



# BEST OF V's AND TROUBLESHOOTING METHODOLOGIES



Anuj Sharma

EMC<sup>2</sup>

## Table of Contents

Introduction .....	3
Storage Design .....	4
VNX Best Practices.....	5
FAST Best Practices.....	7
Enhancing VNX File SNMP Alerting Capabilities in ECC or SNMP Monitoring Tool.....	8
VMAX Best Practices .....	9
FAST VP Best Practices .....	10
DSE Best Practices.....	12
SRDF/A Best Practices .....	13
VPLEX Best Practices.....	16
Performance Troubleshooting Methodologies .....	18
References .....	19
Appendix A .....	19

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**

Cloud is making a significant impact on Enterprises primarily because of scalability, flexibility, and cost benefits. Enterprises are looking for public, private, and hybrid cloud options as per their requirements. Storage is a vital component for a solid cloud foundation whether it's a public or private cloud. The storage selected for their cloud infrastructure that enterprises and service providers choose for their cloud foundation should be agile, flexible, and scalable enough to complement the cloud features. EMC has taken a lead in complementing storage systems with new features that equip storage systems with more scalability, agility, and flexibility. However, along with these features comes more complexity in the environment from the administration point view as virtualization is introduced between the storage and users. Consequently, when performance or technical issues arise there are more places for administrators to look when troubleshooting.

This Knowledge Sharing article will explain the best practices that should be followed while implementing VPLEX, VNX, and VMAX along with the operational best practices to get the best out of the V's. We will also look at the troubleshooting methodologies that will explain that how Storage Admins should look at the performance issue and the different perspectives from which the performance issues can be addressed.

# Storage Design

## Build a Strong Foundation

- Keep in mind that each storage has its own limitation; i.e. Front End Ports, Back End Ports, Cache, and Disk Drives.
- Configure storage with enough components to sustain the workload.
- Bear in mind Little's Law that response time shoots up suddenly after the utilization of equipment reached 70-80%.
- As a rule of thumb, we should design the storage components in a way that they are on average almost 60% utilized.

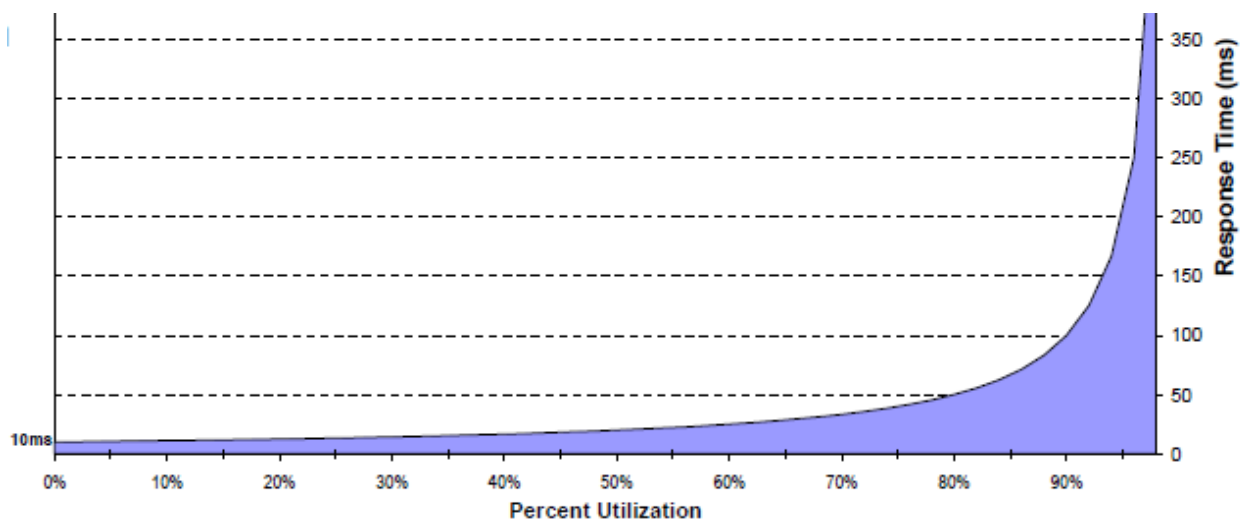


Figure 1: Response Time

Drive specs that can be used while deciding upon the number of disks are shown below.

Drive speed	IOPS (8 KB random mix)	Bandwidth (256 KB sequential mix, RAID 3)
5,400 rpm ATA	50	7 MB/s
7,200 rpm ATA	60	7 MB/s
7,200 rpm SATA	80	8 MB/s
7,200 rpm LCFC	80	8 MB/s
10,000 rpm FC or SAS drives	150	10 MB/s
15,000 rpm FC	180	12 MB/s
EFD, solid state drive	2500	100 MB/s

- Reads Performance across different RAID Protection groups is similar provided the number of disks is the same. Write performance is the one that differs because of read penalty associated with every RAID Protection.

#### **RAID 1: 1 Host Write = 2 Writes**

#### **RAID 5: 1 Host Write = 2 Reads + 2 Writes**

#### **RAID 6: 1 Host Write = 3 Reads + 3 Writes**

- Sequential writes for RAID 5 or RAID 6 can be more efficient than RAID 1
- Meta Volumes should be used for LUNs larger than 256 GB for better performance.

The pointers above are generic storage design pointers that we should consider while deciding upon the storage design. In the following sections we will discuss specific VNX, VPLEX, and VMAX best practices.

