

EMC VNX UNIFIED BEST PRACTICES FOR PERFORMANCE

Applied Best Practices Guide

EMC UNIFIED STORAGE

Abstract

This white paper provides recommended best practices for installing and configuring VNX™ unified systems for best performance.

August 2012

Copyright © 2012 EMC Corporation. All rights reserved.

Published August 2012

EMC believes the information in this publication is accurate as 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.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

EMC, Unisphere, VNX, CLARiiON, and Celerra are registered trademarks in the United States. All other trademarks used herein are the property of their respective owners.

EMC VNX Unified Best Practices for Performance

Applied Best Practices Guide

Part Number 10938

Purpose	5
Audience.....	6
Related documents	6
Chapter 1 System Configuration.....	7
Essential guidelines.....	8
Storage processor cache size.....	8
Storage processor cache page size	8
Storage processor cache watermarks.....	9
Physical placement of drives.....	9
General guidelines	9
Flash drives for FAST Cache.....	9
Flash drives for extreme performance FAST VP tier	9
Hot spares	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
Traditional RAID group considerations	14
Drive count	14
Drive location selection.....	15
Large element size	15
Storage pool considerations	16
Storage pool creation.....	16
Virtual LUN type	17
Virtual LUN ownership considerations.....	17
Storage tiers	18

Considerations for VNX for File	19
Chapter 3 Data Services.....	21
FAST VP.....	22
Pool capacity utilization.....	22
Relocation.....	22
Considerations for VNX for File.....	22
FAST Cache	22
General considerations.....	22
Enabling FAST Cache on a running system	23
Considerations for VNX for File.....	24
Chapter 4 VNX for Block Specific Considerations	25
Availability and connectivity	26
Storage group settings.....	26
Host file system alignment.....	26
VNX for Block features	27
Block replication.....	27
Block compression	27
Chapter 5 VNX for File Specific Considerations	29
File-system creation	30
Traditional RAID group LUNs	30
Storage pool LUNs	30
VNX for File network and protocol considerations	31
Interfaces.....	31
NFS.....	31
VNX for File features.....	31
SnapSure.....	31
IP Replicator.....	31
Deduplication and compression	32
CAVA.....	32
VNX for File application tuning.....	32
VMware/database over NFS.....	32

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 might not be supported by all revisions of the software or hardware 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 might be released to the EMC Online Support website. 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 such as sole-use implementations of backup-to-disk, data stream ingest, and so on, 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

The guidelines contain a few “don’ts”:

- DON’T means: Do not do it; there is some pathological behavior.
- AVOID means: All else being equal, it is generally best not to, but it still is okay to do it

Audience

This document is intended for EMC customers, partners, and employees who are installing and/or configuring VNX 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. If you do not have access to the following content, contact your EMC representative.

- *Introduction to the EMC VNX Series - A Detailed Review – White Paper*
- *EMC Unified Storage System Fundamentals for Performance and Availability – White Paper*
- *EMC VNX Virtual Provisioning – White Paper*
- *EMC FAST VP - A Detailed Review – White Paper*
- *EMC FAST Cache - A Detailed Review – White Paper*
- *Microsoft Exchange 2010: Storage Best Practices and Design Guidance for EMC Storage – White Paper*
- *Deploying EMC VNX Unified Storage Systems for Data Warehouse Applications - Best Practices for Adoption and Deployment – White Paper*
- *EMC Performance for Oracle - EMC VNX, Enterprise Flash Drives, FAST Cache, VMware vSphere – White Paper*
- *Using EMC VNX Storage with VMware vSphere – TechBook*
- *EMC VNX for File on the Enterprise Network – Tech Note*
- *Celerra IP Replicator V2 Best Practices for Performance – Tech Note*
- *EMC CLARiiON Reserved LUN Pool Configuration Considerations - Best Practices Planning*
- *MirrorView KnowledgeBook: Release 30-32 - A Detailed Review*
- *EMC Host Connectivity Guides*
 - Windows/Linux/Solaris/HP-UX/AIX/Tru64/VMware

