

ENABLING SYMMETRIX FOR FAST WITH FTS FOR 3RD PARTY STORAGE

Gaurav Roy
EMC



Table of Contents

Introduction	3
Business Challenges.....	4
EMC VMAX.....	5
VMAX Architecture Design	5
VMAX Key Facts	5
Technology overview	6
Operating Environment.....	6
Device Types in VMAX.....	7
Virtual Provisioning Terminology	9
Fully Automated Storage Tiering	11
FAST VP Architecture	13
FAST device movements.....	16

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

Introduction

Mergers and acquisitions have become norm in today's business environment. There is an increased focus on how to be a market leader in the industry by giving the best to customers in terms of product and after-sales service to ensure that customer satisfaction is met.

EMC introduced the Symmetrix® VMAX® line of products with an aim to become a block storage leader in the market and also leverage them to have the thirst of performance quenched with the add on features to offer on plate.

Some remarkable features of Symmetrix VMAX starting from the global cache to scaling of data center environment based on requirement to having thin provisioned storage to Enterprise Flash Drives (EFD) usage to Data At Rest Encryption (DARE) to Fully Automated Storage Tiering (FAST) to Federated Tiered Storage(FTS) to Federated Live Migrations(FLM) and still the add on features keep increasing with the growing needs of the data center market.

These solutions must be capable of handling today's requirements as well as flexible enough to meet tomorrow's emerging requirements.

Business Challenges

Most customers face similar business-related storage challenges:

- Exponential data growth
- Flat to lower storage budgets (doing more with less)
- Using existing storage with new storage systems (optimum resource utilization)
- Performance/Capacity Tiering for workloads
- Non-disruptive Data Migrations between storage subsystems (technology refresh)
- Shrinking to near-zero Recovery Point Objectives (RPO's)
- Shrinking to near-zero Recovery Time Objectives (RTO's)
- Shrinking to non-existing maintenance windows
- Meeting Increasingly stringent performance Service Level Agreements (SLA's)

This Knowledge Sharing article highlights how to use the FAST feature along with FTS-qualified storage subsystems so that optimum performance can be delivered using existing FTS-qualified vendor storage subsystems.

EMC VMAX

Customers have huge databases and highly intensive performance critical applications along with low workload demanding applications that require an enterprise-class, nonstop storage subsystem. EMC VMAX\VMAXe is a perfect choice as it delivers ultimate performance, scale, and reliability.

We often underutilize some of our existing arrays due to the fact that they suffer from performance-related issues when storage utilization increases. VMAX allows us to scale while adding performance (if needed). As we add storage bays, we can also add processing, cache, and bandwidth to adequately handle the additional storage characteristics. Systems requiring ultimate performance and/or ultimate reliability are placed on the VMAX platform.

VMAX Architecture Design

A single VMAX scales from 1 to 8 engines (4 engines in the case of 10k) and is capable of supporting up to:

- 1,080 drives to 2,400 drives (EFD, SAS, and/or SATA)
- 512GB to 2048GB Cache
- 32 to 192 CPU cores
- 64 to 128 front-end fiber ports
- 32 to 64 FCoE, iSCSI, and FICON ports (excluding 10k which does not support)

VMAX Key Facts

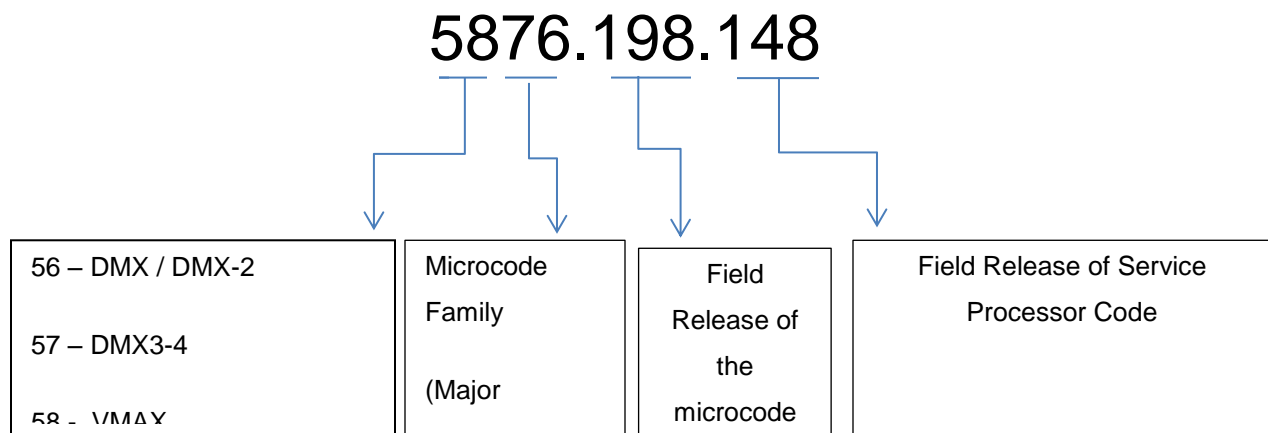
- Ultimate Scale and Performance
- Virtual Provisioning
 - Easy to use and fast storage provisioning capabilities
- Fully Automated Storage Tiering Virtual Provisioning (FAST VP)
 - Automated data migration between different tiers of storage
 - Non-disruptive data migrations within the VMAX
 - Automatically moves data to the proper drive tier (Flash, Fiber Channel, SAS, and SATA drives)
 - Lowers overall storage costs due to increased utilization of 7.2k SATA drives
- Symmetrix Performance Analyzer (SPA)
 - Automated performance monitoring and trending charts
- Symmetrix Management Console (SMC) and Command Line Interface (CLI)
 - Easy to use management consoles

Technology overview

Operating Environment

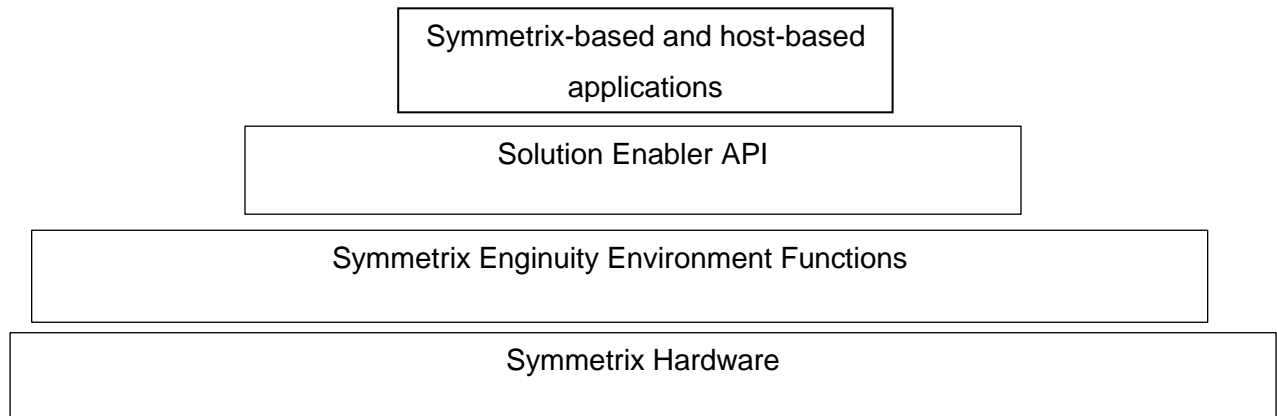
Engenuity is the operating environment in a VMAX .Each processor in each director of an engine is loaded with Engenuity.

Engenuity makes the independent director processors act as one integrated cached disk array. It also supports features such as Symmetrix Remote Data Facility (SRDF[®]) and TimeFinder[®].



To configure the FTS feature, VMAX must be running Engenuity 5876 code level.

The Engenuity operating layers can be presented as:



The initial configurations of the Symmetrix are done by the Symm Win Application running on the service processor of the Symmetrix. The “impl.bin” file contains configuration information of a Symmetrix, such as physical directors, emulation , physical drives, memory, data protection, and so on.

