



# CONFIGURATION AND IMPLEMENTATION OF EMC VNX STORAGE



Octavio Guerra M.  
Solution Architect  
EMC - NOLA  
[octavio.guerra@emc.com](mailto:octavio.guerra@emc.com)

## Table of Contents

Introduction .....	3
Where to Start? .....	4
Customer Feedback.....	5
Looking for Descriptions, Features (old and new), and Further Improvements in VNX Storage..	6
Distribution of Physical Disks in VNX .....	7
What RAID type to use?.....	8
Which is better? Pool / RAID Group in VNX? .....	9
Considerations to create LUNs .....	11
Considerations to Create Reserved LUN Pool .....	11
Considerations for Assigning Ports per Host or Server.....	12
References .....	13
Author's Biography.....	13

Disclaimer: The views, processes, or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.

## Introduction

Frequently, customers who purchase an EMC VNX® storage system have many questions about the best configuration or implementation. Some questions are:

- Is it best to use Pools or RAID Groups?
- Would Thin Provisioning be adequate for my environment?
- What RAID to use: Raid 10, Raid 5, Raid 10?
- Divide or mix the I/O for Block and File?
- How many ports are recommended by host?

This article provides Implementers and Solution Architects with rules, ideas, and recommendations to arrive at the “adequate” configuration in a new VNX, focusing on three things:

1. **Distribution.** Recommendations on how distribute the disks in the VNX storage, based on bus quantity, DAEs, etc. are also a good point to have in count before to configure a new VNX storage.
2. **Performance.** Identifying applications that will reside in the VNX and clues about the I/O are good starting points to best define how to configure Pools / RAID Groups, type of RAID, and even distribution of ports for the hosts.
3. **Availability.** Considerations about how data could be affected in an emergency over the VNX. Critical Data vs. RAID type vs. Quantity of Disks vs. Pool are good points to check out before configuring the VNX.

Additional considerations include space reserved when MirrorView™ or SnapView™ operations are involved in the configuration and quantity of ports per host.

While configuring and implementing VNX storage is a great challenge for implementers and Solution Architects, “adequate” configuration based on recommendations and practical ideas coming from field experience are certainly a great help in getting the best results and the best use of the VNX.



It's recommended to know all security configurations of the VNX. Doing so will remove doubts or assumptions that could generate confusions or misunderstanding when beginning the configuration and implementation phase.

## **Customer Feedback**

Tip: Very Important - Don't Forget.

Remember, at the end, the customer will be directly involved with using the VNX. Take all the time (if the project allows) you need to talk with the customer about these questions:

- What is the use for the VNX storage?
- Is to replace old EMC or third-party storage? If yes, why are you replacing the old storage? Are measures available of the I/O of the actual storage?
- What data could be in the VNX? (Databases, files, Application Data, etc).
- How many servers between Production, Test, and Development environments could be connected to the VNX?
- Will there be replication?
- Are measures available about server / application performance?
- Is the purpose of the new VNX to replace any other storage?

The answers to these questions are a starting point to define the VNX configuration according to the hardware acquired. The answers can also signal an alarm that something could be missing based on customer requirements. Without a doubt, it's one thousand times better to make sure that the "best fit" is determined for the customer environment—or if something additional is needed—before VNX configuration and implementation.

## **Looking for Descriptions, Features (old and new), and Further Improvements in VNX Storage**

Between Flare versions, sometimes new things are added or improvements are done to actual functionalities in the storage that could prove useful in the customer's environment. To stay current, your first stop should be to look in internal sites such as:

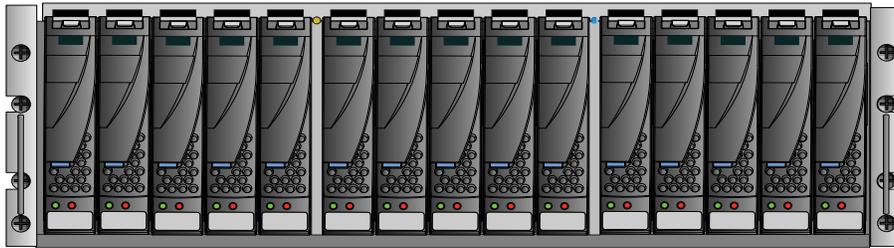
<http://powerlink.emc.com> or <http://support.emc.com>

Also, it's a good idea talk with local co-workers or those from other sites because, again, they could know new things or ideas about the utilities and functionalities in this kind of storage that could be useful. Additionally, if you're not familiar with the VNX model acquired by the customer, read about the characteristics, limitations, restrictions, and functionalities of the Model. It's better to know what you'll be implementing.