### **VNX Best Practices**

- Always ensure that File System alignment is done to avoid disk crossings. We have seen numerous instances where lack of disk alignment leads to serious performance concerns.
- ALUA is the default failover mode for VNX Systems but make sure that the PowerPath® version installed on the host is ALUA aware. PowerPath Versions 5.1 and later are ALUA aware. Refer to the host compatibility matrix to determine if ALUA is supported. EMC primus emc99467 provides more information about the Failover Modes for different operating systems.

- For iSCSI implementations, always use Jumbo Frames and a dedicated Network/VLAN.
- Host should always be zoned to both the SP's of the VNX.
- DAE's should be balanced evenly across the Back-End buses instead of overloading the single bus.
- RAID 5 is the recommended RAID type for SAS disks and RAID 6 is the recommended RAID Type for NL-SAS Disk.
- Avoid placing LUNs on vault drives. It's not recommended to bind LUNs to Vault Drives.
- As per Best Practices, there should be a hot spare configured for every 30 drives.
- It is strongly recommended that all the drives in a RAID group be of the same form factor, type, speed, and capacity.
- As a rule of thumb, storage pools should contain at least four private RAID Groups, meaning that the drives should be added in multiples of four for default drive combinations of a RAID Type. For example, a RAID 5 pool is made up of default 5-drive (4+1) private RAID Groups. Per best practices, the smallest pool would consist of 20 drives (four private RAID groups, each RAID 5 4+1). For RAID 6, the default is 8-drive, i.e. (6+2). For RAID 1/0, initial allocations should be at least 8 drives.
- Pool expansion should also be done in multiples of the default disk drive combinations; for example, in multiples of 5 for a RAID 5 group.
- Striped MetaLUNs are preferred for better performance. However, in cases where storage is urgently needed, concatenation with the same LUN size as the original LUN size can be used. There after multiples of original LUN size.
- Do use the correct RAID level for the pattern of I/O (e.g. An application generating 8KB Random Writes should ideally be using RAID 1/0 LUNs)
- Do not have more than one component LUN from any particular RAID group within the same metaLUN stripe.
- Do use drives of the same capacity and rotational speed in all the RAID groups in the metaLUN. The RAID groups should also contain the same number of drives.
- Do use sets of RAID groups that only contain metaLUNs that are all striped in the same way, if possible. Having standard LUNs in the same RAID group as metaLUN components will lead to some parts of a metaLUN having uneven response times across the metaLUN. The order in which the component LUNs are added can be changed to evenly distribute the file system load (for example, RAID Group 1, 2, 3, and 4; then RAID Group 4, 1, 2, and 3, etc.).

- Do not concatenate stripes with large numbers of drives to components with much fewer drives. This will lead to performance varying dramatically in different parts of the same metaLUN.
- Do name the component LUNs in such a way as they are easy to identify. Numbering the components in a logical order helps to choose the correct RAID group and default SP owner although the component LUNs will be renumbered when the metaLUN is created. The metaLUN will have its own default owner, but choosing the same default owner as all the components avoids the components being reported as being trespassed in some tools.

## **FAST Best Practices**

- FAST Cache drives can sustain very heavy workloads, but if they are all on the same bus, they could completely saturate this bus with I/O. This would especially be an issue if the drives were all on bus 0 because this is used to access the vault drives. Therefore, spread the FAST Cache drives over as many buses as possible.
- The FAST Cache driver has to track every I/O in order to calculate whether a block needs to be promoted to FAST Cache, adding to SP CPU utilization. Disabling FAST Cache for LUNs unlikely to need it will reduce this overhead and thus can improve overall performance levels. This is especially true for secondary mirror and clone destination LUNs, which would gain no significant performance benefit from using FAST Cache. In the case of Thin or Thick LUNs, the FAST Cache would have to be disabled at a Pool level, so it would be best to have separate Pools for LUNs which do not need FAST Cache. This in turn will result in higher performance gain to the LUNs that really do need FAST cache.
- After a few hours, FAST cache will be using nearly all of the available drive capacity of its flash drives. For every new block that is added into FAST cache, another must be removed. These will be the blocks that are the oldest in terms of the most recent access. If there is not much FAST cache capacity (in comparison to LUN capacity), blocks that are frequently accessed at certain times of day will have been removed by the time they are accessed again the next day. Restricting the use of FAST cache to the LUNs or Pools which need it the most can increase the probability of a frequently accessed block still being in FAST cache the next time it is needed.
- The order that the drives are added into FAST Cache is important since that will dictate which drives are Primary and which are Secondary. The first drive added is the first

Primary, the next drive added is its Secondary, the third drive is the Primary for the second mirror and so on. Do not enable FAST Cache on any reserved / private LUNs, apart from metaLUN components. This includes the Reserved LUN Pool, Clone Private LUNs, and the Write Intent Logs.

- Do think about the type of data on a LUN and consider if FAST Cache is needed. For example, log files are generally written and read sequentially across the whole LUN. Therefore, these would not be good candidates for FAST Cache. Avoiding use of FAST Cache on unsuitable LUNs reduces the overhead of tracking I/O for promotion to FAST Cache.
- Do not enable FAST Cache on MirrorView™/S Secondary mirror LUNs or SnapView™ clones. This will involve disabling FAST Cache for Pools which contain these LUNs.
- Do not put all FAST Cache drives on a single bus.

## Enhancing VNX File SNMP Alerting Capabilities in ECC or SNMP Monitoring Tool

- Create config file trap\_eventlog.cfg in the directory /nas/site. Load config file by running the following command.
  - nas\_event -Load /nas/site/trap\_eventlog.cfg
- Check that file was loaded. You'll see trap\_eventlog.cfg in the output from the following command.
  - nas\_event -Load -info
- Confirm that all ECC NAS Agent servers / Monitoring Servers are listed in /nas/site/trap.cfg file.

