

EMC® VNXe3200™ Best Practices For Performance

Applied Best Practices Guide

EMC® Enterprise & Mid-range Systems Division

Abstract

This applied best practices guide provides recommended best practices for installing and configuring VNXe™ systems for best performance.

September, 2014

Copyright © 2014 EMC Corporation. All rights reserved. Published in the USA.

Published September, 2014

EMC believes the information in this publication is accurate 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.

EMC2, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to the technical documentation and advisories section on EMC Online Support.

EMC VNXe Best Practices for Performance

Applied Best Practices Guide

Part Number H13512

Chapter 1	System Configuration.....	7
Essential guidelines		8
Storage Processor cache		8
Physical placement of drives		8
Hot Sparing		9
Availability and connectivity.....		9
Fibre Channel Connectivity.....		9
iSCSI Connectivity.....		10
NAS Connectivity.....		10
Chapter 2	Storage Configuration	11
General considerations		12
Drive type.....		12
Rules of thumb.....		12
RAID level.....		13
Calculating disk IOPS by RAID type.....		13
Storage pool considerations.....		14
Storage pool creation.....		14
Pool capacity considerations		15
Storage tiers		16
Storage object considerations		17
Virtual LUN Creation.....		17
File System Creation.....		17
VMware Datastore Creation.....		17
Chapter 3	Data Services.....	18
FAST VP.....		19
General		19
Tiering Policy.....		19
Data Relocation		19
Pool capacity utilization.....		19

Contents

Multicore FAST Cache	20
General considerations	20
Enabling Multicore FAST Cache on a running system	20
Snapshots.....	21
LUNs	21
File Systems.....	21
Deduplication.....	21
File Systems.....	22
Chapter 4 Application Specific Considerations	23
Block application tuning.....	24
Host file system alignment.....	24
VMware ESX Server with iSCSI Datastore.....	24
Chapter 5 Conclusion	25

As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this guide may not be supported by all revisions of the hardware or software currently in use. For the most up-to-date information on product features, refer to your product release notes.

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 EMC Online Support. Check the website to ensure that you are using the latest version of this document.

Purpose

The *Applied Best Practices Guide* delivers straightforward guidance to the majority of customers using the storage system in a mixed business environment. The focus is on system performance and maximizing the ease of use of the automated storage features, while avoiding mismatches of technology. Some exception cases are addressed in this guide; however, less commonly encountered edge cases are not covered by general guidelines and are addressed in use-case-specific white papers.

Guidelines can and will be broken, appropriately, owing to differing circumstances or requirements. Guidelines must adapt to:

- Different sensitivities toward data integrity
- Different economic sensitivities
- Different problem sets

These guidelines contain a few *DON'T* and *AVOID* recommendations:

- *DON'T* means: Do not do it; there is some pathological behavior
- *AVOID* means: All else being equal, it is recommended not to, but it still acceptable to do it

Audience

This document is intended for EMC customers, partners, and employees who are installing and/or configuring VNXe unified systems. Some familiarity with EMC unified storage systems is assumed.

Related documents

The following documents provide additional, relevant information. Access to these documents is based on your logon credentials. All of the documents can be found on <http://support.emc.com>. If you do not have access to the following content, contact your EMC representative.

Introduction to the EMC VNXe3200 - A Detailed Review – White Paper

EMC VNXe3200 High Availability – A Detailed Review – White Paper

EMC Unisphere for Next-Generation VNXe Series – White Paper

EMC FAST Suite for Next-Generation VNXe Series – White Paper

EMC VNXe3200 Unified Snapshots - A Detailed Review – White Paper

EMC VNXe3200 File Deduplication & Compression – A Detailed Review – White Paper

EMC Next-Generation VNXe Series Capacity and Performance Metrics – VNXe3200 A Detailed Review - White Paper

EMC VNXe3200 - Introduction to SMB 3.0 Support – White Paper

Using a VNXe3200 System with Fibre Channel (FC) or iSCSI LUNs – Manual & Guides

Using a VNXe3200 System with CIFS File Systems – Manual & Guides

Using a VNXe3200 System with NFS File Systems – Manual & Guides

Using a VNXe3200 System with VMware NFS or VMware VMFS – Manual & Guides

Chapter 1 System Configuration

This chapter presents the following topics:

Essential guidelines.....	8
Storage Processor cache.....	8
Physical placement of drives.....	8
Hot Sparing	9
Availability and connectivity.....	9