Chapter 1 System Configuration

This chapter presents the following topics:

Essential guidelines.....	8
Storage processor cache size	8
Storage processor cache page size.....	8
Storage processor cache watermarks	9
Physical placement of drives.....	9
Hot spares	10

Essential guidelines

This guide introduces specific configuration recommendations that enable good performance from a VNX Unified storage system. At the highest level, good performance design follows a few simple rules. The three main tenets of designing a storage system for performance are:

- Distribute load over available hardware resources.
- Design for 70 percent utilization (activity level) for hardware resources.
- AVOID mixing response-time-sensitive I/O with large-block I/O and high-load sequential I/O.

Storage processor cache size

When configuring cache settings for storage processors (SP) on a new VNX system:

- Install add-on software enablers before setting read cache and write cache
- Set read cache to roughly 10 percent of available cache; 200 MB is the recommended minimum and 1024 MB is the recommended maximum
 - Read cache facilitates pre-fetching and does not have to be large. Increase read cache above recommended values only if known to have multiple sequential read-intensive applications.
 - Set read cache the same on both SPA and SPB
- Give all remaining memory to write cache

Storage processor cache page size

Cache page size determines the minimum amount of SP memory used to service a single I/O operation.

- Default of 8 KB is fine for the majority of workloads.
 - In mixed environments, this default provides a good balance.
 - Leave at 8 KB for unified configurations with both Block and File storage and for File-only configurations.
- Increase to maximum 16 KB if large-block I/O size is predominant in the environment.
- With predominant small-block access, like 2 KB and 4 KB database environments, match cache page size to the predominant I/O size.

Storage processor cache watermarks

Cache watermarks control the flushing behavior of write cache.

- Start with low watermark of 60 percent and high watermark of 80 percent.
 - This is the default level and is suitable for the majority of workloads.
- If frequent forced flushing occurs, reduce watermark values.
- Maintain a difference between low and high of about 20 percent.
- AVOID drastic changes to these values unless advised by EMC Support.
- Recommended values for cache settings do not change when VNX for File is involved

Physical placement of drives

General guidelines

When initially placing drives in the array:

- Place highest performing drives in lowest numbered enclosures on each bus.
- Spread each drive type across all available buses.
- AVOID placing drives in the DPE or DAE-OS enclosure (0_0) that will be mirrored with drives in another enclosure. For example, AVOID mirroring a disk in 0_0 with a disk in 1_0.
- AVOID placing drives in the DPE or DAE-OS enclosure (0_0) if they will be in a parity RAID Group where one disk is placed outside of that enclosure.

Flash drives for FAST Cache

When Flash drives will be included for use as FAST™ Cache:

- Place all Flash drives in enclosure 0_0, up to 8 drives. If over 8 drives:
 - Spread Flash drives across all available buses.
 - Mirror drives within an enclosure, to AVOID mirroring across 0_0.

Flash drives for extreme performance FAST VP tier

When Flash drives will be included for use as a FAST VP tier:

- Spread Flash drives across all available buses.
- AVOID using enclosure 0_0.

Hot spares

Hot spares are drives provisioned to automatically replace failing drives or drives that have abruptly failed.

- Allocate one hot spare for every 30 drives.
- Ensure that hot spares for each drive type in the system are available.
 - FLASH must spare for FLASH.
 - SAS must spare for SAS (regardless of rotational speed).
 - NL-SAS must spare for NL-SAS.
- Hot spare capacity should be equal to or larger than the production drives.

Chapter 2 Storage Configuration

This chapter presents the following topics:

General considerations	12
Traditional RAID group considerations	14
Storage pool considerations	16

General considerations

Drive type

Match the appropriate drive type to the expected workload:

- FLASH for extreme performance; best performance for transactional random workloads
- SAS for the general performance tier
- NL-SAS for well-behaved streaming, aging data and 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.