- cat /nas/site/trap.cfg

*#snmp trap configuration file*

*#example:*

*#snmpmanager 128.154.11.20 ; communityname public*

*#snmpmanager host1 ; communityname public*

*snmpmanager 10.175.1.145 ; communityname public*

*snmpmanager 10.175.1.146 ; communityname public*



- If not all of the NAS agents servers for the environment are listed in the trap.cfg file, enter the missing NAS agent servers. The entry in trap.cfg file should be in the same format as below.

- snmpmanager 128.154.11.20 ; communityname xyz\_community

- Confirm that the SNMP services are started for the control stations and Data Movers.
- You can enable SNMP on the Control Station and Data Movers through Celerra<sup>®</sup> Manager. Go to Security> Network Services.
- From root, send a trap by using this syntax:

```
/nas/sbin/nas_snmptrap config_file_path -m /nas/sys/emccelerra.mib -r trap_number -f facility_id -i event_id -s severity_level -d "description".
```

- Use the following example for test.

```
/nas/sbin/nas_snmptrap /nas/site/trap.cfg -m /nas/sys/emccelerra.mib -r 1 -f 64 -i 5 -s 7 -d "test SNMP traps"
```

- If the test SNMP trap did not work, make sure that ports 161 and 162 are open from the control station and the NAS agent servers.

**Refer to Appendix A for the contents of trap\_eventlog.cfg file.**

## **VMAX Best Practices**

- Distribute the workload evenly across the front-end directors.
- Use all of the “zero” ports on all of the directors first and then use the “on” ports on the directors because two active ports on one FA CPU do not generally do more IO/s.
- As per best practice, zone the Host to different Directors as it will provide redundancy. It’s not preferred to zone the Host to a different ports on the same director; host should be zoned to different directors .
- As a general rule, fewer and larger hypervisors will give better overall system performance.
- For LUN size larger than 240 GB, a meta LUN is always preferred.
- RAID 5 3+1 is recommended for EFD, RAID 1 for FC, and 4+2 for RAID 6 for SATA.
- 8 EFD are recommended per engine.
- Try to zone hosts to the same FA Ports that have opposite I/O behavior. For example, one host that requires high bandwidth and another that requires high I/O rate.

## FAST VP Best Practices

- In FAST VP, the default performance window should be left unchanged. However, if there are extended periods of time when the workloads managed by FAST VP are not active, these time periods should be excluded.
- FAST VP best practice recommendation is that, at a minimum, the data movement window should allow data movements for the same period of time that the performance time windows allow data collection. This allows FAST VP to react quicker and more dynamically to any changes in workload that occur on the array. Unless there are specific time periods to avoid data movements—during a backup window, for example—it may be appropriate to set the data movement window to allow FAST VP to perform movements 24 hours a day, every day.
- The best practice recommendation for the WorkLoad Analysis Period is to use the default value of 168 hours (1 week).
- The default value for the Fast VP Relocation Rate is 5. However, the best practice recommendation for the initial deployment of FAST VP is to start with a more conservative value for the relocation rate, perhaps 7 or 8. The reason is that when FAST VP is first enabled, the amount of data to be moved is likely to be greater, compared to when FAST VP has been running for some time.
- The Pool Reserved Capacity (PRC) reserves a percentage of each pool included in a VP tier for non-FAST VP activities. When the percentage of unallocated space in a thin pool is equal to the PRC, FAST VP no longer performs data movements into that pool. The PRC can be set both as a system-wide setting and for each individual pool. If the PRC has not been set for a pool, or the PRC for the pool has been set to NONE, the system-wide setting is used. For the system-wide setting, the best practice recommendation is to use the default value of 10 percent.
- Drive size is not a factor when adding data devices to a thin pool. For example, data devices configured on 300GB FC 15k drives can coexist in a pool with data devices configured on 600GB FC 15k drives. However, when planning to have data devices of different drive sizes exist in the same storage tier, it is recommended to create two separate pools for each drive size, and then combine those two pools into a single tier. The best way to determine appropriate policies for a FAST VP implementation is to examine the workload skew for the application data to be managed by FAST VP. The workload skew defines an asymmetry in data usage over time. This means a small

percentage of the data on the array may be servicing the majority of the workload on the array.

One tool that provides insight into this workload skew is Tier Advisor. While performance and availability requirements ultimately determine the configuration of each tier within the Symmetrix array, it is recommended, as a best practice, to choose RAID 1 or RAID 5 protection on EFDs. The faster rebuild times of EFDs provide higher availability for these protection schemes on that tier. Also, it is recommended to use either RAID 1 or RAID 6 on the SATA tier. This is due to the slower rebuild times of the SATA drives (compared to EFD and FC), and the increased chance of a dual-drive failure, leading to data unavailability with RAID 5 protection.

- For the FC tier, RAID 1 is the recommended protection level. Mirrored data devices on FC pools provide a higher level of performance than both RAID 5 and RAID 6, particularly for write workload. Availability of RAID 1, in regard to a dual-drive failure, is also greater than RAID 5. To obtain the best availability numbers for RAID 1 on FC, using lower capacity drives is recommended.
- Pre-allocated, but unwritten, extents show as inactive and as a result, are demoted to the lowest tier included in the associated FAST VP policy. When these extents are eventually written to, the write performance will be that of the tier to which it has been demoted. A best practice recommendation is to not pre-allocate thin devices managed by FAST VP. Pre-allocation should only be used selectively for those devices that can never tolerate a write failure due to a full pool.
- EMC best practices, for both synchronous and asynchronous modes of SRDF<sup>®</sup> operation, recommend implementing a balanced configuration on both the R1 and R2 Symmetrix<sup>®</sup> arrays. The reason is that in case of failover there is no performance lag encounter at the remote site due to the placement of data extent different than the primary Symmetrix.