The “impl.bin” file can be altered by EMC personnel only. Subsequent changes need to be done online through management tools such as EMC Solution Enabler Interface CLI, EMC Symmetrix Management Console (SMC), EMC Unisphere® for VMAX, EMC ControlCenter®, and ProSphere®.

Device Types in VMAX

- Standard Volumes, or STD, are configured for normal production operations.
- Virtual Devices, or VDEVs, are configured for TimeFinder/Snap local pointer-based replication.
- TDEV devices are configured for Virtual Provisioning.
- R1 and R2 volumes are configured for remote SRDF replication.
- Save Devices are configured for TimeFinder/Snap and SRDF/A DSE.
- Data Devices are configured for Virtual Provisioning.
- Dynamic Reallocation Volumes, or DRV devices, are configured for Symmetrix Optimizer and FAST.

Supported RAID types

RAID 1

- High Performance and availability
- Two hyper volumes on separate physical drives

RAID 5

- Distributed parity and data is spread across all devices in RAID group
- RAID 5 (3+1): Parity and data striped across four hyper volumes
- RAID 5 (7+1): Parity and data striped across eight hyper volumes

RAID 6

- Double distributed parity (diagonal and horizontal) data is spread across all devices in RAID group
- RAID 6 (6+2): Parity and data striped across eight hyper volumes
- RAID 6 (14+2): Parity and data striped across sixteen hyper volumes

RAID 10

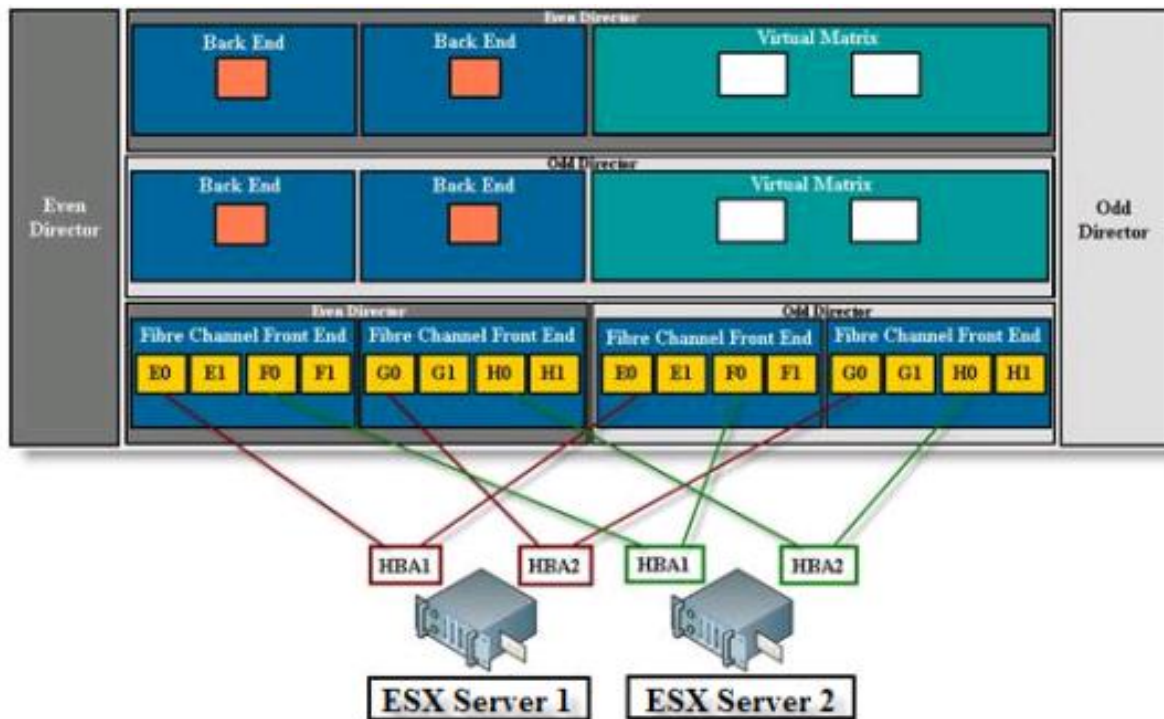
- Combination of RAID 1 (mirrored) and RAID 0 (no protection) used for mainframe environments

RAID 0

- Striping is done with no protection to the data

Zoning Configuration in VMAX

The even and odd director ports are zoned to HBA1 and HBA2 ports of ESX servers as per the diagram below. The red and green colour signify there are two different fabrics through which the SAN is zoned for these servers.



Virtual Provisioning Terminology

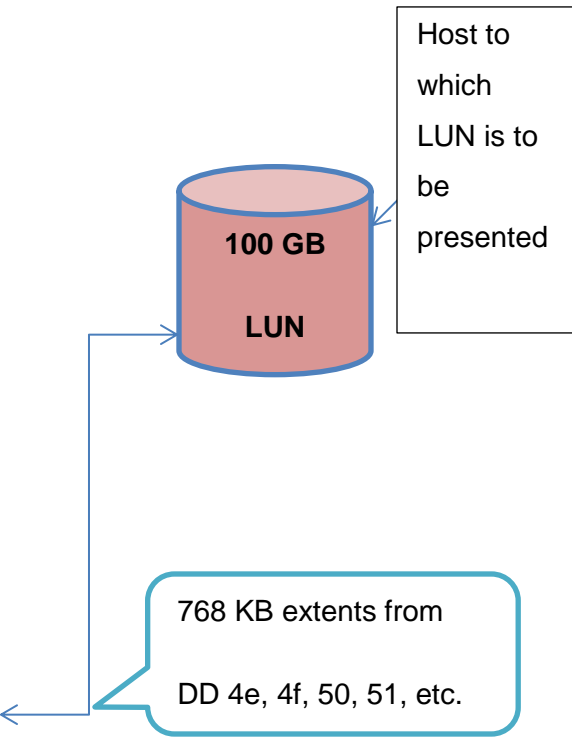
Virtual Provisioning presents an application with more capacity than is physically allocated and in some situations, may provide a more efficient way of allocating capacity for applications that are somewhat predictable in capacity growth patterns.

It makes use of several other terminologies and devices to function properly.

- Data devices (TDAT) consist of the underlying physical disks that will store the actual data. These devices are part of the thin pool and are treated as array devices that are inaccessible to the host.
- Thin devices (TDEV) are devices of a user-defined size that point to the available space on the data devices in a thin pool. TDEVs use the storage space in a pool and must be bound to a thin pool prior to being usable. A TDEV can only belong to one pool at a time and will inherit the RAID protection of the data devices that make up the chosen pool. These are the devices that are presented to the host.
- Thin pool is a grouping of data devices that are used by TDEVs. All devices within the thin pool must be of the same emulation and RAID protection type though sizes can vary. TDEVs can be added to the pool as long as there are free tracks available.

Storage Pool Design example is shown below.

