

FAST Theory and Best Practices for Planning and Performance

Technical Notes

P/N 300-010-342

REV A02

May 2010

This technical note contains information on these topics:

◆ Executive summary	2
◆ Introduction and overview	2
◆ Pertinent terminology	3
◆ Drive technology and tiering.....	4
◆ Understanding workload skew.....	6
◆ Policies and FAST managed objects.....	7
◆ Symmetrix Performance Analyzer and FAST.....	10
◆ Operational controls	24
◆ Workload analysis and FAST algorithms.....	26
◆ Device movement between tiers.....	31
◆ Limitations and restrictions.....	40
◆ Consolidated list of FAST best practices.....	41
◆ Summary and conclusion	44

Executive summary

The EMC® Symmetrix® VMAX™ Series with Enginuity™ is the newest addition to the Symmetrix family, and the first high-end system purpose-built for the virtual data center. The Enginuity operating environment provides the intelligence that controls all components in a Symmetrix VMAX array.

The Enginuity version 5874 Q4'09 Service Release (SR) is the latest Enginuity release that supports the Symmetrix VMAX storage arrays. This Enginuity release features new software functionality such as Fully Automated Storage Tiering (FAST) and enhancements to current products including EMC Virtual Provisioning™ and EMC SRDF®. EMC Symmetrix VMAX FAST technology is designed to automate the allocation and relocation of application data across different performance tiers based on the changes in application performance requirements.

This document describes FAST theory and best practices as necessary to maximize the benefits existing and planned FAST configurations. Features of pertinent Symmetrix Performance Analyzer (SPA) version 2.0 are also described in this document.

Introduction and overview

The EMC Symmetrix VMAX Series with Enginuity is the newest addition to the Symmetrix family, and the first high-end system purpose-built for the virtual data center. Based on the Virtual Matrix Architecture™, Symmetrix VMAX scales performance and capacity to unprecedented levels, delivers nondisruptive operations, and greatly simplifies and automates the management and protection of information.

The Enginuity operating environment provides the intelligence that controls all components in a Symmetrix VMAX array. It coordinates real-time events related to the processing of production data. It applies self-optimizing intelligence to deliver the ultimate performance, availability, and data integrity required in a platform for advanced storage functionality.

The Enginuity version 5874 Q4 '09 SR is the latest Enginuity release supporting the Symmetrix VMAX storage arrays. Enginuity version 5874 features new software functionality such as Fully Automated Storage Tiering (FAST) and enhancements to current products including Virtual Provisioning and SRDF.

EMC Symmetrix VMAX FAST technology, for standard provisioned, non-thin, environments is designed to automate the allocation and relocation of application data across different performance tiers based on the changes in the application performance requirements. FAST helps customers to maximize the benefits of tiered storage within the array by prioritizing and optimizing application performance and reduce labor cost.

Symmetrix VMAX FAST proactively monitors volume workloads and automatically moves the busier volumes to higher performing flash drives, and the slower volumes to higher capacity, lower performance drives (SATA). The promotion/demotion activity is driven by policies assigned to the FAST storage groups and is executed nondisruptively without affecting business continuity and availability.

Though not specifically included in this document, the management and operation of FAST is accomplished via Symmetrix Management Console (SMC) or Solutions Enabler (SYMCLI). The creation of tiers, policies, and storage groups may be accomplished through easy-to-use wizards within SMC. Applicable chargeback information is provided via EMC StorageScope™ within EMC ControlCenter®.

The features of Symmetrix Performance Analyzer (SPA) version 2.0 are described in this document to provide detailed insight into performance trending, forecasting, alerts, and resource utilization. In support of FAST, SPA 2.0 specifically provides detailed information regarding storage groups and Symmetrix tiers within its Snapshot, Trend, and Diagnostic views.

Audience

This technical note is intended for anyone who needs to understand FAST theory and best practices as necessary to achieve the best performance for FAST configurations. Pertinent Symmetrix Performance Analyzer (SPA) version 2.0 features are also covered in this document. This document is specifically targeted at EMC customers, sales, and field technical staff who are either running FAST or are considering FAST for future implementation.

Pertinent terminology

Symmwin Disk Group — Physical disk groups defined by technology, speed, and size within Symmwin.

Storage Type/Symmetrix Tier — Storage Type is a collection of physical disks of the same technology used to implement logical devices with the same protection level. Storage Type is referred to as a Symmetrix Tier on the Symmetrix platform. A tier can be constructed from Symmwin Disk Groups. It may include all disks with the same technology. If the tier type is defined as dynamic, all disks of the same technology are included. Also, when new disks of the same technology are added to the Symmetrix VMAX, they are dynamically added to the appropriate tier.

Best Practice! For FAST environments targeting dynamic Symmetrix tiers, it is possible to start with dynamic tiers or migrate existing static tiers to dynamic. However, be aware, that it is not possible to migrate dynamic tiers back to static.

Storage Group — A storage group is a collection of Symmetrix devices or LUNs. A storage group typically includes all the devices used by one application or by a set of functions of this application. As storage groups are also used for Auto-provisioning, FAST storage groups may coincide with Auto-provisioning storage group or can be a subset or a superset of these groups. Storage groups may also require the assignment of a FAST handling priority.

FAST Policy — A FAST policy is a set of Storage Group/Symmetrix tier rules that apply to associated storage groups.

Association of FAST Policies and Storage Groups — An association defines the relationship between storage groups and FAST policies. Each storage group, which is managed by FAST, must be associated with only one FAST policy. Multiple storage groups may be associated with the same FAST policy.

Drive technology and tiering

Symmetrix supports Enterprise Flash Drives (EFDs) in addition to Fibre Channel (FC) and SATA disks. EFDs deliver much better performance at a higher cost, and therefore, must be used as efficiently as possible. The goal is to keep the EFD devices busy by storing the most frequently accessed data on them. Data access pattern usually change over time and disks that are busy at one time may not be busy at a different time intervals. The goal of FAST is to detect these patterns of data usage and move the appropriate data to the EFD, while moving out data that is no longer as heavily used.

The performance characteristics of a device are determined by the hardware technology of the physical disk on which the device is stored, protection level, and data access pattern. EFDs provide superior performance because they store their data in electronic flash memory and have no physical moving parts. EFDs provide subhalf-millisecond access times for both reads and writes. They provide the best performance when significantly more data is read than written.

FC provides reasonable performance (6–8 ms or 160 IOPS) at a lower cost. SATA is the lowest cost storage with the lowest performance of approximately 12–20 ms; or 80 IOPS. In addition, SATA drives typically store 1 or 1.5 TB, which is 3–20 times as much data as common FC or EFD disks. This has a two fold impact. It lowers the cost per GB and also the physical footprint and power consumption on a per GB basis.

Best Practice! SATA drive technology is suited for data that is not accessed often, or data that is mostly accessed sequentially. It is best for I/O bandwidth to be *shared* by the larger pool of SATA disks in comparison to FC disks. In this way, on an average, each SATA disk will be accessed less often.

A Symmetrix tier is a collection of disks that have similar performance characteristics and are equivalent from the perspective of where to place the data. For example, a Symmetrix that has FAST might have three tiers, one for EFD, one for FC, and one for SATA. Usage of the tiers is governed by FAST policies; for example, FAST may be used to manage only the FC and SATA disk technologies and not include EFDs.

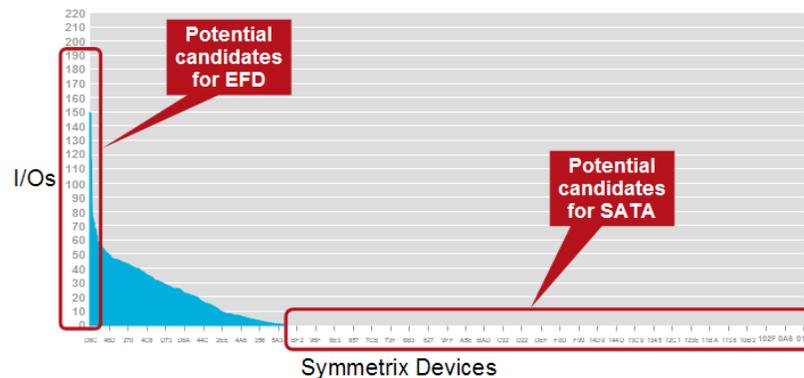
As RAID protection level affects the disk performance, disks of the same technology but with different RAID protection levels must belong to different tiers. For example, a Symmetrix system may have three tiers consisting of:

- RAID 1 15K FC disks
- RAID-5 (3+1) 15K FC disks
- RAID-6 (14+2) SATA disks.

Recent enhancements to the Symmetrix VMAX platform include enhanced Virtual LUN technology which allows FAST to move devices between different protection levels. This feature also enables the movement of devices to the preferred tiers. Virtual LUN move operations are based on Raid Virtual Architecture (RVA) technology.

Understanding workload skew

Skew is a key term for understanding FAST applicability and describes the asymmetry in device usage over time. A small fraction of devices in most cases serve the majority of I/Os; that is, on a Symmetrix there are devices that are used very heavily and others that are used very lightly. Skew among devices may change over time. Devices that are busy at one time period may not be busy at other times.



Best Practice! FAST is most effective when the I/O access patterns are highly skewed.

In cases where the skew is high and the workload is known to change significantly over time, FAST ensures that highly accessed devices will reside on the best performing tier.

In contrast, environments where skew is high and device usage has a relatively low degree of variability over time; users may choose to place the devices statically on the appropriate technology and protection level. In such environments, FAST may not be necessary to automatically optimize the usage of a higher performing tier as no shift in I/O access patterns is expected.

