

**Implementing
Fully Automated Storage Tiering
for Virtual Pools (FASTVP) for
EMC[®] Symmetrix[®] VMAX[™] Series Arrays**

Technical Notes

P/N 300-012-015

REV A04

May, 2012

This technical note contains information on these topics:

- ◆ Executive summary 2
- ◆ Introduction 2
- ◆ Fully Automated Storage Tiering 3
- ◆ FAST VP 6
- ◆ FAST VP architecture 13
- ◆ FAST VP configuration 16
- ◆ FAST VP performance data collection 22
- ◆ FAST VP performance data analysis 26
- ◆ FAST VP data movement 31
- ◆ Advanced FAST VP features 38
- ◆ FAST VP interoperability 42
- ◆ Management interface: Unisphere for VMAX 47
- ◆ Management interface: SYMCLI 111
- ◆ Reporting 157
- ◆ Conclusion 179
- ◆ Appendix A: FAST VP state 180
- ◆ Appendix B: Feature support 182
- ◆ References 183

Executive summary

Information infrastructure must continuously adapt to changing business requirements. EMC® Symmetrix® Fully Automated Storage Tiering for Virtual Pools (FAST VP) automates tiered storage strategies, in Virtual Provisioning™ environments, by easily moving workloads between Symmetrix tiers as performance characteristics change over time. FAST VP performs data movements, improving performance, and reducing costs, all while maintaining vital service levels.

Introduction

EMC Symmetrix VMAX® FAST VP automates the identification of active or inactive application data for the purposes of reallocating that data across different performance/ capacity tiers within an array. FAST VP proactively monitors workloads at both the LUN level and sub-LUN level in order to identify busy data that would benefit from being moved to higher-performing drives. FAST VP also identifies less-busy data that could be moved to higher-capacity drives, without affecting existing performance. This promotion/ demotion activity is based on policies that associate a storage group to multiple drive technologies, or RAID protection schemes, by way of virtual pools, as well as the performance requirements of the application contained within the storage group. Data movement executed during this activity is performed non-disruptively, without affecting business continuity and data availability.

Audience

This technical note provides an in-depth look at the Fully Automated Storage Tiering VP feature for Symmetrix VMAX Family arrays. The intended audience includes system and storage administrators, customers, EMC staff, and partners who need to understand how to implement and leverage FAST VP.

Fully Automated Storage Tiering

Fully Automated Storage Tiering (FAST) automates the identification of active or inactive data for the purposes of relocating application data across different performance/ capacity tiers within an array.

The primary benefits of FAST include:

- ◆ Elimination of manually tiering applications when workload characteristics change over time.
- ◆ Automating the process of identifying data that can benefit from Enterprise Flash Drives (EFDs) or that can be kept on higher-capacity, less-expensive SATA drives without impacting performance.
- ◆ Improving application performance at the same cost, or providing the same application performance at lower cost. Cost is defined as: acquisition (both HW and SW), space/ energy, and management expense.
- ◆ Optimizing and prioritizing business applications, allowing customers to dynamically allocate resources within a single array.
- ◆ Delivering greater flexibility in meeting different price/ performance ratios throughout the lifecycle of the stored information.

The need for FAST

Due to advances in drive technology, and the need for storage consolidation, the number of drive types supported by Symmetrix arrays has grown significantly. These drives span a range of storage-service specializations and cost characteristics that differ greatly.

Several differences exist between the four drive technologies supported by the Symmetrix VMAX Series arrays: Enterprise Flash Drive (EFD), Fibre Channel (FC), Serial Attach SCSI (SAS), and SATA. The primary differences are:

- ◆ Response time
- ◆ Cost per unit of storage capacity
- ◆ Cost per unit of storage request processing

At one extreme are EFDs, which have a very low response time and the ability to handle very high levels of requests, but with a high cost per unit of storage capacity.

At the other extreme are SATA drives, which have a low cost per unit of storage capacity, but high response times and high cost per unit of storage request processing.

Between these two extremes lie Fibre Channel and SAS drives.

Based on the nature of the differences that exist between these four drive types, the following observations can be made regarding the most suitable workload type for each drive:

- ◆ Enterprise Flash Drives: 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 response time provided by the drive, and the low cost per unit of storage request processing, without requiring a lot of storage capacity.
- ◆ SATA drives: SATA drives are suited towards workloads that have a low back-end storage request density.
- ◆ Fibre Channel and Serial Attach SCSI drives: 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.

This disparity in suitable workloads presents both an opportunity and a challenge for storage administrators.

To the degree it can be arranged for storage workloads to be served by the best-suited drive technology, the opportunity exists to improve application performance, reduce hardware acquisition expenses, and reduce operating expenses (including energy costs and space consumption).

The challenge, however, lies in how to realize these benefits without introducing additional administrative overhead and complexity.

The approach taken with FAST is to automate the process of identifying which regions of storage should reside on a given drive technology, and to automatically, and non-disruptively, move storage between tiers to optimize storage resource usage accordingly. This also needs to be done while taking into account optional constraints on tier capacity usage that may be imposed on specific groups of storage devices.

FASTDP and FASTVP

EMC Symmetrix VMAX FAST DP and FAST VP automate the identification of data volumes for the purposes of relocating application data across different performance/ capacity tiers within an array, or to an external array using Federated Tiered Storage (FTS).

Note: Federated Tiered Storage is only available for Open Systems environments. Also, FAST support applies only to FAST VP.

For more information on Federated Tiered Storage, refer to the Design and Implementation *Best Practices for EMC Symmetrix Federated Tiered Storage (FTS)* Technical Note available at [http:// powerlink.emc.com](http://powerlink.emc.com).

FAST DP operates on disk group provisioning Symmetrix volumes. Data movements executed between tiers are performed at the full volume level.

FAST VP operates on Virtual Provisioning thin devices. As a result, data movements can be performed at the sub-LUN level. A single thin device may have extents allocated across multiple thin pools within an array, or on an external array using FTS.

Note: For more information on Virtual Provisioning, refer to the *Best Practices for Fast, Simple Capacity Allocation with EMC Symmetrix Virtual Provisioning Technical Note* available at <http:// powerlink.emc.com>.

Because FAST DP and FAST VP support different device types (disk group provisioning and Virtual Provisioning, respectively), they both can operate simultaneously within a single array. Aside from some shared configuration parameters, the management and operation of each can be considered separately.

The main focus of this document is to discuss the implementation and management of FAST VP in Virtual Provisioning environments.

Note: For more information on FAST, refer to the *Implementing Fully Automated Storage Tiering (FAST) for EMC Symmetrix VMAX Series Arrays Technical Note* available at <http:// powerlink.emc.com>.

FAST managed objects

There are three main elements related to the use of both FAST and FAST VP on Symmetrix VMAX Series arrays. These are:

- ◆ Storage tier: A shared resource with common technologies and RAID protection
- ◆ 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
- ◆ Storage group: A logical grouping of devices for common management

Figure 1 shows the FAST managed objects.

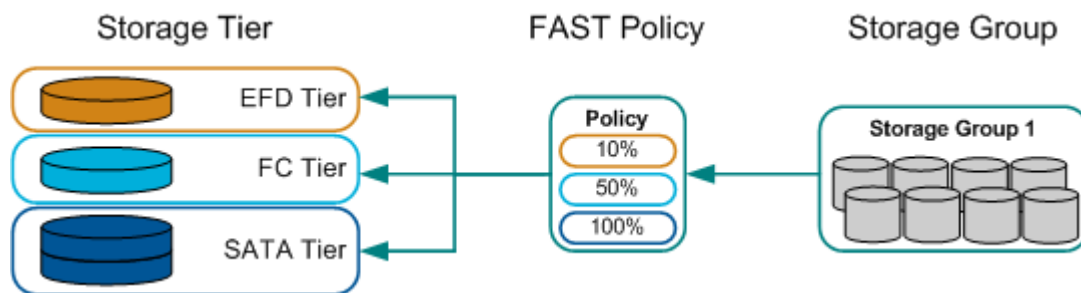


Figure 1. FAST managed objects

Each of the three managed objects can be created and managed by using either Unisphere® for VMAX, Symmetrix Management Console, or the Solutions Enabler Command Line Interface (SYMCLI).

FASTVP

FAST VP automates the identification of thin device extents for the purposes of reallocating application data across different performance tiers within a single array, or to an external array. FAST VP proactively monitors workloads at the LUN level and sub-LUN level in order to identify busy data that would benefit from being moved to higher-performing drives. FAST VP also identifies less-busy data that could be

moved to higher-capacity drives, without existing performance being negatively affected.

For FAST VP to operate, the three storage elements that need to be configured are: Storage tiers, FAST policies, and storage groups.

Storage tiers

A Symmetrix storage tier is a specified set of resources of the same drive technology type (EFD, FC/ SAS, SATA, or external), combined with a given RAID protection type (RAID 1, RAID 5, RAID 6, or unprotected), and the same emulation (FBA or CKD).