100 GB X 24 EFD disks RAID 5(3+1)					
100 GB	100 GB	100 GB	100 GB		2a,30,36,etc
100 GB	100 GB	100 GB	100 GB		2b,31,37,etc
100 GB	100 GB	100 GB	100 GB		2c,32,38,etc
100 GB	100 GB	100 GB	100 GB		2d,33,39,etc
100 GB	100 GB	100 GB	100 GB		2e,34,40,etc
100 GB	100 GB	100 GB	100 GB		2f,35,41,etc
500 GB X 35 10k FC Disks RAID 5(3+1)					
500 GB	500 GB	500 GB	500 GB	500 GB	4e,55,6b,72,etc
500 GB	500 GB	500 GB	500 GB	500 GB	4f,56,6c,73,etc
500 GB	500 GB	500 GB	500 GB	500 GB	50,57,6d,74,etc
500 GB	500 GB	500 GB	500 GB	500 GB	51,58,6e,75,etc
500 GB	500 GB	500 GB	500 GB	500 GB	52,59,6f,76,etc
500 GB	500 GB	500 GB	500 GB	500 GB	53,60,70,77,etc
500 GB	500 GB	500 GB	500 GB	500 GB	54,6a,71,78,etc
1TB X 48 7.2k SATA RAID 6(6+2)					
1 TB	1 TB	1 TB	1 TB	1 TB	79,87,9e,11a,etc
1 TB	1 TB	1 TB	1 TB	1 TB	80,88,9f,11b,etc
1 TB	1 TB	1 TB	1 TB	1 TB	81,89,10a,11c,etc
1 TB	1 TB	1 TB	1 TB	1 TB	82,90,10b,11d,etc
1 TB	1 TB	1 TB	1 TB	1 TB	83,9a,10c,11e,etc
1 TB	1 TB	1 TB	1 TB	1 TB	84,9b,10d,11f,etc
1 TB	1 TB	1 TB	1 TB	1 TB	85,9c,10e,12a,etc
1 TB	1 TB	1 TB	1 TB	1 TB	86,9d,10f,12b,etc



Thin pools are comprised of TDATs, from which 768 KB extents are taken to create a TDEV. The 100 GB TDEV volume is created by taking 768 KB slices from each disk drive until the required size is fulfilled, effectively utilizing all of the drives in that pool.

Fully Automated Storage Tiering

Fully Automated Storage Tiering (FAST™) automates identification of active and inactive data for the purpose of moving application data across different performance or capacity tiers within an array. FAST monitors the I/O activity over user defined windows and performs auto migration of 'hot' data to a higher storage tier and 'cold' data to a lower storage tier.

It Improves application performance at the same cost or provides the same application performance at lower cost. FAST eliminates the need to manually tier applications when workload characteristics change over time.

FAST automates identification of data that should reside on a given drive technology and moves data between tiers automatically and non-disruptively to optimize storage resource usage accordingly.

Need for FAST

Over time, drive types supported by VMAX have become more robust and drive technologies have been made to match the current data center requirements.

Storage drives can be differentiated on the basis of response time, cost per unit of storage capacity, and cost per unit of I/O request processing.

Enterprise flash drives have a very low response time but can handle high volume of I/O requests enabling high cost per unit storage. SATA or NL-SAS drives have low cost per unit of storage with high response times. Between the above two types, there are FC and SAS drives.

Depending on the differences between the drive types, here are a few observations regarding the suitable workload type:

- EFDs are more suited for workloads that have a high back-end random-read storage request density. Such workloads take advantage of both the low service time provided by the drive, and the low cost per unit of storage request processing, without requiring a lot of storage capacity.
- SATA and near-line SAS drives are suited to workloads that have a low back-end storage request density.
- FC and SAS drives are the best drive type for workloads with a back-end storage request density that is neither consistently high nor low.

FAST DP

It operates on disk groups provisioning of Symmetrix volumes. The data movements based on the workload analysis are executed between the tiers and performed on the LUN level.

FAST VP

As the name suggests, FAST VP operates on virtual provisioned thin devices with data movements performed on the sub-LUN or extent level. A single thin LUN will have extents allocated from several data devices in a particular thin pool or multiple thin pools across an array.

FAST Managed Objects

There are three main elements related to FAST on Symmetrix VMAX arrays.

1. Storage Tier

A shared resource created with a name which contains common drive technologies and RAID protection to facilitate data movement from the storage group based on the FAST policy.

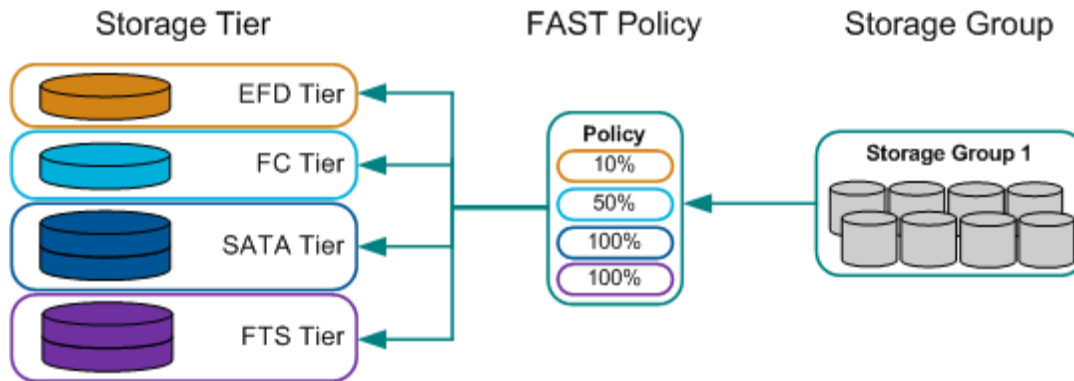
2. FAST policy

A set of tier usage rules that provide guidelines for data placement and movement across Symmetrix tiers to achieve service levels for one or more storage groups.

3. Storage group

A logical grouping of devices for common management across the array for various purposes.

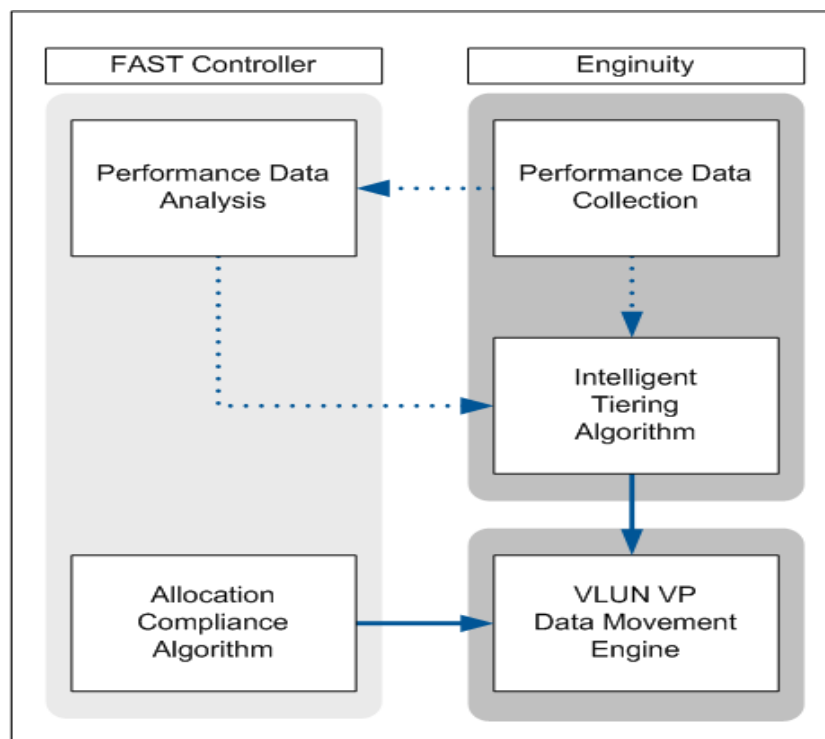
An example of the above components is shown in the figure below.



FAST VP Architecture

Symmetrix Enginuity code is the storage operating environment that controls components within the array. The FAST controller is a service that runs on the service processor.

There are two algorithms—the intelligent-tiering algorithm and the allocation-compliance algorithm—which help the Symmetrix determine appropriate data placement from the storage groups associated on the basis of policy within the storage tiers.



The intelligent-tiering algorithm uses performance data collected by Enginuity as well as supporting calculations performed by the FAST controller to issue data-movement requests to the VLUN VP data-movement engine.

The allocation-compliance algorithm enforces the upper limits of storage capacity that can be used in each storage tier by a given storage group by also issuing data-movement requests to the VLUN VP data-movement engine.

Performance time windows can be defined to specify when the FAST VP controller should collect performance data. Analysis is then performed to determine the appropriate tier for devices. By default, performance data collection occurs 24 hours a day.

FAST VP states

There are five reported states for FAST VP.

Enabled: All FAST VP functions are performed (performance data collection, performance data analysis, data-movement request generation, and data-movement execution).