Rules of thumb for drive performance:

- These are a conservative starting point for sizing, not the absolute maximums.
- IOPS assumes small block random with good response time.
 - Drives are capable of a sustained IOPS workload based on drive type, as follows:

	NL-SAS	SAS 10K	SAS 15K	Flash
IOPS	90	140	180	3500

- For IOPS with spinning drives (SAS and NL-SAS), VNX models scale linearly with additional drives, up to the maximum drive count per model.
- Bandwidth assumes large-block sequential with multiple streams.
 - Spinning drives provide about 30 MB/s of bandwidth capability; Flash drives provide about 100 MB/s.
 - For bandwidth with spinning drives, VNX models scale up to the sweet spot drive count, which varies by model as follows:

	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Drive Count	60	80	120	140	300

RAID level

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

- RAID 1/0 for heavy transactional with high (greater than 25 percent) random write component, and you plan to stay on HDD
- RAID 5 for medium-high performance, general-purpose workloads, and sequential
- RAID 6 for NL-SAS, read-biased workloads, archiving; it provides additional RAID protection to cover longer rebuild times of large drives.

Calculating disk IOPS by RAID type

Block 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 application read I/O = 1 back-end read I/O

For random writes:

- RAID 1/0 - 1 application write I/O = 2 back-end write I/O
- RAID 5 - 1 application write I/O = 4 back-end disk I/O (2 read I/O + 2 write I/O)
- RAID 6 - 1 application write I/O = 6 back-end write I/O (3 read I/O + 3 write I/O)

For sequential writes (doing full-stripe writes):

- RAID 1/0 - 1 application write I/O = 2 back-end write I/O
- RAID 5 - 1 application write I/O = $1 + 1/N$ back-end disk I/O (where N is number of drives in the RAID Group)
- RAID 6 - 1 application write I/O = $1 + 2/N$ back-end disk I/O (where N is number of drives in the RAID Group)

Traditional RAID group considerations

Drive count

RAID Groups can have a maximum count of 16 drives. For parity RAID, the higher the count, the higher the capacity utilization, but with higher risk to availability.

Sequential optimization occurs when the array can perform full-stripe operations normally associated with large-block I/O. Certain drive counts are more likely to enable this behavior, though these optimizations can also occur with non-preferred drive counts.

With a predominance of large-block sequential operations the following applies:

- RAID 5 has a preference of 4+1 or 8+1.
- RAID 6 has a preference of 8+2.
- RAID 1/0 has a preference of 4+4.

Storage pools have different recommended preferred drive counts as described in the [Storage Tiers](#) section.

For predominantly random workloads, use the drive type rule of thumb values and the RAID level IOPS calculations to determine the number of drives needed to meet the expected application workload.

Drive location selection

When selecting the drives to use in a RAID Group, drives can be selected either from a DAE or DAEs on the same bus (horizontal positioning), or from multiple DAEs on different buses (vertical positioning). There are now almost no performance or availability advantages to using vertical positioning; however there can be scenarios where vertical positioning causes availability concerns. Therefore:

- Use the default horizontal positioning method of drive selection when creating RAID Groups.
- If a single DAE does not have enough available drives to create the RAID Group, selecting drives from a second DAE is acceptable.

The most likely scenario where vertical positioning might still be used is when FAST Cache drives are spread across multiple buses. In this scenario, it is very important to avoid creating mirrored pairs where one drive is in enclosure 0_0 and its mirror is in a different enclosure. Mirrored pairs can be specified as follows:

- When using Unisphere®, select the drives one at a time, while holding the Control key, to ensure the order in which the RAID 1/mirrored-pairs are created.
 - For example; Selection 1 and 2 form a RAID Group, Selection 3 and 4 form a RAID Group, and so on.
- When using the CLI, specify the drives in the order in which the mirrored pairs should be created.
 - For example: The 1st and 2nd drives in the command syntax form a RAID Group, the 3rd and 4th drives in the command syntax form a RAID Group, and so on.

(Drive location selection does not apply to storage pools and therefore is not a consideration.)