Note: The unprotected RAID type may only be applied to a tier residing on an FTS-connected storage array.

There are two types of storage tiers: disk group provisioning (DP) and Virtual Provisioning (VP).

Disk group provisioning tiers

The storage tier type used by FAST is called a DP tier. It is defined by combining one or more physical disk groups of the same technology type and a RAID protection type.

Virtual Provisioning tiers

For FAST VP, the storage tier is called a VP tier. When defined, VP tiers contain between one and four thin storage pools (each thin pool must contain data devices of the same RAID protection type, and be configured on the same drive technology).

VP tier characteristics

For use with FAST VP, a thin storage pool must contain data devices configured as the same RAID protection on a single drive technology, and of the same emulation. In the case of FC, SAS, and SATA drives, the rotational speed of the drives must also match. However, two or more thin pools containing data devices configured on rotating drives of different speeds, but the same emulation, may be combined in a single VP tier.

Note: For the purposes of FAST VP, SAS-based thin pools are considered to be FC. As such, they can be combined into the same tier as a FC thin pool.

A thin pool may only belong to one VP tier. Overlapping of pools between tiers is not allowed.

All VP tiers are considered to be static, meaning that thin pools must be explicitly added to the tier. However, if data devices are added to an existing pool, the additional capacity is automatically made available within the VP tier.

A Symmetrix VMAX Series storage array supports up to 256 Symmetrix tiers. Each Symmetrix tier name may contain up to 32 alpha-numeric characters, hyphens (-), and underscores (_). Tier names are not case-sensitive.

FAST policies

A FAST policy groups between one and three tiers and assigns an upper usage limit for each storage tier. The upper limit specifies the maximum amount of capacity of a storage group associated with the policy that can reside on that particular tier.

FAST policies may include storage tiers of only one type, disk group provisioning (DP) or Virtual Provisioning (VP). Similarly, FAST policies may only include storage tiers of a single emulation type, FBA or CKD.

The first tier added to a policy determines the type of tiers that can subsequently be added.

For policies that include VP tiers, the upper capacity usage limit for each storage tier is specified as a percentage of the configured, logical capacity of the associated storage group.

The usage limit for each tier must be between 1 percent and 100 percent. When combined, the upper usage limit for all thin storage tiers in the policy must total at least 100 percent, but may be greater than 100 percent.

Creating a policy with a total upper usage limit greater than 100 percent allows flexibility with the configuration of a storage group. Data may be moved between tiers without necessarily having to move a corresponding amount of other data within the same storage group.

Multiple FAST policies may reuse the same tier, allowing different usage limits to be applied to different storage groups for the same tier.

A Symmetrix VMAX Series storage array supports up to 256 FAST policies. Each FAST policy name may be up to 32 alpha-numeric characters, hyphens (-), and underscores (_). Policy names are not case-sensitive.

Note: FAST VP only performs promotion/ demotion activity between tiers defined on differing drive technologies. RAID protection and drive rotational speed are not considered. As a result, a FAST VP policy should not be created where two or more tiers use the same drive type. For example, a FAST VP policy should not contain two or more FC tiers.

Storage groups

A storage group is a logical collection of Symmetrix devices that are to be managed together. Storage group definitions are shared between FAST and Auto-provisioning Groups. However, a Symmetrix device may only belong to one storage group that is under FAST control.

Storage groups are associated with a FAST policy, thereby defining the VP tiers that data in the storage group can be allocated on.

FAST VP only supports the movement of certain device types. As a result of this, a storage group created for the purposes of FAST VP may not contain the following device types:

- ◆ Diskless
- ◆ CKD EAV
- ◆ IBM i (520-byte block emulation), ICOS, ICL
- ◆ Metadevice members
- ◆ SAVE (SAVDEV)
- ◆ DATA (TDAT)
- ◆ DRV
- ◆ SFS
- ◆ Vault

Note: Support for IBM i D910 emulation (512-byte block) is provided by means of FBA emulation pools, tiers, and FAST VP policies.

A Symmetrix VMAX Series storage array supports up to 8,192 storage groups associated with FAST policies. Storage groups may contain up to 4,096 devices. Each storage group name may be up to 64 alpha-numeric characters, hyphens (-), and underscores (_). Storage group names are not case-sensitive.

FAST policy association

A policy associates a storage group with up to three tiers, and defines the maximum percentage of logical storage capacity in the storage group that can exist in a particular tier.

The same FAST policy may be applied to multiple storage groups. However, a storage group may only be associated with one policy.

FAST VP supports the association of up to 1,000 storage groups with FAST policies containing thin storage tiers.

It is possible to have both disk group provisioning devices and thin devices in the same storage group. However, the storage group can only be associated with one policy. If it is required that both device types be managed by FAST and FAST VP, then separate storage groups need to be created. These storage groups then need to be associated with policies of the appropriate type.

A storage group associated with a FAST policy may only contain thin devices configured for a single emulation (FBA or CKD). No mixing is allowed. Similarly, the emulation of the thin devices in the storage group must match the emulation of the FAST policy the group is being associated with.

Note: When associating a storage group to a policy containing VP tiers, the bound thin devices in the group must be bound to at least one of the thin pools contained within the policy's tiers.

A storage group associated with a FAST VP policy may contain unbound devices. However, those devices may then only be bound to a thin pool contained within the tiers in the policy.

Priority

When a storage group is associated with a FAST policy, a priority value must be assigned to the storage group. This priority value can be between 1 and 3, with 1 being the highest priority. The default is 2.

When multiple storage groups are associated with FAST VP policies, the priority value is used when the data contained in the storage groups is competing for the same resources in one of the associated tiers. Storage groups with a higher priority are given preference when deciding which data needs to be moved to another tier.

Storage group considerations

The type of policy, DP or VP, associated with a storage group determines which devices are managed under that policy. For example, if a storage group is associated with a policy containing VP tiers, then only the thin devices in that group are managed by that policy.

FAST policy configuration

The FAST VP environment can contain multiple thin storage tiers, FAST policies, and storage groups.

Figure 2 shows three storage groups. Each storage group is associated with a policy. These policies associate the storage groups with up to three storage tiers that are defined in the array.

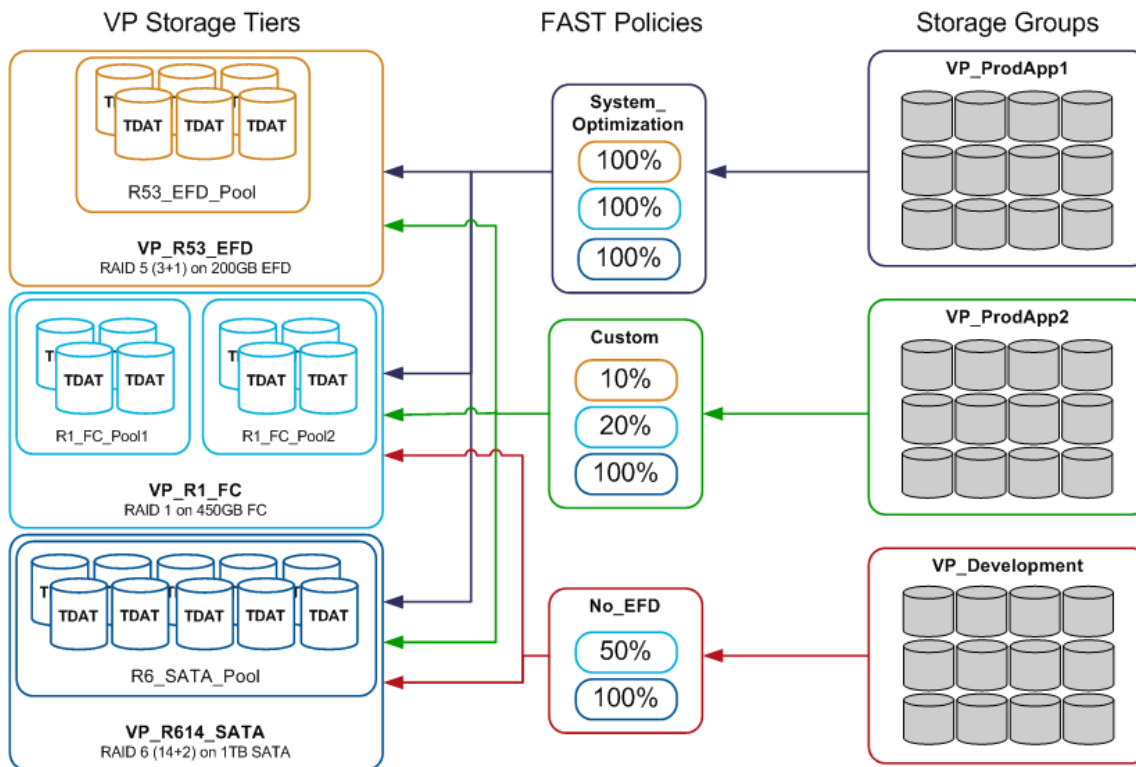


Figure 2. FAST policy association

Based on the System_Optimization policy, FAST VP can place up to 100 percent of the configured, logical capacity of the VP_ProdApp1 storage group in any of the three associated tiers: RAID 5 (3+1) configured on EFD, RAID 1 on FC, or RAID 6 (14+2) on SATA. Such a policy gives the greatest flexibility, as at any given time, all the data can be located on the most appropriate tier.