Best Practice! When migrating to a new configuration or platform targeted for FAST, it is best to plan and configure optimum tiers in the target or upgrade system and to pre-place volumes in these tiers based on performance characteristics.

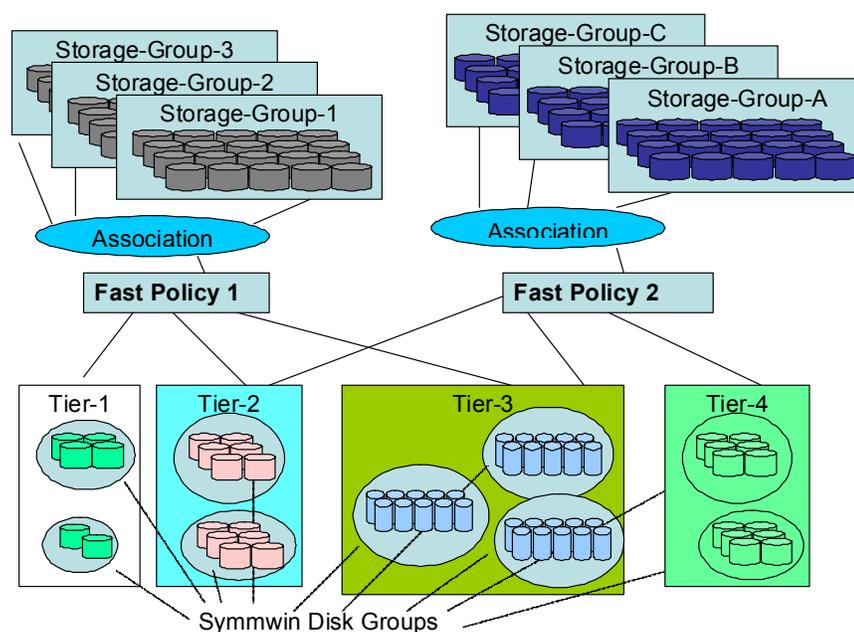
Best Practice! Use Tier Advisor to assist with recommending Storage Types or tiers within the Symmetrix.

Best Practice! Tier Advisor, SymmMerge, and similar tools support configuration planning and are recommended when the user has I/O workload information available. More specifically, SymmMerge is best used to assist with placement or migration planning, while Tier Advisor assists with the sizing of Symmetrix tiers.

Policies and FAST managed objects

The following objects are managed by FAST:

- Symmetrix tiers
- Storage groups
- FAST policies
- Associations of storage groups to FAST policies



As noted previously, Symmetrix tiers are named by collection of physical disks specified by a FAST user that are equivalent in terms of specific performance attributes. In FAST, this includes all the disks of the same technology (EFD, FC, or SATA) and the same RAID protection level. FAST leverages Symmetrix Symmwin disk groups for this purpose. A Symmwin disk group is a collection of disks of similar technology, size, and speed. As a FAST tier is not directly concerned with disk size and speed, a tier can therefore be composed of one or more Symmwin disk

groups.

Storage groups are collections of devices (LUNs) that may represent some logical grouping of the devices usage such as applications. A device, also known as volume or a LUN to the application host may belong to not more than one storage group that is managed by FAST.

A FAST policy is a set of tier usage rules that apply to the associated storage groups. Each FAST policy lists up to three tiers and their corresponding usage rule. Multiple FAST policies may reference the same tiers. Multiple storage groups may share one policy. In this case, the tiers' disks are split between the storage groups.

Usage of the tiers is specified as a percentage of the total storage required by the devices in the storage group. This percentage is actually the upper bound that the storage group can use of the tier. For example, if the FAST policy specifies 20 percent of tier-A, then any of the associated storage groups cannot have more than 20 percent of its storage capacity needs on tier-A.

FAST does not require a complete overlap between the disks on which the devices of a storage group are stored and the tiers in the associated FAST policy. For example, devices of a storage group may be on a Symmwin disk group that is not managed by FAST. In they do not reside in a storage group managed by FAST, these devices will remain where they are and never be moved by FAST. If any of these devices are found to be very busy, as observed by Symmetrix Performance Analyzer (SPA), users may choose to move them manually with Virtual LUN or simply add them to the FAST configuration.

Associations tie the storage groups to FAST policies. Each storage group is associated with only one FAST policy. Multiple storage groups can be associated with the same FAST policy. While two tiers may be configured with the same Symmwin disk group, partial overlapping of Symmwin disk groups in tiers is not allowed.

Multiple storage groups may share the same FAST policy and the storage requests from the underlying tiers may conflict. For this reason, storage groups are assigned priorities—low, medium, and high. Actual tier usage is determined by the priority the storage groups. Those with higher priority are likely to get more of the limited tier capacity than those with lower priority.

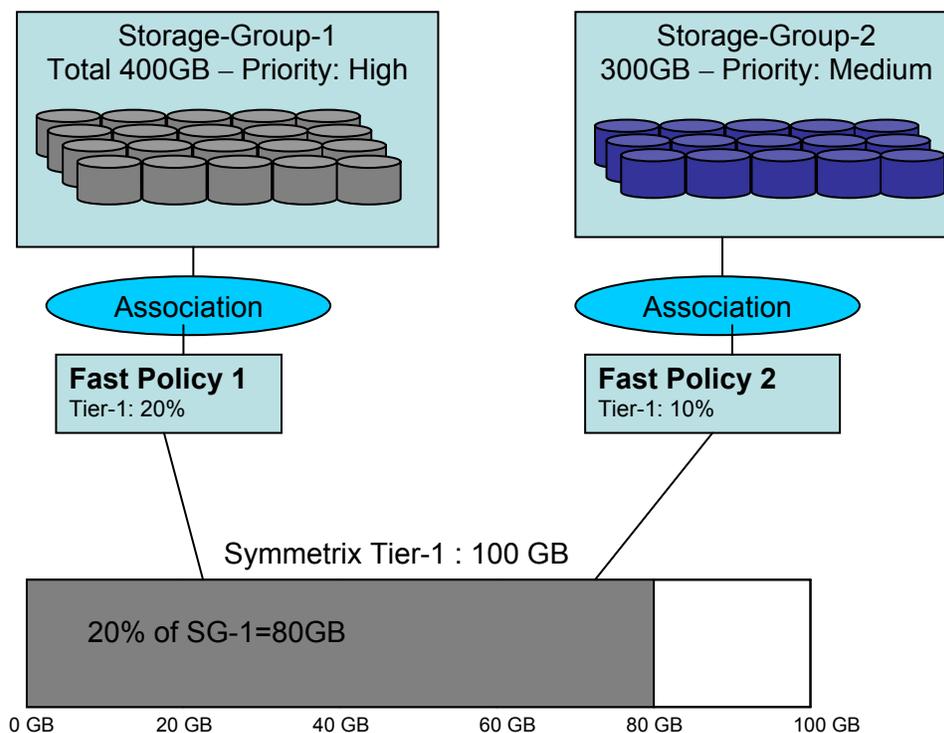
Be aware that a FAST storage configuration can result in oversubscription of a tier. For example:

- Storage group 1 is associated with a FAST policy that states that not more than 20 percent of the storage group capacity can be on EFD.
- Storage group 2 is associated with a FAST policy that states that not more than 10 percent of the storage group capacity can be on EFD.

Both storage groups cannot utilize the maximum EFD that is configured for them. To resolve as to which storage group should be given the preferred storage, the priority of the storage group will be consulted.

In the following diagram:

- Storage-Group-1 is 400 GB and 20 percent of it (80 GB) could be on Tier1.
- Storage-Group-2 is 300 GB and 10 percent of it (30 GB) could be on Tier1.



As Tier-1 is limited to 100 GB in size, both of these storage groups cannot be fully allocated to this tier. As Storage-Group-1 has a higher priority, a *Capacity-based algorithm* will ensure it will receive 80 GB from Tier-1 while Storage-Group-2 will receive the remaining 20 GB.

A FAST policy includes at least one tier and no more than three tiers, each with its assigned maximum capacity percentage. FAST will store devices in a manner which ensures tier usage will not exceed the configured maximum capacity percentage.

For example:

If Storage-Group-A has 100 X 4 GB devices, for a total storage capacity of 400 GB and it is associated with a FAST policy that states that not more than 10 percent will be on Tier-1 and not more than 20 percent will be on Tier-2 then FAST will store no more than 10 devices on Tier-1 and no more than 20 devices on Tier-2. The rest of the devices will be on Tier-3.

It may also be the case that storage groups with the same priority cannot be given the requested storage. In this case, the storage requests are adjusted by employing a capacity adjustment algorithm. This algorithm considers the currently allocated storage and the I/O activity (IOPS).

Symmetrix Performance Analyzer and FAST

SPA version 2.0 product overview

The SPA collects and stores historical performance data for users to analyze and report on workload and resource usage trends for:

- High-frequency data collection
- Root cause analysis
- Alerts and events
- User-definable charts
- Service level management
- Capacity planning

The SPA is a standalone, server-based application that provides a single tool to monitor the workload activity for a number of Symmetrix arrays and FAST configurations within the data center environment.

Although standalone, the SPA application integrates with the existing SMC product and shares as many resources as possible while providing diagnostic, performance, and planning information with easy-to-use graphical data representations.

The SPA version 2.0 interface now includes four views:

Real Time view provides high-frequency data collection and displays the selected metrics as they become available. This view includes user-definable charts for viewing the metrics of Symmetrix arrays and directors (FE, BE, RDF).

