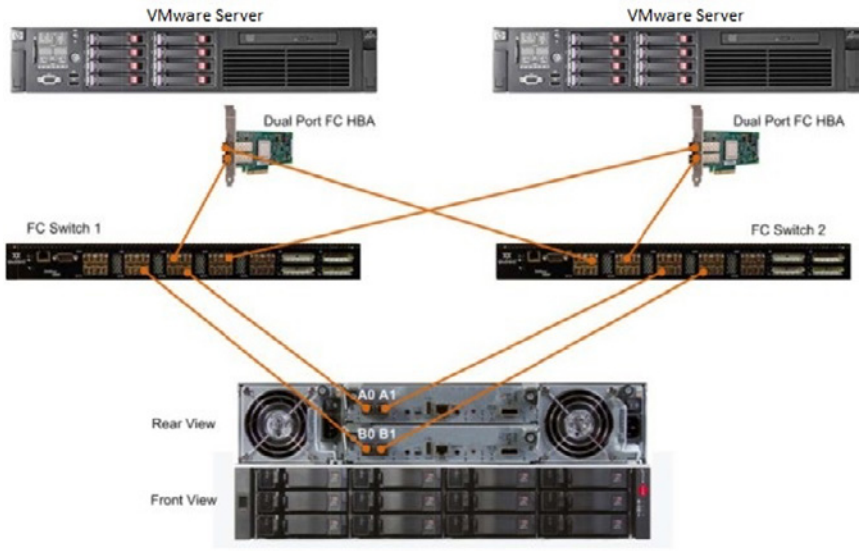

Figure 2. SAN Connectivity Best Practice

Distribute the load among controllers

The DH Exos X series RAID array typically comes with dual active-active controllers. With two controllers, it is important to optimize the processing power of both controllers in order to realize the fastest throughput and I/O to the system. To do this, the RAID groups that are defined on the array should be evenly distributed among the two controllers (each RAID group is assigned to an *owning* controller). For example, if there are 12 drives in the array, then it makes sense to create two RAID groups of six drives each: one RAID group assigned to controller A and the other RAID group assigned to controller B.

Distribute the load among virtual machines

As is demonstrated in the test results presented later in this technology paper, virtualized environments are not optimized for hosting just one virtual machine. While this may seem obvious and intuitive, it bears stating that performance is optimized when it is distributed among several virtual machines in a system.

Understand your multipathing setup

If the host system is connected to the storage system using these best practices, then VMware will be presented with multiple paths to each volume.

In addition, some of the paths will be considered *optimized* or *preferred*, and some of the paths will be considered *non-preferred*. The storage system will process I/O on any of the paths, but it will process I/O faster on the preferred paths. The reasons for having different types of paths stem from the fact that RAID groups are assigned an owning controller, and all I/O to the RAID group passes through the owning controller. Thus, I/O routed directly to the owning controller will be processed very quickly. On the other hand, I/O routed on paths that are connected to the alternate controller will ultimately be passed through the owning controller for processing, which slows things down. Storage devices that distinguish path types as being preferred or non-preferred are known as ALUA (asymmetric logical unit access) compliant devices, as defined in the SCSI 3 (SPC-3) protocol.

All Seagate Exos X series storage products are ALUA-compliant devices. They should be recognized by VMware as such and assigned the default multipathing module for ALUA storage devices. This module is known as *VMW_SATP_ALUA*.

The multipathing module selected for the Exos X array can be verified by examining the ***manage paths*** dialog for the associated devices in the *Configuration->Storage Adapters* section of an ESX server in vCenter. Alternatively, the following command can be issued in the ESX console:

```
esxcli nmp device list
```

The output from this command will list each device path in the system along with associated details. Pay attention to the following entries:

```
Device Display Name: Seagate Fibre Channel Disk... Storage
Array Type: VMW_SATP_ALUA Path Selection Policy: VMW_PSP_MRU
```

The `Device Display Name` will identify the device. The name *Seagate* should be in the output string. The `Storage Array Type` identifies the array as ALUA, active-active, or active-passive. If the policy is not `VMW_SATP_ALUA`, then the policy will need to be specifically set (see below). Finally, the `Path Selection Policy` will identify the load balancing policy.

By default, `VMW_PSP_MRU` (most recently used) is selected. This policy may not provide the best load balancing across the available ports. The policy of `VMW_PSP_RR` (round robin) may be selected instead to improve load balancing. This policy can be easily modified in vCenter for each path in the system. Alternatively, the default policy for the array can be modified to `VMW_PSP_RR` (see below) so that all paths are automatically configured to the new default of round robin, thus eliminating the need for manual adjustments in vCenter.

Data Protection With RAID and ADAPT

Seagate Exos X controllers offer RAID options, as well as Seagate ADAPT for data protection. The level of data protection is set when creating a new disk group.

Creating disk groups

The first step in provisioning storage to ESXi is to create a disk group from a pool of disks on the array. The number of disks and the level of data protection are defined when the disk group is created. A disk group can be created using the Seagate management console or the command line interface.

There are several options available when creating disk groups. In addition to RAID and ADAPT options, there are different linear and virtual storage technologies.

ADAPT and RAID

While RAID and Seagate ADAPT can provide similar levels of data protection, they accomplish this in different ways. RAID is the traditional configuration where the data is striped across a minimum of two disk drives, as well as one or more parity drives.

