



CATE Storage & Backup Engineering
Citi Architecture & Technology Engineering

FAST VP - Directives and Best Practice for VMAX20K & VMAX40K

CSBE-11010-BP

Best Practice

Version 2.5

VERSION CONTROL

Version and Date of Issue	Author	Comment
V1.0 17 th November 2011	Chris Clayton	Production release
V1.1 19th July 2012	Chris Clayton	Changed Best Practice for FC tier changed from RAID6(14+2) to RAID1 for new purchases in line with revised EMC Best Practice. Removed Best Practice when using SE7.2 to manage Time Windows (FAST never deployed into production with SE7.2) Updated CATE recommended values FAST VP parameters in Section 5.5. • FC PRC setting • Relocation rate • Workload period Added Section 5.6 , which covers the iAlert pool monitoring and FAST VP state alerting script. Added commands to Section 6.1 on how to determine the maximum possible size of a migration to the FC tier. Flash erase statements updated with minimum Enginuity levels needed for SecureErase (Section 8.16). Added new Best Practice on decommissioning FAST VP storage (Section 8.17). Updated pool alerting Best Practice (Section 8.8.1.1 & Appendix I) Updated capacity planning section with new Best Practice from EMC on IOPS utilization – rules of thumb (Section 9.2.2.1). Changed standard configurations in Appendix A to use RAID1 for FC tier. Added emergency capacity expansion options into Appendix A for RAID6 configurations. Appendix B updated to reflect change to Best Practice in this version. Appendix H (non-standard configurations) removed.

V1.2 11th February 2013	Chris Clayton	Updated to include support for Enginuity 5876 & SE7.5, including: • FAST support for allocation by policy (Sections 5.2.1, 5.5, 8.8 & Appendix B) • VLUN source pool designation (Section 8.7.1) • Updated entry on monitoring FAST VP engine status using symevent log. • FAST VP Storage Group re-association (Section 8.18) • Document updated throughout to accommodate changes in symcli output in SE7.5 (new or renamed columns).
V1.3 11th February 2013	Carl Williams	Updated to cover adoption of Unisphere for performance analysis in place of SPA. Appendix H added covering Using Unisphere for VMAX with FAST VP
V2.0 4 th April 2013	Chris Clayton	Major update to add support for FAST VP on VMAX40K.
V2.1 12 th April 2013	Chris Clayton	Changed standard for EFD tier protection from RAID5(7+1) to RAID5(3+1) for VMAX40K to protect against loss of an entire engine in a 4-Engine VMAX40K.
V2.2 24 th April 2013	Chris Clayton	Added Best Practice for 2-Tier Gold FAST VP configuration (Appendix H). Numerous other changes made throughout the document to distinguish between 3-tier “Silver” and 2-tier “Gold” configuration best practice.
V2.3 07 th August 2013	Carl Williams	Very Minor update to Drive sizing tables to reflect 100GB EFD drive
V2.4 4 th February 2014	Carl Williams	Edit release – Update to URL’s and applicable SE references
V2.5 15 th May 2014	Carl Williams	Removed SPA references

CONTENTS

Legal Notice	7
Contact	7
Latest Versions	7
1 FAST VP OVERVIEW	8
1.1 Workload skew	9
1.2 Other key FAST concepts	11
2 FAST VP DOCUMENTATION	12
2.1 Nomenclature	13
3 HOW TO USE THIS DOCUMENT	14
4 PLANNING A FAST VP DEPLOYMENT	15
4.1 When should FAST VP be used?	15
4.2 Purchasing requirements and restrictions	15
4.2.1 Hardware purchase	15
4.2.2 Licensing	15
4.3 Supported configuration types	16
4.4 Supported disk drives and RAID types by VMAX model	16
4.4.1 VMAX40K	16
4.4.2 VMAX20K	16
4.5 Supported VMAX FAST VP configurations	17
4.5.1 Introduction	17
4.5.2 CATE standard 3-tier (“Silver”) FAST VP configurations	17
4.5.3 CATE standard 2-tier (“Gold”) FAST VP configurations	18
4.5.4 Bespoke configurations: Analyzing your specific workload	18
4.6 Performance analytics tools for user with FAST VP	18
4.6.1 Unisphere for VMAX	19
4.7 Restrictions and other comments	19
4.7.1 Known FAST VP issues and restrictions	19
4.7.2 TimeFinder/Clone	20
4.7.3 SRDF	20
4.7.4 When should FAST VP not be used?	20
5 3-TIER FAST VP CONFIGURATION BEST PRACTICE	22
5.1 Introduction	22
5.2 3-tier VP pool and Tier Best Practice	22
5.2.1 Selecting and setting the Bound Pool	23
5.2.2 Use of thin device pre-allocation	23
5.3 3-tier FAST VP policy Best Practice	23
5.3.1 Protecting applications from performance degradation	23
5.4 3-tier FAST VP Time Windows Best Practice	24
5.4.1 Managing Time Windows	25
5.5 3-Tier FAST VP control parameters	27
5.6 iAlert : The Intelligent Alerting Script	31
5.6.1 Goals	31
5.6.2 Platform	31
5.6.3 Eliminating False Alerts	32
5.6.4 Architecture	32
5.6.5 Script Location	32
5.6.6 Installation	32
5.6.7 Maintenance	32
6 PLACING APPLICATIONS UNDER FAST VP CONTROL	33
6.1 Existing apps: Tech refresh from a single tier array to a FAST VP array	33
6.2 New application going live on an existing FAST VP array	36

6.3	Application benchmarking (generic)	36
7	CHARGEBACK OF FAST VP MANAGED STORAGE	38
7.1	Provisioned capacity model	38
7.2	Combined capacity & policy chargeback reporting via ECC	38
7.3	Chargeback reporting via local scripts	39
7.4	SANscreen issues	39
8	MANAGING FAST VP	40
8.1	Determining tier residency	40
8.1.1	Current tier residency via CLI	40
8.1.2	Current tier residency via Unisphere	42
8.1.3	Trending via Unisphere	43
8.2	Monitoring data movement	46
8.3	Monitoring performance	47
8.4	Troubleshooting lack of data movement	47
8.5	Audit logging	47
8.6	Dealing with customer storage performance complaints	49
8.7	Rollback and other manual data movement	49
8.7.1	Using Virtual LUN with source under FAST VP control	49
8.8	Dealing with pool exhaustion issues in 5875	52
8.9	Backup/restore of FAST VP configuration	53
8.10	Using Zero Space Reclaim in a FAST VP environment	53
8.11	Copy pace management	54
8.12	Protecting early migrants from artificially high performance	55
8.12.1	Problem statement	55
8.12.2	Best practice	55
8.13	Dealing with irregular workloads	55
8.14	Loss of the Service Processor	56
8.15	Excessive Write Pending	56
8.16	FAST VP EFD drive erasure	57
8.17	Decommissioning storage under FAST VP control	57
8.17.1	Unbinding devices under FAST VP management	57
8.18	FAST VP Storage Group re-association (5876 or later only)	59
9	3-TIER FAST VP CAPACITY PLANNING	61
9.1	Introduction	61
9.2	Determining when expansion is required	62
9.2.1	Storage capacity	62
9.2.2	IOPS capacity	63
9.3	Thin Pool Oversubscription	68
9.3.1	Note for SE7.5 or later	69
10	APPENDIX A – APPROVED 3-TIER FAST VP CONFIGURATIONS	70
10.1	VMAX40K 3-tier (“Silver”) standard configurations	70
10.1.1	3-tier “Silver” configuration emergency Capacity-on-Demand upgrades	72
10.2	VMAX20K Restricted (Legacy) Standard RAID1 FC tier configurations	73
10.2.1	Emergency Capacity-on-Demand upgrades for RAID1 FC tier configs	74
10.3	VMAX20K Restricted (Legacy) RAID6 FC tier configurations	75
10.3.1	Emergency Capacity-on-Demand upgrades for RAID6 FC tier configs	75
11	APPENDIX B - WORKED EXAMPLE OF SETTING UP 3-TIER FAST VP	79
12	APPENDIX C – UPGRADING A SINGLE-TIER VMAX20K VP ARRAY TO FAST VP	82
12.1	Hardware sizing	82
12.2	Placing applications under FAST VP control	82
13	APPENDIX D – WHAT DOES THE POLICY MAX PERCENTAGE MEAN IN PRACTICE?	84
13.1.1	Virtual Provisioning – logical vs. actual	84

14	APPENDIX E – GETTING POLICY INFORMATION VIA SYMCLI	86
15	APPENDIX G – USING UNISPHERE WITH FAST VP	89
15.1	Array Level information	90
15.1.1	Heat Maps	90
15.1.2	Write Pending Count and Percentage	99
15.2	Tier/Pool Level information	101
15.2.1	Tier level information:	101
15.2.2	Pool level information	103
15.3	Storage Group Level information	106
15.4	Device level Information	108
16	APPENDIX H – VMAX40K 2-TIER (GOLD) FAST VP CONFIGURATIONS	109
16.1	Introduction	109
16.2	Supported configurations	109
16.2.1	Gold Configuration capabilities	109
16.3	2-tier specific Best Practice	112
16.4	Placing application under FAST VP control	113
16.5	Worked example of setting up 2-tier FAST VP	114
16.6	Worked example of adding an SSD tier to single tier (SAS) VMAX40K	115

Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP and CITIGROUP are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

Contact

Please address any comments or questions regarding the information contained in this document via email to the “*CATE Storage and Backup Engineering” distribution list or alternatively to dl.CATE.Storage.and.Backup.Engineering.

Latest Versions

It is recommended that the latest version of the document be used from the CSBE website

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/symmetrix.aspx>

Newer versions may include expanded scope, or comments on issues discovered after deployment.

1 FAST VP Overview

The aim of Fully-Automated Storage Tiering (FAST) is to dynamically share VMAX disk resources among applications so that the response time is optimized for the most active and/or critical applications, while less active/critical applications maintain a satisfactory response time.


These disk resources in a FAST-managed array will typically consist of:

Drive type	Performance	Cost	% of array
Enterprise Flash Drives (EFD)	Very High	Very High	Small
Fibre Channel (FC) or SAS	High-Moderate	Moderate	Moderate
SATA	Low	Low	Large

High-performance 2-Tier configurations with only EFD and SAS are also possible, and used in Citi for some franchise critical applications.

The key to the FAST cost model is that the cost of the high-performance Flash drives is more than outweighed by the savings on the slower SATA drives, compared with an all fibre channel configuration of the same average performance. FAST's ability to automatically and transparently use the slower/cheaper SATA storage for less active data drives savings in cost, without compromising the performance of an application.

FAST achieves this by running background algorithms to continuously analyze the utilization (busy rate) of Symmetrix array storage. FAST processes the algorithm data, and plans are generated for moving data between tiers for each application to meet its current performance needs.



There are two flavours of FAST:

FAST a.k.a. FAST V1	Operates on thick (standard) storage, moving data at LUN-level granularity
FAST VP (Virtual Pools)	Operates on thin (virtually provisioned) storage, moving sub-LUN units. Granularity of movement is typically 7680KB (10 thin device extents), but can be as little as a single extent (768KB) in some circumstances.

FAST VP builds incrementally upon the functionality of the original FAST V1, adding finer granularities of both performance measurement and data movement. The data from a single thin device under FAST VP control can be spread across multiple tiers, as the FAST controller is free to relocate small groups of extents of a thin device, based on performance data gathered at the sub-extent level.

This document only covers FAST VP as thinly provisioned storage is now the default for new deployments within Citi. FAST V1 is not being certified by CATE.

FAST VP delivers a reduction in storage costs by combining two STACKED cost reduction technologies, i.e. thin provisioning AND Storage Tiering.



1.1 Workload skew

A key concept with all FAST implementations is that of workload skew. Most people are familiar with the "80:20" rule that states that for many events, roughly 80% of the effects come from 20% of the causes. This is more formally known as the Pareto principle, which is more generally defined as, " $k\%$ is taken by $(100 - k)\%$ of the participants". So in the context of FAST, a skew of 85 means that 85% of the I/O is executed to 15% of the storage.

FAST depends on, and uses the existence of workload skew to move the busiest devices (i.e. some of this busy 15% in the above example) to higher performing tiers (EFD and FC/SAS) and idle devices to lower-cost tiers (SATA).

FAST VP is smarter in that it relocates at the virtual provisioning extent level, rather than an entire Symmetrix device. This results in the effective skew being higher.




FAST VP works at the extent level, resulting in higher effective workload skew

The left side of the above diagram shows a workload with a low skew i.e. 71% of I/O goes to 29% of the storage. The right side shows a workload with a higher skew i.e. 85% of I/O goes to just 15% of the storage. With the higher skew (right), one can have more lower-cost SATA in the configuration in place of FC/SAS drives.

However, these are actually the same workload taken from a Citi array, but on the left the skew is when the workload is looked at from the device level, where as on the right, the skew is at the extent level. This is important because it illustrates how using FAST VP (with its extent-level data movement) shifts the skew so that one can achieve increased cost saves (and/or performance benefits) by changing the relative amount of FC/SAS and SATA targets that are needed.


The following slide from EMC illustrates how different workload skews require different percentages of each storage tier. It is anticipated that most Citi workloads will generally fall into the Moderate to Low FAST VP skew categories.



1.2 Other key FAST concepts

The following diagram shows the relationship between three key FAST-related elements:

- **Symmetrix Tiers** in FAST VP are comprised of one or more virtual (thin) pools. Each tier will typically be built from a different technology type: Flash drives (fast but expensive), the familiar fibre channel drives (or newer SAS alternative in VMAX40K), or low-cost (but slower) SATA drives.
- **FAST policies** define an upper limit for each tier, and this limit determines how much data from a storage group associated with the policy is allowed to reside on that tier (e.g. 50%).
- **Storage Groups** are sets of devices. Storage groups define the devices used by specific applications. Storage groups are associated with FAST policies, and all of the devices in the storage group come under FAST control. The FAST controller can move data from the devices between tiers in the associated policy. Storage groups can easily be moved from policy to policy, with FAST VP making the necessary “policy compliance” data movements automatically. A device must be a member of one and only one Storage Group.



When a storage group is associated with a policy, two types of data movement will start to take place:

Performance moves are requested and executed by the Symmetrix microcode. These data movements will be governed by the workload, but will only be executed within the constraints of the associated FAST policy. That is, a performance movement will not cause a storage group to become non-compliant with its FAST policy. The busiest extents will be migrated to the higher tier, and the least busy to the lowest tier.

Compliance moves are generated by the FAST controller (which runs on the Service Processor), and executed by the microcode. These movements bring the capacity of the storage group back within the boundaries specified by the associated policy. As a consequence, some storage is migrated from FC/SAS tier to SATA tier when a storage group is first brought under FAST VP control in order to bring it into compliance with the policy with which it is associated.

2 FAST VP documentation

The present document focuses primarily on directives and best practice for using FAST VP in Citi. A comprehensive overview of FAST VP and details of actual FAST VP usage (set up procedures, and other command line syntax) is clearly described in the following EMC documents, and is not repeated here.

- **EMC Solutions Enabler Symmetrix Array Controls Version 7.x CLI Product Guides**
 - (Section on “Fully Automated Storage Tiering”)
 - (Section on “Managing time windows”)
- **“Implementing FAST VP for EMC Symmetrix VMAX Series Arrays” Technical Notes**

These two vendor documents must be consulted when using FAST VP (see references below).

The present document assumes familiarity with the FAST Section of the Array Controls CLI Product Guide, as a minimum. The above “Implementing FAST VP” document provides valuable background information on the functioning of FAST VP, as well as detailed examples of FAST VP commands.

This document only focuses on Best Practice that is specific to Citi, or filling in gaps in the current EMC FAST VP documentation.

It is also essential that the CSBE VMAX FAST VP Release Notes (see below) be read in conjunction with this Best Practice guide, as it contains vital information on pre-requisites, and known issues and limitations with the current version of the software.

Document	URL
EMC Solutions Enabler Symmetrix Array Controls Version 7.x CLI Product Guides Section on “Fully Automated Storage Tiering” Section on “Managing time windows” (SE7.3 or later)	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/ArrayControls73GA.pdf https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/ArrayControls75GA.pdf
Implementing FAST VP for EMC Symmetrix VMAX Series Arrays Technical Notes	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/Implementing_FASTVP_TechnicalNotes_.pdf
“FAST VP for EMC Symmetrix VMAX Theory and Best Practices for Planning and Performance” Technical Notes	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/FAST_VP_Theory_BestPractice_Theory_Performance.pdf

VMAX FAST VP Release Notes	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/CSBE-11010-RN.pdf
EMC Symmetrix VMAX Virtual Provisioning Core Policies and Directive	http://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/09058-sg.pdf
VMAX Virtual LUN Directives and Best Practice	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/09033-SG.pdf
Unisphere for VMAX Solutions Guide	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/CSBE-12025-SG_UniVMAX.pdf

2.1 Nomenclature

The following terms are used in both Citi and EMC documentation, and their meaning is clarified below for the reader's convenience.

Term	Tracks	Size	Notes
Track	1	64KB	
Thin device extent	12	768KB	Also known as a chunk or a thin device track group This is the unit of allocation of storage from a thin pool to a thin (VP) device.
Extent group	120	7680KB	10 thin device extents

3 How to use this document

This document is broken down into the following major sections:

- Planning a FAST VP deployment
- FAST VP Configuration Best Practice
- Placing Applications under FAST VP control
- Chargeback of FAST VP managed Storage
- Managing FAST VP
- FAST VP Capacity Planning
- Various appendices with supplementary information

If you are planning a FAST VP implementation, it is recommended that you read the following EMC document sections before consulting this “Directives and Best Practice” guide.

- Implementing Fully Automated Storage Tiering for Virtual Pools (FAST VP) for EMC Symmetrix VMAX™ Series Arrays
 - **Recommended:** Pages 3 to 46, which provides important background information, only some of which is covered in the present document.

If you are going to be configuring and managing FAST VP, you **must** read the following documents before consulting this “Directives and Best Practice” guide.

- **Mandatory:** EMC Solutions Enabler Symmetrix Array Controls CLI Version 7.3 Prodcut Guide
 - Section on “Fully Automated Storage Tiering”
 - Section on “Managing time windows”
- **Recommended:** Implementing Fully Automated Storage Tiering for Virtual Pools (FAST VP) for EMC Symmetrix VMAX™ Series Arrays
 - **Mandatory:** Pages 3 to 46, which provides important background information, only some of which is covered in the present document.
 - **Recommended:** Section on Management interface: SYMCLI (P111 onwards)

The ULRs for these documents are given in the previous section.

4 Planning a FAST VP deployment

4.1 When should FAST VP be used?

3-Tier FAST VP configurations should be used for non-SRDF “Silver” storage requirements in the range 220-320 usable TB. For smaller “Silver” requirements, non-FAST VP configurations should be used.

2-Tier FAST VP configurations should be used for non-SRDF “Gold” storage requirements for ~190TB (see Appendix H for supported configurations and Best Practice for 2-Tier Gold FAST deployments).

FAST VP is not supported in Citi for use in conjunction with SRDF. At the time of writing, there are no plans to support this combination.

4.2 Purchasing requirements and restrictions

4.2.1 Hardware purchase

CTI have elected to purchase FAST VP systems **with fixed standard configurations**, and leveraging a Capacity-on-Demand model for emergency capacity expansion. Both the fixed standard configurations and their associated Capacity-on-demand upgrades (where appropriate) are specified in Appendix A.

4.2.2 Licensing

To use FAST VP, both Symmetrix Optimizer and FAST VP licenses are required. Virtual LUN is included in the Optimizer license. From Jan 2013, all required licenses are included “pre-paid” under the EMC/Citi ELA agreement, which runs for 4 years.

You can check that an array has the appropriate FAST VP and Optimizer licenses with the following command:

```
# symmlmf -sid xxxx list -type sym
```

```
Symmetrix ID: 00019260xxxx
```

Feature Name	Lic	Type	Capacity	Units
SYMM_VMAX_ENGINUITY	EMCLM	R-TB-Non-SATA		0
		R-TB-SATA		0
SYMM_VMAX_FAST_VP	EMCLM	R-TB-Non-SATA		128
		R-TB-SATA		200
SYMM_VMAX_OPTIMIZER	EMCLM	R-TB-Non-SATA		128
		R-TB-SATA		200

Legend:

```
Lic(ense Type):  
  EMCLM = emclm license  
  SE     = se license
```

4.3 Supported configuration types

CATE supported 3-tier FAST configurations for “Silver” storage. The majority of this document is focussed on the more widely deployed 3-tier FAST configurations. A 2-tier “Gold” FAST configuration is also supported for franchise-critical applications only (see Appendix H)..

4.4 Supported disk drives and RAID types by VMAX model

4.4.1 VMAX40K

Tier	Drive Type	Technology	Approx Size(GB) ¹	Rotation speed	RAID Protection	Notes
Top	100GB	Enterprise Flash	98GB	N/A	RAID5(3+1)	Nominally a 100GB Flash drive
Middle	300GB	SAS	288GB	10K	RAID1	Nominally a 300GB SAS drive
Bottom	1TB	SATA	1882GB	7.2K	RAID6(6+2)	Nominally a 1000GB SATA drive

Supported drive types for VMAX40K

4.4.2 VMAX20K

Tier	Drive Type	Technology	Approx Size(GB) ²	Rotation speed	RAID Protection	Notes
Top	Samsung 200GB CTSG200	Enterprise ³ Flash	186GB	N/A	RAID5(7+1)	Nominally a 200GB Flash drive
Middle	Seagate 450GB HUC4515	Fibre Channel	432GB	15K	RAID1 ⁴	Nominally a 450GB FC drive
Bottom	C02THJK	SATA	1882GB	7.2K	RAID6(14+2)	Nominally a 2000GB SATA drive

Supported drive types for VMAX20K

¹ 1TB = 2⁴⁰ bytes in this context. Size is also useable capacity after formatting reduction, but excludes reductions caused by RAID protections.

² 1GB = 1000*MB in this context. Size is also useable capacity after formatting reduction, but excludes reductions caused by RAID protections.