Diagnostic view analyzes high- and low-level metrics. It provides an intuitive root cause analysis engine that provides access to all available metrics and includes composite comparison charts of related metrics. This view includes a "fast lane" to quickly navigate to the root cause of an alert.

Snapshot view displays an analysis of collected data and categorizes the data by physical (array, disk, director) and logical (device group, application) monitored objects.

Trend view displays raw data collected (IOPS, Read Hits, Write Hits). By selecting a long time range (6 months or longer) in the Trend view can best show the consumption of resources over time. This view includes linear projection capability for capacity planning. The following table details the features and metrics specific to each of these views.

Symmetrix Performance Analyzer and FAST

Real Time	Diagnostic	Snapshot	Trend
The real time view provides high frequency data collection and displays the selected metrics as they become available. This view also allows for data saves for a 1-hour time window.	The Diagnostic view analyzes high- and low-level metrics. It provides an intuitive root cause analysis engine that provides access to all available metrics and includes composite comparison charts of related metrics. It also has device-level views, and typically focuses on a short time, such as a couple of hours.	The Snapshot view displays a calculation of the collected data over time and shows the relationship between the physical (array, disk, director) and logical (device group, application) components of the array.	The Trend view can best show the consumption of resources over time. The Trend view also supports linear projection lines, which can be used with user-defined thresholds as the basis for capacity planning.
Data is collected (for a subset of data) at 2 - 5 second intervals. Five minutes will provide 1000 metrics.	Data is collected in 5 minute intervals - 1 hour will have 12 points. One week is saved by default. Policy for saving data can be set by user.	Data is kept for 1 year. Data is collected in 15 minute intervals providing 4 points per hour.	
Provides Key Performance Indicators (KPIs) for:			
All/single Symmetrix array Front-end directors Back-end directors RDF directors	All/single Symmetrix arrays All/single Device Group All/single Storage group All/single front-end director All/single back-end director All/single RDF director All/single Cache partition All/single Symmetrix tier All/single Snap pool All/single Thin pool All/single DSE pool All/single disk group	All/single Symmetrix arrays Masking view Storage groups Port group All/single front-end directors All/single back-end directors All/single RDF directors All/single disks Disk group All/single Snap pools All/single DSE pool All/single Cache partitions All/single RDF/A groups RDF/S group Symmetrix tiers All/single device groups All/single composite group Composite group local Symmetrix array Composite group RA group Composite group RDF group	

Navigating Real Time view

The Real Time view provides the same navigation style as Diagnostic view—tabbed tables. Each Symmetrix array has tabbed information for front-end directors, back-end directors, and RDF directors.

Real Time - 000190300822 Local/FE Directors (4)

Format: Latest Maximum Minimum Average

Time Range: Now 10:49 - 11:49 EDT Capture

Symmetrix ID	Host IOs/sec	Host MBs/sec	FE Reqs/sec	System WP Events/sec	Device WP Events/sec
000190300822	98	0	25	0	0
000192600238	1349	78	2510	0	4

Real Time - 000190300822 Local/FE Directors (4)

Format: Latest Maximum Minimum Average

Time Range: Now 10:50 - 11:50 EDT Capture

Symmetrix Arrays: 000192600238

FE Director | BE Director | RDF Director

ID	Host IOs/sec	% Busy	Host MBs/sec	Reqs/sec	System WP Events/sec	Device WP Events/sec
FA-7E	1353	0	80	2585	0	4
SE-7G	11	0	0	0	0	0
FA-8E	10	0	0	0	0	0
SE-8G	0	0	0	0	0	0
SE-7H	0	0	0	0	0	0
SE-8H	0	0	0	0	0	0

Navigating the Diagnostic view

The Diagnostic view uses table-based navigation. When Diagnostic view is opened, the licensed Symmetrix arrays appear in the table at the top of the window, as shown in the following figure.

Diagnostic - SPA

Time Range: Jul 21, 2009 10:00 - Jul 21, 2009 14:00 GMT

Symmetrix	FE Utilization (%)	BE Utilization (%)	RDF Utilization (%)	Cache Partition(s) Utilization (%)	Disk(s) Utilization (%)	Avg. LRU B F
000192600238	✓ 13	✓ 25	✓ 0	N/A	⚠ 39	
000190300822	✓ 25	● 77	✓ 2	N/A	⚠ 37	

Heat maps for visual cues

Symmetrix Performance Analyzer and FAST

The high-level utilization metrics for each array appear numerically across each table row and visually in the charts at the bottom of the window, called the *Dashboard*, as shown in the following figure:



When a Symmetrix array is double-clicked, a new tab opens for that array, as shown in the following figure.

The screenshot shows a table with several tabs above it. A green arrow points to a new tab labeled '000190300822'. The table below has columns for 'Device Group', 'Storage Group', 'FE Director', 'BE Director', 'RDF Director', 'Cache Partition', 'Snap Pool', 'TP Pool', 'DSE Pool', 'Disk Group', and 'Event'. The table contains two rows of data: 'Exchange' and 'Oracle'.

Device Group	Storage Group	FE Director	BE Director	RDF Director	Cache Partition	Snap Pool	TP Pool	DSE Pool	Disk Group	Event
Id	I/O Rate (per Sec)	Read Response Time (ms)	Write Response Time (ms)	Hit Ratio (%)	Write (%)	Read Miss (%)				
Exchange	2086.2	10.4	2.1	55.3	26.0	60.3				
Oracle	2179.9	19.8	17.4	36.7	25.0	75.9				

Notice the new tab at the top (green arrow) for the Symmetrix array, and the additional tabs that appear across the table. Each tab opens the numerical metrics for the individual objects belonging to that category. For example, *Device Group* lists each device group in the selected Symmetrix array, and the high-level performance indicators for each group. The active object in the tables (top) appears visually in the Dashboard (bottom). They are two representations of the same data.

To navigate the tables above: *double-click* to select, *single-click* to view.

You may only double-click a table row, not a tab. When you double-click an object in the table (row), a new tab appears with the related objects for that tab. In the next example, The *Exchange* device group has been selected (double-clicked).

Device	FE Director	BE Director	Cache Partition	Disk Group	Event		
	Id	I/O Rate (per Sec)	Read Response Time (ms)	Write Response Time (ms)	Hit Ratio (%)	Write (%)	Read Miss (%)
	0420	39.9	6.6	2.3	53.5	26.0	62.1
	041F	39.9	6.6	2.3	53.5	26.0	62.0
	041E	39.9	11.7	2.0	54.4	26.0	61.2
	041D	39.9	8.9	2.1	54.4	26.0	61.3
	041C	39.9	7.0	2.3	53.4	26.0	62.2
	041B	39.9	9.5	2.1	54.4	26.0	61.3
	041A	39.9	9.5	2.1	54.5	26.1	61.2
	0419	39.9	11.7	2.0	54.5	26.0	61.0

Notice the blue triangle to the left of the tab name (green arrows). These are called *bread crumbs* that show you where you have been.

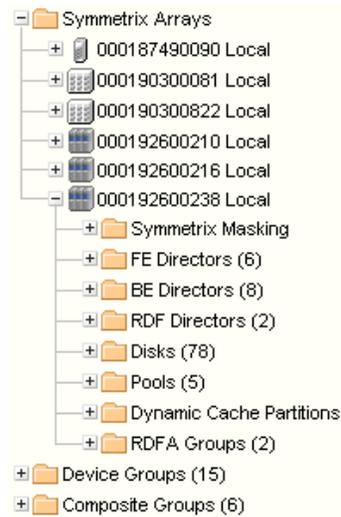
Snapshot and Trend views

The Snapshot view displays an analysis of collected data and categorizes the data by physical (array, disk, director) and logical (device group, application) monitored objects. Snapshot view is the default view for SPA. For example, to view data for all device groups, select the Device Groups folder. To view the data for one specific device group, select the device group name in the navigation tree. The data for the selected object appears in the right side of the SPA window.

Snapshot views that are not percentage-based are dynamic. The axis that shows the data counts, such as I/Os per second, write pending counts, and MBs per second, adjust according to the data.

The Trend view displays raw data collected (IOPS, Read Hits, Write Hits). By selecting a long time range (6 months or longer) in the Trend view can best show the consumption of resources over time. This view includes linear projection capability for capacity planning. Trend views can be opened by selecting the object in the navigation tree while in the Trend view. For example, to view data for all RDF directors, select the RDF Directors folder. To view the data for one specific RDF director, expand the RDF Directors folder and select the director from the navigation tree. The data for the selected object appears in the right side of the SPA window.

The SPA navigation tree displays for the Snapshot and Trend views. The tree contains folders and objects. Each folder contains a collection of objects, such as, *Symmetrix Arrays*, *RDF Directors*, *Device Groups*, and so on. The navigation tree is used to find and select objects for viewing performance details.



All the Snapshot and Trend views can be opened by selecting the object in the navigation tree while in the Snapshot or Trend view.

SPA 2.0 support for FAST

The SPA provides information about FAST storage groups and Symmetrix tiers in the following Snapshot, Trend, and Diagnostic views:

Snapshot view

Best Practice! Use SPA Snapshot views to determine applicable workloads and target Symmetrix tiers for FAST.

The storage group Snapshot view provides the performance data for the selected storage group. It is a high-level description of the workload characteristics over the selected time range.