Large element size

When creating a 4+1 RAID 5 RAID Group, there is an option to select a 1024 block (512KB) element size.

- Use large element size when the predominant workload is large-block random read activity (such as data warehousing).
- AVOID using large element size when the workload is predominantly write-oriented.

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 or RAID Groups.
- Dedicate resources, when you have specific performance goals
- Minimize failure domains
 - Although unlikely, loss of a private RAID group in the pool compromises the total capacity of that pool. It can therefore be desirable to create multiple smaller pools than to use the total capacity available in a single pool
 - As a rule of thumb, consider limiting the pool size to the maximum number of drives that can initially be placed in a pool at creation for a given model.

Model	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Maximum drives in a single pool	71	121	246	496	996
Maximum drives in a pool at creation	20	40	80	120	180

Tiered storage pools have multiple RAID options per tier for preferred type and drive count.

- Use RAID 5 with a preferred count of 4+1 for the best performance versus capacity balance.
 - Using 8+1 improves capacity utilization at the expense of slightly lower availability.
- Use RAID 6 for NL-SAS tier.
 - Options for preferred drive count are 6+2 and 14+2 where 14+2 provides the highest capacity utilization option for a pool, at the expense of slightly lower availability.
- 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 14+2 RAID 6
- AVOID a drive count that results in a small remainder when divided by the preferred drive count specified.

- It is best to maintain a multiple of the preferred drive count for each tier you select (such as 5, 10, 15, and so on, when using RAID 5 4+1).

Virtual LUN type

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

- Thin LUNs are recommended when storage efficiency requirements outweigh performance requirements.
 - Thin LUNs maximize ease-of-use and capacity utilization at some cost to maximum performance.
 - This does not always mean you will get slower response time or fewer IOPS, but the potential of the drives to deliver IOPS to the host is less.
 - When using thin LUNs, adding a small Flash tier to the pool can improve performance.
 - Thin LUN metadata is promoted to the Flash tier.
 - DON'T use thin LUNs with VNX for File; use only thick LUNs when using pool LUNs with File.
- Thin LUNs are recommended when planning to implement VNX Snapshots and/or block compression.
- Thick LUNs are recommended for the highest level of pool-based performance.
 - A thick LUN's performance is comparable to the performance of a traditional LUN and is better than the performance of a thin LUN.

Virtual LUN ownership considerations

When creating LUNs in a storage pool:

- Create at least two LUNs per Storage Pool, to allow ownership by both SPA and SPB.
- Try to evenly balance capacity of the LUNs in a storage pool across SPA and SPB.
- DON'T change ownership of a pool LUN after initial creation.
 - Use LUN migration to move LUNs to their peer SP if required to change ownership.
- Reference the [Considerations for VNX for File](#) section in this chapter if configuring virtual LUNs for File.

Storage tiers

The number of tiers required in a storage pool is influenced by the following:

- Performance requirements
- Capacity requirements
- Knowledge of the I/O

The capacity required for each tier depends on expectations for locality of your active data. As a starting point you can 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.

If you know your active capacity (skew), then you can be more systematic in designing the pool and apportion capacity per tier accordingly.

In summary, follow these guidelines:

- When FAST Cache is available, use a 2-tier pool comprised of SAS and NL-SAS and enable FAST Cache as a cost-effective way of realizing Flash performance without the expense of Flash dedicated to this pool.
 - You can add a Flash tier later if FAST Cache is not capturing your 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 you do not know your skew.
 - You can always expand a tier after initial deployment to effect a change in the capacity distribution.
- 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 an Flash+NL-SAS 2-tier pool if you are unsure of your skew.
 - The SAS tier provides a buffer for active data not captured in the Flash tier, where you can still achieve modest performance, as well as quicker promotion to Flash when relocations occur.
- Add a small Flash tier to a pool with thin LUNs so that metadata is promoted to Flash and overall performance is improved.