³ STEC Enterprise Flash drives must not be used.

⁴ RAID1 is the standard for all new purchases

4.5 Supported VMAX FAST VP configurations

4.5.1 Introduction

Configuring a single-tier storage array (non-FAST VP) is relatively straightforward. One determines the usable capacity required, and then purchases sufficient spindles disks of the required size to meet that requirement (allowing for RAID protection & formatting capacity reductions, as well as hot spares), while also ensuring that disks are purchased in units of one drive per loop.

FAST VP configurations also have the significant addition constraint of needing to have the right relative mix of tiers to match the skew and IOPS requirements of the workload. Hence, the three primary configurations goals that one aims to meet in a FAST VP configuration are:

1. Usable capacity meets arbitrary amount required by customer (desirable outcome)
2. 1 drive per loop Best Practice (constraint)
3. Tier ratios to meet workload and skew requirements (constraint)

It is not possible to meet all three requirements. The second two constraints place significant limits on what useable capacities can be configured. Furthermore, the large size of the SATA disks and the desire to have the majority of the array comprised of this disk type to provide the desired cost savings also results in very large capacity configurations for a relatively modest number of disk spindles, and large steps in capacity between different configurations that meet the two constraints listed above.

As a consequence, FAST VP configurations:

- have very large usable capacities as an entry point
- cannot meet arbitrary usable capacity requirements

4.5.2 CATE standard 3-tier ("Silver") FAST VP configurations

Chris Clayton - 10th April 2013													
VMAX40K	Number of Engines	This Pool Capacity (Eng TB)	EFD drives per loop	SAS drives per loop	SATA drives per loop	EFD drives	SAS drives	SATA drives	Additional Spare Drive Count	Total Drives	% of EFD capacity	% of FC capacity	% of SATA capacity
"New Small" FAST VP config VMAX40K 4-Engine (4 * 96GB) (384 GB total cache) - 100GB 2.5" SSD RAID5(3+1) - 300GB 2.5" 10K SAS RAID1 - 1TB 2.5" SATA RAID6(6+2)	4	216	4	20	6	128	640	192	4*EFD 14*SAS 4*1TB	982	4%	39%	57%
"New Medium" FAST VP config VMAX40K 4-Engine (4 * 96GB) (384 GB total cache) - 100GB 2.5" SSD RAID5(3+1) - 300GB 2.5" 10K SAS RAID1 - 1TB 2.5" SATA RAID6(6+2)	4	315	6	28	9	192	896	288	4*EFD 18*SAS 6*1TB	1404	4%	37%	59%

Appendix A lists further details on the CATE standard 3-tier FAST configurations. The VMAX40K standards should be used for all new installations. The VMAX20K information is only included in Appendix A in support of the existing installed base.

CSBE is recommending deployment of FAST VP-enabled arrays with initially fixed configurations for each capacity (i.e. fixed relative tier ratios) suitable for workloads that are low-to-moderately FAST-friendly to accommodate unknown workloads.

Pros of fixed configurations:

- Much simpler to determine configuration of FAST VP arrays
- Does not require detailed knowledge of target workload

Cons of fixed configurations:

- Fixed capacities. Does not provide the ability to meet arbitrary capacity requirements or workloads with unusual skews. See next section of advice on configuring bespoke configurations to meet extreme skew requirements.

4.5.3 CATE standard 2-tier (“Gold”) FAST VP configurations

VMAX40K 2-tier "GOLD" FAST	Number of Engines	Thin Pool Capacity (Eng TB)	EFD drives per loop	SAS drives per loop	SATA drives per loop	EFD drives	SAS drives	SATA drives	Additional Spare Drive Count	Total Drives	% of EFD capacity	% of FC capacity	% of SATA capacity
"Small Gold" FAST VP config VMAX40K 4-Engine (4 * 192GB) (768GB total cache) - 100GB 2.5" SSD RAID5(3+1) - 300GB 2.5" 10K SAS RAID1	4	189	6	42	0	192	1344	0	4*EFD 28*SAS	1568	6.8%	93.2%	0%

Note that these configurations differ from the 3-tier configurations in that:

- They have no SATA disk
- They have 192GB cache per VMAX40 engine, rather than 96GB (as in 3-tier Silver configurations)

Please see Appendix H for further documentation on 2-tier (Gold) FAST configurations. The above configuration fills a 5-bay VMAX40K configurations and should not be expanded further. The 189TB capacity has been selected to meet failure domain requirements.

4.5.4 Bespoke configurations: Analyzing your specific workload

Use of bespoke configurations is no longer supported by CATE. Use the standard VMAX40K configurations in Appendix A for all new installations.

4.6 Performance analytics tools for user with FAST VP

It is essential that Unisphere for VMAX exists or is rolled out at the same time as FAST VP in order to:

- provide visibility of FAST VP activity within the array
- provide application/server level performance metrics on FAST VP arrays
- provide metrics for FAST VP capacity planning purposes

4.6.1 Unisphere for VMAX

The Unisphere for VMAX product has been introduced by EMC to combining and superseding Symmetrix Performance Advisor (SPA) and Symmetrix Management Console (SMC). Unisphere for VMAX supports the following Symmetrix Enginuity versions:

- Symmetrix VMAX1/VMAX20K series systems running Enginuity 5874 or higher
- Symmetrix VMAX10K/VMAX40K Series systems running Enginuity 5876 or higher

Therefore for new FAST VP deployments on VMAX40K, or array upgrades to Enginuity release 5876, it is recommended Unisphere is rolled out at the same time (along with deployment/upgrade to Solutions Enabler if required) in order to monitor new Enginuity 5876 features.

For further details on Unisphere, please see the **Unisphere for VMAX Solutions Guide** (See References sections for URL)

In addition, for FAST VP-specific Unisphere usage recommendations see the following sections in the present document:

- Section 8: Managing FAST VP
- Section 9: FAST VP Capacity planning
- Appendix G: Using Unisphere with FAST VP

Unisphere presents its data in Realtime, Diagnostic and Historical views.

- Real Time: Collects data between 2 and 5 seconds for a limited group of metrics – the data is available for the previous hour.
- Diagnostic: Collects data every 10 minutes – data can be displayed for the last 1-24 hrs
- Historical: Collects Data every 15 minutes – data can be displayed for between 24hrs – 1yr (albeit with reduced resolution for periods > a few days).

4.7 Restrictions and other comments

4.7.1 Known FAST VP issues and restrictions

As these can vary with time and software version, they are listed in the latest version of the CATE document “**VMAX FAST VP Release Notes**” (see References section for URL). Information on interoperability restrictions between FAST VP and other array functions is also covered in that document.

4.7.2 TimeFinder/Clone

Both the source and target devices of a TimeFinder/Clone session can be managed by FAST VP. However, the source and target will be managed independently, and as such may end up with different extent allocations across tiers.

4.7.3 SRDF

Use of FAST VP with SRDF replicated storage is not currently supported. Support is not currently planned (insufficient demand and technical challenges).

4.7.4 When should FAST VP not be used?

FAST VP is not suitable in the following for certain FAST-unfriendly application workload profiles:

4.7.4.1 Large applications with a workload skew < 80:20.

FAST depends on, and uses the existence of workload skew (see section 1.1) to move the busiest devices (i.e. some of this busy 15% in the above example) to higher performing tiers (EFD and FC/SAS) and idle devices to lower-cost tiers (SATA).

If a large application (multiple TB) has a workload skew of < 80:20 i.e. less than 80% of the IO is going to 20% of the storage, then it will not be suitable for FAST as too much IO will be going to lowest tier SATA.

At present, only EMC can determine the actual workload skew for an array or application using STP data and an internal EMC planning tool (Tier Advisor). Citi has requested this functionality be put into Unisphere, but this is not currently available. Hence, at present, within Citi only discussions with the large application owners about their workload profile can identify an application that might fall into this category. Some warehouse applications may prove to be FAST-unfriendly (skew lower than 80:20).

4.7.4.2 Large applications with irregular workloads:

FAST VP can be configured so that it can “ride through” dips in application workload weekly cycles, including weekends. However, for applications with isolated high I/O peaks which occur either regularly (month/quarter end) or irregularly on a timescale of > 1-2 weeks, then it is possible that the relatively low I/O for the intervening periods will result in the relevant storage being on a lower tier than required to meet the I/O peak requirements.

For small applications (<1TB capacity), do not place the application under FAST VP control, but leave on FC/SAS. The main disadvantage is that it will consume more of the FC/SAS tier than an application under FAST VP control would. Hence, this solution can only be used for a small number of small applications per FAST array. See section 8.13 for further details.

For large (multi-TB) applications with this workload profile, then deployment on a non-FAST array is preferable.

4.7.4.3 Performance sensitive, “read-only” applications with a DR counterpart FAST array

If an application workload profile is mostly reads, then the production array may keep the busiest data in higher tiers through the production host read activity. However, the DR copy is not likely to have any read IO, and in the absence of any replicated write IO, data may be demoted to SATA. This would be an issue following invocation of DR, as performance at the DR site would initially be lower than at the production site.

5 3-Tier FAST VP Configuration Best Practice

For Best Practice for 2-tier FAST VP configurations, please see Appendix H.

5.1 Introduction

This section focuses on Citi Best Practice for configuring thin pools, pool monitoring, thin tiers, FAST policies, time windows and other FAST VP tunable parameters for 3-tier implementations.

Comprehensive documentation on FAST VP configuration and actual FAST VP SYMCLI command syntax is clearly described in the EMC documents referenced in section 2, and is not repeated in this document.

**EMC Solutions Enabler Symmetrix Array Controls Version 7.2 CLI Product Guide
(Section on “Fully Automated Storage Tiering”)
(Section on “Managing time windows” – SE7.3 or later)**

Implementing FAST VP for EMC Symmetrix VMAX Series Arrays Technical Notes

*However, a worked example of setting up FAST VP ***is*** included in Appendix B on the present document.*

5.2 3-tier VP pool and Tier Best Practice

*Existing Citi Best Practice specified in the document “**VMAX Virtual LUN Directives and Best Practice**” (see URL in Documentation Section 2) still applies in FAST VP implementations. That document should be adhered to when configuring VP pools for FAST VP, unless explicitly superseded in the present document.*

For FAST VP one VP pool is required for each tier, built using the supported disk devices and standards listed in the FAST VP release notes. Hence for a 3-Tier FAST VP configuration, a Flash-based pool, a FC or SAS-based pool and a SATA pool are required.

Citi Best Practice is to have:

- ***a single pool for each technology type***
- ***one pool per Tier***

5.2.1 Selecting and setting the Bound Pool

Thin devices have no storage allocated to them when they are created; rather storage is allocated on-demand from a "bound" thin pool.

With FAST VP, the data for a thin device may reside in its bound pool and potentially in one or more other pools. However, new extents (with 5875) are always allocated from the bound pool, even though FAST VP may later subsequently move these extents to other pools. With 5876 EMC have added to capability for extents to be allocated out any pool in the FAST policy that a Storage Group is associated with, rather than just from the bound pool. This is known as "allocation by policy" and is enabled with a command of the form:

```
# symfast -sid <SID> set -control_parms -vp_allocation_by_fp enable
```

Citi Best Practice is for ALL thin devices on a Symmetrix to be initially bound to the FC (VMAX20K) or SAS (VMAX40K) pool. FAST VP can then promote the busiest extents to flash and demote the least busy to SATA, as required.

With 5876 or later, "allocation by policy" should be enabled. This is done at the Symmetrix (FAST VP engine) level.

Note that with 5875 if the bound pool runs out of space, allocations do not automatically come from another pool that a thin device has indirect access to via FAST VP. New extent allocations can only come from the thin pool. Hence it is essential that the bound pool has space available. With 5876 EMC has addressed this restriction thought the "allocation by policy" mechanism mentioned above, although only for storage group that are already under FAST VP control (i.e. associated with a FAST policy)

Note that it is not recommended that the larger SATA pool be used as the bound pool as in a busy production array, FAST VP will not be able to promote this currently active storage (active as it has just been allocated and hence written to) quickly enough to higher tiers, and Write Pending counts can rise very rapidly.

5.2.2 Use of thin device pre-allocation

Citi Best Practice for Virtual Provisioning already recommends a thin device pre-allocation value of 0% (no pre-allocation). However, in a FAST VP environment, this use of pre-allocation is particularly forbidden. Preallocated, but unwritten, extents will show as inactive and as such will be 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 it has been demoted to, and immediate promotion to FC/SAS tier would then be needed.

5.3 3-tier FAST VP policy Best Practice

5.3.1 Protecting applications from performance degradation

There is a concern that early migrants to an array will initially enjoy excessive high-tier residency which is then later reduced as more applications are moved onto this array and pressure on these resources increases.

Furthermore, there is the possibility that a test or development server executing excessive I/O could displace more important production applications from higher tiers.

Both these issues can be addressed by using policies that limit the maximum percentage of an application that can ever reside on the top tiers.

However, attempts to address these issue for an individual application must be balanced with the array-wide requirement not to constrain high IO into the SATA tier (EMC recommend that SATA tier IO should not exceed more that 10% of the total traffic to the array), or overloading the FC/SAS tier.

Best Practice for setting the maximum percentage for each tier is likely to continue to evolve as we gain further experience with FAST VP, but currently it is recommended that the following values be used.

	Flash Tier	FC/SAS Tier	SATA tier
VMAX40K (3-tier “Silver”)			
Production/DEV/UAT policy (VMAX40K)	10%	70%	100%
VMAX20K (3-tier “Silver”)			
Production policy (VMAX20K)	10%	50%	100%
DEV / test policy (VMAX20K) ⁵	10%	50%	100%

Max SG percent settings for policies on 3-tier array with mix of production, development and test servers

For a VMAX20K array supporting either “All production” or “All DEV/UAT”, rather than a mix of the two, then the “Production policy” listed above should be used.

5.4 3-tier FAST VP Time Windows Best Practice

Time windows define when the FAST controller should collect performance information and execute data movement. Time windows are shared between Symmetrix Optimizer and the FAST controller. Time windows are set and managed using: the `syntw` command with SE7.3 & 5875.197 or later.

The following is suggested Best Practice, and instructions for its implementation. This Best Practice is subject to change in the future following feedback from operations teams and as we gain more experience with using the product.

Citi Best Practice Summary

By default, both the performance and data movement windows are open at all times, and these should be left unchanged. However, if there are genuinely extended periods of time when the workloads managed by FAST VP are not active, these time periods can be excluded.

VP Data movement time window:

⁵ Previous VMAX20K Best Practice for DEV/UAT was 2%/45%/100%. The new 10%/70%/100% max SG figures should be used for any new DEV/UAT policies. To minimize work, DEV/UAT policies on existing VMAX20K frames should only be changed to the new standard when an array is re-visited to make other changes, or if the old policy standard is causing any issues.

The data migration time window should be left at the default value (i.e. migrations can occur at all times). If there is concern about the possible impact of data movements occurring during a production workload, then the FAST Relocation Rate can be used to minimize this impact (see later in this section).

Performance time window:

The performance time window is applied system-wide. The default performance time window behavior should be left unchanged (i.e. the performance time window that is open at all times) as multiple applications are active on the array, but active at different times.

If (and only if) there are extended periods of time when there is no activity on the array (which can be determined from ECC Performance Manager or Unisphere data), then you can set the performance time window to exclude these periods. For example, if your Performance Manager data shows that your active period is only 06:00 to 21:00 local time, Monday to Friday, then you can set your performance data collection window to cover only these periods.

5.4.1 Managing Time Windows

Solutions Enabler V7.3 introduced a new command, `symtw`, to manage time windows. This command massively improves on the process and formatting of time windows (previously managed with `symoptmz set time_window`).

Time window management via `symtw` is described clearly in the following EMC documentation:

EMC Solutions Enabler Symmetrix Array Controls Version 7.x CLI Product Guides
Section on “Managing time windows”

However, in brief the minimum setup is:

- 1) Convert any existing time windows to use the `symtw` time window format (if not already done)

```
# symtw convert -sid <SID> -force
```

This need only be done once in the lifetime of the storage array. The command will gracefully exit if it has previously been run.

- 2) Create the default move time window

```
# symtw -sid <SID> add -type MOVE_VP -inclusive -days Sun,Mon,Tue,Wed,Thu,Fri,Sat -  
start_time 00:00 -end_time 24:00  
# symfast -sid <SID> -state -vp list  
Symmetrix ID: XXXXXXXX  
  
FAST VP State                : Enabled  
Reason(s)                    : N/A  
FAST VP Current Activities   : RunningPlan  
FAST VP Perf Time Window     : Open  
FAST VP Move Time Window     : Open
```

Failure to create the move time window will result in the following error:

```
# symfast -sid <SID> list -state  
  
Symmetrix ID: XXXXXXXX  
  
FAST VP State                : Degraded
```

Reason(s)	: FAST VP move time window is not
present or does not extend into the future	
FAST VP Current Activities	: Idle
FAST VP Perf Time Window	: Open
FAST VP Move Time Window	: Closed

The following examples show how to modify the default “always open” time windows for performance data collection with SE7.3 & Enginuity 5875.198 or later. **By default, both the performance and data movement windows are open at all times, and these should be left unchanged. However, if there are genuinely extended periods of time when the workloads managed by FAST VP are not active, these time periods can be excluded.**

Example: If your Performance Manager data shows that your active period is 06:00 to 21:00 local time, Monday to Friday, then you can set your performance data collection window to cover only these periods.

1) Create performance windows that are open from 06:00 to 21:00 Monday to Friday only.

```
# symtw -sid 1380 add -type perf -inclusive -days Sun,Mon,Tue,Wed,Thu,Fri,Sat -start_time
00:00 -end_time 24:00 -nop

# symtw -sid 1380 remove -type perf -inclusive -days Mon,Tue,Wed,Thu,Fri -start_time 00:00 -
end_time 06:00 -nop

# symtw -sid 1380 remove -type perf -inclusive -days Mon,Tue,Wed,Thu,Fri -start_time 21:00 -
end_time 24:00 -nop

# symtw -sid 1380 remove -type perf -inclusive -days Sun,Sat -start_time 00:00 -end_time
24:00 -nop
```

2) Confirm that the time windows are set up correctly. The times shown are for the timezone of the SE server.

```
# symtw list -sid 1380 -type perf
```

Symmetrix ID: 000192601380

Performance Time Windows

Sunday	: None
Monday	: 06:00 - 21:00
Tuesday	: 06:00 - 21:00
Wednesday	: 06:00 - 21:00
Thursday	: 06:00 - 21:00
Friday	: 06:00 - 21:00
Saturday	: None

Exclusive Time Windows (0)

5.5 3-Tier FAST VP control parameters

For Best Practice for 2-tier FAST VP control parameters, please see Appendix H.

These are generally set with the syntax:

```
# symfast -sid SymmID set -control_parms
```

For example

```
# symfast -sid 1380 set -control_parms -vp_data_move_mode Auto
```

See Appendix B for examples of the command lines used to set the FAST VP control parameters to the recommended values.

The following settings are recommended for both 5875 & 5876 (unless stated otherwise).

Control parameter argument	Citi Recommended setting	Is default value?	Comments
-min_perf_period	8 (h)	Yes	See (1) below
-vp_data_move_mode	AUTO	Yes	
-workload_period	120 (h)	No	See (2), (3) & (4) below
-vp_reloc_rate	Initial Values 7 for “new application” arrays or existing arrays being converted from single tier FAST VP. 6 for tech refresh target arrays (a rate of 6 is 17% more aggressive than a rate of 7) If array is mixed usage (new apps and tech refresh) use a rate of 6. Mature array value (existing applications growing organically – no new apps) 5 (A rate of 5 is 40% more aggressive than a rate of 7)	Yes	See (5) below
Pool level reserve capacity. Set with symconfigure (see Appendix B for an example). -pool_resv_cap	Initial Values (Tech Refresh) VMAX20K/40K: EDF pool: 1% VMAX20K: FC Pool: 30% - 35% (see tables below) VMAX40K: SAS Pool: 20% (see tables below) VMAX20K/40K: SATA pool: 1% Reset to 20% after tech refresh migrations complete Mature array values (existing applications growing organically – no new apps) EDF pool: 1% FC/SAS Pool: 20% SATA pool: 1%	No	See (6) below
-vp_allocation_by_fp	Enable (SE7.5 & 5876 or later only)	No	See (7) below

ve good characteristics for the workload. However, when benchmarking new or migrated apps, you need FAST VP to move storage to “final” tiers (based on performance requirements). Hence it is recommended when that benchmarking not commence until the *min_perf_period* has elapsed and FAST VP has able to execute performance-based migrations for the SG⁶. Executing typical I/O during this period is beneficial as it will be used to build up a profile of the application I/O needed to generate migrations plans.

See section 6.3 for further Best Practice on executing meaningful benchmarking in a FAST VP environment.

(2) Choice of workload Analysis Period

From V1.1 of this document, it is being recommended that the workload analysis period be reduced from 7 days (168h) to 5 days (120h) to give FAST VP more opportunity to demote inactive data from the FC/SAS tier to the SATA tier. This is

⁶ Capacity migrations will use metrics that are being collected. So least active will still get moved up to flash and to SATA even though *min_perf_period* has not elapsed.

in response to observations in production of data staying in the FC/SAS tier, rather than being demoted, and leading to high utilization of that tier.

(3) Interaction of WPA and performance time window.

When the Workload Analysis Period (WPA) is set, for example, to 7 days (168 hours), it will take into account 7*24 hours work of *collected* data. If the performance window is set so that there are periods when data is not collected, then a WPA of 7 days will only include data collected when the window is open, and hence will take into account data from more than 1 calendar week ago.

For example, if data is collected Monday-Friday only (24 hours), then on Wednesday 10th, performance data is used from the days marked in light green (only), with a decaying weighting for older data:

Mon	1 st
Tues	2 nd
Wed	3 rd
Thur	4 th
Fri	5 th
Sat	6 th
Sun	7 th
Mon	8 th
Tues	9 th
Wed	10 th

With the Workload Analysis Period set to 1 week, an I/O that has just been received is weighted two times more heavily than an I/O received a week previously.