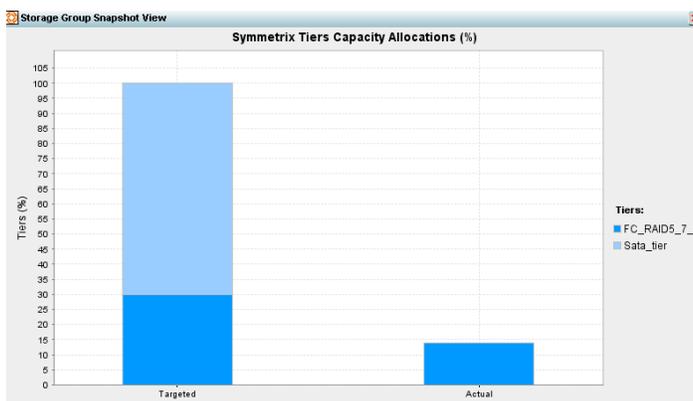
Best Practice! The longer the time range specified in the Storage Group Snapshot view, the more accurate is the view of the workload.

This view provides a profile of the selected storage group and complementary displays of the storage group profile information:

- **Storage Group Profile** — Provides a summary of various activity levels for the storage group such as number of devices, reads and writes per second, read write size, response time, hit ratios, and sequential and random percentage.

Storage Group Profile	
Reads Per Second	841.03
Writes Per Second	1495.12
Read Size	68KB
Write Size	68KB
Read Response Time	11.0ms
Write Response Time	8.9ms
Hit Rate	68.0%
Miss Rate	32.0%
Read Hit	11.0%
Write Hit	100.0%
Sequential Reads	6.7%
Sequential Writes	21.3%
Random Reads	29.3%
Random Writes	42.7%
Number of Devices	28

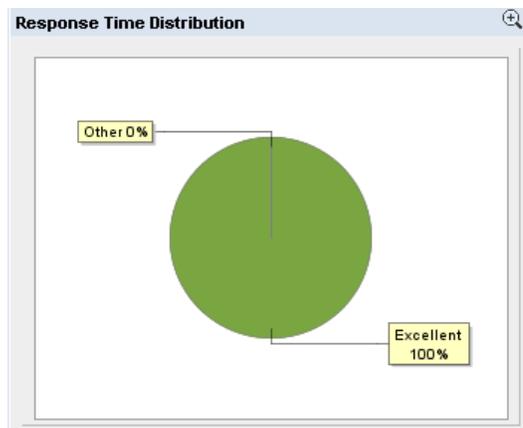
- **Symmetrix Tiers Capacity Allocation (%)** — Shows the targeted and actual allocation percent for each tier associated with the storage group through a FAST policy.



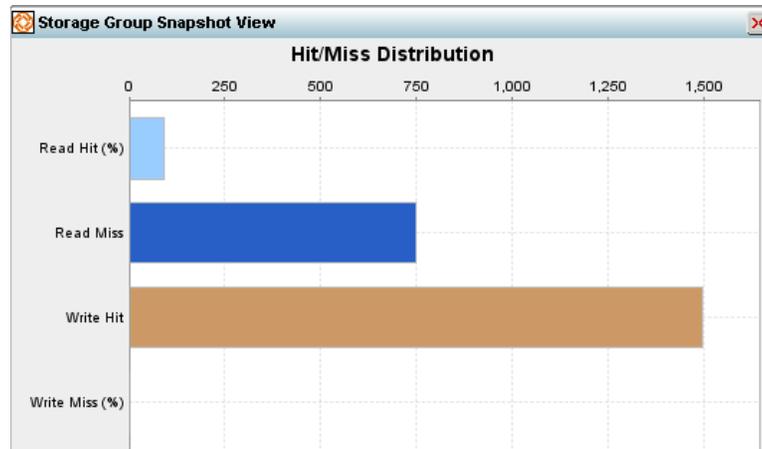
The SPA Snapshot view provides tier allocation information for storage groups that have FAST policies. The preceding example chart shows the tier allocation for a storage group. The storage group policy allows the FAST controller to place up to 30 percent of the storage group devices on FC_RAID_5_7_1 tier, and 70 percent of the storage group devices on SATA_tier, as shown in the Targeted bar. The Actual bar shows that 14 percent of FC_RAID5_7_1 tier contains the storage group data.

Best Practice! Use the SPA 2.0 Snapshot > Storage Group > Symmetrix Tiers Capacity Allocation (%) view to provide tier allocation information for storage groups that have FAST policies.

- **Response Time Distribution** — Displays a graphical representation of the average read and write times. These times are compared to the thresholds that are set for this storage group. If thresholds are not set, the system defaults are used. The following categories are represented: Excellent (green) represents response times which are below the low threshold set. Good (yellow) represents response times that were between the low threshold set and the high threshold set. Fair (red) represents response times that were greater than the highest response time level set.

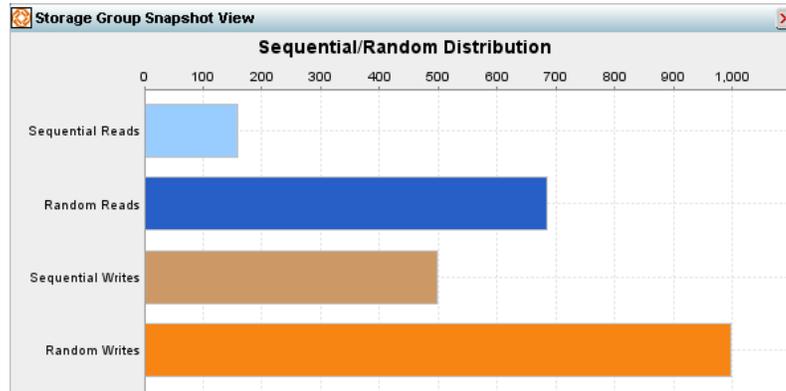


- **Hit/Miss Distribution** — Displays a graphical representation of the number of Read I/Os (hit and miss) and Write I/Os (hit and miss) for the storage group. This is complimentary information to the hit/miss distribution percents in the storage group profile.

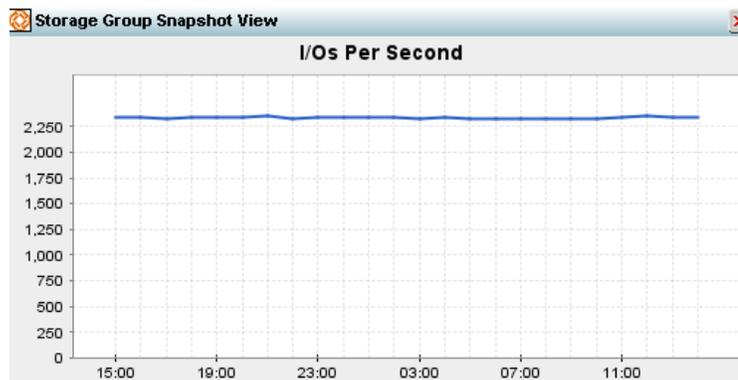


Best Practice! Use the SPA 2.0 Snapshot > Storage Group > Hit/Miss Distribution view to determine applicable workloads and target Symmetrix tiers for FAST. For example, SPA may indicate a significant amount of average read miss activity for a storage group. In this case, the group may be better placed either fully or partly onto an EFD Symmetrix tier as defined within a FAST policy.

- **Sequential/Random Distribution** — Displays a graphical representation of the number of Read I/Os (sequential and random) and Write I/Os (sequential and random) for the storage group. This is complimentary information to the sequential/random distribution percents in the storage group profile.



- I/Os per second** — Provides a high-level overview of the I/O rate pattern for the storage group during the selected time range. This can help identify business cycles of an application. If there is a major high peak or low valley, customize the time window to check that specific time.



Trend view

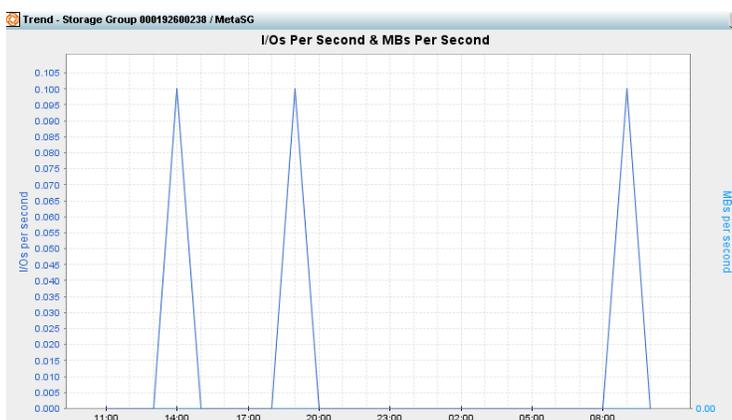
Best Practice! Use SPA Trend views to facilitate comparisons of storage group performance characteristics before and after they have been placed under FAST control.

The Storage Group Trend view provides the performance data for the selected storage group. This Trend view displays raw data collected (IOPS, Read Hits, Write Hits) for a specific storage group. This view includes linear projection capability for capacity planning.

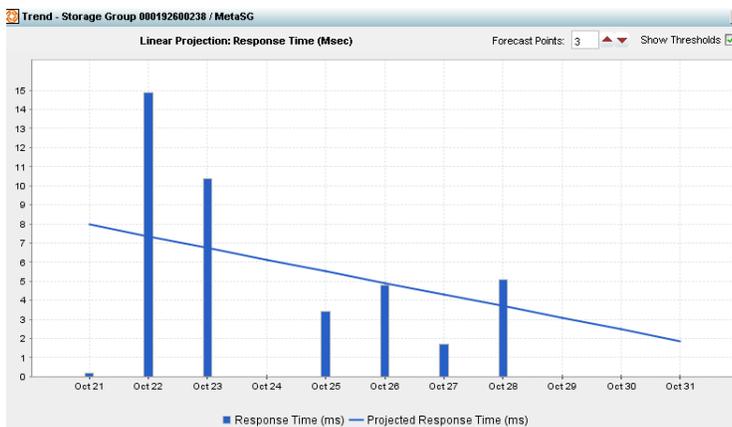
Best Practice! Selecting a long time range (6 months or longer) in the Storage Group Trend view can best compare storage group performance characteristics before and after they have been placed under FAST control.