Note: The default for any newly created policy is to allow 100 percent of the storage group's capacity to be placed on any of the included tiers.

In the case of the Custom policy, 10 percent of the configured capacity of the VP_ProdApp2 storage group can be placed in the EFD tier, 20 percent on the FC tier, and up to 100 percent on the SATA tier.

The No_EFD policy will not move any data to EFD, but does allow 100 percent of the VP_Development storage group's capacity to be on the SATA tier, and up to 50 percent on the FC tier.

If the total for all the tiers combined equals 100 percent and if the thin devices in the storage group are fully allocated, then each tier is utilized as set in the policy. For example, if a policy was set up as 10% EFD, 20% FC, and 70% SATA, then 10 percent of the storage group's capacity will always be located on the EFD tier, and so on for the remaining tiers.

FAST policy compliance

A storage group is considered to be compliant with its associated FAST policy when all data in the storage group is allocated within the bounds of the upper usage limits for each tier contained with the policy.

If all of the data in the storage group is allocated within the tiers contained within the FAST policy, but the allocated capacity in one tier exceeds the upper usage limit for that tier, then the storage group is considered to be non-compliant. In such a case, the FAST controller attempts to correct this non-compliance by relocating data from the VP tier where the usage limit is exceeded to one, or both, of the other tiers contained in the policy. The desired result is to bring the storage group into compliance.

A special case of non-compliance is when some of the data in the storage group is allocated in thin pools not contained within any of the VP tiers within the FAST policy. In this case, the storage group is considered to be non-compliant, and the data considered to be out-of-policy. Again, the FAST controller will attempt to correct this situation by relocating the out-of-policy data to one, two, or all three of the tiers contained in the policy.

FASTVP architecture

There are two components of FAST VP: Symmetrix Engenuity™ and the FAST controller.

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

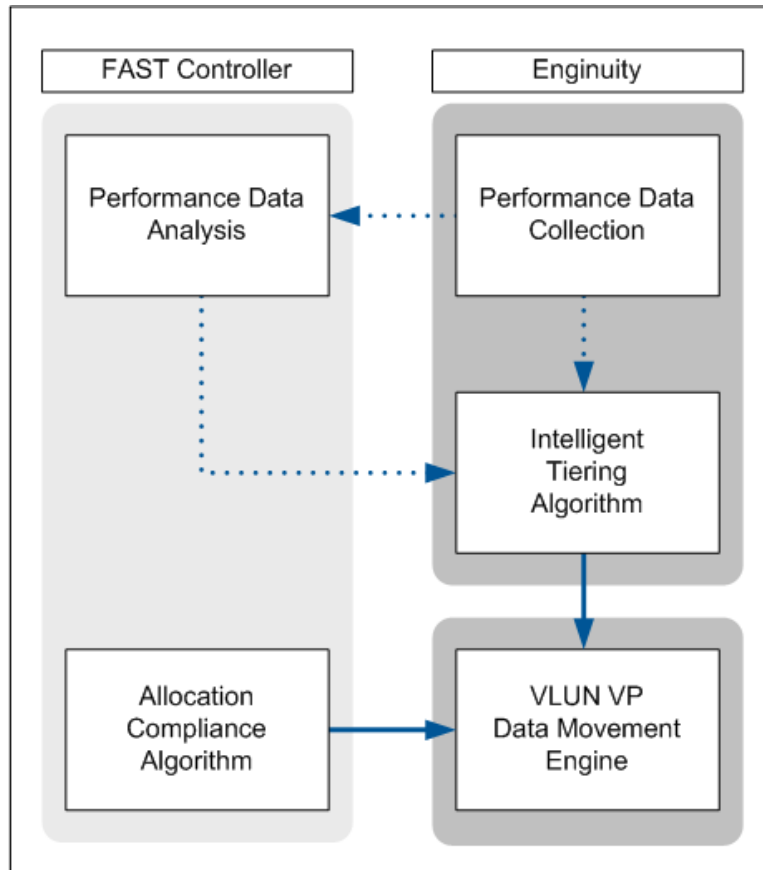


Figure 3. FAST VP components

When FAST VP is active, both components participate in the execution of two algorithms (the intelligent tiering algorithm and the allocation compliance algorithm) to determine appropriate data placement.

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 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.

Data movement windows are used to determine when to execute the data movements necessary to move data between tiers.

Data movements performed by Enginuity are achieved by moving allocated extents between tiers. The size of data movement can be as small as 12 tracks, representing a single allocated thin device extent. More typically, movements are a unit known as an extent group (10 thin device extents), which is 120 tracks in size.

Note: “FAST VP data movement” on page 31 provides more information on the actual data movement.

FAST VP has two modes of operation, Automatic or Off. When operating in the Automatic mode, data analysis and data movements occur continuously during the defined data movement windows. In the Off mode, performance statistics continue to be collected, but no data analysis or data movements take place.

FASTVP state

There are five possible states that FAST VP can be reported in. These are:

- ◆ 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.
- ◆ DisabledwithError: 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.

Note: “Appendix A: FAST VP state” on page 180 provides more information on each of the FAST controller states

FASTVP configuration

FAST VP has multiple configuration parameters that control its behavior. These include time windows that can control when performance data is collected, included in analysis, and when data movements take place. Other settings determine the relevance of historical performance data when analyzed, what percentage of space to reserve in each pool for non-FAST VP activities, and an aggressiveness factor in generating and executing data movement requests.

FASTVP time windows

FAST VP utilizes time windows to define certain behaviors regarding performance data collection and data movement. There are two possible window types:

- ◆ Performance time window
- ◆ Data movement time window

The performance time windows are used to specify when performance data should be collected by Enginuity.

The data movement time windows define when to perform the data relocations necessary to move data between tiers. Separate data movement windows can be defined for full LUN movement, performed by FAST and Optimizer, and sub-LUN data movement, performed by FAST VP.

Both performance time windows and data movement windows may be defined as inclusion or exclusion windows. An inclusion time window indicates that the action should be performed during the defined time window. An exclusion time window indicates that the action should not be performed during the defined time window.

There are two methods for defining time windows, legacy and enhanced. The legacy method uses the Symmetrix Optimizer interface for creating and managing time windows. The enhanced method uses a time window management interface.

On Symmetrix VMAX Series arrays, both the legacy and enhanced methods are supported. However, only one method can be used at a time. The legacy method is the default method, but legacy time windows may be converted to enhanced time windows.

Note: The conversion from legacy to enhanced is a one-way conversion. There is no method for converting from enhanced to legacy.

Legacy time window characteristics

Using the legacy method, both inclusive and exclusive time windows may be defined as periodic or non-periodic. Periodic windows allow a recurrence pattern to be specified as weekly or weekly-by-day. Non-periodic windows are defined to occur only once. Multiple time windows, of both performance and data movement, may be defined.

If multiple legacy time windows of the same type have time ranges that overlap one another, the most recently added time window supersedes the others, including the system default time windows.

All defined time windows apply to all devices configured within the Symmetrix array.

A Symmetrix VMAX Series storage array supports up to 128 defined legacy time windows. Each time window name may be up to 32 alphanumeric characters, hyphens (-), and underscores (_).

Enhanced time window characteristics

Using the enhanced method, inclusive windows are defined by specifying the days of the week and the times in each day during which performance metrics are collected or data is moved. Each time window is defined with a start and end time, in 30-minute increments, and the days

of the week to apply the window.

Exclusive windows are defined as a specific time period during which performance metrics collection or data movement are prevented. Each exclusive window contains a start date and time, as well as an end date and time, and may cover several days. The start and end time must be specified in 30-minute increments.

Multiple inclusive and exclusive enhanced time windows can be defined. The exclusive time windows have the highest priority and override any inclusive time windows they overlap with.

Enhanced time windows have no associated time window name.

Performance time window

The performance time windows are used to identify the business cycle for the Symmetrix array. They specify date and time ranges (past or future) when performance samples should be collected, or not collected, for the purposes of FAST VP performance analysis. The intent of defining performance time windows is to distinguish between the periods of time when the Symmetrix array is idle from the periods when it is active, and to only include performance data collected during the active periods.

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

Note: For legacy time windows, in order to prevent FAST VP from collecting performance statistics on a continuous basis, the first user-defined window should be created to exclude data collection. Inclusive time windows can then be created on top of this exclusive time window.