(4) Closed performance time windows

When the data collection window is closed, any given LUNs performance score will not be changing, as score is determined only from performance information collected when the performance window is open. As scores are not changing, performance-related promotions and demotions are not likely to occur, even if the move window is open (unless there were outstanding performance-related changes due when the performance window closed).

Note: In the future EMC plan to allow different LUNs to use different performance windows. This is needed in situations where a single array is servicing different time zones, and hence different devices are busy at different times.

(5) FAST Relocation Rate (FRR)

The best practice recommendation, for the initial deployment of FAST VP, is to start with a more conservative value for the relocation rate of 7 “new application” arrays . The reason for this 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.

Exceptions:

- For FAST arrays that are target arrays in a tech refresh process, the recommendation is to start with the more aggressive value for the relocation rate of 6. The reason for this is that a high migration rate is needed to ensure migrated data is moved quickly to final tiers after a migration.
- FAST arrays that have a mix of both new applications and tech refresh migrated applications should also have the relocation rate set to 6.

When an array is no longer accepting new applications (“mature array” in the above table), the relocation rate can be set to a more aggressive level of 5. This will allow FAST VP to adjust to small changes in workload more quickly.

(6) Pool Reserve Capacity (PRC)

The PRC reserves a percentage of each pool included in a VP tier for non-FAST VP activities. The purpose of this is to ensure that FAST VP data movements do not fill a thin pool, and subsequently cause (with 5875) a new extent allocation, a result of a host write, to fail. The PRC should be set for each individual pool as per the table above.

When the percentage unallocated space in a thin pool is equal to the PRC, FAST VP will no longer perform data movements into that pool. However, data movements may continue to occur out of the pool to other pools. When the percentage of unallocated space becomes greater than the PRC, FAST VP can begin performing data movements into that pool again.

If host writes or a migration push the FC/SAS pool utilization above (100% - PRC setting for FC/SAS pool), then FAST will actively move data out of the FC/SAS pool to ensure that the free space in the pool is in excess of the PRC limit.

During a tech refresh, the PRC setting needs to be set higher than in a mature array (no new applications being added) to ensure that there is sufficient space in the FC/SAS pool to migrate data to. The following PRC setting are recommended, which should ensure a minimum of 14TB of space being available in the FC/SAS pool as a target for migrations (except for the entry level array), whilst still ensuring that the FC/SAS pool usage does not exceed 85% during the migration itself.

VMAX20K (5875)

FAST VP config	Total Capacity (TB)	FC/SAS tier usable capacity	Tech. refresh. PRC setting	Max FC/SAS pool usage allowed during migration	FC/SAS PRC reserve (TB)	Predicted Minimum Migration Landing space (TB)	FC/SAS space remaining for new allocation (TB)
VMAX20K 2-Engine "Entry"	118	40	35%	85%	14	8	6
VMAX20K 4-Engine "Small"	171	68	35%	85%	24	14	10
VMAX20K 4-Engine "Medium"	297	90	30%	85%	27	14	14

PRC setting during migration for FC pool on VMAX20K running 5875

When a Tech refresh is completed, and no new applications are being added to the array, then for all arrays, the FC PRC setting can be set back to 20%.

As agreed with SAN operations, the 85% limit should be observed for both 5875 & 5876 microcode on VMAX20K. The limit for 5876 code will be increased to 95% once all FAST VP arrays that are migration targets are running 5876.

VMAX40K

As VMAX40K will always be running 5876 and should always have "allocation by policy" enabled, then it has been agreed with CTI SAN Operations that during a migration, the maximum SAS usage can increase temporarily to up to 95%. The Tech Refresh setting of 20% will bring SAS pool utilization back down during the following week, and strive (not guaranteed) to provide the migration landing space capacities shown below.

FAST VP config	Total Capacity (TB)	FC/SAS tier usable capacity	Tech. refresh. PRC setting	Max FC/SAS pool usage allowed during migration	FC/SAS PRC reserve (TB)	Predicted Minimum Migration Landing space (TB)	FC/SAS space remaining for new allocation (TB)
VMAX40K 4-Engine "New Small"	217	84	20%	95%	17	13	4
VMAX40K 4-Engine "New Medium"	317	117	20%	95%	23	18	6

PRC setting during migration for SAS pool on VMAX40K running 5876

(7) Allocation by policy (5876 or later)

What is “allocation by policy”? New allocations for a thin device (resulting from a host write), are satisfied from the most appropriate pool in the FAST VP policy based on available performance metrics. Allocating to the right performance pool saves moving the data at a later time. If no metrics are available a decision tree intelligently provides allocation space removing previous concerns of “bound” pool full status i.e. ***if the bound pool is full, then allocations will take place from one of the other pools⁷, avoiding the issue with 5875 where space might potentially run out in the bound pool causing application failure when there is still space in other pools under FAST VP control.***

VP Allocation by FAST policy: When set to ENABLE, for all thin devices managed by FAST, Enginuity will choose a pool from the policy when making an allocation for the thin device. When set to DISABLE, the allocation will be from the bound pool. The default value for this option is DISABLE.

Note: The -vp_allocation_by_fp option is only supported by Enginuity 5876 and higher. Enginuity 5875 will display N/A for this option with SE7.5.

Note: It will be necessary to manually enable “VP Allocation by FAST policy” after upgrading from 5875 to 5876.

For example:

```
# symfast set -control_parms -vp_allocation_by_fp enable -sid 4261
# symfast list -control_parms -sid 4261

VP Allocation By FAST policy      : Enabled
```

5.6 iAlert : The Intelligent Alerting Script

Pool utilization alerting and FAST VP state alerting should be enabled using the iAlert script described below.

iAlert was developed in response to a major outage in May 2012 where a production Fibre Channel pool on EMC Symmetrix array ran out of capacity. Pool fill alerts were not turned on because there were too many false alerts, making the alerting system nearly useless. Global SAN Operations requested immediate remedy from CATE to prevent similar outages.

The script will also alert if the FAST engine is in DisabledWithError state for an extended period.

5.6.1 Goals

1. Quick to deploy: Should not need additional software on servers
2. Quick to develop: Should be ready in two days
3. Easy to maintain: PERL is already being widely used within SAN Ops, and also the industry in general
4. Reduce (ideally eliminate) false alerts
5. Provide direct alert to migration team in addition to SAN Ops staff

5.6.2 Platform

iAlert is written in PERL. PERL was chosen over Python because at this time PERL is available as a part of the Citi OS install on both Solaris and Linux which are the platforms used for managing the Symmetrix storage arrays. Only core modules are used; again to avoid the task of updating PERL on the servers.

⁷ FAST will try to allocate from a pool that honours compliance with the FAST policy, but can make non-compliant allocations if necessary to ensure an allocation takes place. Allocations will, however, only ever be made from pools in the associated policy.

5.6.3 Eliminating False Alerts

1. SSD and SATA pools are never the BOUND pools and are thus protected from filling up by PRC, and can thus be eliminated from monitoring. SSD pools are designed to run at high utilization and cause false alerts even at 90% threshold
2. Pools that are not oversubscribed can never run out of capacity and can also be eliminated from monitoring
3. Certain pools are driven at high utilization while others are at risk at that utilization level. Thresholds must be configured per pool, rather than at the array level
4. Ideally alert should be generated based on time to pool full, because that is the business metric that drives action such as addition of capacity. For example if a pool is set to alert at 90%, but the consumption rate is 1% a week, it may not be a concern. Also with FAST data is moved out of the pool as well, and in this example if there is negative growth after 90% full there is no need for alert.
5. FAST-VP state alert can sample FAST-VP state across configurable number of script runs to eliminate noise

5.6.4 Architecture

The script runs on every “Control Workstation”. This approach was used as opposed to running it centrally because querying each array using SYMCLI takes time and with hundreds of arrays the script can run for hours delaying the alerts that come out of the script. Centralized script has the advantage of quick updates and troubleshooting and may be possible with parallelizing the PERL code in the future.

5.6.5 Script Location

The latest version of the script is available from

<https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/Symmetrix/ialert.txt>

5.6.6 Installation

The downloaded ialert.txt file must first be renamed to ialert.pl.

The PERL script requires a UNIX OS with PERL 5.8 or better and Solutions Enabler 7.3.1 or better. The script should be installed on every Control Workstation. If there are redundant control workstations the script should only be installed on one. The script will need to be customized before deployment per instructions provided within the script. Once the script is configured, it should be run once to create the per pool threshold configuration file ppconfig.txt. This file should be edited to set up per pool thresholds as desired. The script should be configured to run from CRON at a desired frequency, such as every 30 minutes.

5.6.7 Maintenance

iAlert was written by and is supported by Parul K Jain. Bugs should be reported to Parul K Jain of CATE Storage Engineering, and not to the author of this document. CATE will also release new versions with additional features from time to time. New arrays and pools will automatically show up in ppconfig.txt file. The file should be reviewed periodically to adjust thresholds as needed.

6 Placing applications under FAST VP control

For variations to Best Practice for 2-tier FAST VP data migrations, please see Appendix H.

There are three primary scenarios in which applications may be moved under FAST VP control:

- Existing apps: Tech refresh from a single tier array to a FAST VP array
- New app: new application going live on an existing FAST VP array
- Existing apps: Expansion of an existing “all-FC” VMAX20K array to become a FAST VP array (See Appendix C)

There is some common best practice for all three scenarios (marked as generic), as well as some suggestions that are specific to the particular case. These are all discussed below (or in Appendix C for the third scenario).

Important Note

A device cannot belong to more than one storage group in use by FAST.

Hence, in cluster configurations, a single (shared) storage group (SG) should be used for disks that are to be shared between members. If each host has its own storage group (i.e. the same device may appear in multiple SGs), then it will not be possible to bring all SGs under FAST VP control, and potentially resulting in some devices unique to a storage group not being managed by FAST VP.

Devices that are unique to a host can and should go into the host-specific SG.

6.1 Existing apps: Tech refresh from a single tier array to a FAST VP array

The following is a **high-level overview** of the process for undertaking a tech refresh from a single tier array to a FAST VP array. **Teams executing migrations will need to develop their own more detailed PCMs based on this outline.** Detailed information on how to execute the FAST VP configuration steps can be found in the EMC documentation referenced in Section 2.

Note that migration of applications and placing under FAST VP control should be done in groups, for the reasons described below.

Storage Groups (SGs) should NOT be placed under FAST VP control until after host IO has started on the new storage for that group. This is essential to ensure, that the correct (least active) extents get moved down to SATA when the SG is brought under FAST VP control.

Where possible, the largest applications should be migrated first whilst there is maximum space available in the FC/SAS tier.

- 1) Assess existing workload on array. In this scenario, the workload may already be known. If available, ECC Performance Manager or STP data can be passed to EMC to confirm the workload skew and that the configuration you will be migrating to is suitable for your workload.
- 2) Determine quantity of flash, FC/SAS and SATA disk to configure
 - a. If using a Citi template configuration, it is recommended (but not required) that you ask EMC pre-sales to confirm that it is suitable for your workload. This will require STP or WLA data for the applications to be migrated to be FAST VP array.
 - b. If undertaking a bespoke configuration, EMC will use your workload to determine the appropriate configurations (see Section 4.5.4).
- 3) Install storage array and create thin pools for each
- 4) Enable FAST VP controller
- 5) Optional: Modify default time windows (but only if there is a clear case to exclude certain times – see section 5.4).
- 6) Create storage tiers
- 7) Create storage policies
- 8) Set FAST VP parameters and enable FAST VP on the array
- 9) For each application
 - a. Create TDEVs to be used by application bound to FC/SAS pool
 - b. Create a storage group containing these devices and provision to host with symaccess. In cluster configurations, the same SG should be used when creating views for each cluster member. A device must never be on more than one Storage Group in a FAST VP environment. This is Best Practice with or without FAST VP, but is more important with FAST VP to ensure that you do not have “duplicate” SGs than cannot be brought under FAST VP control.
 - c. Migrate application data to FAST VP array.
 - d. Run Zero Space Reclamation on the devices to which this applications data has been migrated.
 - e. Only associate the storage group with a FAST VP policy **AFTER** the host has cut over to using the storage on the FAST VP array. This is important as the FAST controller needs to be able to collect performance information before it can correctly move data to other tiers. Use the default priority of 2 for all storage groups when associating with a FAST VP policy.

Note that with 5875 initial allocations to TDEVs always come from the bound pool (FC or SAS). The bound pool in a 3-tier configuration will only represent a part of the total storage. Hence, as applications are migrated to the array, care must be taken to ensure that the bound pool does not run out of storage. Note also that even when allocation by policy is enabled on a 5876 array, initial allocations to TDEVs always come from the bound pool until the SG is associated with a FAST policy.

The above limitation and maximum migration size calculation does NOT apply to 2-tier FAST configurations. For 2-tier configurations, there are no limits on the amount of data that can be migrated to the array in a given weekend. See Appendix H for further details.

The maximum size of a migration that can occur to a 3-tier configuration at a given time without exhausting the FC pool can be calculated as follows. As agreed with SAN operations, the 85% limit should be observed for both 5875 & 5876 microcode. The limit for 5876 code will be increased to 95% once all migration target arrays are running 5876.

- a) Determine current FC pool usage. In the example below it is 4797GB (under “Used GB” column).

```
$ symcfg list -sid 425 -pool -thin -detail -gb
    or for SE7.5 or later:
$ symcfg list -sid 425 -pool -thin -detail -gb -mode V73
```

Symmetrix ID: 000192600425

S Y M M E T R I X T H I N P O O L S									
Pool Name	Tech	Dev Emul	Dev Config	Total GBs	Enabled GBs	Used GBs	Used (%)	Subs (%)	Shared GBs
0425_FLASH	EFD	FBA	RAID-5 (7+1)	5136.0	5136.0	1964.6	38	0	0.0
0425_450GB	FC	FBA	RAID-6 (14+2)	32624.6	32624.6	4797.3	14	307	0.0
0425_2TBSATA	SATA	FBA	RAID-6 (14+2)	98114.5	98114.5	17094.3	17	0	0.0
Total GBs				135875.2	135875.2	23856.2	17	74	0.0

- b) Determine maximum FC pool usage allowed to ensure pool is never more than 85% full due to the migration. Multiply the “Enabled GBs” figure for this pool by 0.85.

In the example above, this is 32624GB * 0.85 = 27730GB

The maximum amount of source data that can be migrated in a weekend is the difference between the two figures i.e.

Maximum migration size = “85% full” usage – current FC usage.

In the above example the maximum migration size is 22.9TB (22933GB = 27730 – 4797).

Note that over time as a larger % of the FC tier becomes consumed, then the maximum amount of data that can be migrated to the array in a single weekend will reduce. However, with a PRC setting of 30%-35% (depending on array size – see section 5.5), then for small to large arrays, a minimum 14TB of space should always be available in the FC pool as a landing space for migrations.

Note that if a migration tool is used that can strip out zero space during the migration, this will result in their being less data that needs to actually be written to the FC pool.

Finally, note also that the above max migration size refers to space within the FC pool. In addition, there will be a limit on how much data FAST will be able to move out of the FC pool to other tiers using compliance moves before the next migration weekend. This will typically be 7-8TB. However, performance moves will also be moving data out of the FC pool at a much higher rate.

Pool utilization should be monitored by the migration team (see above command) throughout the migration to ensure that the FC pool utilization does not exceed 85%. If it does, the migration should be terminated, and the target devices unbound to return space consumed by the migration back to the FC pool.

After the migration is completed, run Zero Space Reclamation against the migrated LUNs on the target array, if your migration tool does not already strip out zeros on the fly.

Once the application is running on the new array, FAST VP will relocate allocated extents from FC to other tiers (due to FAST policy compliance driven migrations). Many simultaneous migrations to the FAST VP array (which will cause allocations out of the FC pool) may deplete the FC pool too quickly. Hence, migrations should be done in groups, with applications cutting over to using the new array, being placed under FAST VP control, and the usage of the FC/SAS pool reduced before the next group of applications is migrated.

6.2 New application going live on an existing FAST VP array

The following high-level steps are proposed bringing a new application under FAST VP control. In this case, it is assumed that storage pool, storage tiers and policies have already been created, and FAST VP is configured and enabled, as described earlier in this document.

- 1) Create TDEVs to be used by application bound to FC/SAS pool
 - f. Create a storage group containing these devices. In cluster configurations, the same SG should be used when creating views for each cluster member. **A device must never be on more than one Storage Group in a FAST VP environment.** This is Best Practice with or without FAST VP, but is more important with FAST VP to ensure that you do not have “duplicate” SGs than cannot be brought under FAST VP control.
- 2) Associate storage group with desired FAST VP policy. As there will initially be no storage allocated to these TDEVs, there is no need to wait for the application to be doing IO before putting under FAST VP control.
- 3) Provisioning storage to host

6.3 Application benchmarking (generic)

Great care must be taken when benchmarking applications in an FAST VP environment. This is a common activity in Citi, especially when moving to new storage.


The concern is that unless carefully executed, any benchmarks are meaningless as the responsiveness of the storage is dependent on what tiers the extents that you are accessing are currently on.

Some simple best practice for benchmarking is:

- 1) Benchmarking should be undertaken using the actual application as artificial load generators will usually have a completely different I/O skew compared with a real workload, and hence will result in unreliable results. See Section 1.1 for an explanation of workload skew.
- 2) Initial allocations should be made from the FC/SAS pool. As it is Citi best practice for all TDEVs should be bound to the FC/SAS pool, this will happen automatically. In this way storage will begin its life on medium-to-fast FC/SAS storage.

- 3) A full scale test should be conducted. A small scale test may allow the entire test environment to be resident within higher tiers, potentially giving an artificially low response time. This is because FAST policies which control access to higher tiers apply to all of the storage provisioned to the host (more specifically to storage in the storage group under FAST VP control). See Appendix D for further details of the issue.
- 4) The benchmark workload should be run until the relative percentages of residency of each tier have stabilized. This will give FAST VP an opportunity to promote busy extents (which will have a significant impact on the benchmark) to the appropriate tier to match the I/O requirements of the applications. The final performance figures after stabilization have been run should be use. The SAN support team can determine when stabilization has been reached using Unisphere.

Unisphere: Symmetrix Array → <SID> → Performance → Analyze → Historical (or Diagnostic) → Storage Group → SG by Tier → Plot Allocated Capacity.



Using Unisphere historical/diagnostic view

7 Chargeback of FAST VP managed storage

Storage under FAST VP should be charged back using one of the following two models:

1. **Provisioned capacity model:** Storage capacity (GB) provisioned to storage (i.e. the sum of usable capacity of thin devices, or TDEVs). Chargeback is at a fixed rate per GB of provisioned capacity.
2. **Combined provisioned capacity and policy model:** A combination of:
 - a) Storage provisioned to storage (sum of provisioned capacity of thin devices, or TDEVs)
 - b) The policy associated with the storage group. Potentially different policies could be charged at different rates. Initially all storage would be charged at a single rate, regardless of policy, but the potential would be there to have differing rates in the future.

Chargeback would be against GB in a policy rather than a specific storage tier i.e. based on a right to use a particular amount of a given tier, not actual usage. It is possible to accommodate chargeback on GB within a policy via ECC 6.1 UB9.

At the time of writing, Global SAN operations have elected that FAST VP storage should be charged back using the “Provisioned capacity” model.

7.1 Provisioned capacity model

It should be possible to continue to use existing chargeback mechanisms used for **thinly provisioned storage** without any change for FAST VP managed storage.

The only exception will be when the chargeback rate is determined by drilling down to see what disk type is used to build the underlying thin pool. FAST VP TDEVs are spread over multiple pools, and depending on exactly how you obtain the information on the underlying hardware type (assuming you do), then you may or may not need to modify your chargeback rate determination mechanism.

If instead you use a table that currently maps each array to an array type (and hence chargeback rate), then this mechanism will also work with FAST VP storage.

Contact the author if you wish to discuss possible changes required to your chargeback mechanism.

7.2 Combined capacity & policy chargeback reporting via ECC

SAN operations have indicated that this type of reporting is not required at the present time, and that there should be one chargeback rate regardless of storage type (thick, thin, FAST VP etc).

As a consequence, the material previously in this sub-section has now been moved to Appendix G for preservation for potential future implementation should directives on chargeback change.

7.3 Chargeback reporting via local scripts

If local scripts are still used to compile chargeback data (e.g. as in EMEA), then customers should be charged based on TDEV capacity provisioned to the host.

For combined capacity & policy chargeback reporting (if required in the future), see Appendix E for suggestions on how to get policy information via scripts.

7.4 SANscreen issues

SANscreen 5.1.2 “Volumes” reporting does not support volumes simultaneously resident in multi pools.

SANscreen 6.1.2 is now reporting the BOUND pool (as opposed to a seemingly random pool as it did in 5.1.2). This is the best one can expect, and what ECC also does for multi-pool-resident FAST VP volumes.

8 Managing FAST VP

Full details of the commands to manage FAST VP can be found in the following EMC documentation:

- EMC Solutions Enabler Symmetrix Array Controls Version 7.x CLI Product Guide
 - o Section on “Fully Automated Storage Tiering”
 - o Section on “Managing time windows” (SE7.3 or later)
- Implementing FAST VP for EMC Symmetrix VMAX Series Arrays Technical Notes

See section 2 for URLs for these documents.

This section covered selected FAST VP topics that are either not covered in the EMC documentation, or are of particular note.

Appendix G and the Unisphere for VMAX solutions guide provides additional suggestions on how Unisphere for VMAX (U4V) can be used to help manage FAST VP arrays.

8.1 Determining tier residency

Tier residency refers to the amount of storage in each tier that is used by a storage group. This can be determined in a number of ways.