Storage Group Trend view provides the following data for a storage group:

- **I/Os per second** — Shows the total I/Os per second for the storage group during the selected time range.



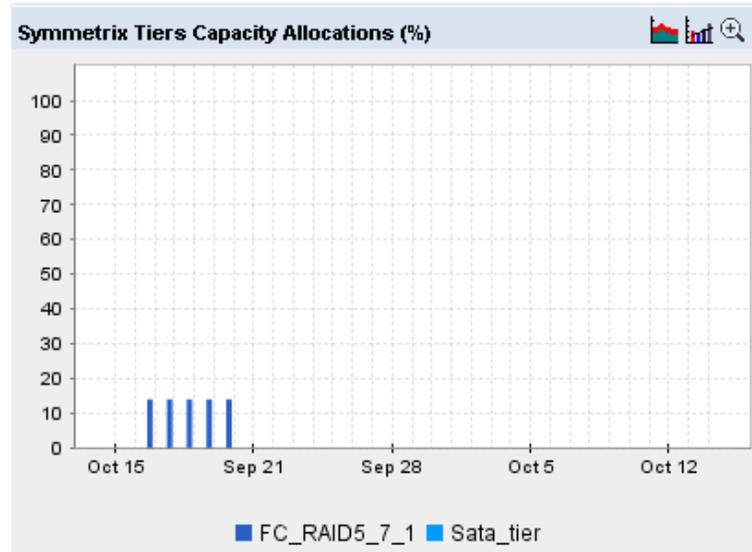
- **Response Time (Msec)** — Shows the response time in milliseconds for the storage group during the selected time range.



- **Read Hit and Write Miss (%)** — Shows the percent of Read Hits and Write Hits for the storage group during the selected time range.



- **Storage Tiers Capacity Allocations (%)** — Shows the capacity allocation percent for each tier associated with the storage group through a FAST policy.



Best Practice! Use the Trend > Storage Tiers Capacity Allocations (%) to show the capacity allocation percent for each tier associated with the storage group through a FAST policy.

Diagnostic view

In addition to alerts and events, Diagnostic view provides information about storage groups and Symmetrix tiers in the following table.

Symmetrix Arrays							
000192600238							
Alert	Device Group	Storage Group	FE Director	BE Director	RDF Director	Cache Partition	Symmetrix Tier
ID	% Busy	IOs/sec	Host MBs/sec	Reads/sec	Avg Read Time (ms)		W
FC_RAID5_7_1	⚠ 50.7	6844.0	275.8	3290.6	7.2		
Fibre_tier	⚠ 50.7	6844.0	275.8	3290.6	7.2		
Sata_tier	✔ 1.5	0.0	0.0	0.0	N/A		

Best Practice! Use the Storage Group and Symmetrix Tier Diagnostic views for critical diagnostic information (alerts and events) regarding storage groups and Symmetrix tiers.

The Storage Group Diagnostic view provides the following data for a storage group:

- **ID** — Shows the name assigned to this storage group.
- **Host IOs/sec** — Provides the number of host read I/O and write I/O operations performed each second by the group.
- **Avg Read Time (ms)** — Shows the average time it takes the group to perform the reads in milliseconds.
- **Avg Write Time (ms)** — Shows the average time it takes the group to perform the writes in milliseconds.
- **% Hit** — Shows the percentage of I/O operations performed by the group, that were immediately satisfied by cache.
- **% Write** — Shows the percent of I/O operations that were writes.
- **% Read Miss** — Shows the percent of Read Miss operations performed each second by the group. A miss occurs when the requested read data is not found in cache or the write operation had to wait while data was destaged from cache to the disks.
- **Capacity (GB)** — Shows the capacity of the storage group in GBs.
- **Number of Members** — Shows the number of devices that comprise this storage group.

In addition, the following information is available for a Symmetrix tier:

- **ID** — Shows name assigned to this storage tier.
- **Disk Busy (%)** — Provides a high-level overview of the utilization (time busy) for the disks that comprise this storage tier during the selected time range.
- **Host IOs/sec** — Shows the total IOPS for the disks that comprise this storage tier during the selected time range.
- **MB Rate** — Shows the total throughput in MBs for the disks that comprise this storage tier during the selected time range.
- **Host Reads/sec** — Shows the average number of host reads performed per second.
- **Read Response Time (ms)** — Shows the average time it takes the Symmetrix array to perform the reads in milliseconds.
- **Host Writes/sec** — Shows the average number of host writes performed per second.
- **Write Response Time (ms)** — Shows the average time it takes the Symmetrix array to perform the writes in milliseconds.
- **Total Capacity (GB)** — Shows the total available capacity in GBs for disks that comprise this storage tier.
- **Used Capacity (%)** — Shows the total capacity used in GBs for disks that comprise this storage tier.
- **Number of members** — Shows the number of members in the tier.
- **Model** — Shows the vendor name and model number.

Operational controls

All FAST controls are set per Symmetrix system and apply to all storage groups (applications) that use this Symmetrix. Some of the parameters used by FAST are also shared with Symmetrix Optimizer. All changes are allowed while FAST is enabled and running.

The following high-level configuration parameters affect FAST:

- **Automatic mode** — Device movement is executed without user approval.
- **User-Approved mode** — Users are presented with a configuration change plan that FAST determined is required to achieve better performance. The user will have to approve this plan in which case all the recommended changes will be executed. If the user declines the plan, no data movement will happen.

Best Practice! FAST and Optimizer share the same approval mode which must be either Automatic or User-Approved. User-Approved mode is recommended for new installs or users unfamiliar with FAST.

Thresholds:

- The maximum number of device movements per day.
- The maximum number of simultaneous device movements.

Workload Analysis settings:

- **Workload Analysis Period** — The period that the analysis should consider. It can be set from 1 hour to 4 weeks.
- **Workload Initial Period** — Elapse time before starting the first analysis. It can be set from 1 hour to 4 weeks but no longer than the Workload Analysis Period.

Scheduling Time Windows:

- **Performance window** — This is the window when statistics are collected.
- **Move window** — This is the time window when device movements are allowed to occur.

Best Practice! Use the workload Trend Views of Symmetrix Performance Analyzer (SPA) to assist with identifying the ideal FAST performance and move time windows based on the sampled I/O activity of FAST targeted volumes. FAST performance windows should ideally coincide with periods of maximum workload; while the FAST move windows should ideally occur during quieter periods of comparatively lower activity.

Swap Only Flag – If the Swap Only flag is set, FAST will not plan any moves. In this case, all device movements are swaps.

Best Practice! FAST device move and swap operations may result in a short term impact to disk response time for a very limited period. In the case of swaps, the impact will be greater due the greater number of devices involved. This impact is not specifically caused by FAST; however, it is a consideration for FAST data movement.

Use unmapped/unmasked devices for swaps — FAST will consider unmapped/unmasked devices that are not in any storage group for swaps. The purpose of this flag is to ensure that devices that are not busy do not occupy high-performance disks.

Workload analysis and FAST algorithms

There are two categories of algorithms available in FAST:

- Performance-based
- Capacity-based

The Capacity-based category is comprised on a single algorithm, while the Performance category is further comprised of the FAST *flash/non-flash promotion demotion algorithm* and the *FC/SATA cross-tier algorithm*.

FAST will choose one of three available algorithms:

- FAST promotion/demotion to and from Flash Drives
- Inter-tier FAST FC/SATA
- FAST capacity based algorithm

The starting time, ending time, and suggestion results of the selected algorithm are recorded within the *configuration change plan*.

Device performance is used to analyze if the devices are located on the right disks. FAST creates periodically the device layout for the best performance, and prepares a list, also known as a configuration change plan, of device movements that will result with this layout.

The goal of the *performance-based algorithms* is to achieve better performance in terms of making best use of faster tiers and balancing the loads within same tiers. The goal of the *capacity-based algorithm* is to eliminate the existing capacity violation by swapping or moving the out-of-tier devices. Intra Symmwin Disk Group load balancing is executed only if the Optimizer is enabled. Each of these algorithms is given a weight and, at every plan creation cycle, one algorithm is chosen for execution.

When FAST is in automatic mode, a plan is created at the start of each configured move window. At this time, one of the algorithms is picked and executed. If the execution of the plan completes before the end of the current move window, another algorithm is picked and executed. The configuration change plan will not include more devices than the configured maximum number of devices to move per day. When FAST is in User Approval mode, a configuration change plan is created every hour. The user can review the last created plan and approve it for execution.

Performance-based algorithms

The goal of the *performance-based inter-tier algorithm* is to make best use of faster disk technologies to achieve increased performance. There are two major disk technologies under FAST management:

- Flash drives
- SATA/FC drives

FAST must model these disk technologies differently due to differing performance behavior. Flash drives are considered the most desired technology and should be populated as much as possible within the FAST policy. On the other hand, Flash drives will be a limited resource and often shared among various storage groups. FAST has been designed to ensure devices serving heavy I/Os will be moved into this technology. The balance of capacity usage across storage groups may then be controlled to ensure one storage group does not consume the entire Flash drive resources.