For each pool LUN:

- Consider dedicated pools or traditional RAID Groups where you have predominantly sequential activity for specific LUNs.
- DON'T use auto-tier for low-skew random workloads where the active dataset would not fit in the highest tier.
 - This might cause excessive tier relocations that may not benefit the active data.

- AVOID using highest available when the LUN capacity exceeds the highest tier.
 - This can affect the overall efficiency of the highest tier to service active data for LUNs running in auto-tier mode.

Considerations for VNX for File

When using Block storage pools with File, use the following recommendations:

- Create a separate pool for File LUNs.
 - AVOID mixing with Block workloads.
- Pre-provision space in the storage pool
 - Pro-actively create LUNs and assign them to File, so that File has available space for file-system creation and extension, checkpoints, and so on.
- Use only thick pool LUNs with File.
 - DON'T use thin LUNs with File.
 - DON'T use compressed LUNs with File.
 - If Virtual Provisioning™ is required for VNX for File, use a thin-enabled file system on traditional or thick LUNs.
- Apply the same tiering policies to all LUNs in the storage pool.
- Allow LUNs to complete the prepare process (thick LUN slice allocation) before adding to the File storage group. Use this command to display the status of the prepare:
 - **naviseccli lun –list -opDetails**

When creating LUNs:

- Create approximately 1 LUN for every 4 drives in the storage pool.
- Create LUNs in even multiples of 10.
 - Number of LUNs = (number of drives in pool divided by 4), rounded up to nearest multiple of 10.
- Make all LUNs the same size.
- Balance LUN ownership across SPA and SPB.

Chapter 3 Data Services

This chapter presents the following topics:

FAST-VP	22
FAST Cache	22

FAST VP

Pool capacity utilization

- Maintain some unallocated capacity within the pool to help with relocation schedules when using FAST-VP.
 - Relocation will reclaim 10 percent free per tier. This space will be used to optimize relocation operations but also helps when new LUNs are being created that want to use the higher tiers.
- This is not a mandatory requirement and does not result in lost capacity.

Relocation

- Schedule relocations for off-hours, so that the primary workload does not contend with relocation activity.
- Enable FAST VP on a pool, even if the pool has only 1 tier, to provide ongoing load balancing of LUNs across available drives.

Considerations for VNX for File

By default, a VNX for File system-defined storage pool is created for every VNX for Block storage pool that contains LUNs available to File. (This is a “mapped storage pool.”)

- All LUNs in a given File storage pool should have the same FAST VP tiering policy.
- Create a user-defined storage pool to separate File LUNs from the same Block storage pool that have different tiering policies.

FAST Cache

FAST Cache is best for small random I/O where data has skew. The higher the locality, the better the FAST Cache benefits.

General considerations

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

- Match the FAST Cache size to the size of active data set.
 - For existing EMC VNX/CX4 customers, EMC Pre-Sales have tools that can help determine active data set size.
- If active dataset size is unknown, size FAST Cache to be 5 percent of your capacity, or make up any shortfall with a Flash tier within storage pools.
- A small FAST Cache can satisfy a high IOPs requirement, but you should consider the ratio of FAST Cache drives to working drives. Large storage pool configurations will distribute I/O across all pool resources. A large pool of HDDs might be able to provide better performance than a few drives of FAST Cache.

Preferred application workloads for 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 FAST Cache for LUNs that are not expected to benefit, such as when:

- The primary workload is sequential.
- The primary workload is large-block I/O.

AVOID enabling FAST Cache for LUNs where the workload is small-block sequential, including:

- Database logs
- Circular logs
- VNX for File SavVol (checkpoint storage)

Enabling FAST Cache on a running system

When adding FAST Cache to a running system, it is recommended to enable FAST Cache on a few LUNs at a time, and then wait until those LUNs have equalized in FAST Cache before adding more LUNs.