Seagate ADAPT is a RAID-based data protection level that maximizes flexibility, provides built-in spare capacity, and allows for very fast rebuilds, large storage pools, and simplified expansion.

ADAPT disk groups use all available space to maintain fault tolerance, and data is spread evenly across all the disks.

ADAPT has several advantages over using RAID.

- Significantly faster rebuild times are possible with ADAPT over RAID.
- ADAPT can use up to 128 drives, while RAID 6 (for example) is limited to 16 drives.
- It is easier to expand ADAPT, with up to 68 disks at a time versus being limited to one to four disks at a time.
- There is no need for spare disks with ADAPT; spare capacity is built in.

When new data is added, new disks are added, or the system will recognize that data is not distributed across disks in a balanced way, and will move the data to maintain balance across the disk group.

Reserving spare capacity for ADAPT disk groups is automatic, since disk space dedicated to sparing is spread across all disks in the system. In the case of a disk failure, data will be moved to many disks in the disk group, allowing for quick rebuilds and minimal disruption to I/O.

About virtual and linear storage

Seagate Exos X arrays use two different storage technologies that share a common user interface. One uses the virtual method while the other one uses the linear method.

Virtual storage is a method of mapping logical storage requests to physical storage (disks). It inserts a layer of virtualization such that logical host I/O requests are mapped onto pages of storage. Each page is then mapped onto physical storage. Within each page, the mapping is linear, but there is no direct relationship between adjacent logical pages and their physical storage.

A page is a range of contiguous LBAs in a disk group, which is one of up to 16 RAID sets that are grouped into a pool. Thus, a virtual volume as seen by a host represents a portion of storage in a pool. Multiple virtual volumes can be created in a pool, sharing its resources. This allows for a high level of flexibility, and the most efficient use of available physical resources.

Some advantages of using virtual storage:

- It allows performance to scale as the number of disks in the pool increases.
- It virtualizes physical storage, allowing volumes to share available resources in a highly efficient way.
- It allows a volume to be comprised of more than 16 disks.

The legacy linear method maps logical host requests directly to physical storage. In some cases, the mapping is 1-to-1, while in most cases the mapping is across groups of physical storage devices, or slices of them. This linear method of mapping is highly efficient. The negative side of linear mapping is lack of flexibility. This makes it difficult to alter the physical layout after it is established.

Virtual storage provides the foundation for data-management features, such as thin provisioning, automated tiered storage, read cache, and the quick rebuild feature. More in-depth information on these technologies is available in the Seagate 5005/4005/3005 Series Storage Management Guide.

To gain maximum flexibility and ease of expansion, it is recommended to use virtual storage.

Select an appropriate storage configuration

An important element of planning storage usage involves selection of RAID groups and RAID levels. While each customer scenario is unique, and trade-offs are necessary to balance performance with capacity, it is understood that RAID 10 will deliver superior I/O performance in random workloads (such as databases), whereas RAID 5 will deliver superior performance in sequential workloads. Given the distributed, virtualized nature of VMware environments, most I/O workloads will be random—regardless of the application.

In addition to RAID selection, a few tuning parameters may be applied to optimize performance in specific environments:

- Chunk Size. A chunk size of 512K has been shown to provide the best performance in standard random and sequential workloads. However, some mixed workloads generate superior results when the chunk size is lowered. Where possible, the customer is encouraged to experiment with a different chunk sizes to determine which works best in their environment.
- Read-Ahead Cache. The read-ahead cache provides significant performance gains in purely sequential workloads. However, in some mixed workloads, this mechanism may not provide optimum results. In largely random workloads, the read-ahead cache should be disabled to ensure the best possible performance.

Take advantage of storage I/O control in vSphere

The Storage I/O Control (SIOC) feature of vSphere provides for policy-based distribution of I/O resources. The benefit of this feature is granular control of bandwidth consumption by applications running on virtual machines. As an example, one virtual machine may be hosting an Oracle database while a second virtual machine hosts a file server. Both virtual machines utilize the same underlying storage system. Therefore, it is possible for the file server to heavily utilize the storage resources while the Oracle database is attempting to perform storage operations. This type of contention can compromise the overall performance of critical applications in a virtualized environment. SIOC can alleviate contention by guaranteeing bandwidth to critical applications.

SIOC is a crucial element in virtual server environments where Tier 1 applications need precedence over lower-priority systems. Even in more modest deployments, SIOC can provide an element of balance among applications, so that no application is starved for access to storage.

Conclusion

The Seagate Exos X storage arrays are an ideal storage solution in VMware computing environments. The Exos X has been demonstrated and certified to work seamlessly with VMware, vSphere servers, and associated virtual machines. In addition to the inherent benefits of SAN based storage, the Exos X storage arrays deliver:

- High Performance. 16G FC, 10G iSCSI, and 12G SAS
- Reliability. Redundant components and paths
- High Availability. Clustered controllers with failover technology
- Scalability and Flexibility. A variety of drive options and configurations

The high-availability features offered by VMware and Seagate unleash the potential of virtualized computing for mission-critical applications. These solutions provide the means by which applications can execute and migrate among various hardware platforms to ensure availability and meet performance requirements.

Additional Resources

VMware has published a Fibre Channel connectivity guide, which expands on some of the topics covered here in a more general fashion. You can refer to the VMware guides at these links:

vCenter Server 6.5 Storage Guide

vCenter Server 6.7 Storage Guide

vCenter Server 7.0 Storage Guide

Ready to Learn More?

Visit us at **seagate.com**

SEAGATE