FAST also incorporates modeling of Flash drive performance for the given I/O. In the traditional Optimizer algorithm, the disk performance modeling is a critical part of optimizer algorithm due to the fact that the target function of the optimization is the disk utilization balance.

However, in FAST algorithm the main goal or target function is to maximize the utilization of Flash drive resource so that the top heavy I/O hitters are moved into Flash drive for better I/O performance.

Flash drives have very specific read/write performance characteristics which differ when their I/O profile contains significant amounts of write activity. Flash drive performance within FAST is estimated through a reference table where measured performance data is organized along three dimensions:

- Write percentage
- Average I/O size in blocks
- Total IOs per seconds

Promotion/demotion to and from EFD

The goal of the EFD promote/demote algorithm is to promote all the busy devices indicated by the weighted read density larger than a threshold to the Flash drives if they are not currently on the Flash drives. A busy device can be promoted to a Flash drive tier either by moving it to a Flash drive directly or by swapping it with a free device on Flash drives.

To determine which devices will benefit most from being stored on EFD, FAST must compare the performance of all devices that are configured in storage groups associated with FAST policies. As EFD exhibits different performance characteristics than HDD disks, FAST evaluates the *EFD performance Score* for each device. To maximize the overall system performance, devices with the highest *EFD performance Score* should be stored on higher performance tiers.

The *EFD performance Score* is a function of Read-Density, Read-Ratio and storage group Priority. Read-Density is derived by dividing the average reads per second by the device size. Read-Ratio is derived by dividing the average reads by the total I/Os (the sum of reads and writes). Devices with highest *EFD performance Score* that are not on EFD are the candidates to be moved there, while devices with low EFD performance Score will be stored on FC or SATA disks.

The algorithm produces a configuration change plan that includes all the moves and swaps that should be executed to improve the system performance. When creating the configuration change plan, FAST prefers moves over swaps. Only if no free devices are available for moves, or when the user configured *SWAP-ONLY*, swaps will be planned. The FAST promotion/demotion to and from EFD algorithm is concerned with device performance (utilization) and takes into account the FAST policy storage allocation rules.

The process to determine candidate devices for the *Promotion/demotion to and from EFD* algorithm is as follows:

1. Build a device list including all the devices in the storage groups with associated FAST policies. Upon users' request, the device list also includes all the unmapped/unmasked devices, which are not in the storage groups, but are on the tiers associated with storage groups. For the metadevices, only the metaheads appear in the device list.
2. Sort all the devices regardless the storage groups and the tiers they belong to according to the metric weighted read density in the descend order.
3. The algorithm first attempts to find all possible moves to promote devices to Flash drives. For the remaining devices on the non-Flash drives, it attempts to find possible swaps to promote busy devices to the Flash drives, and demote free devices to the non-Flash drives. Combine the moving suggestion and the swapping suggestion together and sort them in the descent order according to the source device's value of weighted read density. The first couple of promotions are selected according to the maximum number of promotions allowed each time.

Inter-tier FAST FC/SATA, or non-EFD cross-tier balancing

The FAST FC/SATA cross-tier algorithm is similar to the traditional optimizer algorithm. The goal of this algorithm is to balance the FC and SATA disk utilization. The disk service time (SAT) is used as the metric to indicate the disk utilization. The traditional optimizer utilization calculation is applied to the devices in the storage groups.

Each disk group is executed by this algorithm separately by the user defined priority. For each priority, all devices are sorted by the disk service time (SAT). The algorithm first attempts to move the most utilized devices to less utilized disks. If no disks are available for a move, the algorithm searches for devices that can be swapped.

The process to determine the candidate devices for Inter-tier FAST FC/SATA algorithm is as follows:

1. Sort all the FC and SATA drives in descend order according to the disk service time (SAT).
2. Going from left to pick a *source* drive with the specific storage group priority.
3. Loop through right *target* drives.
4. For each *source* and *target* disk pair:
 - A. Sort the devices on the source disk in the order of high to low according to the device load.
 - B. The algorithm attempts to find moves for the source devices BY using the target disks. The target disks include all the disks in the maximum affinity group with the select *target* disk in step 3. The number of the target disks is up to 16.). After a suitable move is found (the rules are explained next), record the move, and mark all the source disks containing the source device and all the disks in the 16 target disks, which have free space for the source device. They will not be used in the potential moves or swaps.
 - C. If no moves were found, sort the devices on the target disk in the order of low to high, according to the device load. Check each pair of the source and target devices to find a swappable pair (the rules are described next). After a swappable device pair is found, record the swap and continue to process the next source disk. All the disks containing the source device and target device will be marked and will not be in use in the potential moves or swaps.

Capacity-based algorithm

The FAST capacity-based algorithm enforces the FAST policy storage usage rules. The goal of the capacity-based algorithm is to enforce the user defined FAST policy for each storage group. One device is considered to be violating capacity restriction if it is currently out of tier, or in an out-of-capacity limits tier. Storage groups may be in violation of the FAST policy when FAST is first turned on or after configuration changes are made.

The FAST capacity-based algorithm finds all the devices that are in violation of the storage rules and looks for pairs of devices to swap so that after the execution of the swap, both devices will be in compliance with the storage rules.

The capacity-based algorithm attempts to eliminate capacity restriction violation by following operations:

- Swap a pair of out-of-capacity devices, both of them become in-capacity devices
- Move one out-of-capacity device to eliminate the capacity restriction violation
- Swap one out-of-capacity device with one in-capacity device, both of them are in-capacity after the swap

In this way, the *FAST capacity-based algorithm* first attempts to find pairs of devices to swap and when none are found, FAST will attempt to move violating devices to nonconfigured disks. This is done to ensure that after the moves take place, the configuration will be in compliance with the configured storage policy.

The reason for attempting to find swaps before moves is that with swaps the violation can be corrected for two devices at the same time while moves correct the violation of a single device at a time.

Device movement between tiers

When FAST determines that for better performance devices that reside in one tier should be on another tier, it creates the configuration change plan that details how to achieve the preferred configuration by moving the data. Using Virtual LUN technology, FAST either moves devices to unallocated disk space, or swaps pair of devices. In both cases, the device layout changes but the front-end mapping—LUN configuration remains the same. As moves have less overhead, FAST prefers to move devices over swapping them. When the *SWAP ONLY* flag is set, FAST will only swap devices.

Best Practice! If performance issues are encountered as a result of FAST automated move or swap operations, the devices in questions should be manually rolled back to the prior state. This can be accomplished through SymMigrate (moves) or Optimizer (swaps).

In-tier, out of tier — A device is in-tier if the physical location of the device is within the disk groups of the tier and the protection type of the device is compatible with the requirement of the tier. For all other cases, the device is considered as out of tier.

Device move

Devices are moved only when the swap-only flag is not set. Even though device moves allow more flexibility, in FAST, a device can be moved only to unconfigured space. For each device that is being moved, the system identifies unconfigured disk space from the target tier. When moving a device from the source disks to the target disks, FAST will configure the target disks and unconfigure the source disks. This may result in fragmented disks.

One source device can be moved to a set of target disks if all of the following conditions are met:

1. The target disks are in one tier (referred hereon as target tier) of the associated FAST policy of the storage group of the source device.
2. The move of the source device to the target disks will follow the RAID protection rules of the target tier in the associated FAST policy of the source storage group.
3. If the source device is currently out of the target tier, move the source device to the target tier if doing so will not violate the capacity restriction of the target tier.
4. After the move, it will not exceed the target DAs maximum I/O limits and supported target number limits.

Device swap

When FAST determines that a device, hereon referred to it as Device-A, should be on a different tier and no free disk space is available on the target tier, FAST will attempt to find another device, hereone referred to as Device-B, that could be swapped out of the target tier. FAST will then swap the physical hypervisors of device-A with the physical hypervisors of Device-B.

The swap is accomplished through a Device Reallocation Volume (DRV). DRVs have protection level RAID-1 and can reside on any tier. However, the location of the DRV affects the time to complete the swap. The DRV size cannot be smaller than the devices that are being swapped. Therefore, it should be at least as big as the largest device. Swap through a DRV is accomplished through three moves.

Best Practice! The DRV device size must be the same size or larger than the devices that are being swapped.

A source device can be swapped to the location of a target device only if all of following rules are satisfied:

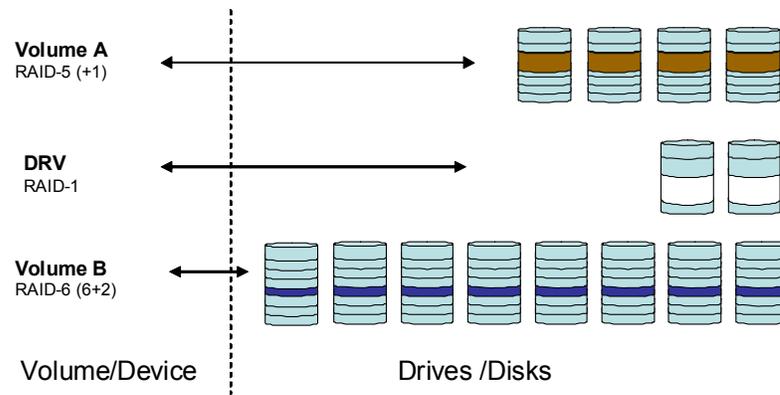
1. The target disks (the disks containing the target device) are in one tier (referred as target tier next) of the associated FAST policy of the storage group of the source device.
2. The swap of the source device to the target device will follow the raid protection rules of the target tier in the associated FAST policy of the source storage group.
3. If the source device is currently out of the target tier, swap the source device to the target tier if doing so will not violate the capacity restriction of the target tier.
4. The source device and target device are in different Symmwin groups or with different protection type.
5. After the swap, it will not exceed the target DAs maximum IOPS limits.