Essential guidelines

This paper introduces specific configuration recommendations that enable good performance from a VNXe3200 storage system. At the highest level, good performance design follows a few simple rules. The main principles of designing a storage system for performance are:

- Flash First – Utilize flash storage for the active dataset to achieve maximum performance
- Distribute the load over available hardware resources
- Design for 70 percent utilization (activity level) for hardware resources
- When utilizing Hard Disk Drives (HDD), *AVOID* mixing response-time-sensitive I/O with large-block I/O or high-load sequential I/O
- Maintain latest released VNXe Operating Environment version

Storage Processor cache

Storage Processor memory configuration is not required. Memory allocation amounts and cache page size are not configurable parameters.

Physical placement of drives

When initially placing drives in the array:

- Spread flash drives across all available buses
- There are no restrictions around using or spanning the DPE

Hot Sparing

Hot sparing is the process of rebuilding a failed drive's data onto a system-selected compatible drive. Any unbound non-system drive can be considered for sparing. When planning Hot Spares consider the following recommendations:

- Plan to reserve at least one of every 30 installed drives of a given type
 - Verify count in the GUI or CLI
 - Storage > Storage Configuration > Spare Disks
 - `uemcli /env/disk -unused show`
 - **Note:** Unbound system drives (DPE Disk 0 through DPE Disk 3) cannot be used as hot spares
- Ensure that unbound drives for each drive type are available
 - SAS Flash / FAST Cache SSD (SLC) must spare for SAS Flash / FAST Cache SSD (SLC)
 - SAS Flash VP / FAST VP SSD (eMLC) must spare for SAS Flash VP / FAST VP SSD (eMLC)
 - SAS must spare for SAS (regardless of rotational speed)
 - NL-SAS must spare for NL-SAS
- The capacity of an unbound drive should be equal to or larger than the provisioned drives for which it will spare

Availability and connectivity

The VNXe3200 storage system offers connectivity to a variety of client operating systems, using multiple protocols, such as FC, iSCSI, NFS, and CIFS. EMC provides connectivity guides with detailed instructions for connecting and provisioning storage via different protocols to the specific host types.

EMC recommends you consult the connectivity documents on <http://support.emc.com>. The host connectivity guides listed below provide detailed configuration guidelines for your operating system:

- *Using a VNXe3200 System with Fibre Channel (FC) or iSCSI LUNs*
- *Using a VNXe3200 System with CIFS File Systems - Manual & Guides*
- *Using a VNXe3200 System with NFS File Systems - Manual & Guides*
- *Using a VNXe3200 System with VMware NFS or VMware VMFS - Manual & Guides*

Fibre Channel Connectivity

Fibre Channel connectivity is facilitated via the FC optical I/O module.

- Use multiple FC I/O ports on each SP, and balance host port connections across FC I/O ports, as host port connections affect the preferred CPU core assignment

- If not connecting all the available FC I/O ports, use the even numbered ports on each FC I/O module before using any odd numbered ports

iSCSI Connectivity

iSCSI connectivity is facilitated via the onboard copper NICs. The NICs will connect at 100Mbps, 1Gbps, or 10Gbps.

- Use 10Gbps for the best performance
- Configure Jumbo Frames (MTU of 9000) on all iSCSI ports
 - **Note:** The entire network infrastructure must also support Jumbo Frames
- When possible, segregate iSCSI traffic onto dedicated storage networks

NAS Connectivity

NAS protocols (NFS and CIFS) are facilitated via the onboard copper NICs. The NICs will connect at 100Mbps, 1Gbps, or 10Gbps.

- Use 10Gbps for the best performance
- Configure Jumbo Frames (MTU of 9000) on all NAS ports
 - **Note:** The entire network infrastructure must also support Jumbo Frames)
- It is recommended to use network trunking and multipathing in order to provide port failover and greater aggregate bandwidth for a NAS server on a single SP
 - Configure Link Aggregation Control Protocol (LACP) across 2 or more ports on a single SP
 - Use LACP instead of EtherChannel
 - Connect matching ports on the 2 SPs in identical configurations
 - This ensures network redundancy across SPs

Chapter 2 Storage Configuration

This chapter presents the following topics:

General considerations	12
Storage pool considerations	14
Storage object considerations.....	17

General considerations

Drive type

Match the appropriate drive type to the expected workload:

Drive type	Workload type
SAS Flash / FAST Cache SSD (SLC)	For extreme performance; these provide the best performance for transactional random workloads, and the lowest write service times. Required for Multicore FAST Cache
SAS Flash VP / FAST VP SSD (eMLC)	For extreme performance FAST VP tier; these are a higher capacity flash option. Not for use with Multicore FAST Cache.
SAS	For general performance tier.
NL-SAS	For less active data, well-behaved streaming data, archive purposes, and backups.