## DSE Best Practices

- Configure DSE on both R1 and R2 Symmetrix system frames.
- Plan for peak workloads.
- Make sure you collect data during peak workloads. If the peak workload occurs during month-end or quarter-end processing, ensure that the peak workload data is used during planning for DSE configuration.
- DSE pool configuration guidelines:
  - The best practice is to spread the DSE pool devices across as many disks as possible. Do not underestimate the bandwidth and throughput demands that DSE place on the disks. When DSE is active, you could have one random write and one random read for every host write. Consider an example where DSE pool devices with RAID 1 protection are used. When DSE is active, a host write load of 1500 writes/sec will result in 3000 backend writes/second and 1500 reads/second to the DSE pool.
  - The minimum number of disks to spread the pool across depends on the peak write rate. For example, if the peak host write rate is 1500 writes/sec and RAID 1 pool devices are used, the DSE pool needs to be spread across about 100 disks.
  - The SymmMerge user data option can model the DSE pool.
  - RAID 1 is the recommended protection type for DSE pool devices. RAID 5 is another option, but it will consume more backend resources. Currently, RAID 1 is the only supported configuration. A RPQ must be submitted for RAID 5.
  - Do not use dedicated disks for the DSE pool. Use only one hyper from each disk for the DSE pool. The other hypervisors can be used for any other purpose.
- The size of the DSE pool depends on:
  - The average write rate when DSE is active.
  - The length of time DSE is active.
- A simple formula for calculating DSE pool size is: Write Rate\*Length of time when SRDF/A will be paged into DSE pool\*64 KB  
*Example: If the write rate is 3000 writes/second and you want to ride over a 1 hour link outage, the size of the pool is  $3000 * 3600 * 64KB = 659 GB$ .*

*This formula assumes that I/Os are smaller than 64 KB. For I/Os larger than 64 KB, adjust accordingly. In the above example, if each write was 128 KB, the size would be 3000\*3600\*128 KB.*

This size is before protection. If the protection is mirrored, you need to double the disk space. In the above example, you will need 659 GBx2 = 1318 GB of disk space. This assumes that there are no rewrites, but it is better to be safe than sorry in this case. Keep the size of the DSE pool reasonable. Typically 2 to 5 times cache size is sufficient for a short outage.

- Ensure sufficient DA and RA CPU resources are available for the DSE task. The DSE task runs on all RA and DA CPUs. Make sure the DA and RA are not more than 50% busy without DSE being active. If they are busier, they may not be able to handle the additional load when DSE is active.
- Use as small a number of DSE pools as possible since this makes the most efficient use of resources and is easiest to manage. The number of RDF groups that can share a DSE pool is limited only by the number of RDF groups that can be created in the system. DSE imposes no limit on the number of groups that can share a pool.
- If multiple groups are assigned to a DSE pool and the need arises to perform paging operations on more than one of those groups at the same time, there will be pool fragmentation. Sometimes, pool fragmentation may impact paging performance.
- If there is 'spare' SRDF/A throughput, the system will return to a normal RPO quicker after a link outage.

## **SRDF/A Best Practices**

- **R2 FRAME:** The R2 Frame should be at least as fast as the R1. This includes: the same amount, size, type of drives, and protection schemes should be used in both the R1 and R2 for the standard volumes. If additional volumes such as BCVs are configured on the R2 side, additional drives and cache should be used. For example, if using RAID 1/0 on the source frame with 15k drives, RAID 1/0 with 15k drives should be used on the target frame. Consideration should be given to segregating standards and BCV volumes onto separate drives.
- The default device write pending limit (amount of cache slots per volume) should be the same or higher in the R2 as in the R1. This may require more physical cache in the R2 than in the R1.

- When defining CLONE on the R2, keep the clone devices on segregated drives and use the pre-copy option.
- QoS with an initial value of 2 can be used to help reduce the copy impact.
- SNAP is not allowed on the R2 volumes.
- **BANDWIDTH:** Sufficient bandwidth needs to be provided to run SRDF/A. Prior to configuring SRDF/A, it is imperative to understand the workload. SRDF/A can sometimes reduce overall bandwidth by 20% over Synchronous SRDF, but it is highly dependent on the workload. It is best to keep this bandwidth reduction in reserve until the actual solution is implemented and the data can be analyzed.
- Symmetrix DMX GigE adapters and some Fibre Channel switches offer compression, but the actual compression values realized are highly dependent on the data and can fluctuate drastically over the business day. For example, certain batch workloads can achieve better compression than online workloads. The actual compression values can be pulled off of the GigE adapters via an Inline or out of the Fibre Channel switches. Compression can help reduce the overall bandwidth required for SRDF/A, but be extremely careful when counting on compression as cycle times can drastically elongate if the compression is not being realized.
- Your support representative can use the SYMMERGE and BCSD tools to model data and correctly size a proper configuration. Bandwidth needs to be at least equal to the average number of writes entering the sub-system. This will not guarantee minimum cycle times.
- If you are targeting minimum cycle times, sufficient bandwidth needs to be configured to handle the peak number of writes entering the system. Keep in mind that we typically are using 10 or 15 minute data to model 30 second cycle times. A sufficient amount of cache must be configured to keep SRDF/A active for the period that that data was collected. In other words, if you model on 15 minute data, you must configure enough cache and bandwidth to keep SRDF/A active for 15 minutes at a minimum. Cycle times may exceed the minimum during this period. Never guarantee minimum cycle times. Required bandwidth must be dedicated to SRDF/A. Do not share bandwidth with network traffic, tape, etc.
- RA COUNT: The correct number of RAs need to be configured. There should be at least N+1 RAs, where N is the number of RAs required, so that a service action can be performed to replace a RA if necessary.