Disabled: Only performance data collection is performed. Data analysis is not performed and data movements are not executed.

Disabling: The FAST controller is transitioning from Enabled to Disabled.

Disabled with Error: The FAST controller has stopped operation due to an internal error. Statistics collection and FAST VP performance data movements continue to be performed. However, FAST VP compliance movements are not performed.

Degraded: FAST VP can perform some or all of its functions. However, it cannot perform each function fully.

FAST VP time windows

There are two time windows when data is collected and moved. Possible window types are:

- Performance Time Window
- Data movement time window

Performance time windows are used to identify the business cycle for the Symmetrix array. They specify date and time ranges (past or future) when samples will be included in, or excluded from, the FAST performance data analysis. The intent of defining performance time windows is to distinguish periods of time when the Symmetrix is idle from periods when the Symmetrix is active, and to only include performance data collected during the active periods.

A default performance time window includes all performance data samples, 24 hours a day, 7 days a week, 365 days a year.

Device movement time windows are used to specify date and time ranges when moves or swaps are allowed, or not allowed, to be performed. While FAST device movements run as low priority tasks on the Symmetrix back end, they can introduce additional processing overhead. Device movement windows can be planned so as to minimize impact on the performance of more critical workloads.

FAST controller settings

Several factors affect the behaviour of the FAST controller. They are:

- **Operating mode**
Two modes of operation, i.e. user-approved mode and automatic mode. In user-approved mode, the FAST controller will generate configuration change plans hourly without any device movement across tiers until approved by the user. In automatic mode, configuration plans will be generated and will have the data movement executed without requiring user intervention.
- **Maximum moves per day**
Controls the maximum number of swaps or moves to be performed by FAST controller during a 24 hour period. Moves affect the back end configurations within Symmetrix so the maximum moves are configured between 2 and 200.
- **Maximum Simultaneous Moves**
Controls the concurrent moves or swaps from source to target pool as per policy associated. This also depends on the dynamic reallocation volumes (DRVs) configured on the array. It can be configured between 2 and 32.
- **Workload Analysis Period**
The amount of historical information maintained by FAST controller which help to determine performance-based data movements collected during the time windows. It can be configured between 1 hour and 4 weeks.

- **Initial Analysis Period**
The minimum amount of samples to be collected initially before generating the first configuration change plan. This can be configured between 1 hour and 4 weeks but should be properly set to capture the typical workload on the array.
- **Move Type**
Determines the type of device movements; only swaps or allow both swaps and moves. Set to disable to allow both swaps and moves.
- **Use host invisible devices**
Determines whether FAST can use configured but unmasked /unmapped devices as target devices during swaps. By default, it is set to disable and prohibits use of configured unmasked/unmapped devices.


FAST device movements

The two methods by which device can be relocated to another tier are:


- **Move:** This occurs when unconfigured space exists in target tier.

Consider an example where, per FAST, the device 100 located originally in RAID 1 Fibre Channel tier needs to be moved to RAID 5 (3+1) EFD tier as per the workload analysis and configuration plan. This can be done in five stages as seen below.

1. Currently, device can be seen to be contained in M1 mirror position before being moved.




2. Move request is submitted and configuration lock is placed to do this. After this, the target RAID group is created in target Symmetrix tier and attached as secondary mirror M2 of the device.

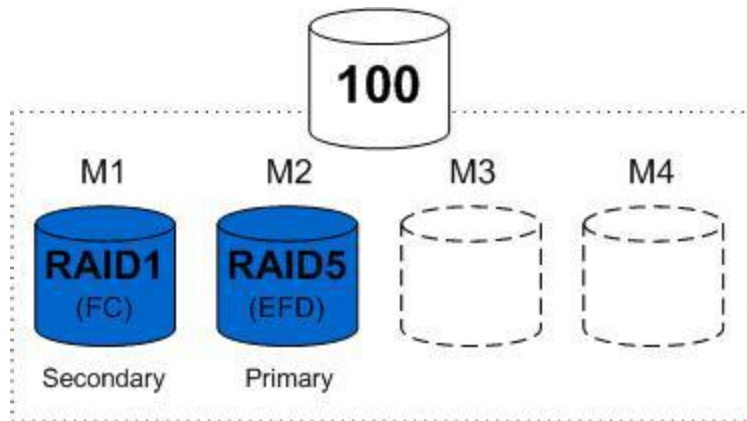


Once done, the secondary RAID group is completely invalidated to be synchronized from primary mirror and configuration lock is released.


3. As the secondary mirror was invalidated when it was associated with the volume, it is then synchronized from the primary mirror of the volume.



4. After synchronization, the lock is acquired once again and the device configuration is updated to swap primary and secondary roles of the mirror positions. Thus, the original RAID 1 mirror is demoted to secondary while new RAID 5 mirror is promoted to primary mirror.



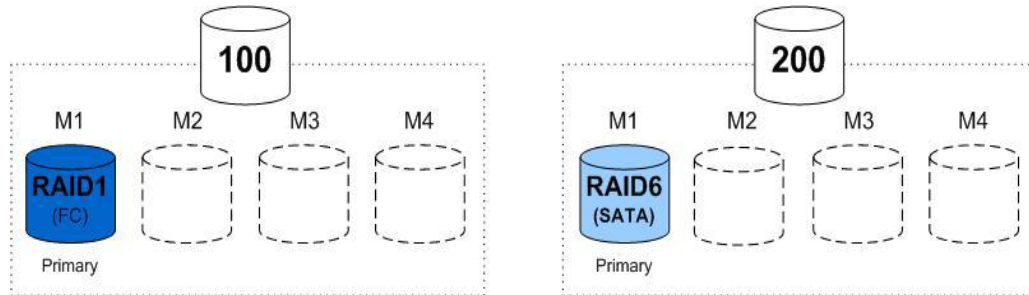
- When the swap is done between mirror roles, the original RAID group is detached from the device and deleted. Also, the space previously used by RAID 1 on Fibre Channel storage is returned to free pool and configuration lock is released.



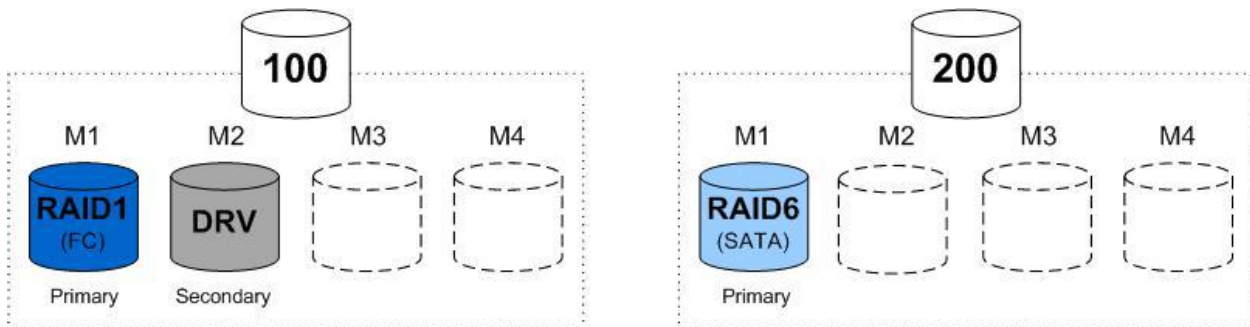
- Swap:** This will happen in cases where there is no unconfigured space in the target tier, resulting in moving a device with some other device in target tier.

Consider a swap example where FAST determined that device 100, which is currently in RAID 1 Fibre Channel, be demoted to RAID 6 (14+2) tier. Simultaneously, device 200 is to be promoted vice versa. This will be done in the eight steps shown below.



1. The devices can be seen as below which need to be swapped as per configuration plan and a swap request is submitted.





2. The configuration lock is placed on an array to perform the swap. Also, a DRV device is chosen and the associated RAID group is detached and associated with one of the devices being swapped.