Rules of thumb

Disk drives are a critical element of unified performance. Use the rule of thumb information to determine the number of drives to use to support the expected workload.

These guidelines are a conservative starting point for sizing, not the absolute maximums.

Rules of thumb (RoT) for drive bandwidth (MB/s):

- Bandwidth assumes multiple large-block sequential streams
- Parity does not count towards host write bandwidth sizing
 - For example; a 4+1 RAID group using SAS 15K to support a sequential write workload is sized at 100 MB/s for the entire RAID group (4*25 MB/s)

Bandwidth	NL-SAS	SAS 10K	SAS 15K	Flash (All)
RoT per drive, Sequential Read	15 MB/s	25 MB/s	30 MB/s	90 MB/s
RoT per drive, Sequential Write	10 MB/s	20 MB/s	25 MB/s	75 MB/s

- For rule of thumb bandwidth sizing with HDD as RAID 5 or RAID 6, VNXe3200 scales up to the sweet spot drive count (drive counts include parity drives)
 - Systems can hit maximum bandwidth with fewer drives when the per-drive performance is exceeding the rule of thumb sizing guidelines

Max Bandwidth	Fibre Channel LUN	iSCSI LUN	NAS file system
Read	150 drives	100 drives	100 drives
Write	120 drives	80 drives	80 drives

Rules of thumb (RoT) for drive throughput (IOPS):

- IOPS assumes small block random with good response time
 - Drives are capable of a sustained IOPS workload based on drive type
 - System drives (DPE Disk 0 through DPE Disk 3) have reduced performance expectations due to the management activities they support; rules of thumb for these drives are adjusted accordingly
 - **Note:** The system drives can be included in storage pools; be aware of the adjusted IOPS, as well as the reduced capacity on these drives
- For rule of thumb IOPS sizing with HDD, the VNXe3200 scales linearly with additional drives, up to the maximum drive count
- To size for host IOPS, you must include the RAID overhead as described in the section [Calculating disk IOPS by RAID type](#)

Throughput	NL-SAS	SAS 10K	SAS 15K	SAS Flash VP (eMLC)	SAS Flash (SLC)
Per drive RoT	90 IOPS	150 IOPS	180 IOPS	3500 IOPS	5000 IOPS
System drive RoT	60 IOPS	100 IOPS	120 IOPS	--	--

RAID level

For best performance from the least number of drives, match the appropriate RAID level with the expected workload:

RAID level	Expected workload
RAID 1/0	Works best for heavy transactional workloads with high (greater than 30 percent) random writes, in a pool with primarily HDDs
RAID 5	Works best for medium to high performance, general-purpose and sequential workloads
RAID 6 for NL-SAS	Works best with read-biased workloads such as archiving and backup to disk RAID 6 provides additional RAID protection to endure longer rebuild times of large drives

Calculating disk IOPS by RAID type

Front-end application workload is translated into a different back-end disk workload based on the RAID type in use.

For reads (no impact of RAID type):

$$1 \text{ application read I/O} = 1 \text{ back-end read I/O}$$

For random writes:

$$\text{RAID 1/0 - } 1 \text{ application write I/O} = 2 \text{ back-end write I/O}$$

$$\text{RAID 5 - } 1 \text{ application write I/O} = 4 \text{ back-end disk I/O (2 read + 2 write)}$$

$$\text{RAID 6 - } 1 \text{ application write I/O} = 6 \text{ back-end disk I/O (3 read + 3 write)}$$

Storage pool considerations

Storage pool creation

Create multiple pools in order to:

- Separate workloads with different I/O profiles
 - Predominantly sequential workloads should be placed in dedicated pools
- Separate pools for block and file
- Dedicate resources, to meet specific performance goals
- Vary pool parameters, such as Multicore FAST Cache enabled/disabled
- Minimize failure domains
 - Although unlikely, loss of a private RAID group in the pool compromises the total capacity of that pool; it may be desirable to create multiple smaller pools rather than use the total capacity available in a single pool