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. Systems should be sized 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:

- Less than 60 percent SP CPU utilization – enable groups of LUNs or one pool at a time; let them equalize in the cache, and ensure that SP CPU utilization is still acceptable before turning on FAST Cache for more LUNs/pools.
- 60-80 percent SP CPU utilization – scale in carefully; enable FAST Cache on one or two LUNs at a time, and verify that SP CPU utilization does not go above 80 percent.
- CPU greater than 80 percent – DON'T activate FAST Cache.

AVOID enabling FAST Cache for a group of LUNs where the aggregate LUN capacity exceeds 20 times the total FAST Cache capacity.

- Enable FAST Cache on a subset of the LUNs first and allow them to equalize before adding the other LUNs.

Note: For storage pools, FAST Cache is a pool-wide feature so you have to enable/disable at the pool level (for all LUNs in the pool).

Considerations for VNX for File

- AVOID enabling FAST Cache for LUNs containing SavVol (checkpoint storage).
 - SavVol activity tends to be small block and sequential, characteristics that do not benefit from FAST Cache.
- When enabling FAST Cache on a Block storage pool with File LUNs, use a separate storage pool or traditional RAID Groups for SavVol storage.

Chapter 4 VNX for Block Specific Considerations

This chapter presents the following topics:

Availability and connectivity	26
VNX for Block features	27

Availability and connectivity

The VNX for Block storage array offers connectivity to a variety of client operating systems, and via multiple protocols, such as FC, FCoE, and iSCSI. EMC provides connectivity guides with detailed instructions for connecting and provisioning storage via different protocols to the specific host types.

It is highly recommended that you consult the EMC Host Connectivity Guide for the host type or types that will be connected to the VNX for Block array for any specific configuration options.

Storage group settings

When registering host HBAs with VNX for Block, make sure to set the appropriate failover mode based on the host type. See the Host Connectivity Guides for details.

Host file system alignment

File-system alignment is covered in detail in the Host Connectivity Guides. In general:

- Windows Server 2008 automatically aligns, as do more recent Linux environments.
- VNX for File automatically aligns when it creates file systems.
- 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.

Alignment practices:

- Use host-based methods to align the file system.
 - DON'T use VNX for Block-based alignment
- EMC recommends aligning the file system with a 1 MB offset.

VNX for Block features

Block replication

For MirrorView™:

- AVOID enabling FAST Cache on MirrorView secondary LUNs.
 - MV/S secondary LUNs replicate only writes from the source and are serviced well by SP cache.
 - MV/A secondary LUNs replicate writes during updates but also incur copy-on-first-write activity. This may incur additional FAST cache promotions that would not lead to performance gain.
- DON'T enable FAST Cache on Write Intent Log (WIL) LUNs since they are already optimized for priority in the storage system's write cache.
 - The WIL exhibits small-block sequential activity not suited for FAST Cache.

For SnapView™:

- Use SAS drives for reserve LUN pool (RLP) configurations, with write cache-enabled LUNs.
 - Match the secondary side RLP configuration to the primary side.
- AVOID configuring RLP on the same drives as the primary and secondary LUNs to avoid drive contention.
- DON'T enable FAST Cache on RLP LUNs.
 - RLP exhibits multiple small-block sequential streams that are not suited for FAST Cache.
- DON'T enable FAST Cache on clone private LUNs (CPL), since they are already optimized for priority in the storage system's write cache.

For RecoverPoint:

- DON'T enable FAST Cache for RecoverPoint journal LUNs.
 - Journals exhibit primarily large-block sequential activity not suited for FAST Cache use.

For VNX snapshots:

- Start with thin LUNs if planning to use VNX snapshots.
 - Traditional or thick LUNs are converted to thin LUNs once a VNX snapshot is created on them.

Block compression

Start with thin LUNs if planning to use Block compression.

- Traditional or thick LUNs are converted to thin LUNs when they are compressed.