3. In the above configuration change, the DRV is synchronized from the original RAID group associated with device 100.
4. After complete synchronization of the DRV device, the configuration lock is placed again. During this change session, the original RAID group of device 100 is detached from the device and associated with the other device in the swap, i.e. device 200. Lock is released after this.



- This will synchronize the newly attached RAID group from device 200's original RAID group.
- The configuration lock is placed again to detach the original RAID group from device 200 and attach it to device 100.



- When the lock is released, the RAID group attached to device 100 is synchronized from the DRV.
- After the synchronization is completed from DRV, the configuration lock is placed again for the last time. During this change, the DRV RAID group is detached from device 100 and re-attached to the DRV device.



Federated Tiered Storage

Federated Tiered Storage (FTS) allows LUNs that exist on external arrays to be used to provide physical storage for Symmetrix VMAX arrays. These LUNs can be used as raw storage space for creation of Symmetrix devices in the same way that internal Symmetrix physical drives are used. The devices configured from external arrays are referred to as eDisks.

Benefits of Federated Tiered Storage

- Scalability of Symmetrix features to existing multi-vendor storage environment
- Allows data mobility and migration between heterogeneous storage arrays and between VMAX and heterogeneous arrays
- The benefits of virtual provisioning along with enterprise replication technologies such as SRDF and TimeFinder and tiering features such as FAST can be used for multi-vendor storage arrays at a lower cost of technology refresh

Components of FTS

FTS is implemented on the Enginuity level and does not require additional Symmetrix hardware. The connectivity is established to external array front-end ports using fibre optic cables.

1. DX Directors

The front-end adapters (FA) from the Symmetrix need to be made as DA eXternal (DX ports) by loading the DA emulation on to them to enable them to act on the external logical units as though they were physical drives. It behaves the same as a DA working on Symmetrix internal backend disks. These directors can only be configured by EMC. Once DX emulation is loaded on the processors, it is completely user-configurable.

2. eDisks

It is the logical representation of an external LUN when it is added into the VMAX configuration.

3. External Disk group

These are the virtual disk groups created by the user to contain eDisks. The external disk group numbers start at 512 onwards. External spindles and internal physical spindles cannot be mixed in a disk group.

4. Virtual RAID group

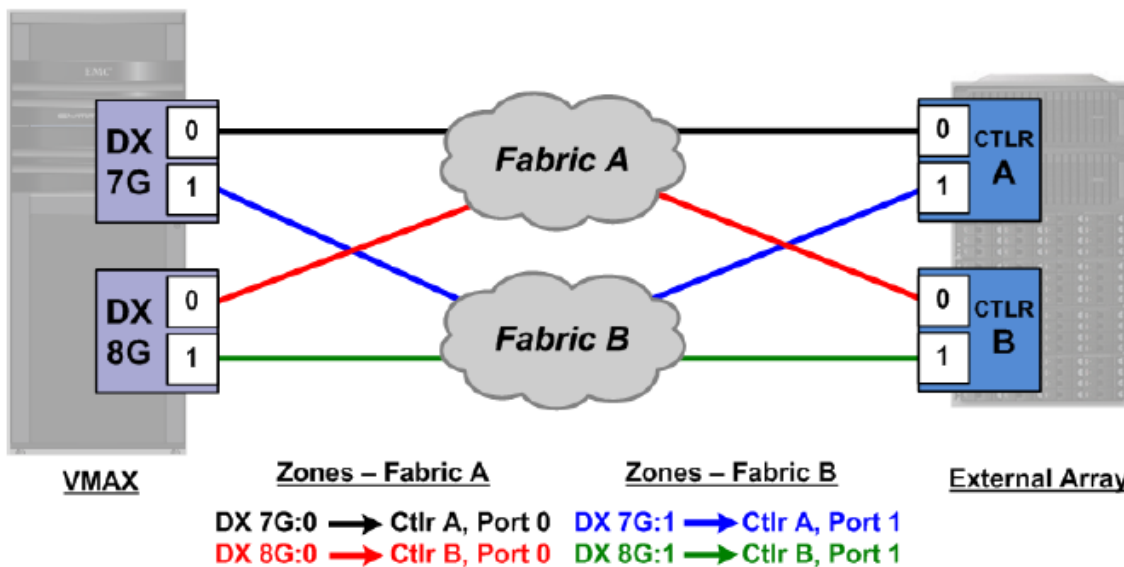
This is an unprotected virtual RAID group created for each eDisk added to the Symmetrix VMAX. It is a virtual RAID group as the eDisk are not protected locally by the VMAX and are dependent on the protection provided by external attached arrays.

FTS configuration

The VMAX should be running on Enginuity code 5876 and above to enable it for FTS. It starts with the creation of the DX directors on the VMAX array, followed by zoning it with the front end of the external array.

The zoning can be done in three ways – single fabric with two external storage ports, dual fabric with two external storage ports, and dual fabric with four external storage ports.

The best practice for redundancy is achieved by using the dual fabrics with four external storage array ports which can be portrayed as shown below.



In the above example, the VMAX array is running DX emulation on odd director 7G and even director 8G which is within engine 4, i.e. the initial engine shipped when VMAX is ordered to EMC.

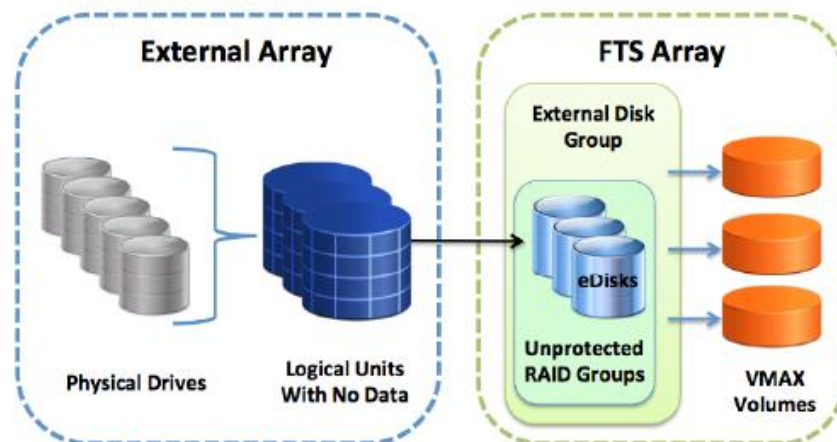
The director ports DX 7G:0, 7G:1 is zoned to external array controller A:0, A:1 via two redundant SAN fabric paths. Similarly, the zoning is done for DX 8G:0, 8G:1 with controller B:0, B:1 .

It is required for the external arrays to be located within the same data center as the VMAX arrays enabled with FTS. However, if the data center is spread across multiple floors in a single building, the external array and VMAX can be in different floors.

The FTS has two modes of operation depending on whether the external LUN is used as a raw storage space or has data that must be preserved.

1. External Provisioning

This allows the user to access the LUNs existing on external storage as raw capacity for new devices within Symmetrix.




2. Encapsulation


In this mode, users can preserve the existing data on external LUNs and access it through Symmetrix volumes. These Symmetrix devices, called encapsulated devices, provide an exact block-by-block match of the external LUN or eDisk. No Volume Table of Contents (VTOC) of the external data will occur.

This mode has two sub-modes; standard encapsulation and virtual provisioning encapsulation.

With standard encapsulation, the eDisk is created and added to an external Disk Group with a Virtualized RAID group. Symmetrix device(s) are created and assigned to the eDisk. If the external LUN is larger than 64 GB, Symmetrix will create a meta volume containing 1 or more Symmetrix devices assigned to the same eDisk. Whether the eDisk contains a single Symmetrix device or a meta device, the represented devices will contain an exact block-by-block representation of the external data contained on the external LUN within the external array.



In virtual provisioning encapsulation, the eDisk is created and added to an external disk group with a virtualized RG. Thin data devices (TDATs) are then created and added to a specified external thin pool. Thin devices will be fully allocated and a one-to-one thin extent to TDAT track is established.