Having information about the hardware and model of VNX, customer answers about its use and the environment where the storage will be, and documentation about functionalities, improvements, characteristics, and limitations is the best starting point to adequately configure and implement the VNX.

## Distribution of Physical Disks in VNX

Distribution of physical disks in the VNX helps to get the best storage performance and efficiency. The VNX typically arrives at the site with the disks distributed in the DAEs in a “serialized” fashion. First, all Flash disks, then all SAS disks and last, all NL-SAS disks. If in the process a DAE is completely filled, the next DAE available is used and so on.



VNX DAE with Disks

Put the different disk models (FAST, SAS, NL-SAS) in all the buses available. For example, if the VNX has two buses, put Flash disks in both buses, as well as SAS disks and NL-SAS disks.

Example:

Customer has 4 DAEs additional to DPE or Enclosure zero (0) and two buses. With the next quantity of disks:

- 10 SAS Flash disks
- 20 SAS disks
- 16 NL-SAS disks

One suggestion to distribute the disks is:

- Five Flash disks for each bus. Two Flash disks in one DAE and three disks in the other one in the same bus.
- Ten SAS disks for each bus. Five SAS disks in each DAE in the same bus.
- Eight NL-SAS disks for each bus. Four NL-SAS in each DAE in the same bus.

**Recommended:** Try to distribute the disks keeping in mind the type of RAID that will be used by model of disk. For example: if you decide use RAID 5 with SAS disks, try to distribute at least 5 disks in each DAE or a multiple of 5.

- Put (faster) Flash disks in the first DAE slots then, in the next slots, the (medium) SAS disks and, in the last slots, the (slower) NL-SAS disks.
- Avoid using the DPE or Enclosure 0 for disks which will be used in Pools or RAID Groups. Leave the DPE for disks that make up the VAULT RAID Group. The only exception is FAST Cache when disk quantity is less than or equal to 8 disks. In that case, use DPE or Enclosure 0.
- Distribute the Spares disks in all the buses available, reserving the last slots in DAE for that. The EMC Unified Best Practices document recommends defining a Spare for every 30 disks, ensuring that there is at least one Spare for each disk model. Note: Never a SAS disk as a spare of NAS disks nor Flash disks of SAS disks.

## What RAID type to use?

The types of RAID supported in VNX are:

RAID 0: Only Stripe with no parity or protection. Offers good performance, but with high risk of data loss if one disk fails.

RAID 1: Is a mirror between only two disks. It means one of the disks is a mirror of the other. Good performance, but is limited to only two disks.

RAID 3: RAID with striping and a dedicated parity disk. Support fail only of one disk. Good performance but is not very widely used. The minimal disks in RAID 3 are three (3).

RAID 10: It is Striping and Mirroring. Good performance to writes. Recommended for high I/O writes. Also, offers good protection because only 50% of the disks in the RAID could be lost. The not so good part is that 50% of RAID space is lost due to mirroring. The minimal disks recommended in RAID 10 are four (4).

RAID 5: It is Striping with One Parity to protect in case of loss of one disk. Suitable performance (not like RAID 10) for writes and works fine with many different I/O profiles. Only one disk can be lost in that kind of RAID. When one disk in the RAID fails, the information is rebuilding from the parity in the other disks to spare disk. Usable space is already 75% of the raw space. The minimal disks in RAID 5 are three (3).

RAID 6: It is striping with double Parity to protect in case of disk failure. Highly recommended with NL-SAS because they are large and slow disks. Occasionally, in the rebuilding process after a fail, a new disk could fail; but the double parity can support the loss of maximum two (2) disks. Usable space is already 60% of the raw space. The minimal disk in RAID 6 are four (4).

The recommended RAID in almost all VNX configurations is RAID 5 for Flash and SATA disks. NL-SAS is really suitable for RAID 6. However, this is where the customer's answers about the I/O profile and data resident in the VNX become important. In some cases, RAID 10 can be used when it's known that high performance in writes is required. If in doubt, use RAID 5 (adequate for Performance vs. Usable Space). In extreme cases, use RAID 5 for NL-SAS (for example, the quantity of disks doesn't permit RAID 6).

## **Which is better? Pool / RAID Group in VNX?**

It's a good question. The answer: It depends.

Pool is the option when:

- FAST VP will be implemented, because this functionality is only supported in Pools.
- Thin Provisioning will be used, because this functionality is only supported in Pools.
- LUN Expanding is required and the customer likes the idea to expand the same LUN and not create more LUNs or add Metaluns. The expansion of the same LUN is only supported in Pools, not RAID Groups.

RAID Groups is the option when:

- Creating LUNs to be used in FILE (NAS part of VNX). It's not a "straitjacket". However, based on field experience, it is adequate to use when NAS is there.
- FAST VP or Thin Provisioning will not be used in the VNX (Definitely the most powerful reason).
- Excellent performance is the goal.
- LUN Expanding based on adding more LUNs to the Server or creating Metaluns is not a problem for the customer.