Best Practice! FAST allows a maximum of 32 simultaneous swaps with a default of eight 8 simultaneous swaps. It is recommended for new FAST implementations (FBA or CKD) to create eight DRVs to match the default number of simultaneous operations; however the maximum number of simultaneous operations will remain limited by the lower value of the following: configured number of simultaneous swaps or the number of available DRVs.

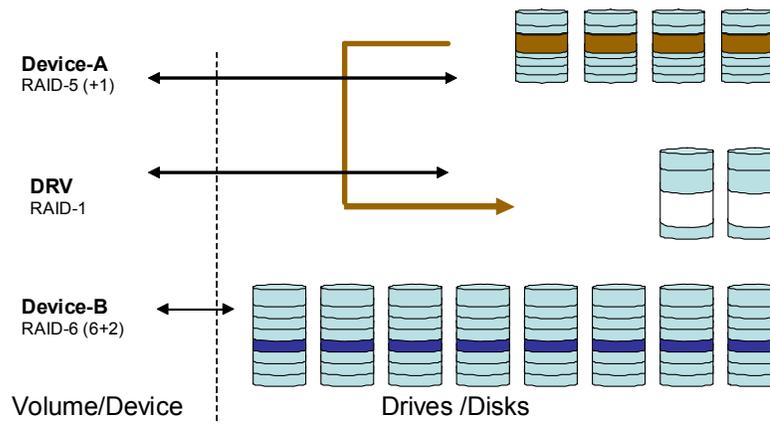
The following illustrates an example of swapping a device (Device-A) from a tier with protection level RAID-5 (3+1) with a device (Device-B) on another tier with protection level RAID-6(6+2). At the end, Device-A will have protection level RAID-6(6+2) and Device-B will have protection level RAID-5 (3+1).

Device movement between tiers

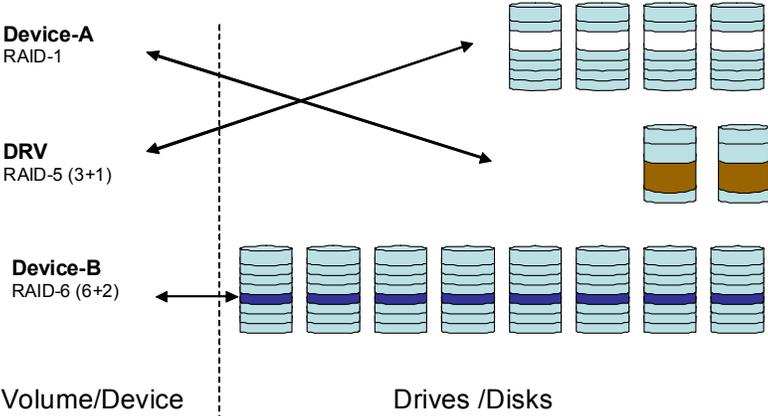
This is the starting configuration before the swap operation is performed.



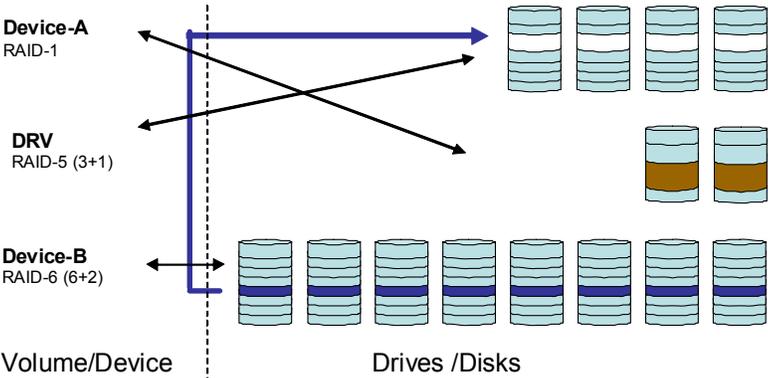
In the first step of the swap operation, Device-A is moved to a DRV.



This is the resulting configuration after Device-A has been moved to a DRV.

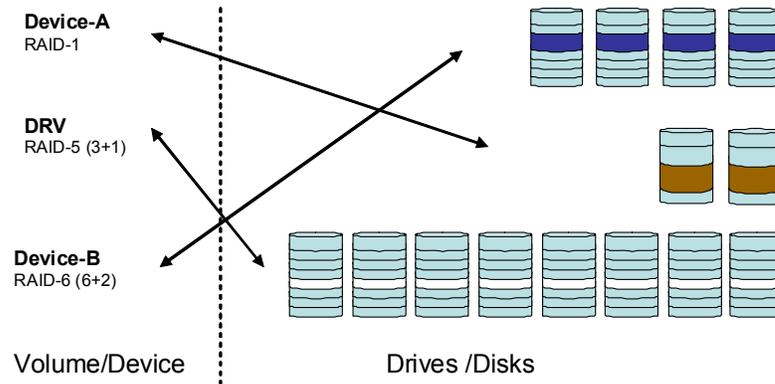


In the next step of the swap operation, Device-B is moved to the DRV

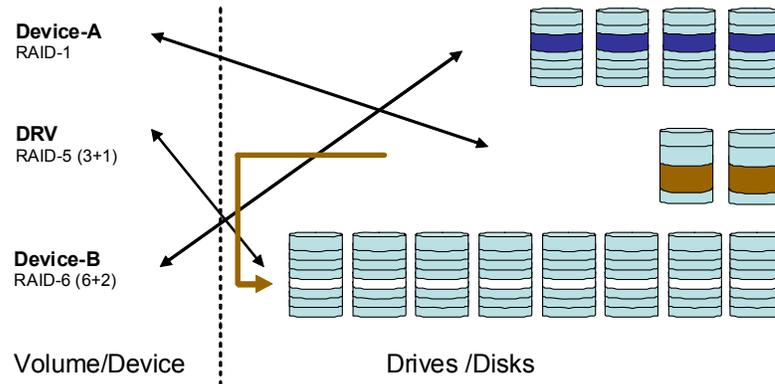


Device movement between tiers

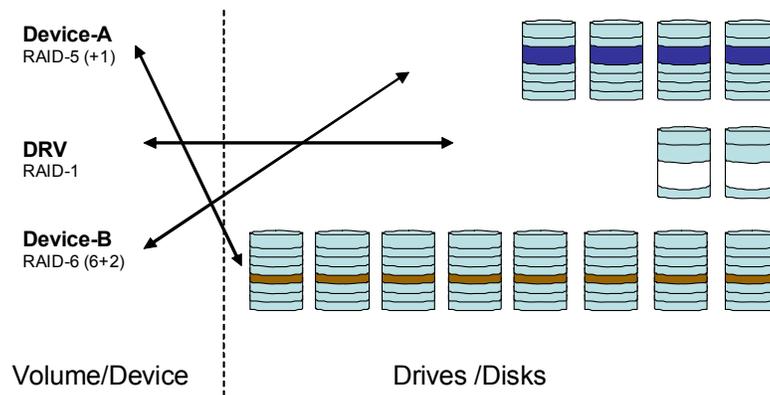
This is the resulting configuration after Device-B has been moved to the DRV.



In the final step of the swap operation after Device-A is moved to the DRV.



This is the final device configuration after the swap operation is complete.



The time required to complete the swap operation depends on the media of the source, target, and the DRVs. Therefore, the DRV should not be on the least performing tier.

Symmetrix movement performance and resources

Data movement can consume large amount of cache. To avoid data movement from taking over a lot of resources, the number of WP tracks that are in the process of being moved (by Virtual LUN) is limited. There is a per-device limit and a system-wide limit. No track copies are scheduled if any of the limits is exceeded. The per-device limit is 126 cache slots. The system limit is 25 percent of the system max limit, which is 80 percent of the number of cache slots available.

At the start of the configured *Move Window*, FAST starts to execute its plan regardless of resource usage at this time. After the device movement starts, it cannot be aborted, paused, or stopped.

Best Practice! Use copy Quality of Service (QOS) to affect the speed of FAST device movement operations within back-end resource constrained environments.

There are two mechanisms to control and isolate the impact of migration processes on the overall system:

Copy QOS — The migration copy rate can be controlled by placing the involved devices in a device group and applying the required QOS level (0 — no delay, 16 — max delay ~ 4 second between copy jobs for the device).

Best Practice! Copy QOS setting must not be configured high enough that it causes FAST to exceed the targeted move window duration. If the device movement duration exceeds 48 hours, the Symmetrix will dial home to initiate a support escalation.

Dynamic Cache Partitioning (DCP) — When a device is assigned to use, a cache partition all I/O related to the device (host I/O as well as copy jobs) will utilize the assigned partition and therefore will be isolated from workload impact on the rest of the system. However, the system WP count limits will still apply.

Best Practice! When using DCP with FAST, it is recommended to assign partitions based on front-end applications rather than back-end disk technology. In addition, also ensure that DCP partitions are configured in relation to FAST policies.

Metavolume considerations

Operations on metadevices are performed on all the metamembers. All the metamembers are either all moved or all swapped. Moving a few, while swapping the others is not allowed.