For enhanced time windows, the default behavior is also to collect performance data, 24 hours a day, 7 days a week. When the first inclusive performance window is created, the default behavior is overridden, and performance metrics are only collected during the defined window.

Data movement time window

Data movement time windows are used to specify date and time ranges when data movements are allowed, or not allowed, to be performed. FAST VP data movements run as low-priority tasks on the Symmetrix back end. They can introduce additional processing overhead on the back end; however, host I/ O should not be impacted. Data movement windows can be planned so they minimize any impact on the performance of other, more critical workloads.

For legacy time windows, a default data movement time window prevents any data movement, 24 hours a day, 7 days a week, 365 days a year.

For enhanced time windows, there is no default data movement window. However, until an inclusive window is defined, the default behavior leads to not allowing any data movement to occur.

FASTVP settings

There are multiple settings that affect the behavior of FAST VP. These include:

- ◆ FAST VP Data Movement Mode
- ◆ Workload Analysis Period
- ◆ Initial Analysis Period
- ◆ Pool Reserved Capacity (PRC)
- ◆ FAST VP Relocation Rate (FRR)
- ◆ VP Allocation by FAST Policy

The following sections describe each of these settings, their effect on the behavior of FAST VP, as well as possible values and the default-setting values.

FASTVP Data Movement Mode

FAST VP, when enabled, may operate in one of two modes, Automatic or Off.

In Automatic mode, a data movement request can be generated to move data based on performance workload. Also, data movement requests based on capacity utilization may be generated. These operations are performed during the periods allowed by the data movement windows.

In Off mode, no data movement requests are generated. As a result, no

data movements occur. However, performance metrics continue to be collected. The default mode is Off.

Workload Analysis Period

The Workload Analysis Period determines the degree to which FAST VP metrics are influenced by recent host activity, and also less-recent host activity, that takes place while the performance time window is considered open.

The longer the time defined in the workload analysis period, the greater the amount of weight assigned to less-recent host activity.

The Workload Analysis Period can be configured to be between 2 hours and 4 weeks. The default is 1 week (7 days).

Note: For more information on the effect of the workload analysis period, refer to “*FAST VP performance data collection*” on page 22.

Initial Analysis Period

The Initial Analysis Period defines the minimum amount of time a thin device should be under FAST VP management before any performance-related data movements should be applied. This period only accounts for time passed while the performance time window is open.

This value should be set to allow sufficient data samples for FAST VP to establish a good characterization of the typical workload on that device. This value allows FAST VP to commence analysis and movement activities on the device, prior to the full Workload Analysis Period elapsing (if so desired).

The initial analysis period can be configured to be between 2 hours and 4 weeks; however, it cannot exceed the Workload Analysis Period. The default is 8 hours.

Pool Reserved Capacity

The Pool Reserved Capacity (PRC) reserves a percentage of each pool included in a VP tier for non-FAST VP activities. The purpose of this is to

ensure that FAST VP data movements do not fill a thin pool, and subsequently cause a new extent allocation to fail, as result of a host write.

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

The PRC can be set both system-wide and for each individual pool. By default, the system-wide setting is applied to all thin pools that have been included in VP tier definitions. However, this can be overridden for each individual pool by using the pool-level setting.

The system-wide PRC can be configured to be between 1 percent and 80 percent. The default system-wide PRC is 10 percent.

The pool-level PRC can be configured to be between 1 percent and 80 percent, or set to NONE. If the PRC is set to NONE, then the system-wide setting is used. The default pool-level PRC is NONE.

FASTVP Relocation Rate

The FAST VP Relocation Rate (FRR) is a quality of service (QoS) setting for FAST VP and affects the aggressiveness of data movement requests generated by FAST VP. This aggressiveness is measured as the amount of data that is requested to be moved at any given time, and the priority given to moving the data between pools.

The FRR can be configured to be between 1 and 10, with 1 being the most aggressive. The default is 5.

<p>Note: The rate at which data is moved between pools can also be controlled by the Symmetrix Quality of VLUN setting.</p>
--

VP Allocation by FAST Policy

The VP allocation by FAST policy feature allows new allocations to come from any of the thin pools included in the FAST VP policy that the thin device is associated with.

With this feature enabled, FAST VP attempts to allocate new writes in the most appropriate tier first, based on available performance metrics. If

no performance metrics are available, an attempt is made to allocate to pool the device is bound to.

If the pool initially chosen to allocate the data is full, FAST VP then looks to other pools contained within the FAST VP policy and allocates from there. As long as there is space available in at least one pool within the policy, all new extent allocations will be successful.

The allocation by policy feature is enabled at the Symmetrix array level and applies to all allocations for all devices managed by FAST VP. The feature is either enabled or disabled. The default setting is disabled. When disabled, new allocations only come from the pool that the thin device is bound to.

Note: For more information on the decision making process of the VP allocation by FAST policy feature, refer to “*VP allocation by FAST policy*” on page 38 in the *Advanced FAST VP features* section.

FASTVP performance data collection

As previously discussed, performance data for use by FAST VP is collected and maintained by Symmetrix Enginuity. This data is then analyzed by the FAST controller and guidelines are generated for the placement of thin device data on the defined VP tiers within the array.

The following sections discuss the collection and decaying of the thin device performance metrics.

Thin device performance collection

Performance metrics collected by FAST VP are measured at both the full LUN and sub-LUN levels for all thin devices associated with a policy.

At the sub-LUN level, each thin device is broken up into multiple regions, known as extent groups and extent group sets. Each thin device is made up of multiple extent group sets which, in turn, contain multiple extent groups.




Figure 4. Thin device performance collection regions

Each extent group is made up of ten contiguous thin device extents. With each thin device extent being twelve tracks in size, an extent group represents 120 tracks of the device.

Each extent group set is made up of forty eight contiguous extent groups, representing 5,760 tracks of the device.

The metrics collected at each sub-LUN level allow FAST VP to make separate data movement requests for each extent group of the device, 120 tracks.

Emulation considerations

FAST VP supports both FBA and CKD emulations. The size of the extent groups and extent group sets are the same size for both emulations, in terms of tracks. However, the tracks sizes are different for each emulation, 64KB for FBA and 57KB for CKD.

The following table shows the relative size of the thin device extent, extent group, and extent group set for both FBA and CKD

Element	Tracks	FBA	CKD
Thin device extent	12	768KB	684KB
FAST VP extent group	120	7,680KB/ 7.5MB	6840KB/ 6.7MB

FAST VP extent group set	5,760	360MB	320.6MB
--------------------------	-------	-------	---------

Performance metrics

When collecting performance data at the LUN level and sub-LUN level for use by FAST VP, Enginuity only collects statistics related to Symmetrix back-end activity that is the result of host I/ O.

The metrics collected are:

- ◆ Read miss
- ◆ Write
- ◆ Prefetch (sequential read)

The read miss metric accounts for each DA read operation that is performed. Reads to areas of a thin device that have not had space allocated in a thin pool are not counted. Also, read hits, which are serviced from cache, are not considered.

Write operations are counted in terms of the number of distinct DA operations that are performed. The metric accounts for when a write is destaged. Write hits, to cache, are not considered.

Writes related to specific RAID protection schemes are not counted. In the case of RAID 1 protected devices, the write I/ O is only counted for one of the mirrors. In the case of RAID 5 and RAID 6 protected devices, parity writes are not counted.

Prefetch operations are accounted for in terms of the number of distinct DA operations performed to prefetch data spanning a FAST VP extent. This metric considers each DA read operation performed as a prefetch operation.

Workload related to internal copy operations, such as drive rebuilds, clone operations, VLUN migrations, or even FAST VP data movements, is not included in the FAST VP metrics.

Note: Performance metrics are only collected during user-defined performance time windows.

I/O activity rates

For each thin device, a short-term and long-term activity rate is maintained at the sub-LUN level. Both activity rates are calculated as weighted moving averages, modeling both recent and less recent workload on each region of the thin device.

As new performance metrics are collected, the previously stored short-term and long-term activity rates are decreased by the application of an exponential decay function. This lessens the influence of previously collected metrics. The newly collected metrics are then added to the decayed rates to calculate the new short-term and long-term activity rates.

The rate of decay applied to the stored I/O rates depends upon the user-defined Workload Analysis Period, with the short-term rate decayed at a faster rate than the long-term rate.

Note: The decay function is only applied to previously collected metrics while the user-defined performance time window is open.

As an example, when calculating the long-term activity rate with the default Workload Analysis Period of 168 hours, an I/O that has just been received is weighted twice as much as an I/O received 24 hours earlier.

Thin device performance score

Based on the performance statistics collected, a prioritized performance score is generated for the extent groups of each thin device. This score is calculated by combining two activity rates maintained in Engenuity (the short-term activity rate and the long-term activity rate) and multiplying that by a factor determined by the priority assigned to the storage group when it was associated with the FAST policy.