The maximum external capacity is determined by the VMAX cache. The maximum number of LUNs per external storage port is 8,000 and 14,000 eDisks per DX director port with a higher limit of 128,000 eDisks per array.

The FTS can only work with FAST VP when using the external provisioning mode of operation. Hence, it cannot be used along with encapsulated devices.

The supported external array list with FTS on VMAX can be found in Appendix B of this article.

FAST VP with FTS set up: Case scenario

The zoning must be done between the DX ports and the external array front end ports.

Assume that DX ports 9F:0, 9F:1 and 10F:0,10F:1 are zoned with front end ports A:0,A:1 and B:0,B:1 in a dual fabric redundant configuration.

The connectivity can be verified by running the commands shown below from the host running Solution Enabler.

```
# symcfg -sid 98 list -DX all
```

```
Symmetrix ID: 000195700398
```

SYMMETRIX			DISK		DIRECTORS	
Ident	Symbolic	Numeric	Slot	Type	Num Of Serviced Hyper Volumes	Status
DX-9F	09F	89	9	EDISK	0	Online
DX-10F	10F	90	10	EDISK	0	Online

This confirms that the directors 9F , 10F are configured as DX ports from engine 5 of the VMAX.

The symsan command can be used to list that LUNs are available on the external storage by using `-sanlun` option and specifying external array port WWN.

```
# symsan -sid 98 -dir 9F -p 0 list -sanluns -wwn 5006048AD5F00BD0

Symmetrix ID:      000195700398
Remote Port WWN:   5006048AD5F00BD0

      ST
      A
      T  Flags  Block  Capacity  LUN  Dev  LUN
DIR:P  E  ICRTHS  Size    (MB)   Num  Num  WWN
-----
09F:0  --  X..F.X   N/A      N/A     0  0000  60060480000190300207000000003000
09F:0  RW  ...F.X   512      61425   0  0043  60060480000190300207533030303433
09F:0  RW  ...F.X   512      61425   0  0044  60060480000190300207533030303434
09F:0  RW  ...F.X   512      61425   0  0045  60060480000190300207533030303435
09F:0  RW  ...F.X   512      61425   0  0046  60060480000190300207533030303436
09F:0  RW  ...F.X   512      61425   0  0047  60060480000190300207533030303437
```

This illustrates that the external array's FA has been mapped to external devices 0043-0047 and are available as eDisks through the DX directors on VMAX array.

The external disk group is then created within the VMAX array with name `ext_prov_1` and disk group number 512 and would be listed as shown below

```
# symdisk -sid 98 list -dskgrp_summary

Symmetrix ID: 000195700398

      Disk Group          Disk          Capacity
-----
Num  Name                Cnt  Flgs  Speed  Size    Total    Free    Actual
-----
10  EFD_R53_DG10          16  IE    0      381470  6103515  9635   6103515
20  FC_R57_DG20           96  IF    15000  558281  53594939  40032238  53594939
21  FC_R1_DG21            80  IF    15000  558281  44662450  21800810  44662450

30  ATA_R614_DG30         32  IS    7200  1823565  58354066  29186746  58354066
512  ext_prov_1            0  N/A   N/A    N/A     0         0         N/A

Total                                162714970  91029429  162714970
```

Hence, the disk group created can be filled with eDisk to enable external provisioning. It will add raw capacity for creating Symmetrix devices as done from any other internal disk drive. The parameter `encapsulate_data` should be set to `NO` so that existing data on the disk is removed.

```
# symconfigure -sid 98 -cmd "add external_disk
wnn=60060480000190300207533030303433 to disk_group=512, encapsulate_data=NO;"
commit -nop
```

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

```
Establishing a configuration change session.....Established.
Processing symmetrix 000195700398
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Queued.
COMMIT requesting required resources.....Obtained.
Step 012 of 060 steps.....Executing.
Step 012 of 060 steps.....Executing.
Step 047 of 118 steps.....Executing.
Step 049 of 118 steps.....Executing.
Step 049 of 118 steps.....Executing.
Step 082 of 118 steps.....Executing.
Step 087 of 118 steps.....Executing.
Step 096 of 119 steps.....Executing.
Step 096 of 119 steps.....Executing.
Step 096 of 119 steps.....Executing.
Step 096 of 119 steps.....Executing.
Local: COMMIT.....Done.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

The disk group 512 would now be listed with a capacity.

```
# symdisk -sid 98 -disk_group 512 list
```

```
Symmetrix ID          : 000195700398
Disks Selected        : 1
Disk Group            : 512
Disk Group Name       : ext_prov_1
Disk Location         : External
Technology            : N/A
Speed (RPM)          : N/A
Form Factor           : N/A
```

Ident	Symb	Int	TID	Vendor	Type	Hypr	Capacity (MB)		
							Total	Free	Actual
DX-9F	09F	-	-	EMC Corp	N/A	0	61425	61425	61425
Total							61425	61425	61425

Similarly, additional external disks can be added as eDisks to the external disk group. Also, additional external disk groups can be created to enhance manageability.

Symmetrix devices can now be configured from the external disk group created. However, protection should be left unprotected as RAID protection would be provided by the external array.

```
# symconfigure -sid 98 -cmd "create dev count=8, size=18414, emulation=FBA,
config=unprotected, dynamic_capability=dyn_rdf, disk_group=512;" commit -nop

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

Establishing a configuration change session.....Established.
Processing symmetrix 000195700398
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Queued.
COMMIT requesting required resources.....Obtained.
Step 012 of 060 steps.....Executing.
Step 046 of 145 steps.....Executing.
Step 054 of 145 steps.....Executing.
Step 061 of 145 steps.....Executing.
Step 082 of 145 steps.....Executing.
Step 100 of 152 steps.....Executing.
Step 127 of 152 steps.....Executing.
Step 129 of 152 steps.....Executing.
Local: COMMIT.....Done.

New symdevs: 0430:0437
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

In the above configuration change, eight devices with cylinder size 18414, emulation FBA, dynamic RDF capability (can be used as R1 or R2 devices) from disk group 512 are created with Symmetrix device names 0430-0437.

As the disk group 512 used here consists of externally provisioned Edisks, they will be non-encapsulated.

The disk group number 513 is created along with an external virtual pool.

```
# symconfigure -sid 98 -cmd "create pool VP_External, type=thin;" commit -nop

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

Establishing a configuration change session.....Established.
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Committing configuration changes.....Reordering.
Creating pools .....Done.
Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.

The configuration change session has successfully completed.
```

The external virtual pool VP_External can now be filled with data devices (TDATs) in enabled state.

```
# symconfigure -sid 98 -cmd "create dev count=8, size=18414, emulation=fba, ,
config=unprotected, attribute=datadev in pool VP_External member_state=enable ,
disk_group=513;" commit -nop
```

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

```
Establishing a configuration change session.....Established.
Processing symmetrix 000195700398
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Queued.
COMMIT requesting required resources.....Obtained.
Step 012 of 060 steps.....Executing.
Step 046 of 145 steps.....Executing.
Step 054 of 145 steps.....Executing.
Step 062 of 145 steps.....Executing.
Step 082 of 145 steps.....Executing.
Step 100 of 152 steps.....Executing.
Step 127 of 152 steps.....Executing.
Step 129 of 152 steps.....Executing.
Step 143 of 152 steps.....Executing.
Local: COMMIT.....Done.
```

```
    New symdevs: 0438:043F [DATA devices]
Terminating the configuration change session.....Done.
Establishing a pool configuration change session.....Established.
Defining pooldev commands.....Defined.
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Locking devices.....Locked.
Committing configuration changes.....Reordering.
Adding pool devs .....Done.
Enabling pool devs .....Done.
Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed

This results in getting data devices 0438-043F in external disk group 513 in enabled state.

These can now be used to provide capacity to the initially cache based thin devices (TDEVs).