8.1.1 Current tier residency via CLI

8.1.1.1 At the Storage Group level

Current usage (in GB) is shown in the “FAST SG Usage” column.

```
# symfast -sid <SID> list -association -demand -sg <SG>
```

For example

```
$ symfast -sid 1380 list -association -demand -sg fastvp02_sangs11
```

```
Symmetrix Id      : 000192601380
```

```
Policy Name       : DEV-UAT
```

```
Storage Group     : fastvp02_sangs11
```

```
Priority          : 3
```

```
Tiers (3)
```

```
{
```

----- Logical Capacities (GB) -----						
Name	Type	Target Prot	Max SG Percent	Max SG Demand	FAST SG Usage	Growth
thin_EFD_400*	VP	R5 (7+1)	4	13	9	+4
thin_FC_450GB	VP	R6 (14+2)	34	115	2	+113
thin_SATA_1TB	VP	R6 (14+2)	100	337	326	+11
Total				-----	-----	

8.1.1.2 At the device level

Current tier residency for individual devices within an SG can be determined with the following command. However, note that the pool name rather than the tier is shown.

```
# symcfg list -tdev -bound -detail -sg <SG> -sid <SID>
```

For example,

```
symcfg list -tdev -bound -detail -sg fastvp02_sangs11 -sid 1380
```

Symmetrix ID: 000192601380

Enabled Capacity (Tracks) : 2602844880

Bound Capacity (Tracks) : 5524200

S Y M M E T R I X T H I N D E V I C E S									

		Flags	Total	Pool	Pool		Total		
Sym	Pool Name	EM	Tracks	Subs (%)	Allocated Tracks (%)		Written Tracks (%)		Status

0331	glab1_g14	F.	552420	0	552432	100	390932	71	Bound
0332	glab1_g14	FX	552420	0	542112	98	245770	44	Bound
	glab1_g10	--	-	-	5760	1	-	-	-
	glab1_g13	--	-	-	4560	1	-	-	-
0333	glab1_g14	FX	552420	0	521472	94	245770	44	Bound
	glab1_g10	--	-	-	23040	4	-	-	-
	glab1_g13	--	-	-	7920	1	-	-	-
0334	glab1_g14	FX	552420	0	517344	94	391382	71	Bound
	glab1_g10	--	-	-	28800	5	-	-	-
	glab1_g13	--	-	-	6288	1	-	-	-

[Rest of output deleted for brevity]

or for SE7.5 or later

Symmetrix ID: 000192601380

Enabled Capacity (Tracks) : 2602844880

Bound Capacity (Tracks) : 5524200


S Y M M E T R I X T H I N D E V I C E S											

	Bound	Flags	Total	Pool	Pool		Total		Compressed		
Sym	Pool Name	ESPT	Tracks	Subs	Allocated		Written		Size/Ratio		
				(%)	Tracks	(%)	Tracks	(%)	Tracks	(%)	
0331	glabl_g13	F..B	552420	0	245784	44	245771	44	245784	0	
	glabl_g10	-.--	-	-	120	0	-	-	120	0	
0332	glabl_g13	F..B	552420	0	245904	45	245771	44	245904	0	
0333	glabl_g13	F..B	552420	0	552420	100	552391	100	552420	0	
0334	glabl_g13	F..B	552420	0	245712	44	245771	44	245712	0	
	glabl_g10	-.--	-	-	192	0	-	-	192	0	

[Rest of output deleted for brevity]

8.1.2 Current tier residency via Unisphere

A high-level tier demand report is available via the FAST Dashboard.
Symmetrix Array → <SID> → Storage → FAST →




The graph panel shows representations of the used and free space for the ENTIRE tier against ALL policies.

The chart includes markers for the following metrics:

- Max SG Demand = Combined Maximum that ALL storage groups/Policies could move into the tier.
- Available to FAST = The amount of storage available for FAST operations on the tier.

A slightly more detailed report can be found within the FAST Policy details windows, which can be reached via:
Symmetrix Array → <SID> → Storage → FAST → Double click on any of the policies




Unisphere also provides a FAST compliance report which provides similar metrics to

```
# symfast -sid <sid> list -association -demand -sg <SG>
```

Where: Symmetrix Array → <SID> → Storage → Storage Group → double click on the storage group to view details.

Current usage (in GB) is shown in the “FAST SG Used” column and space available for growth shown in “Growth” column.




8.1.3 Trending via Unisphere

8.1.3.1 At the storage group level

Tier residency is visible at the storage group level within the predefined FAST by Storage Group dashboard and can be displayed in diagnostic or historical format via:


Symmetrix Array → <SID> → Performance → Monitor → EMC Defined Dashboards → FAST by Storage Group → Select Desired SG



Several other useful canned graphs/reports are provided for side by side comparison which can assist event/symptom correlation:

Hence, one can:

1. Identify if there have been response time changes
2. Easily correlate changes in tier residency (if any) that might have impacted response time



8.1.3.2 At the pool level

Overall Pool capacity & Pool IO can be plotted via:

Symmetrix Array → <SID> → Performance → Analyze → Historical (or Diagnostic) → Virtual Pool tier → Select pool/Pools → Plot required metrics



Note: Unisphere does not display a high-level **combined** pool versus host view for IO/MB per second, but both can be plotted separately and displayed via a personalised dashboard. However, you need to know the names of all storage groups under FAST control to plot a combined Host IO/Sec graph, and you can only plot against a single pool at a time. This is not that useful if wishing to see an overall view specific to FAST especially on an array with potentially several hundred Storage groups.

8.2 Monitoring data movement

With the CLI, data movement between tiers can most easily be detected at the Storage Group (SG) level by repeated running of the command

```
# symfast -sid <SID> list -association -demand -sg <SG> -mb
```

and observing changes in tier residency. This reports to the nearest MB. The Growth column does NOT report changes in tier residency, but the amount that usage of that tier can grow before the SG reaches the maximum associated with the FAST VP policy associated with the SG.


For individual devices, use the command (described earlier)

```
# symcfg list -tdev -bound -detail -sg <SG> -sid <SID>
```

In Unisphere, graphs of ingress tracks (tracks entering tier) and egress tracks (tracks exiting tier) metrics can be plotted for either a specific thin tier or as a combined/side by side view.

Where: Symmetrix Array → <SID> → Performance → Analyze → Historical (or Diagnostic) → Virtual Pool Tier → Plot required metrics

Ingress/egress – combined by tier



Ingress/egress – Single Tier



8.3 Monitoring performance

Unisphere can be used to track host IOPS and throughput, as well as response time for the Storage Group. These can also be tracked against tier residency at the Storage Group level. See section 8.1.3 for example screenshots.

8.4 Troubleshooting lack of data movement

There are a number of reasons why FAST VP may not move data between tiers.

- FAST VP does not deem it necessary to move any extents between tiers i.e. they are all on the most appropriate tier already.
- FAST VP may not be enabled, or the FAST VP data movement window may not currently be open. Check this with the command:

```
# symfast list -state -sid 1380 -vp

Symmetrix ID: 000192601380

FAST VP State           : Enabled
Reason(s)               : N/A
FAST VP Current Activities : RunningPlan
FAST VP Perf Time Window : Open
FAST VP Move Time Window : Open
```

The above output show that FAST VP is enabled and the “Move Time Window” is Open. The “RunningPlan” state also shows that FAST VP is actively moving data now.

- c) The devices in the Storage Group of interest have been pinned (pinning stops FAST from moving devices). Pinned devices can be listed (at the array level) with the command:

```
# symdev list -pinned -sid 1380
```

The symdev command supports use of the `-sg` qualifier, so that the pinned status of devices in just a specific SG can be determined.

- d) The `min_perf_time` may not have elapsed since the Storage Group of interest was brought under FAST VP control.

8.5 Audit logging

FAST VP is continuously moving small chunks of storage between tiers in response to changing IO needs. Individual movements are not tracked.

Human initiated FAST-related activities are, however, logged in the Symmetrix Audit log.

```
# symaudit list -sid <SID> -function_class FAST, OPTMZR, Migrate -start_date MM/DD/YY
```

will show a summary FAST-related activities such as:

- Create/delete a tier; add/remove disk groups

- Create/delete a policy; add/remove tiers; rename a policy; change a policy tier capacity
- Associate/disassociate a storage group to/from a policy; change the priority of a storage group association
- Enable/disable the FAST controller; approve/decline a plan
- Set FAST control parameters

For example

```
# symaudit list -sid 1380 -function_class FAST, OPTMZR, Migrate -start_date 09/02/10
```

A U D I T L O G D A T A

Symmetrix ID : 000192601380

Record Number	Date	Time	Application	Host	Function Class	Action Code
735	11/08/10	04:35:23	SYMFAST	namgdse14	Fast	Create
736	11/08/10	04:35:23	SYMFAST	namgdse14	Fast	Create
1346	11/09/10	16:12:07	SYMFAST	namgdse14	Fast	Create
1347	11/09/10	16:12:07	SYMFAST	namgdse14	Fast	Create
1350	11/09/10	16:14:23	SYMFAST	namgdse14	Fast	Add
1351	11/09/10	16:14:23	SYMFAST	namgdse14	Fast	Add
1352	11/09/10	16:14:59	SYMFAST	namgdse14	Fast	Add
1353	11/09/10	16:14:59	SYMFAST	namgdse14	Fast	Add
1354	11/09/10	16:15:18	SYMFAST	namgdse14	Fast	Add
1355	11/09/10	16:15:18	SYMFAST	namgdse14	Fast	Add
1356	11/09/10	16:16:57	SYMFAST	namgdse14	Fast	Assoc
1357	11/09/10	16:16:57	SYMFAST	namgdse14	Fast	Assoc

Further information on an individual record (using number from above output) can be obtained as follows:

```
# symaudit list -sid <SID> -record_num <rec_num> -n 1 -v
```

This more detailed output includes (but is not limited to):

1. Time
2. Hostname
3. Username
4. Text, with information on the command line issued

An example of the informational Text is shown below:

```
Starting a FAST 'ADD_SYMMETRIX_TIER' operation. Symm=000192601380, Policy_Name=Silver10, Symmetrix_Tier=[FC_300GB_10, 40%SG].
```

Also, if configured as per Citi standards, the file /var/symapi/logs/symevents.logN will also contain symaudit events, with the same information as above.

Further details on FAST VP administration event codes can be found in ***the Implementing FAST VP for EMC Symmetrix VMAX Series Arrays Technical Notes*** Tech Note (see references in section 2 of this document).

8.6 Dealing with customer storage performance complaints

As FAST VP will be pro-actively moving application data between tiers, there is potential for customers to suspect that FAST VP is causing them a performance issue. The following advice is offered to deal with such issues.

1. Use Unisphere Historical or diagnostic data to check response time, host IOPS and tier residency before and after the any alleged performance issues. ECC Performance Manager data may also be of assistance for tracking performance over time.
2. When an application owner is concerned that FAST VP is changing the I/O response time for his/her application, device pinning can be used temporarily to prevent further movement during the investigation. For example to pin all devices in an SG:

```
# symmsg -sid <SID> -sg <SG> pin
```

Device pinning can also be done at the device level

```
# symdev -sid <SID> pin -range 1EA:1EF
```

or for an arbitrary list of devices listed in a file

```
# symdev -sid <SID> pin -file pin_devices.txt
```

Only pin devices when performance is currently acceptable to the customer.

3. One can permanently remove an application's storage from FAST VP control by disassociating a storage group from a FAST VP policy. Note that this will erase all performance metrics collected by FAST VP of this Storage Group. Hence, a Storage Group should only be disassociated if there is no intention to place it under FAST VP control again in the near future, or if you want to deliberately erase previous performance data.

```
# symfast -sid <SID> -fp_name <Policy name> -sg <SG> disassociate
```

In addition, Virtual LUN can be used to non-disruptively, manually move some or all data associated with a storage group to a specific tier (pool) e.g. to a FC/SAS pool. See next section for further details of using Virtual LUN.

8.7 Rollback and other manual data movement

With FAST VP there is no concept of rolling back changes that it has made. The only option to “reverse” an unwanted change is to use Virtual LUN to migrate either all the devices in the Storage Group, or just specific devices, into the FC/SAS pool. After the migration, the migrated thin devices will be wholly resident on and bound to the FC/SAS pool.

Migrating using Virtual LUN is a non-disruptively process (storage is on-line to host at all times).

Virtual LUN operations are described in the GDSE document “**VMAX Virtual LUN Directives and Best Practice**”. See References section for URL. A recipe to move a Storage Group under FAST VP control is shown below.

8.7.1 Using Virtual LUN with source under FAST VP control

If Virtual LUN is being used to migrate thin devices in a storage group under FAST VP control to a specific FAST VP tier, then it is necessary first to identify the appropriate target pool before executing the migration.

1) Identify the storage tiers associated with this storage group

```
# symfast -sid 1380 show -association -sg fastvp02_sangsl1
```

```
Tiers(3)
{
-----
Tier Name                                Type      Max SG O   Target      Flgs
                                Percent C Tech Protection C
-----
thin_EFD_400GB                        VP          10 I EFD   RAID-5(7+1) .
thin_FC_600GB                          VP          50 I FC    RAID-6(14+2) .
thin_SATA_1TB                          VP         100 I SATA  RAID-6(14+2) .
}
```

The Flgs column appears with SE7.5 and later only.

2) Identify the pool(s) that comprise the desired target tier. For example, if one wanted to migrate this storage group to thin_SATA_2TB

```
# symtier show -tier_name thin_SATA_2TB -sid 1380
```

```
Thin Pools(1)
{
-----
Logical Capacities (GB)
Pool Name    Dev    Enabled    Free    Used    Full
Emul
-----
glab1_g14   FBA      98115     94942    3173     3
Total
98115     94942    3173
}
```

Here we see that the pool associated with tier thin_SATA_2TB is glab1_g14.

3) Now that the pool to which the migration is to be conducted has been identified, once need to check that this pool has free capacity

```
# symcfg list -sid 1380 -pool -thin
```

```

S Y M M E T R I X   P O O L S
-----
Pool      Ty  p Dev  Dev    Total  Enabled  Used    Free    Full    Sta
Name      p  e Emul Config Tracks Tracks Tracks Tracks (%)   t
-----
glab1_g10 TH FBA RAID-5(7+1) 167934768 167934768 391992 167542776 0 Ena
glab1_g11 TH FBA RAID-6(14+2) 118199520 118199520 4788960 113410560 4 Ena
glab1_g13 TH FBA RAID-6(14+2) 709197120 709197120 44734488 664462632 6 Ena
glab1_g14 TH FBA RAID-6(14+2) 1607513472 1607513472 52184736 1555328736 3 Ena
Total
Tracks      2602844880 2602844880 102100176 2500744704 3

```

or for SE7.5 or later

```
# symcfg list -sid 1380 -pool -thin 1380
```

S Y M M E T R I X P O O L S							
Pool Name	Flags PTECSL	Dev Config	Usable Tracks	Free Tracks	Used Tracks	Full (%)	Comp (%)
glab1_g10	TEF-EI	RAID-5 (7+1)	167934768	167219580	715188	0	0
glab1_g11	TFF-EI	RAID-6 (14+2)	118199520	4533084	113666436	96	0
glab1_g13	TFF-EI	RAID-6 (14+2)	709197120	595978224	113218896	15	0
glab1_g14	TSF-EI	RAID-6 (14+2)	1607513472	1578995196	28518276	1	0
Total Tracks			2602844880	2346726084	256118796	9	0

4) Optionally pin the devices in the storage group.

If you wish to ensure that the thin devices remain in the target pool in their entirety after the migration is completed, and that FAST VP does not immediately start to relocated thin pool extents to pools in higher or lower tiers, then you should “pin” the storage group before executing the migration. Pinning stops FAST VP from relocating any extents used by pinned devices.

```
# symmsg -sid 1380 -sg fastvp02_sangs11 pin -nop
```

Note that it is also possible to pin individual devices with the `symdev` command.

5) Finally, execute the migration. Note the use of the `-tgt_pool` flag. The process will also bind that thin device to that pool, if not already bound to the target pool.

```
# symmigrate -name fastvp02_sangs11 -sg fastvp02_sangs11 -tgt_pool -pool glab1_g14 -sid 1380 establish -nop
```

Note that migrations to a thin pool can also be executed on source devices listed in a file, or in a device group.

Additional functionality with 5876 & SE7.5 (or later)

In 5875, use of the Virtual LUN migration software in a FAST VP environment only allowed migration of ALL storage in the LUN to a target pool (tier), and also resulted in that LUN being rebound to the target tier, which was not always desirable. In 5876, it is now possible to only migrate storage owned by a LUN that is in a specific pool (tier) to a target pool. For example, one could migrate just those parts on the LUN that are in the SATA pool to the FC/SAS pool.

Furthermore, migrations of just a single tier of a LUN do not result in the LUN being rebound to the target pool.

An example of source pool designation is shown below. In this example, the addition of the `-src_pool` option ensures that only the extents in storage group `fastvp02_sangs11` that are in pool `glab1_g13` are migrated to target pool `glab1_g14`. Furthermore, use of the `src_pool` option in this migration does not result in the devices in storage group `fastvp02_sangs11` being rebound to the target pool `glab1_g14`.

```
# symmigrate -name fastvp02_sangs11 -sg fastvp02_sangs11 -src_pool glab1_g13 -tgt_pool -pool glab1_g14 -sid 1380 establish -nop
```

8.7.1.1 Sources not under FAST VP management

For migration of devices currently outside of FAST VP management, then one simply needs to select the target pool, check that it has free capacity and execute the `symmigrate` command as shown in the previous example. There is no need to pin devices, or to identify relationships between tiers and pools.

8.8 Dealing with pool exhaustion issues in 5875

FAST VP only supports a single tier per pool (c.w. with HDS who support multi-tier pools in their HDT product). Furthermore, for 5875, new allocations are made from one specific pool, known as the bound pool. Hence, multiple pools are required, potentially leading to pool exhaustion when there is still space in other pools which are accessible via the same policy, as there is presently no automatic allocation out of the other tiers within the policy.

EMC have addressed this issue in 5876 via the “allocation by policy” feature described earlier in this document. For 5875, however, several workarounds are possible:

- Reducing future demand on bound pool

- It is possible to manually **rebind** thin devices (TDEVs) to another pool if the original bound pool is becoming full. All future allocations will come out of the new bound pool. Note that the approach does not involve any movement of existing data. Threshold monitoring of thin pool capacity is already part of Citi Best Practice. The FC (bound) pool should be monitored. FAST VP will never move data into this pool when more than 80% used when using Citi Best Practice of setting the `pool_resv_cap` to 20% (see Section 5.5). However, you can expect FAST VP to fill the pool up to 80% utilization.

The pool that a TDEV is bound to can be changed with the `symconfigure` command syntax:

```
# symconfigure -sid <SID> -cmd "rebind tdev <SymDevName[:SymDevName] | in SG SgName> to pool PoolName;" commit -nop
```

For example to rebind all devices in storage group “fastvp02_sangs11” to pool `glab1_g13` on Symm 1380:

```
# symconfigure -sid 1380 -cmd "rebind tdev in SG fastvp02_sangs11 to pool glab1_g13;" commit -nop
```

- Reducing current usage of bound pool

- One can use manual **Virtual LUN migration** both to change the bound pool and move existing allocations out of the current bound pool to a new pool to free up space in the original bound pool. Note that all future allocations from thin pools for migrated devices will now come from the new bound pool.
- After migrating storage to a new pool (to free up storage in the source pool), one optionally can rebind the devices back to the original pool to ensure that any future allocations come from the original pool.

Note that if you use Virtual LUN to migrate (and rebind) to a SATA pool, then all new allocations will be from SATA pool, which may have a negative impact on performance of active applications.

Citi did consider making the large SATA pool the bound pool (where all new allocations come from) and rely on FAST VP to promote active extents to higher tiers. While this would reduce the probability of running out of space in the bound pool, it would also:

- Place additional load on the array as a significant % of new allocations will need to be on FC/**SAS**.

- Give initially poor performance until FAST VP is able to promote busy newly allocated extents.

The above are considered to be emergency processes as they potentially result in TDEVs that are not bound to the FC/SAS pool.

8.8.1.1 Pool utilization alerting

Pool alerting should be set up using the CATE iAlert script (see Appendix I)

Note that alerting should **not** be set up of the EFD tier as very high utilization is expected and desirable, and there is no risk with high utilization, as devices should not be bound to this pool.

8.8.1.2 Other Pool exhaustion issues

No checks are made for oversubscription at the time when a storage group is associated with a policy. A thin pool in a tier in an associated FAST policy becomes a shared resource; its capacity is no longer used only by the bound thin devices but also by thin devices under FAST control.

Hence, even a pool that is not oversubscribed might potentially become full due to a combination of allocations for bound devices and FAST moving chunks into the pool.

To address this, EMC have mechanism know as the Pool Reserve Capacity (PRC). This is percentage of the capacity of each thin pool that will be reserved for non-FAST activities (i.e. used for allocation of new extents). If the free space in a given thin pool (as a percentage of pool-enabled capacity) falls below the PRC, the FAST controller does not move any more chunks into that pool. To move any new chunks to the pool, the FAST controller must first move some chunks from that pool to another pool to free up space. See Section 5.4 for Best Practice for setting the PRC value for each pool.

8.9 Backup/restore of FAST VP configuration

Backups of all FAST data – tiers, policies, associations – are performed nightly on the Service Processor and can be restored by EMC personnel in the case of catastrophic failure. The backup/restore facility is not accessible via symcli.

8.10 Using Zero Space Reclaim in a FAST VP environment

At the time of writing, Zero Space Reclamation (ZSR) does function on FAST VP devices that have some extents currently being migrated.

In order to run ZSR on a storage group (SG) that is under FAST VP control, it is necessary to pin the devices in the SG for the duration of the running of ZSR.