Storage pools have multiple RAID options per tier for preferred type and drive count

- Use RAID 5 with a preferred drive count of 4+1 for the best performance versus capacity balance
 - Using 8+1 or 12+1 improves capacity utilization at the expense of reduced availability
- Use RAID 6 for NL-SAS tier
 - Preferred drive counts of 6+2, 8+2, or 10+2 provide the best performance versus capacity balance
 - Using 14+2 provides the highest capacity utilization option for a pool, at the expense of slightly lower availability and performance
- Use RAID 1/0 when a high random write rate (> 30%) is expected with HDD
 - For best possible performance with RAID 1/0, use the largest available preferred drive count (i.e., 4+4 > 3+3 > 2+2, etc.)
- Consider the following rule of thumb for tier construction:
 - Extreme performance flash tier: 4+1 RAID 5
 - Performance SAS tier: 4+1 or 8+1 RAID 5
 - Capacity NL-SAS tier: 6+2 or 8+2 RAID 6

Recommendations for creating and expanding storage pools:

- When creating a pool, it is best to specify a multiple of the preferred drive count for each tier you select
 - For example, when using RAID 5 4+1, specify a drive count of 5, 10, 15, etc.
- It is best to maintain the same capacity and rotational speed of all drives within a single tier of a given pool
 - For example, *AVOID* mixing 600GB 10K SAS drives in the same pool with 300GB 15K SAS drives; instead, split them into 2 different pools

- Within a given pool, use all of the same flash technology for the extreme performance tier
- When expanding pools, use a multiple of the preferred drive count already in use for the tier being expanded

Pool capacity considerations

EMC recommends leaving free space in the storage pool, to accommodate data services.

Note: The pool can still be oversubscribed above 100% of actual capacity. The values here refer to actual physical space that is not used in the pool.

- When using Snapshots, at least 5% free space is needed; EMC recommends maintaining about 10% to buffer snapped writes
- When using FAST VP, EMC recommends maintaining at least 10% free to accommodate the quickest rebalancing
- When using FAST VP and Snapshots together in a pool, a total of 10% free will meet the requirements of both

Note: By default, the VNXe3200 will begin issuing alerts when more than 70% of available capacity has been subscribed.

File systems share space in the storage pool with their Snapshots, and also with Block LUNs if the pool is shared.

- *DON'T* oversubscribe space in a storage pool that contains file systems
- Ensure that the storage pool has sufficient capacity to cover the maximum size of all file systems, plus the capacity needed for any Snapshots, plus the maximum size of all Block LUNs

Storage tiers

Skew is when a small percentage of the total storage capacity in a storage system is the target for the majority of the IOPS served by the system. It is the locality of active data within the total storage capacity. For instance, in a payroll system, the current month's data will be highly active, year-to-date data will be moderately active, and the data for previous years will be mostly inactive.

Storage tiers improve system performance by providing faster drives to handle the more active data, while keeping the less active data on inexpensive capacity drives.

The number of tiers required in a storage pool is influenced by performance requirements, capacity requirements, and the knowledge of the skew between active and inactive capacity. Best performance is achieved when the entire active dataset can be contained within the capacity of the Extreme Performance (flash) and Performance (SAS) tiers.

If the active capacity (skew) is known, the capacity per tier should be sized accordingly. Using the payroll system example again, if 7 years of data are kept in the storage system, less than 20% of the data is expected to be highly active (1 year out of 7).

If the active capacity is not known, consider capacity per tier of 5 percent flash, 20 percent SAS, and 75 percent NL-SAS. This works on the assumption that less than 25 percent of the used capacity will be active, and infrequent relocations from the lowest tier will occur.

Follow these general guidelines:

- When Multicore FAST Cache is available, use a 2-tier pool comprised of SAS and NL-SAS. Enable Multicore FAST Cache as a cost-effective way of realizing flash performance without dedicating flash to this pool
 - Flash tier can be added later if Multicore FAST Cache is not fully capturing the active data
- For a 3-tier pool, start with 5 percent flash, 20 percent SAS, and 75 percent NL-SAS for capacity per tier if skew is not known
 - Tiers can be expanded after initial deployment to effect a change in the capacity distribution if needed
- Use a 2-tier pool comprised of flash and SAS as an effective way of providing consistently good performance.
 - NL-SAS can be added later if capacity growth and aged data require it
- *AVOID* using a 2-tier pool of flash and NL-SAS if there is uncertainty about the active data fitting in the flash tier
 - The SAS tier provides a buffer for active data not captured in the flash tier; the SAS tier still provides modest performance, as well as quicker promotion to flash when relocations occur
- Add a flash tier to a pool with thin LUNs so that metadata is promoted to flash and overall performance is improved

Storage object considerations

Virtual LUN Creation

LUNs can be created as either thick (fully allocated) or thin (virtually provisioned).