- Synchronous groups and SRDF/A groups should be segregated onto their own physical adapters. *Do not mix Synchronous and SRDF/A on the same adapters.* Directors supporting SRDF/A should not be shared with any other SRDF solution.
- Caution! When moving from a Synchronous solution to SRDF/A, in many cases we have seen the bandwidth and adapter utilization increase as a result of the overall response time to the system decreasing.
- **MONITORING:** SRDF/A should be monitored during the initial roll-out to ensure that all components were properly sized and configured. Data needs to be collected via STP or WLA and then run through the tools again to verify the initial projections were correct. STP includes SRDF/A statistics, which can be very beneficial.
- Do not forget that Mainframe MSC customers have a way to monitor for issues; SCF1562I and SCF1563I messages. These will tell if they are experiencing transmit or restore issues. The messages will also tell which box is the issue.
- The SYMSTAT commands were specifically created for monitoring open systems SRDF/A but, when issued from the Service Processor on the Symmetrix, can be quite informative regardless of whether it is mainframe or open systems.
  - There are three options:
    1. Cycle
    2. Requests
    3. Cache
- Using different combinations of the three options can help determine what caused the CACA and you can even prevent a drop by monitoring the cache utilization closely. SRDF/A should be monitored regularly to look for workload changes and to predict increases in cache or bandwidth due to growth.
- **VERIFICATION:** The network should always be verified to ensure that the projected amount of bandwidth is configured. STP or WLA should be collected during the initial Adaptive Copy Synchronization to ensure that the required bandwidth is configured and that the network runs error free. Compression ratios should also be checked either at the switches or on the GigE adapters to verify that the correct numbers were used.
- **Upgrade or Reconfiguration:** Always re-evaluate the SRDF/A solution prior to upgrades or reconfigurations. This includes drive upgrades, adding volumes to the SRDF/A links, or changing the front end connectivity—for example, changing ESCON to FICON.

- **Starting SRDF/A:** SRDF/A activation is considerate of cache utilization. SRDF/A will capture a delta set of writes and send them in cycles across the link. In addition to the new writes, SRDF/A will include up to 30,000 invalid tracks per cycle. This is a design feature and the 30,000 track value was chosen to prevent cache from being flooded by the invalid tracks. Therefore, EMC generally recommends as a best practice to synchronize the boxes in Adaptive Copy Disk mode to below 30,000 invalid tracks before activating SRDF/A. This will ensure that SRDF/A will become secondary consistent within a few cycles.
- SRDF/A will activate with many more than 30,000 invalid tracks and in fact, some customers choose to activate SRDF/A when they have thousands or millions of invalid tracks. While this is allowed, only a maximum of 30,000 invalid tracks will be sent with each SRDF/A cycle. As a result, it will take many cycles before the frames are secondary consistent.

## **VPLEX Best Practices**

- Before implementing VPLEX<sup>®</sup>, decide on the back-end storage configuration and follow it throughout the storage lifecycle. The size of the LUN's allocated to the VPLEX should be cited and discussed. For better performance and management, try to allocate LUN's of the same size to the VPLEX. For example, we can present LUN's of 250 GB for all FC RAID 1 configurations and 500 GB for all SATA RAID 6.
- Fabric Zoning should consist of a set of zones, each with a single initiator and up to 16 targets.
- Avoid incorrect FC port speed between the fabric and VPLEX.
- If more than one engine is available, spread I/O paths across engines as well as directors.
- Use the same FE/BE ports on each director to avoid confusion; that is, B0-FC00 and A0-FC00
- Metadata volumes should be on storage volumes that have underlying redundant properties such as RAID 1 or RAID 5.
- Metadata volumes should be mirrored between different storage arrays.
- EMC recommends that you create your metadata volume from at least two 78GB storage volumes from two different arrays to provide data protection and redundancy.
- It is recommended when initially setting up the cluster that only the storage for the metadata volume be presented to the VPLEX. This is particularly important if hosts and



storage volumes will be migrating from an existing environment to a VPLEX environment.

- It is recommended that you back up the metadata volume periodically and before an upgrade.
- Host Multipathing should be set up for adaptive and not round-robin.
- Avoid multipathing software that does excessive round-robin and splits I/O.
- A one-to-one mapping of storage view and host is recommended.
- Single Initiator Zoning should be used and uniform zoning naming convention should be followed across the sites.
- Single Storage View should be created for each host to avoid confusion.

## Performance Troubleshooting Methodologies



### Data Collection ASK QUESTIONS?

- Is there a recent change in environment that has caused this issue?
- What application is running on the host?
- Is there a specific time window during which the performance issue is encountered?
- What is the activity running on the host when the performance issue occurs?
- Ask for the host grab?
- Collect Switch Logs?