For example, if SG fastvp02_sangs11 contains devices 0331:033A

- 1) Pin the devices in this SG either by SG name or device range.

```
# symmsg -sid 1380 -sg fastvp02_sangs11 pin
```

```
or
# symdev -sid 1380 -range 0331:033A pin -nop
```

2) Run zero space reclaim against these devices

```
# symconfigure -sid 1380 -cmd "free tdev 0331:033A start_cyl=0 end_cyl=last_cyl
type=zero;" commit -nop
```

3) Monitor the state of the reclaiming devices

```
# symcfg -sid 1380 -devs 0331:033A verify -tdev -bound -i 60
```

4) When the state of these devices has changed from reclaiming to bound, unpin the devices.

```
# symmsg -sid 1380 -sg fastvp02_sangs11 unpin
or
# symdev -sid 1380 -range 0331:033A unpin -nop
```

8.11 Copy pace management

The system-wide aggressiveness of FAST VP copying is controlled using the FAST VP parameter `-vp_reloc_rate` discussed earlier in this document.

If a storage group needs to have migrations suspended temporarily for it, best practice is to pin the devices, as described elsewhere in this document.

If necessary, the copy pace use by FAST and Virtual LUN (VLUN) can also be set for individual devices or storage groups. This is achieved using the `symqos symcli` command and the QoS type "VLUN".

```
# symqos set VLUN pace <value> -sg <SG> -sid <SID>
```

This is not recommended Best Practice, but the capability is mentioned here should a situation arise where a particular group needs to have it copy pace reduced (but not completely suspended).

For example, to reduce the copy pace from the default value of 0 (fastest) to 3:

```
# symqos set VLUN pace 3 -sg fastvp02_sangs11 -sid 1380
# symqos list -sg fastvp02_sangs11 -sid 1380
```

```
Storage Group (SG) Name: fastvp02_sangs11
SG's Symmetrix ID      : 000192601380
```

Device Name		Copy Pace				
Sym Physical	Config	BCV	RDF	MIR	CLN	VLN
331 Not Visible	TDEV	0	0	0	0	3
332 Not Visible	TDEV	0	0	0	0	3

333 Not Visible	TDEV	0	0	0	0	3
334 Not Visible	TDEV	0	0	0	0	3
335 Not Visible	TDEV	0	0	0	0	3
336 Not Visible	TDEV	0	0	0	0	3
337 Not Visible	TDEV	0	0	0	0	3
338 Not Visible	TDEV	0	0	0	0	3
339 Not Visible	TDEV	0	0	0	0	3
33A Not Visible	TDEV	0	0	0	0	3

8.12 Protecting early migrants from artificially high performance

8.12.1 Problem statement

With all shared resources, there is the potential for the first users to have greater access to the resources initially than they will have once there is more competition for them when other applications have moved onto the array.

FAST VP promotes storage associated with a given application to higher tiers based in part on relative IOPS compared with other application workloads. Hence, when there are fewer applications, a given application has more change of getting access to the top tiers.

8.12.2 Best practice

Policies which place an artificial limit on the maximum tier residency for the higher tiers can help limit maximum access to the top tiers from day 1. In section 5.3, the default Citi Production policy is defined as 10% EFD / 50% FC/SAS / 100% SATA Hence, early adopters will be limited in terms of the maximum flash usage.

8.13 Dealing with irregular workloads

FAST VP can be configured so that it can “ride through” dips in application workload weekly cycles, including weekends.

However, for applications with isolated high I/O peaks with occur either regularly (month/quarter end) or irregularly on a timescale of > 1-2 weeks, then it is possible that the relatively low I/O for the intervening periods will result in the relevant storage being on a lower tier than required to meet the I/O peak requirements.

For such applications there are several possible solutions:

- **Recommendation.** Do not place the application under FAST VP control, but leave on FC/SAS.
 - Safest option. Main disadvantage is that it will consume more of the FC/SAS tier than an application under FAST VP control would.
- Pin the devices following one of the isolated high I/O peaks after FAST VP has been able to move the storage to the correct tiers.
 - Main disadvantage is that FAST VP will not be able to promote SATA storage that becomes active after the peak.

Other more manual alternatives would include using VLUN to move entire TDEVs to specific tiers in advance of the I/O peak, but this is not considered an acceptable solution from an operational perspective.

8.14 Loss of the Service Processor

The FAST controller runs on the Service Processor (SP). If the SP dies:

1. Production of FAST VP plans for compliance moves will cease immediately. Any compliance data movement in flight will continue as these are controlled by the DAs.
2. No new promotion and demotion thresholds will be created by the SP (usually created every 10 minutes).
 - a. The last generated promotion and demotion thresholds will continue to be used for four hours as the DAs determine required performance-driven moves.
 - b. After four hours, a conservative set of backup thresholds will be used. These are more conservative than the thresholds created by SP using actual performance data and reduce the risk of FAST VP making inappropriate performance movement. As a consequence, there will be degradation in the quality of the moves that FAST VP is making.
 - c. After four days, all FAST VP performance moves will stop if the SP has not been replaced.

Note that performance data is not kept on the SP any more and so changing the SP does not result in the loss of performance metrics.

The FAST State will appear as degraded in the output of a `symfast list -state` command. The reason for the degraded state will also be shown, and in this case would be *“Timed out attempting to communicate with the FAST controller”*

With SE7.5 or higher, it is possible to set up symlive alerting to trigger if the FAST VP controller goes into degraded state. An alert with the event code 1508 is generated that states both the current and previous state for the FAST VP engine.

```
[fmt=evt] [evtid=1508] [date=2012-10-22T06:19:30] [symid=000192601380] [sev=info]  
= The FAST VP controller has switched to state: Degraded (was: Enabled).
```

This alert is available with both 5875 & 5876 microcode.

See the CATE Solutions Enabler Solutions Guide on details of how to set up event alerting.

8.15 Excessive Write Pending

In a situation where there is excessive IO to storage on SATA, it is possible for the array Write Pending (WP) count to increase as the backend is unable to keep up.

If the write pending count reaches 60 percent of the write pending limit, FAST VP data movements will stop. As the write pending count decreases below this level, data movements will automatically restart.

If your Symmetrix alerting is correctly setup as per the Best Practice in the CATE Solutions Enabler Solution Guide, then you will also pick up WP alerts via the symlive. Symmetrix Event Code 1243 is for “Write pending data is now using nn percent of the cache”. However, at the time of writing, the first alert occurs when nn = 90%, which is not helpful for identifying WP issues in a timely manner.

The current WP count and %WP can also be observed in Unisphere (see Section 15.1 for an example). If it is > 60%, then you can infer that FAST VP data movement has paused.

Note that when data movement is paused due to WP levels, `symfast list -state` will still report FAST VP state as Enabled, albeit as

FAST VP Current Activities: Idle

8.16 FAST VP EFD drive erasure

Unlike FC, SAS or SATA drives, flash drives cannot be erased by simply overwriting existing data. This issue is understood by EMC and their SecureErase leverages the firmware of the flash drives (both the STEC and Samsung) to execute the erasure of data for them. SecureErase can be used for both Single drive and Full frame erasure.

The Engenuity minimum levels that fully support disk erase for the Samsung disks are 5875.249.188 & 5876.82.57.

At the time of writing (Sept 2012) CTI is still determining if it wishes to adopt SecureErase as the global standard for Symmetrix drive erasure for single disk erasure. The alternative is to use the CATE Host Security Engineering WipeDrive PRO standard.

8.17 Decommissioning storage under FAST VP control

The process for decommissioning storage from a host will be very similar to your existing process for thin devices, with the additional step of disassociating the storage group from its FAST VP policy. As with all thin devices that are being reused, the devices must be unbound from and then rebound to the FC/SAS pool. An example of unbinding and rebinding is shown in the following sub-section.

- 1) Disassociate Storage Group from policy
- 2) Unmasked devices from host (remove view)
- 3) Unbind devices to avoid any new application server seeing data from old server.
- 4) Delete storage group
- 5) Delete port group and zoning for server (if appropriate)

The unbound devices will need to be rebound to the FC/SAS pool before they are re-allocated to a new host (see below).

8.17.1 Unbinding devices under FAST VP management

There are no- FAST VP specific issues with unbinding and rebinding devices when re-allocating storage to new servers.

Although a device may have storage allocated to it from multiple pools, it is only bound to a single pool. Hence the process of binding and rebinding remains the same. After rebinding, the device will only have a small (or no) allocated from the bound pool.

The following example from the CATE engineering lab shows the end-to-end process for device 0385, which is initially using storage from pools glab1_g13 (bound pool), glab1_g10 & glab1_g14.

```
# symcfg list -tdev -bound -detail -sg fastvp_SPA -sid 1380
```

```
Symmetrix ID: 000192601380
```

```
Enabled Capacity (Tracks) : 2602844880
```

```
Bound Capacity (Tracks) : 2209680
```

S Y M M E T R I X T H I N D E V I C E S

Sym	Pool Name	Flags EMSP	Total Tracks	Pool Subs (%)	Pool Allocated Tracks (%)	Total Written Tracks (%)	Status
0385	glab1_g13	FX..	552420	0	0 0	0 0	Bound
	glab1_g10	--.	-	-	208956 38	- -	
	glab1_g14	--.	-	-	67272 12	- -	
0386	glab1_g13	FX..	552420	0	228240 41	161176 29	Bound
	glab1_g14	--.	-	-	48000 9	- -	
0387	glab1_g13	FX..	552420	0	276216 50	0 0	Bound
	glab1_g14	--.	-	-	12 0	- -	
0388	glab1_g13	FX..	552420	0	0 0	0 0	Bound
	glab1_g14	--.	-	-	276228 50	- -	
Total							
Tracks			2209680	0	1104924 0	161176 0	

or for SE7.5 or later

Symmetrix ID: 000192601380

Enabled Capacity (Tracks) : 2602844880
Bound Capacity (Tracks) : 2209680

S Y M M E T R I X T H I N D E V I C E S

Sym	Bound Pool Name	Flags ESPT	Total Tracks	Pool Subs (%)	Pool Allocated Tracks (%)	Total Written Tracks (%)	Compressed Size/Ratio Tracks (%)
0385	glab1_g13	F..B	552420	0	239976 43	239954 43	239976 0
0386	glab1_g13	F..B	552420	0	228240 41	161176 29	228240 0
	glab1_g14	-.--	-	-	48000 9	- -	48000 0
0387	glab1_g13	F..B	552420	0	503508 91	432521 78	503508 0
	glab1_g14	-.--	-	-	12 0	- -	12 0
0388	glab1_g13	F..B	552420	0	0 0	0 0	0 0
	glab1_g14	-.--	-	-	276228 50	- -	276228 0
Total							
Tracks			2209680	0	1295964 0	833651 0	1295964 0

symdev not_ready 0385 -sid 1380

Execute a 'Not Ready' Device operation for device '0385' (y/[n]) ? y

'Not Ready' Device operation successfully completed for the device.

symconfigure -sid 1380 -cmd "unbind tdev 0385 from pool glab1_g13;" commit -nop

A Configuration Change operation is in progress. Please wait...

```
Establishing a configuration change session.....Established.
Processing symmetrix 000192601380
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Locking devices.....Locked.
Committing configuration changes.....Started.
Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

symcfg list -tdev -bound -detail -sg fastvp_SPA -sid 1380

Symmetrix ID: 000192601380

Enabled Capacity (Tracks) : 2602844880
 Bound Capacity (Tracks) : 1657260

SYMMETRIX THIN DEVICES

Sym	Pool Name	Flags EMSP	Total Tracks	Pool Subs (%)	Pool Allocated Tracks (%)	Total Written Tracks (%)	Status
0386	glabl_g13	FX..	552420	0	228240 41	161176 29	Bound
	glabl_g14	--.	-	-	48000 9	- -	
0387	glabl_g13	FX..	552420	0	276216 50	0 0	Bound
	glabl_g14	--.	-	-	12 0	- -	
0388	glabl_g13	FX..	552420	0	0 0	0 0	Bound
	glabl_g14	--.	-	-	276228 50	- -	
Total							
Tracks			1657260	0	828696 0	161176 0	

(SE7.5 output not shown in the interest of brevity)

```
# symconfigure -sid 1380 -cmd "bind tdev 0385 to pool glabl_g13 preallocate size = 0MB;" commit -nop
```

A Configuration Change operation is in progress. Please wait...

```
Establishing a configuration change session.....Established.
Processing symmetrix 000192601380
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Locking devices.....Locked.
Committing configuration changes.....Started.
Binding devices.....Done.
Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

```
# symcfg list -tdev -bound -detail -sg fastvp_SPA -sid 1380
```

Symmetrix ID: 000192601380

Enabled Capacity (Tracks) : 2602844880
 Bound Capacity (Tracks) : 2209680

SYMMETRIX THIN DEVICES

Sym	Pool Name	Flags EMSP	Total Tracks	Pool Subs (%)	Pool Allocated Tracks (%)	Total Written Tracks (%)	Status
0385	glabl_g13	F...	552420	0	12 0	0 0	Bound
0386	glabl_g13	FX..	552420	0	228240 41	161176 29	Bound
	glabl_g14	--.	-	-	48000 9	- -	
0387	glabl_g13	FX..	552420	0	276216 50	0 0	Bound
	glabl_g14	--.	-	-	12 0	- -	
0388	glabl_g13	FX..	552420	0	0 0	0 0	Bound
	glabl_g14	--.	-	-	276228 50	- -	
Total							
Tracks			2209680	0	828708 0	161176 0	

(SE7.5 output not shown in the interest of brevity)

8.18 FAST VP Storage Group re-association (5876 or later only)

In 5875, if a storage group (SG) is disassociated from a FAST VP policy, then all of its performance information is lost.

In 5876, performance data is retained if the SG is associated with another policy (or re-associated with the original policy if dis-associated by accident) within ~2 hours.

Consequently, there is no mechanism to remove performance data for an SG in 5876, not that this is expected to be necessary in a production environment.

9 3-tier FAST VP Capacity planning

Note: The 2-tier 189TB “Gold” FAST VP standard configuration is installed with a fixed capacity that fills all 4-storage bays. Hence expansion of this configurations is not possible. See Appendix H for further details.

9.1 Introduction

An existing 3-tier FAST VP array may require additional capacity. This is more complex than just adding more disks to a single-tier array as:

- 1) Disk will need to be added to more than one tier, and the correct ratios must be determined / maintained.
- 2) The need to add drives in increments of 1 drive per loop (e.g. 16 for 2-engine array, 32 for 4-engine array etc) will restrict the granularity of the upgrades that can be achieved.
- 3) There are maximum limits on number of flash drives per loop.
- 4) Once must consider if there is sufficient “capacity” in terms of both:
 - a. Storage capacity (GB)
 - b. IOPS capacity

It may not be appropriate to just add storage to each of the tiers in the same ratios as currently installed. This is because one tier may be overloaded from an I/O perspective, and it may be necessary to add disproportionately more storage to a higher tier, to allow more busy extents to be promoted.

In addition, the age of the array and its status within its depreciation cycle may also weigh against upgrading an existing array.

Capacity upgrades for FAST VP arrays beyond a given configuration should be delivered where possible by expanding the array to match the next higher capacity Standard FAST VP configuration listed in Appendix A.

CATE and CTI have also agreed that emergency increases in FAST VP capacity can be met via an emergency “capacity-on-demand” upgrade to the current configuration. These upgrades are also described in Appendix A.

The processes described in this section can be used to check capacity and IOPS utilization of your FAST VP environment prior to an upgrade.

9.2 Determining when expansion is required

Expansion (additional drives) may be required when:

- 1) **Storage Capacity:** The FAST VP environment overall storage capacity (sum of all pool that FAST VP uses) is running low.
- 2) **IOPS Capacity:** There is insufficient storage capacity in higher tiers to meet the I/O profile (specifically the skew of the I/O profile). For example, extents that ideally should be on FC/SAS tier are running on SATA due to lack of capacity in the FC/SAS tier.

9.2.1 Storage capacity

Pool utilization can be tracked in a number of ways:

1. Existing regional processes already in place for thin pools, based on either ECC or scripts.
2. Using symcli commands and logging the results

```
$ symcfg list -sid 425 -pool -thin -detail -gb
```

```
Symmetrix ID: 000192600425
```

S Y M M E T R I X T H I N P O O L S									
Pool Name	Tech	Dev Emul	Dev Config	Total GBs	Enabled GBs	Used GBs	Used (%)	Subs (%)	Shared GBs
0425_FLASH	EFD	FBA	RAID-5 (7+1)	5136.0	5136.0	1964.6	38	0	0.0
0425_450GB	FC	FBA	RAID-6 (14+2)	32624.6	32624.6	4797.3	14	307	0.0
0425_2TBSATA	SATA	FBA	RAID-6 (14+2)	98114.5	98114.5	17094.3	17	0	0.0
Total GBs				135875.2	135875.2	23856.2	17	74	0.0


In the above example, we can see overall consumption of the pool capacity across all three tiers is 17% (bottom of the “Used % column”). Individual pool consumption is also shown.

With SE7.5 or later, the above command will show different columns (e.g. “Used(%)” is renamed to “Full(%)” and there is an additional “Comp(%)” column).

3. Unisphere - use either Historical or Diagnostic view: Trends pool average capacity allocation (total and allocated) for all pools or a single thin pool.


All thin pools on an array (very helpful when all pool are used by FAST VP)

Where: Symmetrix Array → <SID> → Performance → Analyze → Historical (or Diagnostic) → TP Pool → <Select ALL pools to display total combined or single pool> → Plot required metrics as either time series or snap shot graphs.



Trend in allocated storage for all pools

- Individual thin pools (potentially misleading – see below)



Trend in allocated storage for a single FC pool

The key pool to track from a long term storage capacity planning perspective is of the pool supporting the lowest tier (SATA). All storage groups can access as much SATA as they need as Citi policies always have a 100% limit on the SATA tier.


Note that high utilization of the flash and FC/**SAS** tiers is desirable as it shows that these higher cost resources are being used as fully as possible.

9.2.2 IOPS capacity


At the time of writing, IOPS Utilization by tier does not yet appear to be available in Unisphere. U4V does not provide trending information on IOPS utilization per tier as a percentage of the capability of that tier. However, trending of % utilization is available for disk groups. These usually map to tier in a simple way and their use is a workable alternative when there is only one diskgroup per tier.

Symmetrix Array → <SID> → Performance → Analyze → Historical (or Diagnostic) → Disk Group → Select ALL → Plot either Snap shot or Timescale for IO/Sec

Note: Hover the cursor over the disk group to display %




The trend view also shows Average disk utilization for a disk group.



Alternatively, the snapshot view for all Disks for a given array will show you both:

- Utilization by (Disk) Group
- Response time by (Disk) Group

averages for the snapshot period selected.



Solutions Enabler can also be used to map disk groups to technology type. In this case, disk group 013 is SATA.

```
symdisk -sid 4261 list -dskgrp_summary
```


Symmetrix ID: 000192604261

Disk Group		Disk			Capacity			
Num	Name	Flgs		Speed (RPM)	Size (MB)	Total (MB)	Free (MB)	Actual (MB)
		Cnt	LT					
10	DISK_GROUP_010	32	IE	0	190782	6105031	333713	6105031
11	DISK_GROUP_011	32	IF	15000	558281	17864980	14659176	17864980
12	DISK_GROUP_012	96	IF	15000	558281	53594940	844403	53594940
13	DISK_GROUP_013	64	IS	7200	1823565	116708132	494424	116708132
Total						194273082	16331717	194273082

Alternatively, the predefined Unisphere “ALL Disk Groups” dashboard will provide snapshots for ALL Disk groups for a given array showing:

- IOs/Sec by Disk Group
- Avg Response time by Disk Group
- % Used Capacity by Disk Group
- % Busy by Disk Group


Where: Symmetrix Array → <SID> → Performance → Monitor → EMC Dashboards → ALL Disk Groups



For individual disk groups you can use the predefined Disk Group dashboard, and select the disk group accordingly, this can be displayed in diagnostic or historical views showing:

- % Busy
- MB Written & Read/Sec
- Writes & Reads/Sec
- Total & Used Capacity

Where: Symmetrix Array → <SID> → Performance → Monitor → EMC Dashboards → Disk Group



9.2.2.1 IOPS utilization – rules of thumb

At the time of writing, the tools required for end users to model the load on each tier prior to planning the upgrade are not yet available from EMC. Instead, EMC are able to take an existing workload, and model that workload against different configurations and determine disk utilization of each tier using a tool called Tier Advisor. This tool is not available to customers.

To optimize the tiered storage configuration EMC will use Tier Advisor to attempt to maximize the use of the capacity available in each tier but limit the maximum throughput utilization to ensure a high performance of the overall configuration.

Tier	Desired use of the storage capacity	Desired maximum use of the total throughput
Flash	100%	80 - 90%
Fibre	100%	40 - 50%
SATA	100%	0 - 10%

Important rules of thumb:

- Backend SATA utilization should be < 50%. SATA response time degrades quickly at higher utilization levels.
- Backend FC/SAS utilization should be < 50%. If FAST cannot get a 50% improvement in response time though data relocation to a higher tier, it will not move the data. Hence data can get stuck down in SATA tier if the FC/SAS tier is excessively busy.
- EFD. Utilization should be as high as possible, May not reach 80% if the skew of the workload is not high enough.

At the time of writing, it is not know if the CATE Standard Configurations will provide the correct mix of tier capacities to support all mixes of typical Citi workloads without resulting in excessive of a specific tier.

The current IOPS utilization of a disk group can be determined using one of the method(s) described above and compared against the above table to determine if you have excessively high utilization in any of your tiers.

9.3 Thin Pool Oversubscription

How do we meaningfully track thin pool subscription in a FAST VP environment?

Today we can measure how subscribed a pool is by considering the total size of the TDEVs bound to a pool, and the total capacity of the DATADEVs in the pool (maximum capacity that can be allocated to the TDEVs).

Pool subscription level = (Sum TDEVs bound to pool) / (Sum DATADEVs in the pool).

With FAST VP, we still need to measure subscription levels, but now pool subscription levels are less meaningful as data on a TDEV bound to a specific pool can actually reside in multiple pools.

Instead the subscription formula to be used should consider subscription at the array or FAST VP level.

FAST VP subscription level = (Sum of bound TDEVs) / (Sum DATADEVs in ALL pools)

The above assumes one pool for each technology type (EFD, FC/SAS & SATA), and FAST policies have give access to all three pools.