```
# symconfigure -sid 98 -cmd "create dev count=16, size=18414, emulation=fba,
config=TDEV, binding to pool= VP_External;" commit -nop
```

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

```

Establishing a configuration change session.....Established.
Processing symmetrix 000195700398
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Started.
Committing configuration changes.....Queued.
COMMIT requesting required resources.....Obtained.
Step 011 of 060 steps.....Executing.
Step 046 of 145 steps.....Executing.
Step 054 of 145 steps.....Executing.
Step 070 of 145 steps.....Executing.
Step 092 of 145 steps.....Executing.
Step 124 of 149 steps.....Executing.
Step 126 of 149 steps.....Executing.
Step 132 of 149 steps.....Executing.
Local: COMMIT.....Done.
Binding devices.....Done.

  New symdevs: 0448:0457 [TDEVs]
Terminating the configuration change session.....Done.

```

The thin devices 0448-0457 are then created and bound to pool VP_External.

Suppose these virtually provisioned thin devices need to be presented to host Win_01 with initiator port wwns as 10:00:90:fa:08:15:4a and 10:00:90:fa:08:15:4c.

The Symmetrix can be checked for existing Port groups configured by the command below.

```

symaccess -sid 0398 list -type port

Symmetrix ID          : 000195700398
Port Group Name
-----
host_1e0_2e0_1f0_2f0_PG
host_1e0_2e0_PG
host_1e1_2e1_1f1_2f1_PG
host_1g0_2g0_PG
host_1g1_2g1_PG
host_3e0_4e0_3f0_4f0_PG

```

The pre-configured thin pools can be listed along with the external virtual pool configured.

```
symcfg -sid 0398 list -type pool -thin -detail -gb
```

```
Symmetrix ID: 000195700398
```

SYMMETRIX THIN POOLS											
Pool Name	Flags	Dev PTECSL	Config	Total GBs	Usable GBs	Free GBs	Used GBs	Full (%)	Subs (%)	Comp (%)	Shared GBs
Std_RAID6_TP	TSFDEI	RAID-6(6+2)		146930.7	146930.7	146930.7	0.1	0	0	0	0.0
Mid_RAID5_TP	TFEDEI	RAID-5(3+1)		22537.9	22537.9	22537.9	0.0	0	0	0	0.0
EFD_RAID5_TP	TEFDEI	RAID-5(3+1)		2200.2	2200.2	2200.2	0.0	0	0	0	0.0
VP_External	T---EX			2764.0	2764.0	2764.0	0.0	0	0	0	0.0
Total				171668.8	171668.8	171668.8	0.1	0	0	0	0.0
GBs											

Legend:

```
(P)ool Type:
  S = Snap, R = Rdfa DSE T = Thin
(T)echnology:
  S = SATA, F = Fibre Channel, E = Enterprise Flash Drive, M = Mixed, - = N/A
Dev (E)mulation:
  F = FBA, A = AS400, 8 = CKD3380, 9 = CKD3390, - = N/A
(C)ompression:
  E = Enabled, D = Disabled, N = Enabling, S = Disabling, - = N/A
(S)tate:
  E = Enabled, D = Disabled, B = Balancing
Disk (L)ocation:
  I = Internal, X = External, M = Mixed, - = N/A
```

The VP_External thin pool contains TDATs from eDisks that are configured from externally provisioned space of an external array. The FAST control parameters can be verified on the Symmetrix array as shown below.

```
symfast -sid 0398 list -control_parms
```

```
Symmetrix ID: 000195700398
```

Optimizer and FAST Control Parameters:

```
Data Movement Mode           : N/A
Max Simultaneous Device Moves : N/A
Max Device Moves Per Day      : N/A
```

Optimizer, FAST and FAST VP Control Parameters:

```
Min Initial workload Period(hrs) : 24
workload Analysis Period(hrs)    : 168
```

FAST Control Parameters:

```
Swap Not Visible Devices      : N/A
Allow Only Swap                : N/A
```

FAST VP Control Parameters:

```
FAST VP Data Movement Mode    : AUTO
FAST VP Data Relocation Rate   : 8
Thin Pool Reserved Capacity(%) : 10
VP Allocation By FAST policy   : Enabled
FAST VP Time to Compress       : 40 days
FAST VP Compression Rate       : 8
```

In this example, the initial workload period has been set to 24 hours and workload analysis period is 168 hours. Data movement mode has been set to automatic with VP allocation by FAST policy as enabled. Current state of the FAST VP can also be verified.

```
symfast -sid 0398 list -state
Symmetrix ID: 000195700398
FAST State           : N/A
Reason(s)           : N/A
FAST Current Activities : N/A
FAST Perf Time Window : N/A
FAST Move Time Window : N/A

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

FAST VP is enabled with open performance time window and move time window. Symmetrix tiers can be listed and the detailed capacity of the tiers can be found as per below.

```
symtier -sid 0398 list -v
Symmetrix ID       : 000195700398
Tier Name          : High_Tier_VP
Tier Type          : VP
Disk Location      : Internal
Technology         : EFD
Target Protection  : RAID-5(3+1)
Emulation          : FBA
Include Type       : Static

Thin Pools(1)
{
-----
          Logical Capacities (GB)
-----
Pool Name   Dev  Enabled   Free   Used   Full
-----
EFD_RAID5_TP FBA    2200    2200     0     0
-----
Total      2200    2200     0
}
}
```



```
Tier Name      : Mid_Tier_VP
Tier Type     : VP
Disk Location  : Internal
Technology    : FC
Target Protection : RAID-5(3+1)
Emulation     : FBA
Include Type  : Static
```

```
Thin Pools(1)
```

```
{
-----

```

Pool Name	Dev Emul	Logical Capacities (GB)			Full (%)
		Enabled	Free	Used	
Mid_RAID5_TP	FBA	22538	22538	0	0
Total		22538	22538	0	

```
-----
}
```

```
Tier Name      : Std_Tier_VP
Tier Type     : VP
Disk Location  : Internal
Technology    : SATA
Target Protection : RAID-6(6+2)
Emulation     : FBA
Include Type  : Static
```

```
Thin Pools(1)
```

```
{
-----

```

Pool Name	Dev Emul	Logical Capacities (GB)			Full (%)
		Enabled	Free	Used	
Std_RAID6_TP	FBA	146931	146931	0	0
Total		146931	146931	0	

```
-----
}
```

```
Legend:
```

```
Tier Type      : DP = Disk Group Provisioning, VP = Virtual Pools
```

The FAST Policies can be set and verified as per below. In this example, three FAST policies are set – All-Tier-Pol_VP, High-Tier-Pol_VP, and Low-Tier-Pol_VP which include all internal Symmetrix tiers with different settings. All-Tier-Pol_VP enables the associated storage groups to be moved across all the tiers which have EFD, FC, and SATA drive technologies. This can be considered standard policy if workload is predicted to be normal by the user. High-Tier-Pol_VP policy is mainly aimed to accommodate application data predicted to be business critical and don't deem movement to low tier SATA drive under any circumstances. Low-Tier-Pol_VP policy is suitable for low to medium predicted storage groups.

symfast -sid 0398 list -fp -v

Symmetrix ID : 000195700398

Policy Name : All-Tier-Pol_VP
Emulation : FBA

Tiers(3)

{

Tier Name	Type	Max Percent	SG	L O C	Tech	Target Protection	Flgs C
High_Tier_VP	VP	100	I	EFD	RAID-5(3+1)	.	
Mid_Tier_VP	VP	100	I	FC	RAID-5(3+1)	.	
Std_Tier_VP	VP	100	I	SATA	RAID-6(6+2)	.	

}

No Storage Groups associated with Policy All-Tier-Pol_VP

Policy Name : High-Tier-Pol_VP
Emulation : FBA

Tiers(2)

{

Tier Name	Type	Max Percent	SG	L O C	Tech	Target Protection	Flgs C
High_Tier_VP	VP	100	I	EFD	RAID-5(3+1)	.	
Mid_Tier_VP	VP	100	I	FC	RAID-5(3+1)	.	