Then FAST VP places those extent groups with the highest performance score onto the higher-performing tiers within the associated policy.

Storage group association

The subLUN performance metrics previously discussed are only

collected for thin devices contained within storage groups associated with a FAST VP policy.

Disassociating a storage group from a FAST policy, or removing devices from the storage group, causes Enginuity to stop collecting performance metrics for those devices, for the purposes of FAST VP.

If the thin device is associated again to a FAST policy, the initial analysis period must pass for that device before data movements can be performed once again.

Reassociation

A storage group may be moved from one policy to another, without interruption to ongoing FAST VP management, by reassociating the storage group to a new policy.

When a reassociation is performed, all performance metrics previously collected for the devices in the storage group are maintained within Enginuity. Also, all attributes of the association, for example priority, are maintained for the association to the new policy.

Cache consumption

In order to maintain the sub-LUN level metrics collected by Enginuity, the Symmetrix allocates one cache slot for each thin device that is under FAST VP control.

When managing FBAS metadevices, cache slots are allocated for both the meta head and for each of the meta members.

Note: Each cache slot on a Symmetrix VMAX Series array is one track in size.

FASTVP performance data analysis

FAST VP uses two distinct algorithms, one performance-oriented and one capacity allocation-oriented, in order to determine the appropriate tier a device should belong to. These algorithms are:

- ◆ Intelligent tiering algorithm

- ◆ Allocation compliance algorithm

The intelligent tiering algorithm considers the performance metrics of all thin devices under FAST VP control, and determines the appropriate tier for each extent group.

The allocation compliance algorithm is used to enforce the per-tier storage capacity usage limits.

The following sections provide additional data on each of the algorithms.

Intelligent tiering algorithm

The goal of the intelligent tiering algorithm is to use the performance metrics collected at the sub-LUN level to determine which tier each extent group should reside in and to submit the needed data movements to the Virtual LUN (VLUN) VP data movement engine.

The determination of which extent groups need to be moved is performed by a task that runs within the Symmetrix array.

Algorithm structure

The intelligent tiering algorithm is structured into two components:

- ◆ A main component that executes within Enginuity
- ◆ A secondary, supporting, component that executes within the FAST controller on the service processor

The main component assesses whether extent groups need to be moved in order to optimize the use of the FAST VP storage tiers. If so, the required data movement requests are issued to the VLUN VP data movement engine.

When determining the appropriate tier for each extent group, the main component makes use of both the FAST VP metrics, previously discussed, and supporting calculations performed by the secondary component on the service processor.

Promotion-and-demotion thresholds

The supporting component of the intelligent-tiering algorithm, running on the service processor, does not make explicit data movement requests. Rather, it produces a set of promotion-and-demotion thresholds that define recommendations for data placement across thin tiers.

In a multi-tier configuration, there are up to six potential thresholds

generated:

- ◆ EFD promotion threshold
- ◆ FC promotion threshold
- ◆ SATA promotion threshold
- ◆ FC demotion threshold
- ◆ SATA demotion threshold
- ◆ FTS demotion threshold

However, as a FAST VP policy can only contain a maximum of three tiers, each extent group is only compared to up to four thresholds.

Note: There is no EFD demotion threshold, as data is never demoted to EFD. Similarly, there is no FTS promotion threshold, as data is never promoted to FTS.

The purpose of these thresholds is to allow Enginuity to determine and schedule specific data movements that are needed to optimize the allocation of thin devices under FAST VP control.

The thresholds are calculated using the performance metrics collected and the capacity available in each tier. As a result, the thresholds are dynamic and may increase or decrease as the workload on the FAST VP managed devices increases or decreases.

When calculating the thresholds, the goal is to maximize the utilization of the highest performing tier in the FAST VP policy, while also proactively demoting inactive or lightly accessed data to the most cost-effective tier.

Intelligent tiering data movement requests

The size of the data movement requests generated by the intelligent tiering algorithm depends foremost on the amount of capacity that it deems not to be on the appropriate tier. It also depends on the size of the VMAX Series array.

Each back-end disk adapter (DA) in the array is responsible for executing data movements. The more DAs there are available, the more data that can be moved at a single time, and, therefore, the larger the request size.

The FAST VP relocation rate does not affect the size of the request generated, but it does influence the pace at which the requests are executed and data moved.

Intelligent tiering algorithm execution

Operating within Enginuity, the main component of the intelligent tiering algorithm runs continuously during open data movement windows, when FAST is enabled and the FAST VP operating mode is Automatic. As a result, performance-related data movements can occur continuously during an open data movement window.

The supporting component, operating within the FAST controller on the service processor, runs every 10 minutes. It recalculates the promotion-and-demotion thresholds based on changes in workload on the FAST VP managed devices.

Allocation compliance algorithm

The goal of the allocation compliance algorithm is to detect and correct situations where the allocated capacity for a particular storage group within a thin storage tier exceeds the maximum capacity allowed by the associated FAST policy.

Policy compliance

A storage group is considered to be in compliance with its associated FAST policy when the configured capacity of the thin devices in the storage group is located on tiers defined in the policy, and when the usage of each tier is within the upper limits of the tier usage limits specified in the policy.

Compliance violations may occur for multiple reasons, including:

- ◆ New extent allocations performed for thin devices managed by FAST VP
- ◆ Changes made to the upper usage limits for a VP tier in a FAST policy
- ◆ Adding thin devices to a storage group that are themselves out of compliance

- ◆ Manual VLUN VP migrations of thin devices

Compliance data movement requests

When a compliance violation exists, the algorithm generates a data movement request to return the allocations within the required limits. This request explicitly indicates which thin device extents should be moved, and the specific thin pools they should be moved to.

The size of the data movement request depends on the amount of capacity that is currently out of compliance, but also on the user-defined relocation rate. The maximum size of request that can be generated by the allocation compliance algorithm is 10 GB worth of data movements.

When the relocation rate is set to anything other than 1, the FAST controller divides 10 GB by the relocation rate to determine the new maximum. For example, if the relocation rate is set to 2, the maximum request size is 5 GB; if it is 10, the maximum size is 1 GB.

Note: The allocation compliance algorithm is not designed to be a tool for bulk data movement. To move large amounts of data from one pool to another, for a given set of devices, VLUN VP should be used. VLUN VP allows both a source pool and a target pool to be specified to zero in on the exact data to be migrated.

Intelligent tiering algorithm coordination

The compliance algorithm attempts to minimize the amount of movements performed to correct compliance that may, in turn, generate movements performed by the intelligent tiering algorithm.

Compliance violations can be corrected by coordinating the movement requests with the analysis performed by the intelligent tiering algorithm to determine the most appropriate extents to be moved to the most appropriate tier.

Allocation compliance algorithm execution

The compliance algorithm runs every 10 minutes during open data movement windows, when FAST is enabled, and the FAST VP operating mode is Automatic.

FASTVP data movement

There are two types of data movement that can occur under FAST VP: Intelligent tiering algorithm-related movements and allocation compliance algorithm-related movements. Both types of data movement occur only during user-defined data movement windows.

Intelligent tiering algorithm-related movements are requested and executed by Enginuity. These data movements are governed by the workload on each extent group, but are executed only within the constraints of the associated FAST policy. This is to say that performance movements will not cause a storage group to become non-compliant with its FAST policy.

Allocation compliance algorithm-related movements are generated by the FAST controller, and executed by Enginuity. These movements bring the capacity of the storage group back within the boundaries specified by the associated policy. Performance information from the intelligent tiering algorithm is used to determine the more appropriate sub-extents to move when restoring compliance.

Data movement engine

Data movements executed by FAST VP are performed by the VLUN VP data movement engine. They involve moving thin device extents between thin pools within the array.

Extents are moved by way of a move process only; extents are not swapped between pools.

The movement of extents, or extent groups, does not change the thin device binding information. The thin device remains bound to the pool it was originally bound to. New allocations for the thin device, as the result of host writes, continue to come from the bound pool, unless VP allocation by FAST VP is enabled.

Movement considerations

To complete a move, the following must hold true:

- ◆ The FAST VP operating mode must be Automatic.
- ◆ The VP data movement window must be open.
- ◆ The thin device affected must not be pinned.
- ◆ There must be sufficient unallocated space in the thin pools included in the destination tier to accommodate the data being moved.
- ◆ The destination tier must contain at least one thin pool that has not exceeded the pool reserved capacity (PRC).

Note: If the selected destination tier contains only pools that have reached the PRC limit, then an alternative tier may be considered by the movement task.

Other movement considerations include:

- ◆ Only extents that are allocated are moved.
- ◆ No back-end configuration changes are performed during a FAST VP data movement, and, as a result, no configuration locks are held during the process.
- ◆ As swaps are not performed, there is no requirement for any swap space, such as DRVs, to facilitate data movement.

Data movement process

The following section details the process followed during a FAST VP data movement.