Any reports written to report on FAST VP subscription level must use this formula instead. At the time of writing, EMC have not yet included in ECC out-of-the-box reporting on FAST VP subscription level.

Information on the usage and subscription can be obtained via Solutions Enabler. For example:

```
$ symcfg list -sid 425 -pool -thin -detail -gb
```

Symmetrix ID: 000192600425

S Y M M E T R I X T H I N P O O L S									
Pool Name	Tech	Dev Emul	Dev Config	Total GBs	Enabled GBs	Used GBs	Used (%)	Subs (%)	Shared GBs
0425_FLASH	EFD	FBA	RAID-5 (7+1)	5136.0	5136.0	1964.6	38	0	0.0
0425_450GB	FC	FBA	RAID-6 (14+2)	32624.6	32624.6	4797.3	14	307	0.0
0425_2TBSATA	SATA	FBA	RAID-6 (14+2)	98114.5	98114.5	17094.3	17	0	0.0
Total GBs				135875.2	135875.2	23856.2	17	74	0.0

or for SE7.5 or later (to keep the same output format):

```
$ symcfg list -sid 425 -pool -thin -detail -gb -mode V73
```

However, care must be taken when interpreting these figures in a FAST VP environment for two reasons:

1. The subscription rate is based on the capacity of the thin devices (TDEVs) that are bound to the pool, and not the capacity of TDEVs that are provisioned to hosts (unless TDEVs are only bound to a pool at the time of provisioning).
2. The above command reports a subscription rate based on the capacity of the bound pool (FC or SAS). While this is technically correct, in a FAST VP environment where data can move into all of the pools, the FAST VP subscription rate should be based on capacity of all of the pools accessible to FAST VP.

Hence, the high 307% “Pool subscription %” figure for the FC/SAS tier should be ignored, and instead, the 74% figure at the bottom of the “Subs (%)” column used instead as the “FAST VP subscription %” figure.

Note that this last figure is only correct if all pools in the list are managed via FAST VP. The “FAST VP subscription %” in other circumstances can also be calculated as follows:

$$\begin{aligned} \text{FAST VP Sub. \%} &= \frac{(\text{FC/SAS Sub \%} * \text{FC enabled GB})^8}{(\text{EFD} + \text{FC/SAS} + \text{SATA enabled GB} \text{ *under FAST control*})} \\ &= \frac{(307\%) * 32624.6}{(135875.2)} = 74\% \end{aligned}$$

Again, this is the subscription rate based on all bound TDEVs, not just those that are already provisioned to hosts, unless TDEVs are only bound to a pool at the time of provisioning.

9.3.1 Note for SE7.5 or later

The output for the command

```
# symcfg list -sid 425 -pool -thin -detail -gb
```

has changed considerably between SE7.3 & SE7.5 and above. In the SE7.5 and above output, the following changes have occurred:

SE7.3 Output Column Heading	SE7.5 Output Column Heading
Total GBs	Total GBs
Enabled GBs	Usable GBs
Used GBs	Used GBs
N/A	Free GBs
Used (%)	Full (%)
Subs (%)	Subs (%)
N/A	Comp (%)

The old format can be retained by using the “-mode V73” option on the command line:

```
$ symcfg list -sid 425 -pool -thin -detail -gb -mode V73
```

⁸ FC Sub % * FC Enabled GB = Capacity of TDEVs bound to pool

10 Appendix A – Approved 3-tier FAST VP configurations

This section has been updated to include new VMAX40K configurations that should be used for all new 3-tier “Silver” storage FAST VP deployments. 2-tier VMAX40K “Gold” storage standards and best practice are describe in Appendix H.

3-tier VMAX20K configurations are also still described as they are still supported on existing VMAX systems installed prior to the release of this document (March 2013).

10.1 VMAX40K 3-tier (“Silver”) standard configurations

The table in this subsection summarizes configurations that are suitable when the skew of the workload is not known in advance. The table defines:

- Number of engines
- Total usable capacity in engineering TB (1TB = 2⁴⁰ bytes)
- RAID type for each tier
- Number of drives per loop for each tier
- Number of drives required for each type (excluding spares)
- Pool capacities
- Spare drive requirements
- Total drive count
- Percentage of total usable capacity by tier.

This table can be used for both initial configurations, or to plan an upgrade from “new small” to “new medium” by finding the configuration that has the required total capacity and adding the required number drives per loop to reach the new configuration.

Chris Clayton - 10th April 2013													
VMAX40K	Number of Engines	Thin Pool Capacity (Eng TB)	EFD drives per loop	SAS drives per loop	SATA drives per loop	EFD drives	SAS drives	SATA drives	Additional Spare Drive Count	Total Drives	% of EFD capacity	% of FC capacity	% of SATA capacity
"New Small" FAST VP config VMAX40K 4-Engine (4 * 96GB) (384 GB total cache) - 100GB 2.5" SSD RAID5(3+1) - 300GB 2.5" 10K SAS RAID1 - 1TB 2.5" SATA RAID6(6+2)	4	216	4	20	6	128	640	192	4*EFD 14*SAS 4*1TB	982	4%	39%	57%
"New Medium" FAST VP config VMAX40K 4-Engine (4 * 96GB) (384 GB total cache) - 100GB 2.5" SSD RAID5(3+1) - 300GB 2.5" 10K SAS RAID1 - 1TB 2.5" SATA RAID6(6+2)	4	315	6	28	9	192	896	288	4*EFD 18*SAS 6*1TB	1404	4%	37%	59%

- 100GB SSD drives are used to allow an even distribution across all director CPUs. RAID6(3+1) should be used so that the pool can survive the failure of an entire VMAX engine. This replaces the previous recommendation to use RAID6(7+1) following loss of an entire engine in NAM.
- 300GB 10K SAS RAID1 is required for in order to meet the the workload IOPS requirements for the middle tier.
- RAID6(6+2) is used for the SATA tier as this offers optimized writes, which minimize RAID6 overheads for sequential writes. 1TB drives have replaced use of the 2TB drives in the SATA tier in order to double the IOPS in this tier, which is sometime subject due heavier than expected IO due to infrequent database dumps.

The above configurations have been designed in conjunction with EMC to support the following workload profile:

Assumed workload parameters	
Written TB:	Assume maximum 85% of total usable pool storage allocated from pool
Host IOPS/TB:	400* (TB here = TB allocated from pool)
Maximum total Host IOPS:	400 x 0.85 * total usable TB i.e. • 74K host IOPS for “New Small” configuration • 108K host IOPS for “New Medium” configuration
Sub-LUN skew:	85/15 i.e. 85% of IO goes to 15% of storage capacity Config 50 (“Silver” FAST Standard with RAID5(3+1) SSD – SSD 4.1% of usable capacity): • 42% to SSD • 55% to SAS • 3% to SATA
Read hits	30%
Random reads	15%
Random writes	45%
Sequential writes	10%
Random write IO size	8KB
Sequential write IO size	64KB

10.1.1 3-tier “Silver” configuration emergency Capacity-on-Demand upgrades

The following table shows supported emergency upgrade configurations, if desired. **The addition of SATA drives is potentially may introduce performance issues than expanding the SAS tier and should only be undertaken once you have first confirmed that your SATA tier is not already overloaded.** EMC have recommended that if the SATA tier IOPS %utilization is presently greater or equal to 50% *OR* if the SAS tier is > 50%, then you should upgrade the SAS tier. This can be checked using Unisphere (see Section 9.2.2.1 for details).

For the “New small” emergency upgrades, you either add 128*300GB 10K SAS drives **OR** add 32*1TB SATA drives, not both.

VMAX40K	Number of Engines	Thin Pool Capacity (Eng TB)	EFD drives per loop	SAS drives per loop	SATA drives per loop	EFD drives	SAS drives	SATA drives	Additional Spare Drive Count	Total Drives	% of EFD capacity	% of FC capacity	% of SATA capacity
"New Small" FAST VP config VMAX40K 4-Engine (4 * 96GB) (384 GB total cache) - 100GB 2.5" SSD RAID5(3+1) - 300GB 2.5" 10K SAS RAID1 - 1TB 2.5" SATA RAID6(6+2)	4	216	4	20	6	128	640	192	4*EFD 14*SAS 4*1TB	982	4%	39%	57%
Emergency SAS Expansion for "New medium" configuration. Add 128 300GB 10K SAS drives (plus 2 additional spare drives) (+17TB ; +8%)	4	233	4	24	6	128	768	192	4*EFD 16*SAS 4*1TB	1112	4%	43%	53%
Emergency SATA Expansion for "New Small" configuration. Add 32 1TB SATA drives. (plus 2 additional spare drives) (+21TB ; +9%)	4	236	4	20	7	128	640	224	4*EFD 14*SAS 6*1TB	1016	4%	35%	61%
"New Medium" FAST VP config VMAX40K 4-Engine (4 * 96GB) (384 GB total cache) - 100GB 2.5" SSD RAID5(3+1) - 300GB 2.5" 10K SAS RAID1 - 1TB 2.5" SATA RAID6(6+2)	4	315	6	28	9	192	896	288	4*EFD 18*SAS 6*1TB	1404	4%	37%	59%
Emergency SAS Expansion for "New medium" configuration. Add 160 300GB SAS drives (plus 4 additional spare drives) (+21TB ; +6.6%)	4	336	6	33	9	192	1056	288	4*EFD 22*SAS 6*1TB	1568	4%	41%	55%
Emergency SATA Expansion for "New medium" configuration. Add 32 1TB SATA drives (plus 2 additional spare drives) (+21TB ; +6.5%)	4	336	6	28	10	192	896	320	4*EFD 18*SAS 8*1TB	1438	4%	35%	61%

10.2 VMAX20K Restricted (Legacy) Standard RAID1 FC tier configurations

The following configurations are NOT supported for new installations. They remain in the documentation to cover expansion of existing deployments only.

Valid VMAX20K FAST VP configurations															
Chris Clayton 20/06/22															
	Number of Engines	Thin Pool Capacity (TB)	200GB EFD drives per loop	450GB B FC drives per loop	2TB SATA drives per loop	200GB EFD drives	450GB FC drives	2TB SATA drives	Additional Spare Drive Count	Total Drives	Cabs	% of EFD capacity	% of FC capacity	% of SATA capacity	Notes
2-Engine Entry	2	117	2	12	3	32	192	48	1 * EFD 6 * FC 2 * SATA	281	3	4.4%	32.9%	62.7%	Entry-level configuration
4-Engine Small	4	168	1	10	2	32	320	64	1 * EFD 8 * FC 2 * SATA	427	5	3.1%	38.4%	58.5%	Entry-level configuration
4-Engine Medium	4	297	2	14	4	64	448	128	2 * EFD 10 * FC 4 * SATA	656	5	3.5%	30.4%	66.2%	Recommended config. 1 system plus 4 storage bays

CATE-certified FAST VP RAID1 FC tier configurations

Notes on the above table

1. The **2-Engine Entry** configuration should only be used when there are no plans to expand the array beyond this initial size.
2. For configurations that are planned to be expanded, the **4-Engine Small** configuration (or larger) should be installed.
3. If an unplanned expansion of the **2-Engine Entry** configuration is required, then it should be upgraded to the **4-Engine Medium** configuration.
4. When upgrading **any 4-Engine configuration**, add storage to the unit to match a configuration lower in the table.
5. If a capacity greater than a 297TB pool is required, then a new FAST VP instance should be installed to ensure that the average amount of storage per segment (FA pair) does not grow too large. A 4-engine array can have a maximum of 32 segments.
6. The **4-Engine large** configuration is not currently supported with a RAID1 FC tier as this would require 6 bays, which is not supported within Citi for a 4-Engine configuration.
7. **6-Engine configurations** are not currently supported on the recommendation from the SAN hardware implementation team due to space planning considerations (minimum footprint of 5 bays).

These standard configurations are:

- Large usable capacity configurations (driven by use of SATA drives)
- Fixed usable capacities, with large steps in capacity between configurations.

- Meet Citi and EMC best practice on deploying integer numbers of drives per loop
- Meet Citi and EMC best practice on minimum (for SATA and FC) and maximum (for Flash) drives per loop

The above template configurations are not optimized for a specific workload. Instead, they have the following characteristics:

- The ability to deal with a range of I/O distribution skews (85 or greater) without significant performance degradation for any skew that is low-to-moderate (85) or higher. The relative response time for configurations with a lower percentage of FC compared with SATA degrades more quickly as the skew decreases (as I/Os that would otherwise be on FC have to reside on SATA) (see graph titled “Reponse Time degradation with low % FC” in section 4.5.3 for an example of the impact caused by skimping on the FC tier %).
- The ability to deal with a large range of IOPS.

In exceptional circumstances where the capacity-on-demand configurations cannot be used, a bespoke configuration is required. See Appendix H for details. At present, approval is required for the author for all bespoke configurations.

10.2.1 Emergency Capacity-on-Demand upgrades for RAID1 FC tier configs

The following two subsections show supported emergency upgrade configurations. **The addition of SATA drives is more likely to introduce performance issues than expanding the FC tier and should only be undertaken once you have first confirmed that your SATA tier is not already overloaded.** EMC have recommended that if the SATA tier IOPS %utilization is presently greater or equal to 50% *OR* if the FC tier is > 50%, then you should upgrade the FC tier. This can be checked using Unisphere (see Section 9.2.2 for details).

	Number of Engines	Thin Pool Capacity (TB)	200GB EFD drives per loop	450GB 8 FC drives per loop	2TB SATA drives per loop	200GB EFD drives	450GB FC drives	2TB SATA drives	Additional Spare Drive Count	Total Drives	Cabs	% of EFD capacity	% of FC capacity	% of SATA capacity	Notes
2-Engine Entry	2	117	2	12	3	32	192	48	1 * EFD 6 * FC 2 * SATA	281	3	4.4%	32.9%	62.7%	Entry-level configuration.
Emergency CoD upgrade for config above (+21%/+25TB)	2	142	2	12	4	32	192	64	1 * EFD 4 * FC 4 * SATA	297	3	3.6%	27.2%	69.2%	Adding 1 SATA drive per loop
4-Engine Small	4	168	1	10	2	32	320	64	1 * EFD 8 * FC 2 * SATA	427	5	3.1%	38.4%	58.5%	Entry-level configuration.
Emergency CoD upgrade for config above (+29%/+49TB)	4	217	1	10	3	32	320	96	1 * EFD 8 * FC 2 * SATA	459	5	2.4%	29.7%	67.9%	Adding 1 SATA drive per loop
4-Engine Medium	4	297	2	14	4	64	448	128	2 * EFD 10 * FC	656	5	3.5%	30.4%	66.2%	Recommended config. 1 system plus 4 storage bays
Emergency CoD upgrade for above config (+17%/+49TB)	4	346	2	14	5	64	448	160	2 * EFD 10 * FC 4 * SATA	688	5	3.0%	26.1%	71.0%	Adding 1 SATA drive per loop.

Low cost Emergency Capacity-on-Demand SATA upgrades for RAID1 FC tier configs

Valid VMAX20K FAST VP configurations, with emergency CoD FC expansion options															
Chris Clayton 20/06/12															
	Number of Engines	Thin Pool Capacity (TB)	200GB EFD drives per loop	450GB B FC drives per loop	2TB SATA drives per loop	200GB EFD drives	450GB FC drives	2TB SATA drives	Additional Spare Drive Count	Total Drives	Cabs	% of EFD capacity	% of FC capacity	% of SATA capacity	Notes
2-Engine Entry (RAID1 FC)	2	117	2	12	3	32	192	48	1 * EFD 6 * FC 2 * SATA	281	3	4.4%	32.9%	62.7%	Entry-level configuration.
Emergency CoD upgrade for config above (+16%/+197B)	2	137	2	18	3	32	288	48	1 * EFD 6 * FC 2 * SATA	377	3	3.8%	42.4%	53.9%	Adding 6 FC drives per loop
4-Engine Small (RAID1 FC)	4	168	1	10	2	32	320	64	1 * EFD 8 * FC 2 * SATA	427	5	3.1%	38.4%	58.5%	Entry-level configuration.
Emergency CoD upgrade for config above (+15%/+267B)	4	194	1	14	2	32	448	64	1 * EFD 10 * FC 2 * SATA	557	5	2.7%	46.6%	50.7%	Adding 4 FC drives per loop
4-Engine Medium (RAID1 FC)	4	297	2	14	4	64	448	128	2 * EFD 10 * FC 4 * SATA	656	5	3.5%	30.4%	66.2%	Recommended config. 1 system plus 4 storage bays
Emergency CoD upgrade for config above (+15%/+457B)	4	342	2	21	4	64	672	128	2 * EFD 14 * FC 4 * SATA	884	5	3.0%	39.6%	57.4%	Adding 7 FC drives per loop

Higher cost Emergency Capacity-on-Demand FC upgrades for RAID1 FC tier configs

10.3 VMAX20K Restricted (Legacy) RAID6 FC tier configurations

The following configurations are NOT supported for new installations. They remain in the documentation to cover expansion of existing deployments only.

Valid VMAX CoD FAST VP configurations																
Chris Clayton 26/10/11																
Configuration Name	Number of Engines	Thin Pool Capacity (TB)	Raw capacity (TB)	200GB EFD drives per loop	450GB B FC drives per loop	2TB SATA drives per loop	200GB EFD drives	450GB FC drives	2TB SATA drives	Thin Pool Capacity (TB)	Additional Spare Drive Count	Total Drives	% of EFD capacity	% of FC capacity	% of SATA capacity	Notes
2-Engine Entry	2	118	135	2	7	3	32	112	48	118	1 * EFD 6 * FC 2 * SATA	201	4.3%	33.4%	62.3%	Entry-level configuration.
4-Engine Small	4	171	195	1	6	2	32	192	64	171	1 * EFD 6 * FC 2 * SATA	297	3.0%	39.6%	57.4%	Entry-level configuration.
4-Engine Medium	4	297	339	2	8	4	64	256	128	297	2 * EFD 6 * FC 4 * SATA	460	3.5%	30.4%	66.2%	Recommended configuration.
4-Engine Large	4	445	509	3	12	6	96	384	192	445	2 * EFD 8 * FC 4 * SATA	686	3.5%	30.4%	66.2%	Recommended configuration.

Restricted Legacy FAST VP RAID6 FC tier configurations

10.3.1 Emergency Capacity-on-Demand upgrades for RAID6 FC tier configs

The following two subsections show supported emergency upgrade configurations. **CATE recommends emergency expansion via the addition of 450GB FC drives**, but will also support the more cost effective expansion via 2TB SATA drives, as requested by CTI senior management. **However, it must be stressed that expansion via the addition of SATA drives is more likely to introduce performance issues than expanding**

the FC tier and should only be undertaken once you have first confirmed that your SATA tier is not already overloaded. EMC have recommended that if the SATA tier IOPS %utilization is presently greater or equal to 50% *OR* if the FC tier is > 50%, then you should upgrade the FC tier. This can be checked using SPA (see Section 9.2.2 for details).

10.3.1.1 Option A – Only add 450GB FC drives for emergency expansion

This set of options only requires the addition of FC drives, typically the number needed to match the FC drive requirement of the next largest standard configuration.

Advantages

1. Can leverage 450GB FC spares pool used for expanding non-FAST VP arrays.
2. Will increase average IOPS/GB capability of array. Will not introduce performance issues.
3. Small capacity expansions are possible (< 15%) c.w. Option B (20-30%).
4. Expands the most critical pool (FC) where all new write are allocated from with Enginuity 5875.

Disadvantages

5. Higher cost per GB than Option B

Valid VMAX FAST VP configurations, with emergency CoD FC (only) expansion options															
Chris Clayton/06/01/12															
	Number of Engines	Thin Pool Capacity (TB)	Raw capacity (TB)	200GB EFD drives per loop	450GB FC drives per loop	2TB SATA drives per loop	200GB EFD drives	450GB FC drives	2TB SATA drives	Additional Spare Drive Count	Total Drives	% of EFD capacity	% of FC capacity	% of SATA capacity	Notes
2-Engine Entry	2	118	135	2	7	3	32	112	48	1 * EFD 6 * FC 2 * SATA	201	4.3%	33.4%	62.3%	Entry-level configuration.
Emergency CoD upgrade for config above (+14%/+17TB)	2	135	154	2	10	3	32	160	48	1 * EFD 6 * FC 2 * SATA	249	3.8%	41.7%	54.5%	Adding 3 FC drives per loop
4-Engine Small	4	171	195	1	6	2	32	192	64	1 * EFD 6 * FC 2 * SATA	297	3.0%	39.6%	57.4%	Entry-level configuration.
Emergency CoD upgrade for config above (+13%/+23TB)	4	194	221	1	8	2	32	256	64	1 * EFD 6 * FC 4 * SATA	363	2.7%	46.6%	50.7%	Adding 2 FC drives per loop
4-Engine Medium	4	297	339	2	8	4	64	256	128	2 * EFD 6 * FC 4 * SATA	460	3.5%	30.4%	66.2%	Recommended config. 1 system plus 4 storage bays
Emergency CoD upgrade for config above (+15%/+45TB)	4	342	391	2	12	4	64	384	128	2 * EFD 6 * FC 4 * SATA	588	3.0%	39.6%	57.4%	Adding 4 FC drives per loop
4-Engine Large	4	445	509	3	12	6	96	384	192	2 * EFD 8 * FC 4 * SATA	684	3.5%	30.4%	66.2%	Recommended config. 1 system plus 4 storage bays
Emergency CoD upgrade for config above (+10%/+45TB)	4	490	560	3	16	6	96	512	192	2 * EFD 8 * FC 6 * SATA	816	3.1%	36.8%	60.1%	Adding 4 FC drives per loop

Emergency Capacity-on-Demand FC upgrades for RAID6 FC tier configs

10.3.1.2 Option B: Add primarily 2TB SATA drives

These options primarily add SATA drives, although FC also needed for one configuration.

Advantages

1. Lower cost per GB than Option A.

Disadvantages

2. Requires a 2TB SATA pool just for FAST VP arrays. FC spares also needed for emergency upgrades to 4-Engine medium upgrade.
3. Will **DECREASE** average IOPS/GB capability of array. More likely to introduce performance issues than Option A, depending on array workload.
4. Adding even the minimum amount of SATA of just 1 drive per loop (except for the largest array) gives you more than required upgrade (20-29%), but uses more cost-efficient disks.