}

No Storage Groups associated with Policy High-Tier-Pol_VP

Policy Name : Low-Tier-Pol_VP
Emulation : FBA

Tiers(2)

{

Tier Name	Type	Max Percent	SG	L O C	Tech	Target Protection	Flgs C
Mid_Tier_VP	VP	100	I	FC	RAID-5(3+1)	.	
Std_Tier_VP	VP	100	I	SATA	RAID-6(6+2)	.	

}

No Storage Groups associated with Policy Low-Tier-Pol_VP

Legend:

Tier Type : DP = Disk Group Provisioning, VP = Virtual Pools
Disk (Loc)ation : I = Internal, X = External

Flgs:

(C)ompression : X = Compression Capable, . = Not Compression Capable

The host can now be provisioned with LUNs which are from eDisks, namely 0448-0457.

This can then be mapped and masked using the auto provisioning groups within Symmetrix.

```
Symmetrix ID           : 000195700398
Masking View Name     : win_01_MV
Last updated at      : 07:14:52 AM on Sat Nov 10,2013
Initiator Group Name  : win_01_IG

  Host Initiators
  {
    WWN   : 100090fa08154a  [alias: 100090fa08154a/100090fa08154a]
    WWN   : 100090fa08154c  [alias: 100090fa08154c/100090fa08154c]
  }

Port Group Name       : host_1e0_2e0_PG

  Director Identification
  {
    FA-1E:0
    FA-2E:0
  }

Storage Group Name    : win_01_SG

  Number of Storage Groups : 0
  Storage Group Names     : None

Sym
Dev  Dir:P  Physical Device Name  Host Lun  Attr  Cap(MB)
-----
0448  01E:0  Not Visible           0        512
      02E:0  Not Visible           0
0449  01E:0  Not Visible           1        512
      02E:0  Not Visible           1
0450  01E:0  Not Visible           2        512
      02E:0  Not Visible           2
0451  01E:0  Not Visible           3       51200
      02E:0  Not Visible           3
0452  01E:0  Not Visible           4       51200
      02E:0  Not Visible           4
0453  01E:0  Not Visible           5       51200
      02E:0  Not Visible           5
-----
Total Capacity                155136
```

The Storage group can then be associated to a FAST policy, for example, All-Tier-Pol_VP when predicted workload is considered to be normal by the below command.

```
symfast -sid 0398 -fp_name All-Tier-Pol_VP associate -sg Win_01_SG -priority 2
```

The associated storage group can be further verified as per below.

```
symmsg -sid 0398 show win_01_SG
Name: win_01_SG
Symmetrix ID      : 000195700398
Last updated at   : Sat Nov 10 14:19:36 2013
Masking Views     : Yes
FAST Policy       : Yes

Devices (6):
{
-----
Sym  Device Config  Sts  Cap
Dev  Pdev Name
-----
0448 N/A          TDEV  RW   512
0449 N/A          TDEV  RW   512
0450 N/A          TDEV  RW   512
0451 N/A          TDEV  RW  51200
0452 N/A          TDEV  RW  51200
0453 N/A          TDEV  RW  51200
}

```

This completes the illustration of FAST along with FTS-provisioned external storage from existing storage arrays. The scenario can be rolled across multi-vendor storage arrays currently qualified by the EMC Simple Support Matrix (Appendix B) to scale the Symmetrix features for other arrays.

This would also account for cost-neutral firms to buy at least a minimum configuration Symmetrix VMAX array to add the performance and robustness to their storage environment in a minimized fashion with future prospects for complete technology refresh if required.

Appendix A – VMAX Engine FA Layout

VMAX Engine FA Ports

Directors	Even									Odd							
Engine	1	2	3	4	5	6	7	8	Cache	9	10	11	12	13	14	15	16
Engine 8	16E-0	16E-1	16F-0	16F-1	16G-0	16G-1	16H-0	16H-1		15E-0	15E-1	15F-0	15F-1	15G-0	15G-1	15H-0	15H-1
Engine 7	14E-0	14E-1	14F-0	14F-1	14G-0	14G-1	14H-0	14H-1		13E-0	13E-1	13F-0	13F-1	13G-0	13G-1	13H-0	13H-1
Engine 6	12E-0	12E-1	12F-0	12F-1	12G-0	12G-1	12H-0	12H-1		11E-0	11E-1	11F-0	11F-1	11G-0	11G-1	11H-0	11H-1
Engine 5	10E-0	10E-1	10F-0	10F-1	10G-0	10G-1	10H-0	10H-1		9E-0	9E-1	9F-0	9F-1	9G-0	9G-1	9H-0	9H-1
Engine 4	8E-0	8E-1	8F-0	8F-1	8G-0	8G-1	8H-0	8H-1		7E-0	7E-1	7F-0	7F-1	7G-0	7G-1	7H-0	7H-1
Engine 3	6E-0	6E-1	6F-0	6F-1	6G-0	6G-1	6H-0	6H-1		5E-0	5E-1	5F-0	5F-1	5G-0	5G-1	5H-0	5H-1
Engine 2	4E-0	4E-1	4F-0	4F-1	4G-0	4G-1	4H-0	4H-1		3E-0	3E-1	3F-0	3F-1	3G-0	3G-1	3H-0	3H-1
Engine 1	2E-0	2E-1	2F-0	2F-1	2G-0	2G-1	2H-0	2H-1	1E-0	1E-1	1F-0	1F-1	1G-0	1G-1	1H-0	1H-1	

Appendix B – Simple support matrix for EMC Federated Tiered Storage

FTS Array — EMC Enginuity Release	Qualified External Arrays	
Minimum Enginuity version 5876.82.57 ^{a, b, c, d}	EMC	Symmetrix VMAX 40K VMAX 20K/VMAX, VMAX 10K/ VMAXe, VMAX SE
		Symmetrix DMX™
		VNX® Series
		CLARiiON® CX4 and CX3 Series
	HDS	VSP (Virtual Storage Platform)
		USP-V (Universal Storage Platform)
		USP-VM
		TagmaStore USP 1100
		TagmaStore USP 600
		TagmaStore USP 100
	HP	XP P9500
		XP 24000
		XP 20000
		XP 12000
IBM	DS8700 ^f	
	XIV	
SUN	StorEdge 9990V	
Minimum Enginuity version 5876.159.102 ^{a, b, c, d}	All of the above-listed models are supported. In addition, the following models are also supported.	
	EMC	Symmetrix VMAX 10K xxx987xxxx and VMAX 10K xxx959xxxx
	HDS	AMS 1000 (Adaptable Modular Storage) ^e
		AMS 500 ^e
	HP	EVA 8100 (Enterprise Virtualization Array)
		EVA 8000
		EVA 6100
		EVA 6000
		EVA 4100
		EVA 4000
	IBM	DS8800 ^f
		N7900
	NetApp	FAS3200
FAS6080		
FAS6200		
Minimum Enginuity version 5876.229.145	All of the above-listed models are supported. In addition, the following models are also supported.	
	IBM	DS8300 ^f
		DS6xxx Series ^f
		DS5xxx/4xxx Series ^g
	SUN	STK 6x40 ^g
		FLX 380 family ^g

This table provides information for EMC Symmetrix® Federated Tiered Storage for EMC Symmetrix VMAX 40K, VMAX 20K/VMAX, VMAX 10K xxx987xxxx, VMAX 10K xxx959xxxx, VMAXe, and VMAX SE and minimum EMC Enginuity versions.

Appendix C – References and Further Reading

Symmetrix Foundation Student Guide

EMC Symmetrix VMAX Product Guide

Implementing Fully Automated Storage Tiering for EMC Symmetrix VMAX Series Arrays

Implementing Fully Automated Storage Tiering for Virtual Pools (FAST VP) for EMC Symmetrix VMAX Family Arrays

Implementing Fully Automated Storage Tiering (FAST) for EMC Symmetrix VMAX Series Arrays

EMC Simple Support Matrix EMC Federated Tiered Storage

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.