In the following illustration, device 100 is bound to an FC thin pool. The device is associated with a FAST policy that also contains tiers with an EFD pool and a SATA pool.

Over time, FAST VP has determined that several of the devices extents need to be moved to the other tiers within the policy.

As data movements can occur continuously during the thin data movement time window, a single device's allocations may change multiple times while it is actively managed by FAST VP.

Initial allocation

Figure 5 shows the thin device with its initial binding and allocation on a FC tier. The additional tiers the device is associated with as part of the FAST policy are also shown.

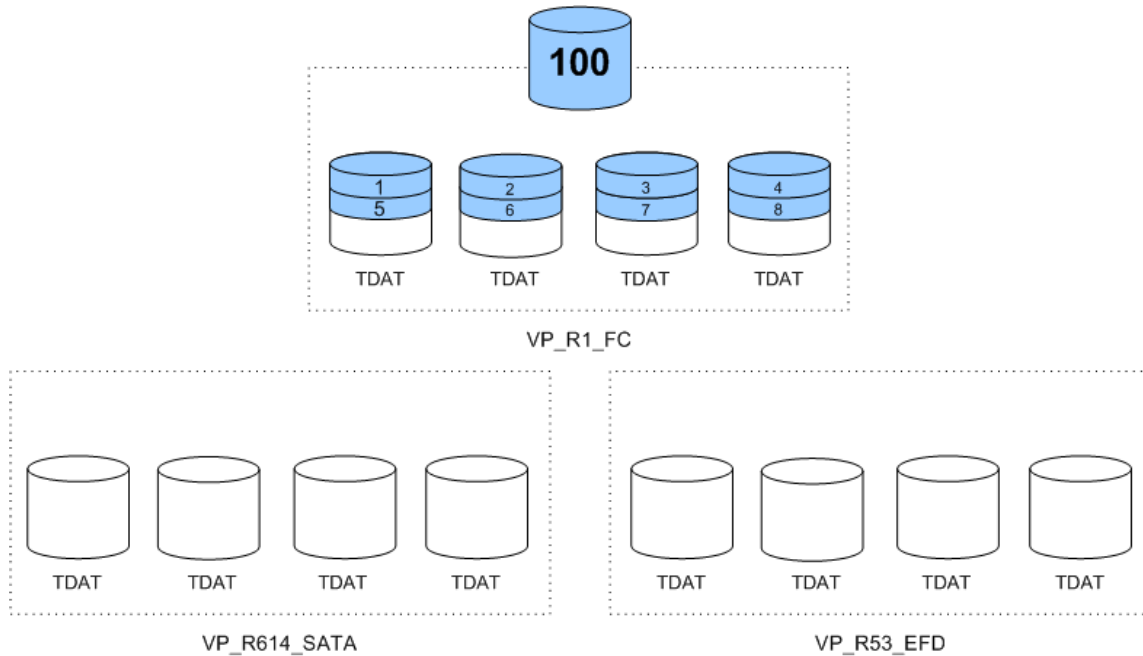


Figure 5. FAST VP data movement: Initial allocation and associated tiers

Extent movements

After the initial analysis period has passed, the performance metrics collected for device 100 are analyzed by the FAST controller. Based on this performance analysis, it is determined that several of the allocated extents need to be moved to the other tiers in the FAST policy.

Due to a higher level of activity, the extents labeled 1 and 6 are promoted to the EFD tier. Meanwhile, the extents labeled 2, 3, 7, and 8 are demoted to SATA, as they were determined to be less active.

The data movements are queued up on each DA, and the data is transferred. When the data transfer is complete, the space originally consumed by the extents in the FC tier is deallocated and is reported as

free space in that tier, as shown in Figure 6.




Figure 6. FAST VP data movement: Extent relocations

Even though data has been moved to other tiers within the array, the thin device remains bound to the pool it was originally bound to, which is contained in the FC tier.

New host writes

By default, new allocations that occur as a result of host writes come from the pool to which the thin device is bound.

Note: The default behavior of allocating new extents from the pool to which the device is bound may be changed. See “VP allocation by FAST policy” on page 38 in the *Advanced FAST VP features* section.

Figure 7 shows additional extents 9 and 10 that have been allocated in the pool.

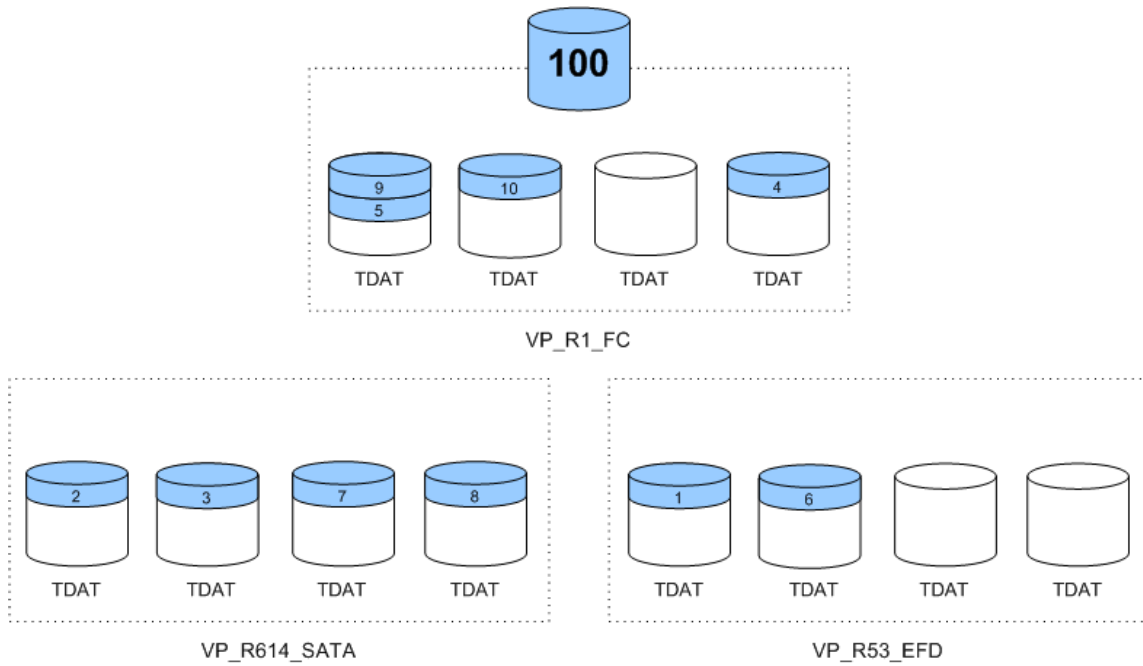


Figure 7. FAST VP data movements: New host writes

Continuous operations

Over longer periods of time, new data is generated, causing more allocations. Also, data access patterns may change, causing additional promotions and demotions to be performed.

Figure 8 shows a snapshot of device 100 and its data allocation across all three tiers.




Figure 8. FAST VP data movement: Continuous operations

Data movement control

Several mechanisms exist for controlling the movement of data related to thin devices under FAST VP control.

Device pinning

In order to prevent FAST VP changing the current tiering allocation of a thin device, a feature called device pinning may be used. Pinning a device locks all extent allocations for the device in their current locations, and prevents FAST VP from relocating them.

Any new allocations performed for a pinned device comes from the thin pool the device is bound to. These allocations are not moved by FAST VP.

The thin device needs to be unpinned in order to re-enable data movement.

Note: While a device is pinned, performance statistics continue to be collected for that device.

FAST policy disassociation

Aside from pinning a device to prevent a thin device's allocation from being changed by FAST VP, the device may be removed from the storage group that is associated to a FAST policy. Removing the device from the storage group disassociates it from the FAST policy and removes the device from the control of FAST VP. As a result, no further data movements are performed for that device. Allocated data previously moved by FAST VP remains in its current locations, and is not automatically returned to the thin device's bound pool.

Note: Disassociating a device from a FAST policy prevents performance statistics from being collected for that device, and all previously collected metrics are discarded. If the thin device is associated again to a FAST policy, the initial analysis period must pass for that device before data movements can once again be performed.

Changing the quality of service

If the performed data movements by FAST VP are impacting other applications or replication tasks within the Symmetrix array, Quality of Service (QoS) tools may be used to change the pace at which data is moved. Slowing down the FAST VP data movements gives higher priority to the other tasks running on the Symmetrix back end.

The VLUN copy pace can be set between 0 and 16 inclusively, with 0 being the fastest and 16 being the slowest. The default QoS value is 0.

Setting the pace value to 16 gives the FAST VP data movements the lowest priority on the array. Once the unrelated copy tasks have completed, the pace setting can be reset to 0.

Note: QoS can be used to change the copy pace for individual devices under FAST VP control. The FAST Relocation Rate can be used to change the copy pace for all devices under FAST VP control.