Moving a metadevice

If the number of metamembers of the source device is N, FAST must locate N unconfigured disk spaces. When moving to FC or SATA disks, these N unconfigured disk spaces must reside on different disks. When moving to EFD, these N unconfigured disk spaces may reside on less than N disks, which means that multiple metamembers reside on a single disk.

Swapping a metadvice

If the number of metamembers of the source device is N , FAST must locate N destination drives or metamembers that are available for swapping. Examples of options for swaps of a source metadvice with N members are as follows:

- A destination meta with N members
- A destination meta with M members and a destination meta with $N-M$ members
- N (nonmeta) devices
- A destination meta with M members and $N-M$ (non-meta) devices
- Multiple metadvice with total number of metamembers M and $N-M$ (non-) devices

If a target device is a metadvice, all its members must be swapped.

Note: Multiple metamembers of a device that resides on EFD can be on the same physical disk. However, each of the metamembers of a device that resides on FC or SATA must reside on distinct physicals. This restriction may make the swapping of a metadvice out of EFD impossible.

If there are more metamembers than DRVs, then metavolumes are moved in multiple stages. It should be noted that during a multi-stage move, metamembers may have different protection levels.

For a given source metadvice, FAST will loop through the target devices until it identifies swap pairs for all the source metamembers or the target device list is exhausted.

The procedure to find swaps for source metamembers is as follows:

1. If the target device is not a metamember and the target device is swappable with one unused source metamember, perform the swap and mark the selected source metamembers as used.
2. If the target device is a metamember, and FAST can find swaps for the entire target metamembers with unused source metamembers, perform the swaps and mark the selected source metamembers as used.
3. If the target device is a metamember and the target metadvice has the same number of members as the source metadvice, FAST attempts to swap the entire source metadvice and the entire target metadvice. If this is successful FAST records the swaps and is done with the source metadvice.

Device size consideration

Best Practice! For optimal use by FAST, Symmetrix devices (including metadvice) should be relatively small and uniform in size. Ideally, devices should be an order of magnitude less than available Flash drive capacity. As a general rule, size devices under FAST to ensure at least 10 devices are able to fit onto the available Flash drives. The more candidate devices which may fit onto Flash drives; the most optimally FAST will be able to exploit the workload skew.

Limitations and restrictions

The following considerations apply to FAST:

- FAST management of Virtual Provisioning devices is not supported, that is, thin devices and DATA devices.
- FAST management of EMC TimeFinder®/Snap devices is not supported, that is, VDEV and SAVE devices.
- FAST management of CKD EAV, or IBM i (AS400) devices is not supported.
- At least one DRV of each emulation must be configured in the array for FAST to perform swap operations.

Consolidated list of FAST best practices

The following list represents the FAST best practices, which were presented throughout this document, consolidated in a single location. The associated section provides additional details regarding a specific best practice.

Pertinent terminology:

- For FAST environments targeting dynamic Symmetrix tiers, it is possible to start with dynamic tiers or migrate existing static tiers to dynamic. However, be aware that it will not be possible to convert dynamic tiers back to static.

Drive technology and tiering:

- SATA drive technology is suited for data that is not accessed often, or data that is accessed mostly sequentially.
- It is best for I/O bandwidth to be shared by the larger pool of SATA disks (in comparison to FC disks). In this way, on an average, each SATA disk will be accessed less often.

Understanding workload skew:

- FAST is most effective when the I/O access patterns are highly skewed.
- When migrating to a new configuration or platform targeted for FAST, it is best to plan and configure optimum tiers in the target or upgrade system and to pre-place volumes in these tiers based on their performance characteristics.
- Use Tier Advisor to assist with recommending Storage Types or tiers within the Symmetrix.
- Tier Advisor, SymmMerge, and similar tools support configuration planning and are recommended when the user has I/O workload information available.
- SymmMerge is best used to assist with placement or migration planning, while Tier Advisor assists with the sizing of Symmetrix tiers.

SPA and FAST:

- Use SPA Snapshot views to determine applicable workloads and target Symmetrix tiers for FAST.
- The longer the time range specified in the Storage Group Snapshot

view, the more accurate the view of the workload will be.

- Use the SPA 2.0 Snapshot > Storage Group > Symmetrix Tiers Capacity Allocation (%) view to provide tier allocation information for storage groups that have FAST policies.
- Use the SPA 2.0 Snapshot > Storage Group > Hit/Miss Distribution view to determine applicable workloads and target Symmetrix tiers for FAST. For example, SPA may indicate a significant amount of average read miss activity for a storage group. In this case, the group may be better placed either fully or partly onto an EFD Symmetrix tier as defined within a FAST policy.
- Use SPA Trend views to facilitate comparisons of storage group performance characteristics before and after they have been placed under FAST control.
- Select a long time range (6 months or longer) in the Storage Group Trend view to compare storage group performance characteristics before and after they have been placed under FAST control.
- Use the Trend > Storage Tiers Capacity Allocations (%) to show the capacity allocation percent for each tier associated with the storage group through a FAST policy.
- Use the Storage Group and Symmetrix Tier Diagnostic views for critical diagnostic information (alerts and events) regarding storage groups and Symmetrix tiers.

Operational controls:

- FAST and Optimizer share the same approval mode which must be either Automatic or User-Approved. User-Approved mode is recommended for new installs or users unfamiliar with FAST.
- FAST device move and swap operations may result in a short term impact to disk response time for a very limited period. In the case of swaps, the impact will be greater due the greater number of devices involved. This impact is not specifically caused by FAST; however, it is a consideration for FAST data movement.
- Use the workload Trend Views of Symmetrix Performance Analyzer (SPA) to assist with identifying the ideal FAST performance and move time windows based on the sampled I/O activity of FAST targeted volumes. FAST performance windows should ideally coincide with periods of maximum activity; while the FAST move windows should ideally occur during quieter periods of comparatively lower activity.

Device movement between tiers:

- If performance issues are encountered as a result of FAST automated move or swap operations, the devices in questions should be manually rolled back to the prior state. This can be accomplished through SymMigrate (moves) or Optimizer (swaps).
- The DRV device size must be the same size or larger then the devices that are being swapped.
- FAST allows a maximum of 32 simultaneous swaps with a default of eight 8 simultaneous swaps. It is recommended for new FAST implementations (FBA or CKD) to create eight DRVs to match the default number of simultaneous operations; however the maximum number of simultaneous operations will remain limited by the lower value of the following: configured number of simultaneous swaps or the number of available DRVs.
- It is not recommended in CKD environments to configure less than four DRVs.
- Use copy Quality of Service (QOS) to affect the speed of FAST device movement operations within back-end resource constrained environments.
- Copy QOS setting must not be configured high enough that it causes FAST to exceed the targeted move window duration. If the device movement duration exceeds 48 hours, the Symmetrix will dial home to initiate a support escalation.
- When using Dynamic Cache Partitioning (DCP) with FAST, it is recommended to assign partitions based on front-end applications rather than back-end disk technology. In addition, also ensure that DCP partitions are configured in relation to FAST policies.
- For optimal use by FAST, Symmetrix devices (including metadevices) should relatively small and uniform in size. Ideally, devices should be an order of magnitude less than available Flash drive capacity. As a general rule, size devices under FAST to ensure at least 10 devices are able to fit onto the available Flash drives. The more candidate devices which may fit onto Flash drives; the most optimally FAST will be able to exploit the workload skew.

Summary and conclusion

The EMC Symmetrix VMAX Series with Enginuity is the newest addition to the Symmetrix family, and the first high-end system purpose-built for the virtual data center. Based on the Virtual Matrix Architecture, Symmetrix VMAX scales performance and capacity to unprecedented levels, delivers nondisruptive operations, and greatly simplifies and automates the management and protection of information.

The Enginuity operating environment provides the intelligence that controls all components in a Symmetrix VMAX array. It coordinates real-time events related to the processing of production data. It applies self-optimizing intelligence to deliver the ultimate performance, availability, and data integrity required in a platform for advanced storage functionality.

The Enginuity version 5874 Q4 '09 SR is the latest Enginuity release supporting the Symmetrix VMAX storage arrays. Enginuity version 5874 features new software functionality such as Fully Automated Storage Tiering (FAST) and enhancements to current products including Virtual Provisioning and SRDF.

EMC Symmetrix VMAX FAST technology is designed to automate the allocation and relocation of application data across different performance tiers based on the changes in application performance requirements. FAST helps customers maximize the benefits of tiered storage within the array by prioritizing and optimizing application performance and reduce labor cost.

Symmetrix VMAX FAST Standard Provision environments will proactively monitor volume workloads and automatically move busier volumes to higher performing flash drives, and the slower volumes to higher capacity, lower performance drives (SATA). The promotion/demotion activity will be driven by policies assigned to the FAST storage groups and will be executed nondisruptively, without affecting business continuity and availability.

Management and operation of FAST is accomplished via Symmetrix Management Console (SMC) or Solutions Enabler (SYMCLI). The creation of tiers, policies, and storage groups may also be accomplished through easy-to-use wizards within SMC. Applicable chargeback information is provided via EMC StorageScope within EMC ControlCenter.

SPA version 2.0 provides detailed insight into performance trending, forecasting, alerts and resource utilization. In support of FAST, SPA version 2.0 specifically provides detailed information regarding storage groups and Symmetrix tiers within its snapshot, trend, and diagnostic views.

Copyright © 2009, 2010 EMC Corporation. All Rights Reserved.

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

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

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

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

For the most up-to-date regulatory document for your product line, go to the Technical Documentation and Advisories section on EMC Powerlink.

All other trademarks used herein are the property of their respective owners.