### Analysis

- Analyze the Host Grabs and Switch Logs for errors on the HBA or Switch Ports. A faulty HBA Port or Switch Port can cause performance issues.
- Once you find that host HBA and Switch Logs are fine, analyze the Storage Performance using WLA, Symmerge for VMAX/DMX, or Analyzer for VNX for the time window for which the problem was reported.
- Check the FE Port Utilization, LUN Utilization, Cache Utilization, and BE Utilization for that time window.
- If a Front End Port is highly utilized, hosts can be migrated to other FE Ports.
- In cases of high cache utilization, enable FAST Cache in VNX or adjust the FAST VP policy for that Host to accommodate more data on EFD.
- High LUN utilization means that the current LUN doesn't have sufficient RAID Geometry to meet the I/O requirement. Either the FAST VP policy can be modified or the LUN can be migrated to another RAID Geometry after calculating the peak IOP requirements.
- If all vital parameters of the storage are OK, if there is a VPLEX between the Host and Storage, the VPLEX is acting as a bottleneck. If the host is zoned to more VPLEX FE Ports and the problem persists, contact VPLEX Support for analyzing the logs.
- If no issue is detected at HBA, Switch, VPLEX, or Storage means there is some issue with the application itself and the application vendor needs to be contacted.

### MANTRA

- If there was no performance issue occurring earlier, what can be the sudden cause of the issue? There might have been some change in the environment that caused this and we have to find that change.

## References

- [www.emc.com](http://www.emc.com)
- [powerlink.emc.com](http://powerlink.emc.com)

## Appendix A

# Includes dispositions for alerts for home status page

#

# Note that events that require the use of the threshold

# features cannot be grouped together with a range since the

# threshold count would apply to all events in that range

# MasterControl - CS\_PLATFORM - File system error

facilitypolicy 6:129, 6

disposition range=4-4 threshold=1 rearm=100000 resetafter=86400, trap  
"/nas/site/trap.cfg 2"

disposition range=7-9, trap "/nas/site/trap.cfg 2"

disposition range=10-10, trap "/nas/site/trap.cfg 2"

disposition range=12-13, trap "/nas/site/trap.cfg 2"

disposition range=15-16, trap "/nas/site/trap.cfg 2"

disposition range=17-20, trap "/nas/site/trap.cfg 2"

# CFS - DART

facilitypolicy 1:27, 6

disposition range=1-1, trap "/nas/site/trap.cfg 2"

disposition range=4-4, trap "/nas/site/trap.cfg 2"

# CFS - CS\_PLATFORM

facilitypolicy 6:27, 6

disposition range=101-101, trap "/nas/site/trap.cfg 2"

disposition range=103-115, trap "/nas/site/trap.cfg 2"

disposition range=123-125, trap "/nas/site/trap.cfg 2"

disposition range=6-6,

disposition range=6-6, trap "/nas/site/trap.cfg 2"

disposition range=8-10, trap "/nas/site/trap.cfg 2"

disposition range=12-15, trap "/nas/site/trap.cfg 2"

# IP - DART - Duplicate address

facilitypolicy 1:43, 3

disposition range=1-1 threshold=1 rearm=100000 resetafter=21600, trap "/nas/site/trap.cfg 2"

disposition range=2-2 threshold=1 rearm=100000 resetafter=21600, trap "/nas/site/trap.cfg 2"

# VC - DART - No Virus Checker server; stop CIFS

facilitypolicy 1:81, 6

disposition range=2-3, trap "/nas/site/trap.cfg 2"

disposition range=4-4 threshold=1 rearm=100000 resetafter=21600, trap "/nas/site/trap.cfg 2"

disposition range=5-5 threshold=1 rearm=100000 resetafter=21600, trap "/nas/site/trap.cfg 2"

disposition range=6-6 threshold=1 rearm=100000 resetafter=21600, trap "/nas/site/trap.cfg 2"

disposition range=7-7, trap "/nas/site/trap.cfg 2"

disposition range=9-9, trap "/nas/site/trap.cfg 2"

disposition range=12-15, trap "/nas/site/trap.cfg 2"

disposition range=19-19, trap "/nas/site/trap.cfg 2"

disposition range=20-20 threshold=1 rearm=100000 resetafter=21600, trap "/nas/site/trap.cfg 2"

disposition range=21-27, trap "/nas/site/trap.cfg 2"

disposition range=28-28 threshold=1 rearm=100000 resetafter=21600, trap "/nas/site/trap.cfg 2"

# BoxMonitor - CS\_PLATFORM - Various Hardware Problems

facilitypolicy 6:131, 6

disposition range=1-1, trap "/nas/site/trap.cfg 2"

disposition range=2-2, trap "/nas/site/trap.cfg 2"

disposition range=3-15, trap "/nas/site/trap.cfg 2"

disposition range=18-21, trap "/nas/site/trap.cfg 2"