Valid VMAX FAST VP configurations, with emergency CoD SATA & FC expansion options															
Chris Clayton 06/01/12															
	Number of Engines	Thin Pool Capacity (TB)	Raw capacity (TB)	200GB EFD drives per loop	450GB FC drives per loop	2TB SATA drives per loop	200GB EFD drives	450GB FC drives	2TB SATA drives	Additional Spare Drive Count	Total Drives	% of EFD capacity	% of FC capacity	% of SATA capacity	Notes
2-Engine Entry	2	118	135	2	7	3	32	112	48	1 * EFD 6 * FC 2 * SATA	201	4.3%	33.4%	62.3%	Entry-level configuration.
Emergency CoD upgrade for config above (+21%/+257B)	2	143	163	2	7	4	32	112	64	1 * EFD 6 * FC 2 * SATA	217	3.6%	27.6%	68.8%	Adding 1 SATA drive per loop.
4-Engine Small	4	171	195	1	6	2	32	192	64	1 * EFD 6 * FC 2 * SATA	297	3.0%	39.6%	57.4%	Entry-level configuration.
Emergency CoD upgrade for config above (+29%/+497B)	4	220	251	1	6	3	32	192	96	1 * EFD 6 * FC 4 * SATA	331	2.3%	30.7%	66.9%	Adding 1 SATA drive per loop.
4-Engine Medium	4	297	339	2	8	4	64	256	128	2 * EFD 6 * FC 4 * SATA	460	3.5%	30.4%	66.2%	Recommended config. 1 system plus 4 storage bays.
Emergency CoD upgrade for above config (+20%/+607B)	4	357	408	2	9	5	64	288	160	2 * EFD 6 * FC 4 * SATA	524	2.9%	28.4%	68.7%	Adding 1 FC drive & 1 SATA drive per loop.
4-Engine Large	4	445	509	3	12	6	96	384	192	2 * EFD 8 * FC 4 * SATA	684	3.5%	30.4%	66.2%	Recommended config. 1 system plus 4 storage bays.
Emergency CoD upgrade for above config (+11%/+497B)	4	494	565	3	12	7	96	384	224	2 * EFD 8 * FC 6 * SATA	720	3.1%	27.4%	69.5%	Adding 1 SATA drive per loop.

Emergency Capacity-on-Demand SATA & FC upgrades for RAID6 FC tier configs

10.3.1.3 Cost comparison of Options A & B (Legacy configurations)

A comparison of the cost of the two emergency upgrade option sets that are most likely to occur are shown in the table below, using pricing as of 6th Jan 2012.

The analysis is just for pool drive hardware. Price analysis excludes and additional extra spare drives or licensing costs, which would likely be very similar for either option.

Both Global VPA and NA ELA pricing (VPA – 10%) is shown.

Note that when the array is latter upgraded for increased capacity to the next configuration down the table (e.g. from 4-Eng Small to 4-Eng Medium), any saves obtained via an Option B emergency upgrade will be negated after the capacity upgrade to the Medium configuration are completed i.e. the total cost is the same cost in the end regardless of whether you go with Option A or B for the emergency upgrade.

Upgrade option	Config being upgraded	# 450GB FC Drives	# 2TB SATA Drives	Pool size increase (TB)	VPA FC \$/MB	VPA SATA \$/MB	Upgrade VPA cost	Upgrade ELA cost
A	4-Eng Small	64	0	23	0.0074	0.0017	\$213,120	\$191,808
B	4-Eng Small	0	32	45	0.0074	0.0017	\$108,800	\$97,920
A	4-Eng Medium	128	0	49	0.0069	0.0014	\$397,440	\$357,696
B	4-Eng Medium	32	32	60	0.0069	0.0014	\$188,960	\$170,064

11 Appendix B - Worked example of setting up 3-tier FAST VP

The following VMAX20K example demonstrates the main steps for setting up 3-tier FAST VP. It assumes that the thin pools are already configured, but does also show the creation of devices, masking views and storage groups for completeness.

Note that the naming conventions do not reflect best practice, but are just those used in the lab when generating this example.

VP pool names:

- EFD pool: glab1_g10
- FC pool: glab1_g13
- SATA pool: glab1_g14

Storage Tier names:

- EFD tier: thin_EFD_200GB
- FC tier: thin_FC_450GB
- SATA tier: thin_SATA_2TB

Production policy max SG percentages:

- EFD: 10%
- FC: 50%
- SATA: 100%

Enable FAST VP controller (one-off task)

```
symfast -sid 1380 enable -vp
```

Create Storage Tiers (one-off task)

```
symtier -sid 1380 create -name thin_EFD_200GB -tgt_raid5 -tgt_prot 7+1 -tech EFD -thin -pool glab1_g10
symtier -sid 1380 create -name thin_FC_450GB -tgt_raid1 -tech FC -thin -pool glab1_g13
symtier -sid 1380 create -name thin_SATA_2TB -tgt_raid6 -tgt_prot 14+2 -tech SATA -thin -pool glab1_g14
```

Create "Production" policy (one-off task)

```
symfast -sid 1380 -fp create -name Production
symfast -sid 1380 -fp -fp_name Production add -tier_name thin_EFD_200GB -max_sg_percent 10
symfast -sid 1380 -fp -fp_name Production add -tier_name thin_FC_450GB -max_sg_percent 50
symfast -sid 1380 -fp -fp_name Production add -tier_name thin_SATA_2TB -max_sg_percent 100
symfast -sid 1380 show -fp_name Production
```

Modify time windows

Note: Best Practice is to use default time 7x24 windows unless there is reason to believe that there are extended periods of time when the workloads managed by FAST VP are not active.

The examples below show the case where applications are only active Monday to Friday between 06:00 to 21:00.

```
# symtw -sid 1380 add -type move_vp -inclusive -days Sun,Mon,Tue,Wed,Thu,Fri,Sat -start_time 00:00 -end_time 24:00 -nop

# symtw -sid 1380 add -type perf -inclusive -days Sun,Mon,Tue,Wed,Thu,Fri,Sat -start_time 00:00 -end_time 24:00 -nop

# symtw -sid 1380 remove -type perf -inclusive -days Mon,Tue,Wed,Thu,Fri -start_time 00:00 -end_time 06:00 -nop

# symtw -sid 1380 remove -type perf -inclusive -days Mon,Tue,Wed,Thu,Fri -start_time 21:00 -end_time 24:00 -nop
```

```
# symtw -sid 1380 remove -type perf -inclusive -days Sun,Sat -start_time 00:00 -end_time 24:00 -nop
```

Set other FAST VP parameters

```
# symfast -sid 1380 set -control_parms -min_perf_period 8
# symfast -sid 1380 set -control_parms -vp_data_move_mode AUTO
# symfast -sid 1380 set -control_parms -workload_period 120
# symfast -sid 1380 set -control_parms -vp_reloc_rate 6

# symconfigure -sid 1380 -cmd "set pool glab1_g10, type=thin, pool_resv_cap=1;" commit
# symconfigure -sid 1380 -cmd "set pool glab1_g13, type=thin, pool_resv_cap=30;" commit
# symconfigure -sid 1380 -cmd "set pool glab1_g14, type=thin, pool_resv_cap=1;" commit
```

The following command to enable FAST VP allocation by policy only applies to SE7.5 & 5876 or later.

```
# symfast -sid 1380 set -control_parms -vp_allocation_by_fp enable
```

Create devices bound to FC tier pool (if not already created)

```
symconfigure -sid 1380 -cmd "create dev count=60, size=34526MB, emulation=FBA, config=TDEV, binding to
pool=glab1_g13, preallocate size=0 MB, dynamic_capability=dyn_rdf;" commit -nop
```

Create Storage Groups (if not already created)

```
symaccess -sid 1380 create -name fastvp01_sangsl -type storage devs 0313:031C
symaccess -sid 1380 create -name fastvp02_sangsl -type storage devs 031D:0326

symaccess -sid 1380 create -name fastvp01_sangsl1 -type storage devs 0327:0330
symaccess -sid 1380 create -name fastvp02_sangsl1 -type storage devs 0331:033A
```

Create Initiator groups (if required)

```
symaccess -sid 1380 create -name ig_sangsl -type initiator -wwn 10000000c9368305
symaccess -sid 1380 add -name ig_sangsl -type initiator -wwn 10000000c93dd604

symaccess -sid 1380 create -name ig_sangsl1 -type initiator -wwn 10000000C9517FF5
symaccess -sid 1380 add -name ig_sangsl1 -type initiator -wwn 10000000C95172CF
```

Create port groups (if required)

```
symaccess -sid 1380 create -name pg_sangsl -type port -dirport 7E:1,10F:1
symaccess -sid 1380 create -name pg_sangsl1 -type port -dirport 9E:1,8F:1
```

Now create the masking views (if required)

```
symaccess -sid 1380 create view -name view_fastvp01_sangsl -sg fastvp01_sangsl -ig ig_sangsl -pg pg_sangsl
symaccess -sid 1380 create view -name view_fastvp02_sangsl -sg fastvp02_sangsl -ig ig_sangsl -pg pg_sangsl

symaccess -sid 1380 create view -name view_fastvp01_sangsl1 -sg fastvp01_sangsl1 -ig ig_sangsl1 -pg
pg_sangsl1
symaccess -sid 1380 create view -name view_fastvp02_sangsl1 -sg fastvp02_sangsl1 -ig ig_sangsl1 -pg
pg_sangsl1
```

Associate these storage groups with the policy to bring under FAST VP control

```
symfast -sid 1380 -fp_name Production associate -sg fastvp01_sangsl -priority 2
symfast -sid 1380 -fp_name Production associate -sg fastvp02_sangsl -priority 2

symfast -sid 1380 -fp_name Production associate -sg fastvp01_sangsl1 -priority 2
symfast -sid 1380 -fp_name Production associate -sg fastvp02_sangsl1 -priority 2
```

The priority setting of the storage group decides which group gets to use a resource when there is competition for that same resource. Priority 1 has the highest priority. The default of 2 should usually be used, as shown here.

At this point, the storage groups are under FAST VP control. Once I/O begins to the devices, and allocations occur, FAST VP will begin to migrate the thin extents from the FC tier to other tiers to:

a) Ensure compliance with the policy with the storage group is associated

b) To optimize tier placement based on actual I/O and hence required performance. Performance migrations will not take place until the min_perf_period has elapsed after the storage group is associated with a policy.

12 Appendix C – Upgrading a single-tier VMAX20K VP array to FAST VP

12.1 Hardware sizing

The scenario covers the situation where an existing “all-FC” array using thin provisioning is FAST-VP enable by adding flash and SATA tiers.

Note, however, that existing arrays are only likely to be expanded from single to multiple tier (FAST VP) configurations if ≤ 1 year old, as older arrays may be considered to be too advanced in the depreciation cycle

In this upgrade scenario, use of one of the standard FAST VP templates may not be suitable as the FC tier will likely have a different number of spindles and/or a different drive type.

In this case, it will likely be necessary to determine a bespoke configuration (see section 4.5.4) with EMC. Fortunately, some of the target applications will already be running on the array and their I/O profiles can be used for the modeling of the tiers.

Again, any configurations that EMC suggest must still comply with CSBE VMAX configuration best practice. For example, drives should be configured in increments of 1 drive per loop (e.g. 16 for 2-engine array, 32 for 4-engine array etc).

12.2 Placing applications under FAST VP control

Existing arrays are only likely to be expanded from single to multiple tier (FAST VP) configurations if ≤ 1 year old (due to reduced lifetime of new storage in an existing array). However, where this does occur, the following high-level steps are proposed for bringing existing applications under FAST VP control. In this case, it is assumed that the TDEVs and related storage groups already exist for existing applications on the array.

- 1) Assess existing workload on array. In this scenario, the workload is already well known as it is already on the array. ECC Performance Manager or STP data should be passed to EMC to confirm the workload skew of your workload.
- 2) Determine quantity of flash and SATA disk to purchase and configure. A bespoke configuration is needed in this case and EMC will use your workload to determine the appropriate configuration (see Section 4.5.4).
- 3) Install additional storage tiers and create thin pools for flash and SATA
- 4) Enable FAST VP controller
- 5) Optional: Modify default time windows (but only if there is a clear case to exclude certain times – see section 5.4).
- 6) Create storage tiers
- 7) Create storage policies
- 8) Set FAST VP parameters and enable FAST VP
- 9) Associate storage groups one at a time with their storage policies. As applications will be active, applications should be brought under FAST VP one at a time, waiting for the tier residency for one to stabilize (if time

permits) before starting on the next. This is to minimize the amount of data migration that is occurring on the array due to compliance moves, and reduce response time impact.

13 Appendix D – What does the policy max percentage mean in practice?

This innocuous sounding question requires some explanation in a virtually provisioned, over-subscribed environment.

13.1.1 Virtual Provisioning – logical vs. actual

The *MaxSgPercent* specified in a FAST policy indicates the upper limit of how much of the total storage group's logical capacity can be allocated to a specific tier.

The key point here is that the percentage applies to the logical capacity and not the actual amount of storage allocated out of the pools.

For example, if the policy has an upper limit of 10% for flash, and 100GB is provisioned to a host, but only 50GB of this is actually allocated out of thin pools, then the maximum amount of storage that could reside on the flash tier is 10GB (10% of the logical 100GB, but 20% of the 50GB of actual storage “in use” i.e. allocated out of the thin pools).

Hence, as application usage of storage out of the thin pool grows, the % (not absolute amount) of “in use” (allocated out of the thin pool) storage that is allowed in the flash tier may reduce. This could impact small-scale benchmarking activities (see section 6.3), or performance during the infancy of the lifecycle of an application.

Application 1

Provisioned	400GB
Used from pools	100GB
Flash capacity allowed	40GB (10% of logical capacity; 40% of currently used capacity)

Application 2

Provisioned	400GB
Used from pools	300GB
Flash capacity allowed	40GB (10% of logical capacity; ~13% of currently used capacity)

Note that the skew of the workload is a very important factor here. FAST VP will only move storage to the flash tier if very busy. Hence, the maximum flash capacity allowed will not always be consumed.

13.1.1.1 Over-subscription

The situation is potentially even more confusing in VP environments where over-subscription is used.

For example, consider the following situation taking multiple applications together. All applications have their own SG associated with a policy with 10%

Total size of pools:	400TB
Total provisioned storage:	640TB (i.e. 60% over-subscribed)
Total used from pools:	160TB

In total, 640TB is allocated, and hence 64TB (10% of total logical) is permitted to be allocated in flash.

If, for example, only 4% of the installed usable storage is flash, this would equate to 16TB in this instance i.e. a maximum possible demand on the flash storage of 4X that installed (and not the 2.5X that one might expect from comparing the 10% flash in the policy to the 4% flash installed by capacity).

14 Appendix E – Getting policy information via symcli

Note that FAST VP will not always be enable for all storage groups on a FAST VP array. It is switched on at the storage group level. In cases where there is no policy associated with a storage group, the hosts should just be marked as “Local RAID6” (in EMEA) and the Level (Silver or Gold) determined using whatever mechanism you use for non-FASTVP thin storage on VMAX.

How to determine if a device is under FAST VP control or not

- 1) Look up the storage group associated with the device
- 2) Determine if that storage group is under FAST VP control, and if so which policy.

There are two approaches to this:

- a) Check if there is a policy associated with an SG and only if confirmed that there is do you go and get the name
- b) Just try to get the name of the policy associated with the SG, and deal with the error if there is no policy.

The second is quicker as fewer steps. I will cover that option first.

`$ symfast -sid <SID> show -association -sg <SG>`

For example:

```
$ symfast -sid 1380 show -association -sg fastvp_SPA
```

```
Symmetrix ID          : 000192601380
```

```
Storage Group         : fastvp_SPA
```

```
Thin Devices(4)
```

```
{
-----

```

Sym	Flags P	Dev Emul	Total Tracks	Bound Pool Name	Allocated Tracks
0385	Y	FBA	552420	glab1_g10	276228
0386	Y	FBA	552420	glab1_g13	276228
0387	Y	FBA	552420	glab1_g13	276228
0388	Y	FBA	552420	glab1_g13	276228
Total			-----		-----
Tracks			2209680		1104912
GBs			135		67

```

}
```

```
Policy Name           : Production
```

```
Priority               : 3
```

```
Tiers(3)
```

```
{
-----

```

Tier Name	Type	Max Percent	SG Tech	Target Protection
thin_EFD_200GB	VP	10	EFD	RAID-5 (7+1)
thin_FC_450GB	VP	50	FC	RAID-6 (14+2)

```

}
```

```
thin_SATA_1TB          VP          100 SATA RAID-6(14+2)
}
```

If you run this command against an SG that is NOT associated with a policy, then you get:

```
$ symfast -sid 1380 show -association -sg sg_sangs4
```

```
Symmetrix ID          : 000192601380
```

```
No matching FAST association found
```

In such cases, the hosts should just be marked as “Local RAID6” (in EMEA) and the Level (Silver or Gold) determined using whatever mechanism you use for non-FASTVP thin storage on VMAX.

You can also get a list of all storage group that have associations with FAST VP policies, and the associated policy with the command

```
symfast -sid 1380 list -association
```

```
Symmetrix ID          : 000192601380
```

Storage Group Name	Policy Name	Pri
fastvp01_namgdse14	Production	1
fastvp01_sangs1	Production	1
fastvp01_sangs11	Production	1
fastvp02_namgdse14	Production	2
fastvp02_sangs1	Production	2
fastvp02_sangs11	SilverVP	3
fastvp_SPA	Production	3

Note that this does not list SGs that have no policy associated with them. The above policy names are just examples I the lab, not recommendations!!

Now onto the other option, when you check if there is a policy first before getting the name of the policy.

The following will give you a yes/no as to whether there is a FAST policy associated with a group or not, as well as a list of devices in that group:

```
$ symmsg show <SG> -sid <SYMMID>
```

For example:

```
$ symmsg show fastvp_SPA -sid 1380
```

```
Name: fastvp_SPA
```

```
Symmetrix ID          : 000192601380
Last updated at       : Thu Nov 18 03:02:43 2010
Masking Views         : No
FAST Policy         : Yes
```

```
Devices (4):
```

Sym	Device	Cap
Dev	Pdev Name	Config
		Sts
		(MB)

```

0385    N/A                TDEV          RW    34526
0386    N/A                TDEV          RW    34526
0387    N/A                TDEV          RW    34526
0388    N/A                TDEV          RW    34526
}

```

Note that FAST VP will not always be enable for all storage groups on a FAST VP array. It is switched on at the storage group level. In such cases, the hosts should just be marked as “Local RAID6” and the Level (Silver or Gold) determined using whatever mechanism you use for non-FASTVP thin storage on VMAX.

Alternatively, the following command may (or may not) also be of use to you. It lists ALL storage groups on an array, and if they are associated with a FAST VP policy or not.

```
$ symmsg list -sid 1380
```

S T O R A G E G R O U P S

```
Symmetrix ID:          000192601380
```

Storage Group Name	Number Devices	Use F M
-----	-----	-----
16test	16	. .
fastvp01_namgdse14	10	X X
fastvp01_sangs1	10	X X
fastvp01_sangs11	10	X X
fastvp02_namgdse14	10	X X
fastvp02_sangs1	10	X X
fastvp02_sangs11	10	X X
fastvp_SPA	4	X .
sangs11_sg	36	. X
sg-F5-DL385-X1	60	. X
sg_c14-p560-viocl4	6	. X
sg_g9_p550_vio_R2	2	. X
sg_gdse03	8	. X
sg_namgdse06	3	. X
sg_namgdse12	3	. X
sg_namgdse14	8	. X
sg_sangs4	2	. X

Legend:

Use Flags:

(F)ast Policy, X = Associated with Fast Policy, . = N/A

(M)asking View, X = Contained in Mask View(s), . = N/A

Once you have determined that a storage group is under FAST VP control, then you need to get the name of the policy, as shown above.

15 Appendix G – Using Unisphere with FAST VP

The following are suggestions on how to use Unisphere to help with management of FAST VP. These are based on the original metrics previous captured in SPA, and in addition to details shown within section 8.

For supplementary information please refer to the **Unisphere for VMAX Solutions Guide** (See References sections for URL)

Notes:

- Not all functionality from SPA has been ported in to Unisphere there are some instances where equivalents are not yet available or multiple ways to display similar/same information.
- Where SPA graphed using Snap shot and Trend views, Unisphere provides equivalents in Diagnostic and Historical views.
 - Real Time: Collects data between 2 and 5 seconds for a limited group of metrics – the data is available for the previous hour.
 - Diagnostic: Collects data every 5-10 minutes – data can be displayed for the last 1-24 hrs
 - Historical: Collects Data every 15 minutes – data can be displayed for 24hrs – 1yr (albeit with reduced granularity for periods > a few days).
- Unisphere for VMAX does not save **individual** disk statistics for historical and diagnostic views; instead it will save some statistics on activity and capacity breakdown by 'buckets' of disks – please see section 17.1.1.5 for more details.

15.1 Array Level information

15.1.1 Heat Maps

Provides a color coded graphical % utilization for Front End, Backend ports and other components depending on type of chart chosen.

Where: Symmetrix Array → <SID> → Performance → Monitor → EMC Defined Dashboards → Heatmap

15.1.1.1 Heat Map Type: Realtime



Displays Front end, Back End Director and disk Utilization. Cache % WP Used is also shown. The heatmap is updated dynamically.




Double Click Cache '%WP Used' status bar to display time series chart.