Operating mode

To stop data movements for all thin devices under FAST VP control, the operating mode may be set to Off. While off, performance statistics continue to be collected by Enginuity. However, no data movement requests are generated by Enginuity or the FAST controller.

Advanced FASTVP features

Along with the features already described, advanced features exist that further improve the ease of management and usability of FAST VP.

These advanced features include:

- ◆ VP allocation by FAST policy: Allows thin device allocations to come from thin pools other than the pool the devices are bound to.
- ◆ FAST VP SRDF® coordination: Allows promotion-and-demotion decisions of data belonging to a thin R2 device to account for workload on the corresponding thin R1 device.

The following sections describe each of these advanced features in detail.

VP allocation by FAST policy

By default, new extent allocations generated by writes to a thin device come from the thin pool the device is bound to. This behavior has two potential consequences. First, data allocated from the bound pool may be promoted or demoted shortly after the allocation that causes an additional movement of data on the back end. Second, should the bound pool fill up, a write generating a new allocation to the pool may fail.

The VP allocation by FAST policy feature addresses these potential consequences for devices managed by FAST VP. When the feature is enabled, new allocations can come from any of the thin pools included in the FAST VP policy that the thin device is associated with. (Allocations are not limited to just the bound pool.)

The allocation by policy feature is a setting that is either enabled or disabled across the entire Symmetrix array. The default setting is

disabled. When disabled, new allocations come only from the pool the thin device is bound to.

When this feature is enabled, FAST VP attempts, based on available performance metrics, to direct the allocation to the most appropriate tier. If the selected pool is full, an alternate pool is chosen from the other thin pools included within the FAST VP policy.

Note: The VP allocation by FAST policy requires a minimum of Engenuity 5876.

The following section describes the decision-making criteria for selecting the thin pool to allocate from.

Allocation by policy decision criteria

As previously mentioned, FAST VP collects performance metrics at three distinct levels for a thin device: Full LUN, extent group set, and extent group. Performance metrics already collected for each extent group set can be applied to newly allocated data within each of those sets.

When allocation by policy is enabled on a Symmetrix array, FAST VP first checks for the existence of performance metrics on the set for which the new allocation is generated.

If available, the performance metrics are used to determine the most appropriate tier, from a workload perspective, from which to allocate the extent. Once the tier has been selected, a pool within the tier is selected and the allocation is attempted.

If there are no performance metrics available for the extent group set being written to (this would happen when the entire set is unallocated), FAST VP attempts to allocate new extents from the pool to which the device is bound.

In either of these cases, should the selected pool be full, an alternate pool within the policy is chosen.

To determine the alternate pool, the tiers are sorted in ascending order of the maximum storage group percentage values within the policy. Each tier is then checked for available space within each pool in the tier, and the allocation attempted. Should the selected pool be full, FAST VP moves onto the next pool within the tier (if it exists), or the next tier on

the list.

If a selected tier contains multiple pools, the pools are sorted into ascending order of unallocated capacity within the pool.

This continues until the allocation is successful, or until a determination is made that all pools within the policy are full. As long as there is space available in at least one of the pools within the policy, all new extent allocations will be successful.

Note: In the case the device generating the allocation is pinned, the existence of performance metrics is ignored. FAST VP first attempts to allocate from the bound pool. If the bound pool is full, an alternate pool within the policy is chosen.

FAST VP SRDF coordination

FAST VP has no restrictions in its ability to manage SRDF devices. However, it must be considered that FAST VP data movements are restricted to the array upon which FAST VP is operating. By default, however, there is no coordination of data movements between the source and target arrays. FAST VP acts independently on both.

While an R1 device typically undergoes a read-and-write workload mix, the corresponding R2 device only sees a write workload. (Reads against the R1 are not propagated across the link.) A consequence of this is that the R2 device data may not be located on the same tier as the related data on the R1 device.

FAST VP SRDF coordination allows R1 performance metrics to be used when making promotion-and-demotion decisions for data belonging to an R2 device.

SRDF coordination is enabled or disabled by the storage group associated with a FAST VP policy. The default state is disabled.

Note: FAST VP SRDF coordination requires a minimum of Engenuity 5876 on both the local and remote Symmetrix arrays.

This feature can be enabled when the storage group is associated with a policy. It can also be enabled by modifying an existing association.

You only need to enable this feature on a storage group containing R1 devices. If it is enabled on a storage group containing R2 devices, the setting has no effect. The only case that this would not be true is if an SRDF swap operation occurs and converts the R2 devices to R1 devices.

FAST VP SRDF coordination is supported for single and concurrent SRDF pairings (R1 and R11 devices) in any mode of operation:

Synchronous, asynchronous, or adaptive copy.

FAST VP SRDF coordination is not supported for SRDF/ Star, SRDF/ EDP, or Cascaded SRDF, including R21 and R22 devices.

Effective performance score

When enabled on a storage group, the collected FAST VP performance metrics for R1 devices are periodically transmitted across the SRDF link to the corresponding R2 devices.

Note: Performance metrics are only sent for R1 devices when their corresponding R2 devices are being managed by FAST VP.

On the R2 device, the R1 performance metrics are merged with the actual R2 metrics. This creates an effective performance score for data on the R2 devices.

When promotion-and-demotion decisions are made on the remote array, the effective score is used for the R2 data, thereby allowing the R1 workload to influence the movement of the R2 data. Data that is heavily read on the R1 device is likely to be promoted to the higher tiers in the policy that the R2 device is associated with.

Note: The R2 device's effective performance score also influences the calculation of promotion-and-demotion thresholds on the remote array.

SRDF coordination considerations

When the SRDF link between the R1 and R2 devices is not ready, the R1 performance metrics are not transmitted to the R2 device. When the link is restored, the metrics are transmitted again.

The SRDF link is considered to be not ready when the SRDF pair state is in one of the following states:

- ◆ Split
- ◆ Suspended
- ◆ Failedover
- ◆ Partitioned

During the period that the metrics are not being sent, previously received R1 metrics are decayed in a manner similar to normally collected performance data.

Note: For more information on decaying of performance metrics, see “*I/O activity rates*” on page 25.

If an SRDF personality swap operation is performed, performance metrics are only transmitted when SRDF coordination is enabled on the storage group containing the new R1 devices.

Note: After a swap operation, performance metrics are automatically transmitted from the new R1 devices, as long as SRDF coordination had been previously enabled while the storage group contained R2 devices.

FASTVP interoperability

FAST VP is fully interoperable with all Symmetrix replication technologies: EMC SRDF, EMC TimeFinder®/ Clone, TimeFinder/ Snap, and Open Replicator. Any active replication on a Symmetrix device remains intact while data from that device is being moved. Similarly, all incremental relationships are maintained for the moved or swapped devices.

FAST VP also operates alongside Symmetrix features such as Symmetrix Optimizer, Dynamic Cache Partitioning, and Auto-provisioning Groups.

SRDF

Thin SRDF devices, R1 or R2, can be associated with a FAST policy. Extents of SRDF devices can be moved between tiers while the devices are being actively replicated, in either synchronous or asynchronous modes.

Note: For more information on using FAST VP with SRDF, see “*FAST VP SRDF coordination*” on page 40 in the *Advanced FAST VP features* section.

TimeFinder/Clone

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

TimeFinder/Snap

The source device in a TimeFinder/ Snap session can be managed by FAST VP. However, the target device (VDEV) may not be brought under FAST VP control.

TimeFinder VP Snap

The source device in a TimeFinder VP Snap session can be managed by FAST VP. Target devices may also be managed by FAST VP, however, extent allocations that are shared by multiple target devices are not moved.

Open Replicator for Symmetrix

The control device in an Open Replicator session, push or pull, can have extents moved by FAST VP.

Virtual Provisioning

Each thin device, whether under FAST VP control or not, may only be bound to a single thin pool. All host-write-generated allocations, or user-requested pre-allocations, are performed on this pool. FAST VP data movements do not change the binding information for a thin device.

It is possible to change the binding information for a thin device without changing any of the current extent allocations for the device. However,

when rebinding a device that is under FAST VP control, the thin pool the device is being re-bound to must belong to one of the VP tiers contained in the policy with which the device is associated.

Virtual Provisioning space reclamation

Space reclamation may be run against a thin device under FAST VP control. However, during the space reclamation process, the sub-LUN performance metrics are not updated, and no data movements are performed.

Note: If FAST VP is actively moving extents of a device, a request to reclaim space on that device will fail. Prior to issuing the space reclamation task, the device should first be pinned. This suspends any active FAST VP data movements for the device and allows the request to succeed.

Virtual Provisioning T10 unmap

Unmap commands can be issued to thin devices under FAST VP control.

The T10 SCSI unmap command for thin devices advises a target thin device that a range of blocks are no longer in use. If this range covers a full thin device extent, then that extent can be deallocated and the free space is returned to the pool.