disposition range=24-34, trap "/nas/site/trap.cfg 2"  
disposition range=38-39, trap "/nas/site/trap.cfg 2"  
disposition range=43-43, trap "/nas/site/trap.cfg 2"  
disposition range=52-52, trap "/nas/site/trap.cfg 2"  
disposition range=58-58 threshold=1 rearm=1000 resetafter=86400, trap  
"/nas/site/trap.cfg 2"  
disposition range=62-67 threshold=1 rearm=1000 resetafter=86400, trap  
"/nas/site/trap.cfg 2"  
disposition range=70-70 threshold=1 rearm=1000 resetafter=86400, trap  
"/nas/site/trap.cfg 2"  
disposition range=71-73, trap "/nas/site/trap.cfg 2"  
disposition range=77-78, trap "/nas/site/trap.cfg 2"  
disposition range=80-91, trap "/nas/site/trap.cfg 2"  
disposition range=92-92 severity=1-1, trap "/nas/site/trap.cfg 2"  
disposition range=100-115 threshold=1 rearm=1000 resetafter=86400, trap  
"/nas/site/trap.cfg 2"  
disposition range=117-128 severity=2-2 threshold=1 rearm=1000 resetafter=86400, trap  
"/nas/site/trap.cfg 2"  
disposition range=150-152, trap "/nas/site/trap.cfg 2"  
disposition range=200-215 threshold=1 rearm=1000 resetafter=86400, trap  
"/nas/site/trap.cfg 2"  
disposition range=300-315 threshold=1 rearm=1000 resetafter=86400, trap  
"/nas/site/trap.cfg 2"  
disposition range=500-500 severity=1-2, trap "/nas/site/trap.cfg 2"  
disposition range=502-502 severity=1-2, trap "/nas/site/trap.cfg 2"  
disposition range=504-504 severity=1-2, trap "/nas/site/trap.cfg 2"  
disposition range=506-506 severity=1-2, trap "/nas/site/trap.cfg 2"  
disposition range=508-508 severity=1-2, trap "/nas/site/trap.cfg 2"  
disposition range=510-510 severity=1-2, trap "/nas/site/trap.cfg 2"  
disposition range=512-512 severity=1-2, trap "/nas/site/trap.cfg 2"  
disposition range=514-515 severity=1-2, trap "/nas/site/trap.cfg 2"  
disposition range=517-517 severity=1-2, trap "/nas/site/trap.cfg 2"

```
disposition range=519-519 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=521-521 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=523-523 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=524-524 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=527-527 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=536-536, trap "/nas/site/trap.cfg 2"
disposition range=538-538 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=541-542 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=544-544 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=546-546 severity=1-2, trap "/nas/site/trap.cfg 2"
disposition range=550-557 severity=2-2, trap "/nas/site/trap.cfg 2"
disposition range=558-558 severity=2-2 threshold=1 rearm=100000 resetafter=1800,
trap "/nas/site/trap.cfg 2"
```

#### # JServer - CS\_PLATFORM - Various events

facilitypolicy 6:135, 4

```
disposition severity=0-4 range=101-101, trap "/nas/site/trap.cfg 2"
disposition severity=0-4 range=201-800, trap "/nas/site/trap.cfg 2"
disposition severity=0-4 range=1001-1001, trap "/nas/site/trap.cfg 2"
```

#### # UFS - DART - Quota events

facilitypolicy 1:64, 4

```
disposition severity=0-4 range=4-5, trap "/nas/site/trap.cfg 2"
disposition severity=0-4 range=7-9, trap "/nas/site/trap.cfg 2"
disposition range=11-11, trap "/nas/site/trap.cfg 2"
disposition range=15-15 threshold=1 rearm=100000 resetafter=86400, trap
"/nas/site/trap.cfg 2"
disposition severity=0-4 range=16-16, trap "/nas/site/trap.cfg 2"
disposition severity=0-4 range=41-41, trap "/nas/site/trap.cfg 2"
disposition severity=0-4 range=46-47, trap "/nas/site/trap.cfg 2"
```

#### # SVFS - DART - SavVol problems

facilitypolicy 1:70, 4

disposition severity=0-4 range=2-2, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=5-9, trap "/nas/site/trap.cfg 2"

# SVFS - CS\_PLATFORM - Checkpoint problems

facilitypolicy 6:70, 4

disposition severity=0-4 range=10-21, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=29-29, trap "/nas/site/trap.cfg 2"

# NASDB - CS\_PLATFORM - Schedule problems

facilitypolicy 6:137, 3

disposition severity=0-4 range=202-202, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=211-212, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=214-214, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=307-307, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=311-311, trap "/nas/site/trap.cfg 2"

# NASDB - CS\_CORE - NASDB problems

facilitypolicy 2:137, 4

disposition severity=4-4 range=2-2, trap "/nas/site/trap.cfg 2"

# VRPL - DART - Replication V1 problems

facilitypolicy 1:77, 3

disposition severity=0-4 range=1-2, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=6-6, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=8-8, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=10-10, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=14-14, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=17-18, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=22-23, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=27-27, trap "/nas/site/trap.cfg 2"

disposition severity=0-4 range=30-30, trap "/nas/site/trap.cfg 2"

# XLT - DART - XLT file corruption

facilitypolicy 1:72,2

disposition severity=0-2 range=1-2, trap "/nas/site/trap.cfg 2"

# UPS - CS\_PLATFORM - Hammerhead UPS events

facilitypolicy 6:140,6

disposition range=1-1000 severity=0-4, trap "/nas/site/trap.cfg 2"

# LOGCOLLECT - CS\_PLATFORM - Automatic Log collect/transfer

facilitypolicy 6:141,6

disposition range=3-11 severity=4-4, trap "/nas/site/trap.cfg 2"

disposition range=101-109 severity=4-4, trap "/nas/site/trap.cfg 2"

disposition range=306-312 severity=4-4, trap "/nas/site/trap.cfg 2"

disposition range=400-400 severity=4-4, trap "/nas/site/trap.cfg 2"

# Checkup - CS\_PLATFORM

facilitypolicy 6:142,6

disposition range=112-113 severity=3-4, trap "/nas/site/trap.cfg 2"

disposition range=122-123 severity=3-4, trap "/nas/site/trap.cfg 2"

# ADMIN - CS\_PLATFORM - Output from dskMon

facilitypolicy 6:24,6

disposition range=16-19 severity=0-4, trap "/nas/site/trap.cfg 2"

# EventLogger - CS\_PLATFORM - Pre-CCMD Alerts

facilitypolicy 6:130,4

disposition range=50-54 severity=0-4, trap "/nas/site/trap.cfg 2"