- Thick LUNs (without Snapshots) are recommended for the highest level of pool-based performance
 - A thick LUN's performance can be better than the performance of a thin LUN
- Thin LUNs are recommended when storage efficiency requirements outweigh performance requirements
 - When using thin LUNs, adding a flash tier to the pool can improve performance
 - Thin LUN metadata can be promoted to the flash tier when FAST VP is enabled
- Thin LUNs are recommended when implementing Snapshots on Block LUNs

File System Creation

When provisioning file systems for File access, a thick storage object is automatically provisioned from the storage pool to hold the file system.

- Create a separate pool for File Systems
 - *AVOID* mixing with Block workloads
 - *DON'T* oversubscribe a storage pool that contains File Systems
- Create thin (virtually provisioned) file systems
 - Thin file systems provide better capacity utilization from the pool, especially if Snapshots will be taken
 - Thin file systems provide better performance when used with FAST VP tiering

When creating NAS servers

- Create at least one NAS server on each SP
- Balance file system access across NAS servers on SPA and SPB
 - A file system is only accessed from a single SP

VMware Datastore Creation

VMware Datastores are created as either NFS or VMFS datastores.

- When creating NFS datastores, observe the best practice recommendations for [File System Creation](#)
- When creating VMFS datastores, observe the best practice recommendations for [Virtual LUN Creation](#)

Chapter 3 Data Services

This chapter presents the following topic:

- FAST VP 19
- Multicore FAST Cache..... 20
- Snapshots 21
- Deduplication 21

FAST VP

General

FAST VP moves data between tiers in a pool based on the performance needs of the data. Construct the pool such that each tier will provide consistent performance.

- Use consistent drive technology for each tier within a single pool
 - Same flash drive technology and drive size for the extreme performance tier
 - Same SAS RPM and drive size for the performance tier
 - Same NL-SAS drive size for the capacity tier

Tiering Policy

For most tiered pools, using auto-tier will provide the best performance.

- *DON'T* use auto-tier for LUNs with low-skew random workloads where the active dataset will not fit in the highest tier
 - This might cause excessive tier relocations that may not benefit the active data
- *AVOID* using highest-available for a LUN whose capacity exceeds 90% the highest tier capacity
 - This can affect the overall efficiency of the highest tier to service active data for LUNs running in auto-tier mode
- *AVOID* using lowest-available with thin LUNs, as this will force the metadata into the lowest tier as well

Data Relocation

Relocation is the process of moving pool data slices across tiers, or within the same tier, to move hot data to higher performing drives, or to balance underlying drive utilization. Relocation can occur as part of a FAST VP scheduled relocation, as an automated relocation after a storage pool expansion, or as a result of manually requested relocation.

- Enable FAST VP on a pool, even if the pool only contains a single tier, to provide ongoing load balancing across available drives based on slice temperature and capacity utilization
- Schedule relocations for off-hours, so that relocation activity does not contend with the primary workload
- Schedule relocations to run before or during backup windows, so that the relocations are based on the primary workload activity

Pool capacity utilization

FAST VP requires unallocated space within the pool to accommodate data relocations.

- EMC recommends leaving 10% free space in storage pools when FAST VP is enabled
 - Relocation will attempt to reclaim 10 percent free per tier

- Free space is used to optimize relocation operations
- Free space is used for new allocations to thin LUNs and file systems
- Free space is used to support Snapshot schedules

Multicore FAST Cache

Multicore FAST Cache is best for small random I/O where the active data has skew. The higher the skew, the greater the benefit that Multicore FAST Cache provides. Multicore FAST Cache also adapts quickly to changes in locality.

General considerations

EMC recommends first utilizing available flash drives for Multicore FAST Cache, which can globally benefit all LUNs in the storage system. Supplement performance as needed with additional flash drives in storage pool tiers.