In some scenarios, the Mix (Pools and RAID Groups) is a good idea. For example:

- Pools to use with I/O profiles where FAST VP is required.
- Pools to use Thin Provisioning.
- RAID Groups to use in NAS or for environments where Dedicated I/O profile is required.

- RAID Groups where space is required for Reserved LUN Pool due to Snapshot use or MirrorView Async implementations.

Keep in mind that a RAID Group can support a maximum 16 disks. For the Pools, see Table 2.

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

Table 2

- If possible, create a minimum of two (2) Pools in the VNX. Though creating just one Pool would be sub-optimal, if the Quantity disks / DAEs / RAID don't permit any other option, create just one Pool.
- The mix of disk models in Pools with FAST VP recommended are:
  - Flash Disks + SAS Disks
  - Flash Disks + SAS Disks + NL-SAS Disks
  - SAS Disks + NL-SAS Disks

The mix Flash Disks + NL-SAS Disks is not recommended.

- The RAID configurations suggested in a Pool are:

For Flash and SAS Drives:

RAID 10 - Quantity of Disks: 4 (2+2)

RAID 5 - Quantity of Disks: 5 (4+1) or 9 (8+1)

For NL-SAS Disks:

RAID 6 - Quantity of Disks: 8 (6+2) or 16 (14+2)

## Considerations to create LUNs

- Just create Thin LUNs in Pools when performance is not the priority.
- In cases where you don't want to use Dedicated Pools or RAID Groups for customer environments, create LUNs distributed over all Pools and RAID Groups in the VNX. The goal is get a balanced load in the storage.
- Use at least two RAID Groups to create the LUNs for File (NAS) in the VNX.
- Keep an equal number of LUNs in each SP.
- When a LUN is created in a Pool and is subdivided in Sub-LUNs of 1 GB to enable FAST VP functionality, do movements between Tiers.
- Only create LUNs in the Vault RAID Group to be used, such as repository, swap areas, or OS volumes.

## Considerations to Create Reserved LUN Pool

The Reserved LUN Pool is an area used by the VNX to create "Snapshots" of LUNs or to be used like a "buffer" when Replication using MirrorView Asynchronous is used.

Suggested considerations when this area needs to be created:

- If possible, if there is availability, create a simple Pool or RAID Group to be used in creating the LUN for the Reserved LUN Pool.
- When there is no dedicated Pool or RAID Group for the LUNs used in the Reserved LUN Pool, create and distribute them in all the Pools or RAID Groups available.
- Always use the same disk model for LUNs of Reserved LUN Pool in both VNX (Source or Target).
- Create double the number of LUNs for the Reserved LUN Pool based in the quantity of LUNs to be replicated. For example: If 10 LUNs will be replicated, then create 20 LUNs for the Reserved LUN Pool.
- The size of the Reserved LUN Pool, in general, could be estimated at 20% of the total size of data replicated. Example: if 10 LUNs will be replicated and total data size is already 800 GB, the size of the RLP will be 160 GB divided into 20 LUNs.

## Considerations for Assigning Ports per Host or Server

The following recommendations are the result of field experience regarding the quantity of ports in VNX Storage vs. Servers:

- Never assign just one port to any server. Always assign at least two Ports (one for each SP) from the VNX to the Server.
- Use Native Multipath or EMC PowerPath® from the server to optimize efficient use of all paths toward the VNX.
- Distribute the ports for servers, keeping in mind to not use the same ports for servers with high I/O.
- Once assigned, monitor the ports for usage. As a rule, check that the usage or load doesn't exceed 50% in the port. If the load is more than 50%, re-distribute ports for the servers.

## References

- EMC VNX Unified Best Practices for Performance Applied Bests Practices Guide
- White Paper - MirrorView Knowledge Book
- VNX Virtual Provisioning
- Managing Statistics for VNX

## Author's Biography

Octavio Guerra is a Solutions Architect in the EMC Global Products Practice, supporting the EMC Storage product portfolio, Avamar® and Data Domain® products (parts of BRS portfolio), EMC Connectrix® and SAN Portfolio as well. He has been working in the IT Industry for 14 years and at EMC for 3 years. He has experience with Unix, Linux, Windows, Fibre Channel, TCP/IP and SCSI environments. Octavio has a degree in Computer Engineering and is a Certified EMC Proven Professional Symmetrix® Expert with experience in field implementations with EMC CLARiiON®/VNX, Symmetrix, and VMAX®. He is located in Bogota, Colombia and works on projects covering EMC – NOLA (North – Latin American) Region.

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.