# CommandService - CS\_CORE

facilitypolicy 2:9,4

disposition range=3-3 severity=0-4, trap "/nas/site/trap.cfg 2"

disposition range=6-9 severity=0-4, trap "/nas/site/trap.cfg 2"

disposition range=11-18 severity=0-4, trap "/nas/site/trap.cfg 2"

# DBMS - CS\_CORE

facilitypolicy 2:122,4

disposition range=1-5 severity=0-4, trap "/nas/site/trap.cfg 2"



# DBMS - DART

facilitypolicy 1:122,4

disposition range=1-1 severity=0-4 threshold=1 rearm=100000 resetafter=21600, trap  
"/nas/site/trap.cfg 2"

disposition range=2-3 severity=0-4, trap "/nas/site/trap.cfg 2"

# REP - CS\_CORE - Replication V2

facilitypolicy 2:108,4

disposition range=2-24 severity=0-4, trap "/nas/site/trap.cfg 2"

# REP - DART - Replication V2

facilitypolicy 1:108,4

disposition range=1-12 severity=0-4, trap "/nas/site/trap.cfg 2"

disposition range=14-15 severity=0-4, trap "/nas/site/trap.cfg 2"

disposition range=22-22 severity=0-4, trap "/nas/site/trap.cfg 2"

disposition range=24-25 severity=0-4, trap "/nas/site/trap.cfg 2"

disposition range=28-29 severity=0-4, trap "/nas/site/trap.cfg 2"

disposition range=40-45 severity=0-4, trap "/nas/site/trap.cfg 2"

disposition range=47-54 severity=0-4, trap "/nas/site/trap.cfg 2"

# CEPP - DART

facilitypolicy 1:146,4

disposition range=2-2 severity=0-4, trap "/nas/site/trap.cfg 2"

# DNS - DART

facilitypolicy 1:118,4

disposition range=1-1 severity=0-4 threshold=1 rearm=100000 resetafter=21600, trap  
"/nas/site/trap.cfg 2"

# LOCK - DART

facilitypolicy 1:68,4

disposition range=1-2 severity=0-4, trap "/nas/site/trap.cfg 2"

# NETLIB - DART

facilitypolicy 1:73,4

```

        disposition range=1-1 severity=0-4, trap "/nas/site/trap.cfg 2"
# SECMAP - DART
facilitypolicy 1:115,4
        disposition range=3-3 severity=0-4 threshold=1 rearm=100000 resetafter=21600, trap
"/nas/site/trap.cfg 2"
        disposition range=4-4 severity=0-4 threshold=1 rearm=100000 resetafter=21600, trap
"/nas/site/trap.cfg 2"
# USRMAP - DART
facilitypolicy 1:93,4
        disposition range=6-6 severity=0-4, trap "/nas/site/trap.cfg 2"
        disposition range=10-10 severity=0-4 threshold=1 rearm=100000 resetafter=21600, trap
"/nas/site/trap.cfg 2"
        disposition range=11-11 severity=0-4, trap "/nas/site/trap.cfg 2"
# WINS - DART
facilitypolicy 1:117,4
        disposition range=1-1 severity=0-4, trap "/nas/site/trap.cfg 2"
#
# DART:CAM
#
facilitypolicy 1:26, 4
        disposition range=7-7 severity=0-3, trap "/nas/site/trap.cfg 2"
#
# DART:FSTOOLS
#
facilitypolicy 1:40, 7
        disposition range=1-1, trap "/nas/site/trap.cfg 2"
#
# DART:NFS
#
facilitypolicy1:52, 7

```

```
    disposition range=1-3, trap "/nas/site/trap.cfg 2"
#
# DART:STORAGE
#
facilitypolicy 1:58, 7
    disposition range=3-4, trap "/nas/site/trap.cfg 2"
    disposition range=38-38, trap "/nas/site/trap.cfg 2"
#
# DART:CHAMII - CHAMIIENCMON Data mover enclosure events
#
facilitypolicy 1:86, 7
    disposition range=6-6 severity=1-2, trap "/nas/site/trap.cfg 2"
    disposition range=16-16 severity=1-2, trap "/nas/site/trap.cfg 2"
    disposition range=21-21 severity=1-2, trap "/nas/site/trap.cfg 2"
    disposition range=23-23 severity=1-2, trap "/nas/site/trap.cfg 2"
    disposition range=25-26 severity=1-2, trap "/nas/site/trap.cfg 2"
    disposition range=31-33 severity=1-2, trap "/nas/site/trap.cfg 2"
#
# CS_PLATFORM:EmailUser
#
facilitypolicy 6:144, 7
    disposition range=1-1 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
    disposition range=2-2 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
    disposition range=3-3 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
    disposition range=4-4 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
    disposition range=5-5 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
    disposition range=6-6 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
    disposition range=7-7 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
    disposition range=8-8 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
```

```
        disposition range=9-9 threshold=1 rearm=1000 resetafter=86400, trap "/nas/site/trap.cfg 2"
#
# DART:Deduplication
#
facilitypolicy 1:148, 7
        disposition range=1-3, trap "/nas/site/trap.cfg 2"
#
# Notification definitions for all storage related events.
# CS_PLATFORM - NaviEventMonitor - Severity 6
#
facilitypolicy 6:138, 6
        disposition range=4-4, trap "/nas/site/trap.cfg 2"
        disposition range=10-10, trap "/nas/site/trap.cfg 2"
        disposition range=100-101, trap "/nas/site/trap.cfg 2"
        disposition severity=0-4 range=201-203,trap "/nas/site/trap.cfg 2"
```

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.