Chapter 5 VNX for File Specific Considerations

This chapter presents the following topics:

File-system creation	30
Traditional RAID group LUNs.....	30
Storage pool LUNs	30
VNX for File network and protocol considerations	31
VNX for File features.....	31
VNX for File application tuning	32

File-system creation

Guidelines for creating file systems:

- Balance backup and restore window requirements with capacity requirements when deciding on file-system sizes.
- Use Automatic Volume Management (AVM) to create file systems.
 - AVM creates well-designed file-system layouts for most situations.
- Manual Volume Management can be used to create file systems with specific attributes for specialized workloads.
 - For metadata-heavy workloads, stripe across an odd number of disk volumes (dVols).
 - For sequential access on striped volumes:
 - AVOID wide-striping across a large number of dVols.
 - Match File stripe size to the Block LUN full-stripe width for best sequential write performance.
- AVOID using file-system auto extension.

Traditional RAID group LUNs

When creating LUNs for VNX for File from RAID Groups, consider the following:

- Preferred RAID Group sizes:
 - RAID 1/0: 1+1
 - RAID 5: 4+1, 8+1
 - RAID 6: 8+2, 4+2
- When creating volumes manually:
 - Stripe across 5 LUNs from different RAID Groups.
 - Use a stripe size of 262144 (256KB).
 - Balance SP ownership.

Storage pool LUNs

When creating LUNs for VNX for File from Block storage pools, consider the following:

- When creating volumes manually:
 - Stripe across 5 LUNs from the same pool.
 - Use a stripe size of 262144 (256KB).
 - Balance SP ownership as well as possible.

VNX for File network and protocol considerations

Interfaces

10GbE:

- Best for bandwidth-sensitive applications
- Configure the first NIC from different SLICs before moving to the second NIC on the same SLIC.
- Configure multiple IP addresses (virtual interfaces) per NIC.
- Use Jumbo frames.

1GbE:

- Use Jumbo frames when possible

Trunking and multipathing:

- Use LACP instead of EtherChannel.
- Configure Failsafe Networking (FSN) to avoid a network single point of failure.

NFS

For bandwidth-sensitive applications:

- Increase the value of param `nfs v3xfersize` to `131072` (128K).
- Negotiate larger NFS transfer sizes with `mount -o rsize=X,wsiz=X`.
- If mounting multiple exports, use a different Data Mover IP address for each mount.

VNX for File features

SnapSure

If using SnapSure™ to create user checkpoints of the primary file system:

- Size disk layout for PFS to include copy-on-first-write activity.
- Size layout for SavVol based on expected user load to the checkpoint file systems.
- Place SavVol on separate disks when possible.
- DON'T disable SmartSnap traversal.
- AVOID enabling FAST Cache on SavVol.

IP Replicator

- Size disk layout for the primary file system to include copy-on-first-write and replication transfer activity.
- Place the secondary file system on the same drive type as the primary file system.

- It is okay to place SavVol on NL-SAS for replication.
 - If user checkpoints are also enabled on the primary file system, then consider user load from the checkpoints to determine whether NL-SAS will still be adequate for SavVol.
- Use 1GbE links for Replicator interconnects when traversing WAN links.

Deduplication and compression

If using VNX for File deduplication and compression:

- Target deep compression at inactive files.
- DON'T enable CIFS compression on busy Data Movers.
 - CIFS compression occurs in-line with the host write and can delay response time if the compression activity is throttled.

CAVA

- Always consult with the antivirus vendor for their best practices.
- CIFS should be completely configured, tested, and working before setting up virus checker.
- Antivirus servers should be strictly dedicated for CAVA use only.
- Ensure that the number CIFS threads used are greater than virus checker threads.
- Exclude real-time network scanning of compressed and archive files.
- Set file mask to include only file types recommended for scanning by the antivirus provider.
 - DON'T include *.*.
- Disable virus scanning during migrations.

VNX for File application tuning

VMware/database over NFS

- Mount file systems on the Data Mover using Direct Writes Enabled (uncached mount option).
- If application cannot or will not send direct I/O, enable the param **file asyncUncachedOpt**.

Conclusion

This best practices guide provides configuration and usage recommendations for VNX unified 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 fundamentals guide, or any of the papers in the related documents section.