Preferred application workloads for Multicore FAST Cache:

- Small-block random I/O applications with high locality
- High frequency of access to the same data
- Systems where current performance is limited by HDD capability, not SP capability

AVOID enabling Multicore FAST Cache for pools that are not expected to benefit, such as when:

- The primary workload is sequential
- The primary workload is large-block I/O
- The primary workload is small-block sequential, like database logs or circular logs

Enabling Multicore FAST Cache on a running system

When adding Multicore FAST Cache to a running system, it is recommended to enable Multicore FAST Cache one pool at a time. Wait until the LUNs and/or file systems in that pool have reached steady state in Multicore FAST Cache before enabling more pools.

Multicore FAST Cache can improve overall system performance if the current bottleneck is drive-related, but boosting the IOPS will result in greater CPU utilization on the SPs. Generally, EMC recommends sizing systems so that the maximum sustained utilization is 70 percent.

On an existing system, check the SP CPU utilization of the system, and then proceed as follows:

CPU utilization	Recommended sizing approach
Less than 60%	Enable one pool at a time; let it reach steady state in Multicore FAST Cache, and ensure that SP CPU utilization is still acceptable before enabling Multicore FAST Cache on more pools

60-80%SP	Scale in carefully; enable Multicore FAST Cache on one pool with the smallest capacity, and verify that SP CPU utilization does not go above 80 percent
Greater than 80%	<i>DON'T</i> activate Multicore FAST Cache
Note: For storage pools, Multicore FAST Cache is a pool-wide feature so you have to enable/disable at the pool level (for all objects in the pool).	

Snapshots

Snapshots are used to take point-in-time checkpoints of LUNs and file systems.

LUNs

When using Snapshots with LUNs:

- Start with thin LUNs, to provide the most optimal pool capacity utilization
- Plan for the deletion of snapshots
 - Whenever possible, schedule the deletion of Snapshots during non-peak hours of operation
 - If snapshots must be deleted during peak periods of array activity, lessen the impact by reducing the number of concurrent Snapshot deletes (for example, stagger the delete operations over several hours, instead of all at once)
 - *DON'T* delete the last snapshot of a Thick LUN, if you intend to create another snapshot immediately after deleting the last snapshot
 - Create the new snapshot before deleting the older snapshot
 - Deleting the last snapshot of a Thick LUN will undo the thin conversion, which would then be reconverted for the new snapshot

File Systems

When using Snapshots with file systems:

- Start with thin file systems, to provide the most optimal pool capacity utilization
- Migrate data into a new file system before enabling Snapshots on the file system
- *DON'T* delete the last snapshot of a file system, if you intend to create another snapshot immediately after deleting the last snapshot
 - Create the new snapshot before deleting the older snapshot
 - Deleting the last snapshot of a file system will undo the thin conversion, which would then be reconverted for the new snapshot

Deduplication

File Systems

If using file system deduplication and compression:

- Enable deduplication and compression on a file system and allow the initial scan to complete before enabling Snapshots
- Use file extension and path filtering to exclude large files that will not compress or deduplicate well, such as JPG or MP3 files

Chapter 4 Application Specific Considerations

This chapter presents the following topics:

Block application tuning 24

Block application tuning

Host file system alignment

File system alignment is covered in detail in the host connectivity guides on <http://support.emc.com>. In general:

- Windows Server 2008 and later automatically align
- Recent Linux operating systems automatically align

When provisioning LUNs for older Windows and Linux operating systems that use a 63-block header, the host file system needs to be aligned manually. Follow these alignment practices:

- Use host-based methods to align the file system
- EMC recommends aligning the file system with a 1 MB offset

VMware ESX Server with iSCSI Datastore

When provisioning iSCSI LUNs to an ESX Server:

- Configure ESX to use “round robin” path settings
 - Configure ESX NMP IOPs setting of 1 instead of the default of 1000
- Use Jumbo Frames (9000 MTU) with iSCSI, including all host NICs, network equipment, and array ports
- Use a physically segregated network when possible
- Disable **Delayed Ack** for iSCSI storage adapters and targets.
 - For further detail, see VMware Knowledge Base article 1002598
 - <http://kb.vmware.com/kb/1002598>

Chapter 5 Conclusion

This best practices guide provides configuration and usage recommendations for VNXe3200 systems in general usage cases.

For detailed discussion of the reasoning or methodology behind these recommendations, or for additional guidance around more specific use cases, see the documents the [related documents](#) section.