Double Click specific Front End director to show % Busy




Double Click specific Back End directors to show % Busy




15.1.1.2 Heat Map Type: Diagnostic




Displays Front and Back End Utilization. Also displays Director Port Utilization by 'Disk Bucket' (See 17.1.1.5)




Double Click specific Front End director to show % Busy per port versus % Busy for the array



Double Click specific Back End directors to show % Busy per port versus % Busy for the array (area chart style shown as an example)



Double Click Specific Directors to show Disk Distribution/% utilization




15.1.1.3 Heat Map Type: Historical




Displays Front and Back End Utilization

– Also displays Disk Group Utilization by 'Disk Bucket' (See 15.1.1.5)



Double Click specific Front End director to show % Busy per port versus % Busy for the array



Double Click specific Back End directors to show % Busy per port versus % Busy for the array (stacked line chart style shown as an example).

15.1.1.5 Disk Utilization Buckets


Unisphere for VMAX does not save **individual** disk statistics for historical and diagnostic views; instead it will save some statistics on activity and capacity breakdown by 'buckets' of disks.

Unisphere for VMAX creates 10 buckets for each configured disk group. The buckets are based on disk utilization (% Busy) where each bucket summarizes (or averages) key metrics of disks that their utilization in the last sample put them in the bucket.

- In diagnostic heatmaps, this is relevant to the 'Director Utilization Distribution' for the ports, which are displayed by utilization buckets.

Note: A text summary is shown when hovering the cursor over the director port


In the example below of the 7 disks accessed through director 7B port 0, 2 disks were utilized between 0-10 percent during the last sample period.




- In Historical Heatmaps – this is relevant to the 'Disk Group Utilization Distribution', which is displayed by utilization buckets

Note: A text summary is shown when hovering the cursor over the director port

In this example below 1% of disks in DISK_GROUP_012 were utilized between 11-20 percent utilized during the last sample period.



You can double click into the group to display all buckets and change chart type if required.




Note, on drilling into the disk groups, the bucket names displayed in a different format than shown on previous summaries – these granular bucket names are derived from the disk group name and the utilization range and prefixed by 'DG_'.

e.g. In the heatmap screen shot the text summary displays 'Utilization Bucket 20', on drilling into the disk group this is displayed as 'DG_12_11_20' which equates to Disk Group 12 statistics for the 'bucket' of disks that on average are 11-20% utilized during the last sample period.

Bucket name	Bucket Utilization Range	Heatmap Name
DG <DG> 00 10	00 <= X < 10	Utilization Bucket 10
DG <DG> 11 20	10 <= X < 20	Utilization Bucket 20
DG <DG> 21 30	20 <= X < 30	Utilization Bucket 30
.....		
DG <DG> 91 100	90 <= X <= 100	Utilization Bucket 100

As covered in section 17.1.1.4, from double clicking on a disk group, you can then 'Navigate to Analyze', which will provide the ability to graph additional metrics for the disk utilization buckets. (defined in the table below)



Metric Name	Description
Number of Disks	Number of disks in this bucket
Total Capacity (GB)	Total capacity of disks in this bucket
Used Capacity (GB)	Used capacity of disks in this bucket
Avg Response Time (ms)	Weighted average response time (read + writes) for disks in this bucket
Avg Read Response Time (ms)	Average read response time
Avg Write Response Time (ms)	Average write response time
IOs/sec	Average number of disk IOs of disks in this bucket
Reads/sec	Average number of disk Reads of disks in this bucket
Writes/sec	Average number of disk Writes of disks in this bucket
MBs/sec	Average number of MBs read and written to disks in this bucket
MBs Read/sec	Average number of MBs read to disks in this bucket
MBs Written/sec	Average number of MBs written to disks in this bucket
Avg Busy	Average %Busy of disks in this bucket
Avg Queue Depth	Average queue depth of disks in this bucket

15.1.2 Write Pending Count and Percentage

As shown from the real time heat map, you can drill in from that front pane to display the cache write pending %.

Alternatively the EMC Predefined **Array** Dashboard can provide this metric, along side


- BE Read & Writes/Sec
- Host MB Read & Writes/sec
- Host Reads & Writes/sec

Where: Symmetrix Array → <SID> → Performance → Monitor → EMC Defined Dashboards → Array




An alternative is to plot the required information and create a customized dashboard allowing you to add other relevant metrics should you wish to – e.g. Array Write Pending Count with System Max Write Pending Limit etc, etc.

Where: Symmetrix Array → <SID> → Performance → Analyze → Diagnostic (Or Historical) → select Symmetrix and plot array metric



Any of the graphed metrics can be saved to form a personalized dashboard.




15.2 Tier/Pool Level information

15.2.1 Tier level information:

Predefined **FAST by Tier** Dashboard will show the following parameters **by Tier selected**:


- Read/Write Response Time
- Ingress/Egress Tracks
- IO Density (by SG)
- BE requests (by SG)
- BE MB's Transferred (by SG)
- Allocated Tier Capacity (by SG)

Where: Symmetrix Array → <SID> → Performance → Monitor → EMC Defined Dashboards → FAST by Tier → Select Desired tier, Type and Data period to display



Also variety of metrics can be produced to create personalized dashboards

Where: Symmetrix Array → <SID> → Performance → Analyze → Diagnostic (Or Historical) → Virtual Pool Tier → Select desired tiers/Metric to plot




Note: For a quick view the Tiers Dashboard provides a capacity overview

Where: Symmetrix Array → <SID> → Storage → Tiers

Name	Type	Technology	Disk Location	Emulation	Used Capacity	Capacity (GB)	Protection
EFD_200GB	Virtual Pool	EFD	Internal	FBA	1 %	4855	RAID-5(7+1)
FC_600GB	Virtual Pool	FC	Internal	FBA	19 %	43899	RAID-6(14+2)
SATA_2TB	Virtual Pool	SATA	Internal	FBA	3 %	97711	RAID-6(14+2)
small_tier	Virtual Pool	FC	Internal	FBA	0 %	135	RAID-1

You can double click into each Tier for a more granular view




15.2.2 Pool level information

Predefined **All TP Pools** Dashboard will show the following parameters as a snapshot view:




- ALL Thin Pool Total Capacities
- ALL Thin Pool Used Capacities
- ALL Thin Pool Read/Write response times
- ALL Thin Pool Back End Read/Write requests Per Second

Where: Symmetrix Array → <SID> → Performance → monitor → EMC Defined Dashboards → All TP Pools → Select Type and Data period to Display



A variety of metrics can be produced to create personalized dashboards.


Where: Symmetrix Array → <SID> → Performance → Analyze → Diagnostic (Or Historical) → TP Pool → Select desired Thin Pool/Metric to plot



Note: For a quick view the System Dashboard provides a capacity overview showing:

- Physical Capacity – Used%/Available GB
- Virtual Capacity – Used%/Available GB
- Thin Pool – Used%/Available GB


Where: Symmetrix Array → <SID> → System → Dashboard



From here you can double click into to the **thin pools** dashboard to see more details.

000192604261 > Storage > Thin Pools						
Thin Pools						
Name	Technology	Configuration	Emulation	Allocated Capacity	Enabled Capacity (GB)	
glab261_10	EFD	RAID-5 (7 + 1)	FBA	1 %	4855.23	
glab261_11	FC	RAID-6 (14 + 2)	FBA	87 %	539.47	
glab261_12	FC	RAID-6 (14 + 2)	FBA	19 %	43899.18	
glab261_13	SATA	RAID-6 (14 + 2)	FBA	3 %	97711.08	
small_pool	FC	2-Way Mir	FBA	0 %	134.87	

At either dashboard (System or Thin Pool) you can double click into each pool for a more granular view




15.3 Storage Group Level information

Predefined FAST by Storage Group Dashboard will show the following parameters as a time series view:


- Read & Write Response Times
- Host IO & Back End requests/Sec
- Back End MB's Transferred /Sec (by Tier)
- Back End Requests (by Tier)
- IO Density (by Tier)
- Allocated Capacity (by Tier – shown in previous slides)

Where: Symmetrix Array → <SID> → Performance → Monitor → EMC Defined Dashboards → FAST by Storage Group → Select Desired SG and Type to plot



If Snap shot views are required you can plot these and create a personalized dashboard.

Where: Symmetrix Array → <SID> → Performance → Analyze → Diagnostic (Or Historical) → Storage Group → Double click required Storage Group → Storage Group by Tier → Select Tier/Metric to plot



15.4 Device level Information

If required device level information can be viewed/plotted by Storage group

Where: Symmetrix Array → <SID> → Performance → Analyze → Storage Group → Double Click SG → Select Devices to plot → Select desired tiers/Metric to plot



16 Appendix H – VMAX40K 2-tier (Gold) FAST VP configurations

16.1 Introduction

For VMAX40K 2-tier FAST VP configurations, use the VMAX40K 3-tier Best Practice described in the main body of this document, **unless over-riding Best Practice is provided below.**

2-tier (EFD & SAS) FAST VP configurations should only be used to provide non-SRDF “Gold” storage for franchise critical applications.

16.2 Supported configurations

VMAX40K 2-tier "GOLD" FAST	Number of Engines	Thin Pool Capacity (Eng TB)	EFD drives per loop	SAS drives per loop	SATA drives per loop	EFD drives	SAS drives	SATA drives	Additional Spare Drive Count	Total Drives	% of EFD capacity	% of FC capacity	% of SATA capacity
"Small Gold" FAST VP config VMAX40K 4-Engine (4 * 192GB) (768GB total cache) - 100GB 2.5" SSD RAID5(3+1) - 300GB 2.5" 10K SAS RAID1	4	189	6	42	0	192	1344	0	4*EFD 28*SAS	1568	6.8%	93.2%	0%

Note that these configurations differ from the 3-tier configurations in that:

- They have no SATA disk
- They have 192GB cache per VMAX40 engine, rather than 96GB (as in 3-tier Silver configurations)

Please see Appendix H for further documentation on 2-tier (Gold) FAST configurations. . The above configuration fills a 5-bay VMAX40K configurations and should not be expanded further. The 189TB capacity has been selected to meet failure domain requirements. Furthermore, there is no possibility of emergency expansion should the SAS pool start to show very high levels of utilization. Instead, it will be necessary to run zero-space reclaim and/or migrate some applications off of the array.

16.2.1 Gold Configuration capabilities

16.2.1.1 Resilience for RAID1 configuration

The mean time between pool failures for the 1344 SAS disk R1 thin pool is 107,020 years.


16.2.1.2 Workload profile use for performance analysis

Assumed workload parameters	
Written TB:	Assume maximum 85% of total usable pool storage allocated from pool
Sub-LUN skew:	85/15 i.e. 85% of IO goes to 15% of storage capacity For 6.8% SSD & 93.2% SAS split by usable capacity, and for a nominal workload skew of 85 (85% of IO to 15% of storage) the IO split by tier is: ~59% IOs to SSD ~41% IOS to SAS
Read hits	30%
Random reads	15%
Random writes	45%
Sequential writes	10%
Random write IO size	8KB
Sequential write IO size	64KB


16.2.1.3 Symmerge analysis with above workload profile

- At 5ms of latency R1 SAS configuration shows 185,000 IOPS. The accompanying graph shows the System RT (green) and BE CPU (pink) plots intersecting at 190K IOPS. This indicates that both disk and BE director resource utilization are approaching maximum levels in parallel.


Response Time depending on FrontEnd IOs (alpha)



Response Time depending on FrontEnd IOs (alpha)



- The R1 SAS configuration provides a higher level of resource utilization efficiency; where, both the disks and director (controllers) are reaching maximum utilization in parallel.



- The Gold configuration shows 76% SSD and 62% SAS drive utilization at maximum levels.
- BE loops will support a consistent 4926 IOPS for a maximum utilization of 94% & SSD drives at 76% utilized and the SAS disks at 62% utilized.

16.3 2-tier specific Best Practice

The 3-tier Best Practice in the main body of this document (Sections 5) applies to 2-tier configurations unless superseded below.

- The bound pool should be the SAS tier
- There should be one Flash and one SAS tier.
- “Allocation by policy” should be enabled (see Section 5.2.1)
- The following max SG percent setting should be used for the FAST policy:

	Flash Tier	FC/SAS Tier
VMAX40K (2-tier “Gold”)		
Gold_Production policy	20%	100%

Max SG percent settings for policies on 3-tier array

Max_SG for Flash has been increased from 10% to 20% cw. Silver array policy as Gold array has 7%, rather than 3-4% in Silver arrays.

- The following FAST VP control parameters should be used for VMAX40K 2-tier FAST configurations:

Control parameter argument	Citi Recommended setting	Is the default value?	Is the same as 3-tier?
-min_perf_period	8 (h)	Yes	Yes
-vp_data_move_mode	AUTO	Yes	Yes
-workload_period	120 (h)	No	Yes
-vp_reloc_rate	Initial Values 7 for “new application” arrays or existing arrays being converted from single SAS tier FAST VP. 6 for tech refresh target arrays (a rate of 6 is 17% more aggressive than a rate of 7) If array is mixed usage (new apps and tech refresh) use a rate of 6. Mature array value (existing applications growing organically – no new apps being added) 5 (A rate of 5 is 40% more aggressive than a rate of 7)	Yes	Yes
Pool level reserve capacity. Set with symconfigure (see Appendix B for an example). -pool_resv_cap	- VMAX40K: EDF pool: 1% VMAX40K: SAS Pool: 1% (20-30% PRC for FC/SAS tier needed for 3-tier configuration does not apply to 2-tier configs. See next subsection for further details)	No	No!
-vp_allocation_by_fp	Enable	No	Yes

See Section 5.3 for further details on each control parameter.

16.4 Planning a migration under FAST

ST VP control

The 3-tier Best Practice in the main body of this document (Section 6) applies to 2-tier configurations unless superseded below.

Migration to a 2-tier FAST configuration is easier than to a 3-tier configurations for the following reasons:

- A 2-tier configuration is best considered as a single-tier SAS configuration with some additional (7% by capacity) flash drives.
- The bound pool (from which most storage allocations will take place) is very large (~93% of the usable capacity of the array).
- Is the lowest tier and hence will always have the most free space.
- Has high-performance.

Hence, there is no concept of a FC/SAS migration landing space as with 3-tier configurations, **and no limits on the amount of data that can be migrated to the array on a given weekend.**

16.5 Worked example of setting up 2-tier FAST VP

The following VMAX40K example demonstrates the main steps for setting up 2-tier FAST VP. It assumes that the thin pools are already configured, but does also show the creation of devices, masking views and storage groups for completeness.

Note that the tech type for SAS drives is “FC”.

VP pool names:

- EFD pool: EFD100GB
- SAS pool: SAS300GB

Storage Tier names:

- EFD tier: thin_EFD_100GB
- SAS tier: thin_SAS_300GB

Production policy max SG percentages:

- EFD: 20%
- SAS: 100%

Enable FAST VP controller (one-off task)

```
symfast -sid 1234 enable -vp
```

Create Storage Tiers (one-off task)

```
symtier -sid 1234 create -name thin_EFD_100GB -tgt_raid5 -tgt_prot 3+1 -tech EFD -thin -pool EFD100GB
symtier -sid 1234 create -name thin_SAS_300GB -tgt_raid1 -tech FC -thin -pool SAS300GB
```

Create “Production” policy (one-off task)

```
symfast -sid 1234 -fp create -name Production
symfast -sid 1234 -fp -fp_name Production add -tier_name thin_EFD_200GB -max_sg_percent 20
symfast -sid 1234 -fp -fp_name Production add -tier_name thin_SAS_300GB -max_sg_percent 100
symfast -sid 1234 show -fp_name Production
```

Modify time windows

Note: Best Practice is to use default time 7x24 windows unless there is reason to believe that there are extended periods of time when the workloads managed by FAST VP are not active.

Set other FAST VP parameters

```
symfast -sid 1234 set -control_parms -min_perf_period 8
symfast -sid 1234 set -control_parms -vp_data_move_mode AUTO
symfast -sid 1234 set -control_parms -workload_period 120
symfast -sid 1234 set -control_parms -vp_reloc_rate 6

symconfigure -sid 1234 -cmd "set pool EFD100GB, type=thin, pool_resv_cap=1;" commit
symconfigure -sid 1234 -cmd "set pool SAS300GB, type=thin, pool_resv_cap=1;" commit
```

The following command to enable FAST VP allocation by policy

```
symfast -sid 1234 set -control_parms -vp_allocation_by_fp enable
```

Create devices bound to SAS tier pool (if not already created)

```
symconfigure -sid 1234 -cmd "create dev count=60, size=34526MB, emulation=FBA, config=TDEV, binding to pool= SAS300GB, preallocate size=0 M;" commit -nop
```

Create Storage Groups (if not already created)

```
symaccess -sid 1234 create -name fastvp01_sangs1 -type storage devs 0313:031C
symaccess -sid 1234 create -name fastvp02_sangs1 -type storage devs 031D:0326

symaccess -sid 1234 create -name fastvp01_sangs11 -type storage devs 0327:0330
symaccess -sid 1234 create -name fastvp02_sangs11 -type storage devs 0331:033A
```

Create Initiator groups (if required)

```
symaccess -sid 1234 create -name ig_sangs1 -type initiator -wwn 10000000c9368305
symaccess -sid 1234 add -name ig_sangs1 -type initiator -wwn 10000000c93dd604

symaccess -sid 1234 create -name ig_sangs11 -type initiator -wwn 10000000C9517FF5
symaccess -sid 1234 add -name ig_sangs11 -type initiator -wwn 10000000C95172CF
```

Create port groups (if required)

```
symaccess -sid 1234 create -name pg_sangs1 -type port -dirport 7E:1,10F:1
symaccess -sid 1234 create -name pg_sangs11 -type port -dirport 9E:1,8F:1
```

Now create the masking views (if required)

```
symaccess -sid 1234 create view -name view_fastvp01_sangs1 -sg fastvp01_sangs1 -ig ig_sangs1 -pg pg_sangs1
symaccess -sid 1234 create view -name view_fastvp02_sangs1 -sg fastvp02_sangs1 -ig ig_sangs1 -pg pg_sangs1

symaccess -sid 1234 create view -name view_fastvp01_sangs11 -sg fastvp01_sangs11 -ig ig_sangs11 -pg pg_sangs11
symaccess -sid 1234 create view -name view_fastvp02_sangs11 -sg fastvp02_sangs11 -ig ig_sangs11 -pg pg_sangs11
```

Associate these storage groups with the policy to bring under FAST VP control

```
symfast -sid 1234 -fp_name Production associate -sg fastvp01_sangs1 -priority 2
symfast -sid 1234 -fp_name Production associate -sg fastvp02_sangs1 -priority 2

symfast -sid 1234 -fp_name Production associate -sg fastvp01_sangs11 -priority 2
symfast -sid 1234 -fp_name Production associate -sg fastvp02_sangs11 -priority 2
```

The priority setting of the storage group decides which group gets to use a resource when there is competition for that same resource. Priority 1 has the highest priority. The default of 2 should usually be used, as shown here.

At this point, the storage groups are under FAST VP control. Once I/O begins to the devices, and allocations occur, FAST VP will begin to migrate the thin extents from the SAS tier to flash tier to optimize tier placement based on actual I/O and hence required performance. Performance migrations will not take place until the min_perf_period has elapsed after the storage group is associated with a policy.

16.6 Worked example of adding an SSD tier to single tier (SAS) VMAX40K

The following example demonstrates the main steps for converting a single tier (non-FAST) SAS VMAX40K configuration to a 2-tier FAST VP configuration by adding an SSD tier. It assumes that the thin pools are already configured. Furthermore, it does not show the creation of devices, masking views and storage groups, as these are presumed to have been previously created when running as a single tier array. It does show how to associate these storage groups with the FAST policy.

Note that the tech type for SAS drives is "FC".

VP pool names:

- EFD pool: EFD100GB
- SAS pool: SAS300GB

Storage Tier names:

- EFD tier: thin_EFD_100GB
- SAS tier: thin_SAS_300GB

Production policy max SG percentages:

- EFD: 20%
- SAS: 100%

Enable FAST VP controller (one-off task)

```
symfast -sid 1234 enable -vp
```

Create Storage Tiers (one-off task)

```
symtier -sid 1234 create -name thin_EFD_100GB -tgt_raid5 -tgt_prot 3+1 -tech EFD -thin -pool EFD100GB
symtier -sid 1234 create -name thin_SAS_300GB -tgt_raid1 -tech FC -thin -pool SAS300GB
```

Create “Production” policy (one-off task)

```
symfast -sid 1234 -fp create -name Production
symfast -sid 1234 -fp -fp_name Production add -tier_name thin_EFD_200GB -max_sg_percent 20
symfast -sid 1234 -fp -fp_name Production add -tier_name thin_SAS_300GB -max_sg_percent 100
symfast -sid 1234 show -fp_name Production
```

Modify time windows

Note: Best Practice is to use default time 7x24 windows unless there is reason to believe that there are extended periods of time when the workloads managed by FAST VP are not active.

Set other FAST VP parameters

```
# symfast -sid 1234 set -control_parms -min_perf_period 8
# symfast -sid 1234 set -control_parms -vp_data_move_mode AUTO
# symfast -sid 1234 set -control_parms -workload_period 120
# symfast -sid 1234 set -control_parms -vp_reloc_rate 6

# symconfigure -sid 1234 -cmd "set pool EFD100GB, type=thin, pool_resv_cap=1;" commit
# symconfigure -sid 1234 -cmd "set pool SAS300GB, type=thin, pool_resv_cap=1;" commit
```

The following command to enable FAST VP allocation by policy

```
# symfast -sid 1234 set -control_parms -vp_allocation_by_fp enable
```

Associate these storage groups with the policy to bring under FAST VP control

```
symfast -sid 1234 -fp_name Production associate -sg fastvp01_sangsl -priority 2
symfast -sid 1234 -fp_name Production associate -sg fastvp02_sangsl -priority 2
etc etc
```

The priority setting of the storage group decides which group gets to use a resource when there is competition for that same resource. Priority 1 has the highest priority. The default of 2 should usually be used, as shown here.

At this point, the storage groups are under FAST VP control. Once I/O begins to the devices, and allocations occur, FAST VP will begin to migrate the thin extents from the SAS tier to other tiers to optimize tier placement based on actual I/O and hence required performance. Performance migrations will not take place until the min_perf_period has elapsed after the storage group is associated with a policy.