If the unmap command range covers only some tracks in an extent, those tracks are marked Never Written by Host (NWBH). The extent is not deallocated. However, those tracks do not have to be retrieved from disk should a read request be performed. Instead, the Symmetrix array immediately returns all zeros.

Note: The T10 SCSI unmap command is only supported in Open Systems (FBA) environments.

Virtual Provisioning pool management

Data devices may be added to or removed from a thin pool that is included in the FAST VP tier. Data movements related to FAST VP, into

or out of the thin pool, continue while the data devices are being modified.

In the case of adding data devices to a thin pool, pool rebalancing may be run. Similarly, when disabling and removing data devices from the pool, they drain their allocated tracks to other enabled data devices in the pool.

While running either data device draining or automated pool rebalancing on a thin pool that is included in a VP tier is allowed, both processes may affect performance of FAST VP data movements.

Virtual LUN VP Mobility

A thin device under FAST VP control may be migrated using VLUN VP. Such a migration results in all allocated extents of the device being moved to a single thin pool.

While the migration is in progress, no FAST VP-related data movements are performed. Once the migration is complete, all allocated extents of the thin device will be available to be retired.

To prevent the migrated device from being retired by FAST VP immediately following the migration, it is recommended that the device first be pinned. The device can later be unpinned to re-enable FAST VP-related data movements.

FAST

Both FAST and FAST VP may coexist within a single Symmetrix array. FAST only performs full-device movements of non-thin devices. As a result, there is no impact to FAST VP's management of thin devices.

Both FAST and FAST VP share some configuration parameters. These are:

- ◆ Workload Analysis Period
- ◆ Initial Analysis Period
- ◆ Performance Time Windows

Symmetrix Optimizer

Symmetrix Optimizer operates only on non-thin devices. As such, there is no impact on FAST VP's management of thin devices.

Both Optimizer and FAST VP share some configuration parameters.

These are:

- ◆ Workload Analysis Period
- ◆ Initial Analysis Period
- ◆ Performance Time Windows

Dynamic Cache Partitioning

Dynamic Cache Partitioning (DCP) can be used to isolate storage handling of different applications. As data movements use the same cache partition as the application, movements of data on behalf of one application do not affect the performance of applications that are not sharing the same cache partition.

FBA Auto-provisioning Groups

Storage groups created for the purposes of Auto-provisioning FBA devices may also be used for FAST VP. However, while a device may be contained in multiple storage groups for the purposes of Auto-provisioning, the device may only be contained in one storage group that is associated with a FAST policy (DP or VP).

If a storage group contains a mix of device types, thin and non-thin, only the devices matching the type of FAST policy the group is associated with are managed by FAST.

If it is intended that both device types in an Auto-provisioning storage group be managed by FAST and FAST VP, then separate storage groups need to be created. A storage group with the non-thin devices may then be associated with a policy containing DP tiers. A separate storage group containing the thin devices needs to be associated with a policy containing VP tiers.

If separate storage groups are created for the purposes of applying separate FAST policies, then these groups can be added to a parent storage group, using the cascaded SG feature. A masking view can then be applied to the parent SG, provisioning both sets of devices.

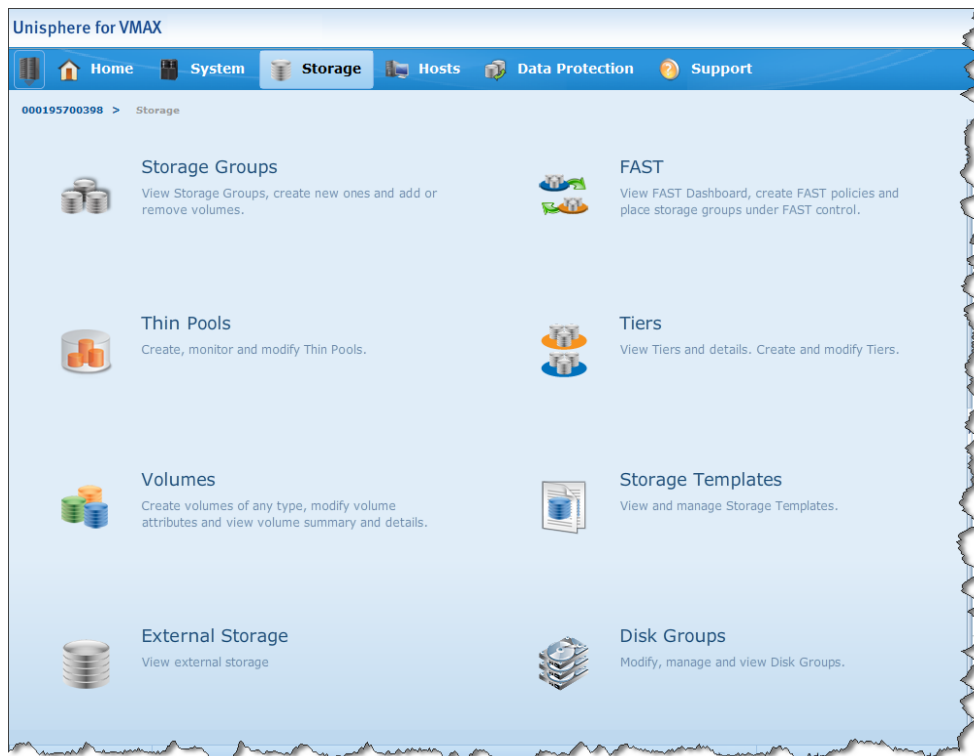
Note: FAST policies may only be associated to storage groups containing devices. A parent SG, containing other storage groups, cannot be associated to a FAST policy.

Management interface: Unisphere for VMAX

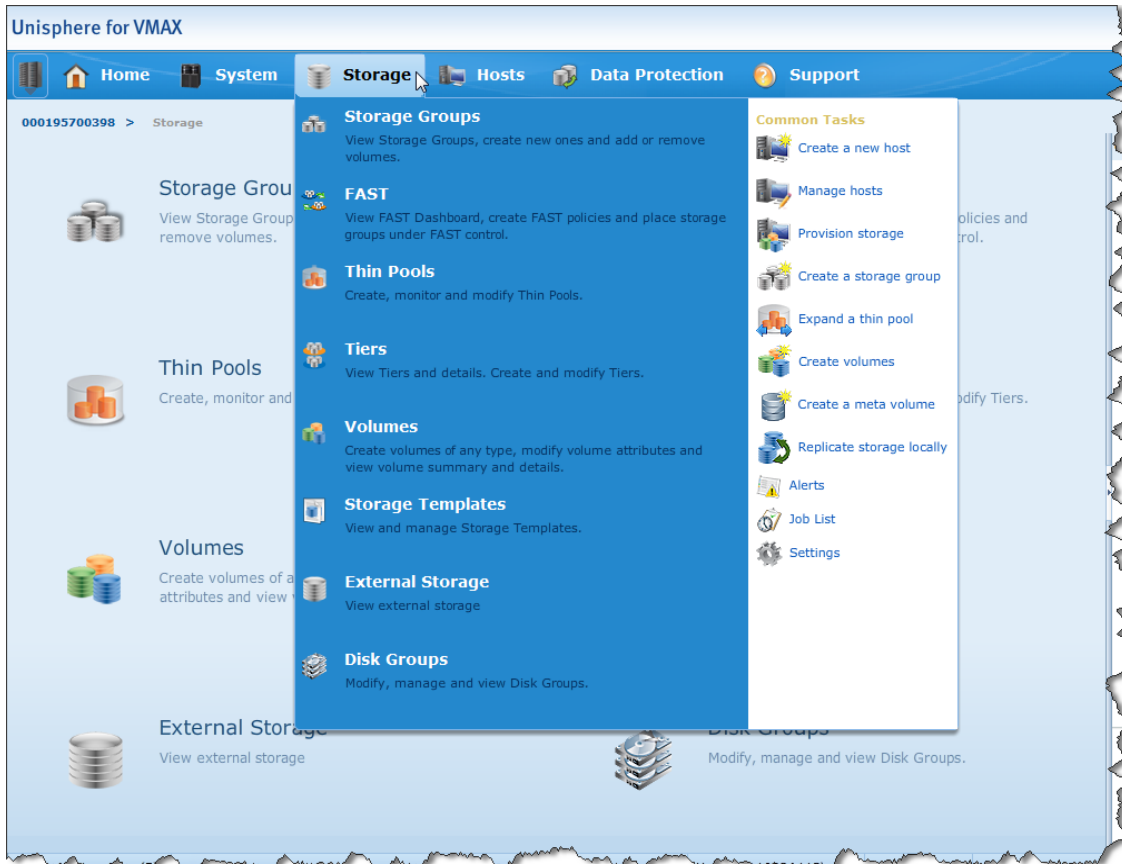
Unisphere for VMAX provides features and functionality for managing FAST VP in both Open Systems and mainframe environments. All FAST VP managed objects and associations, configuration settings, and time windows may be created and managed.

All FAST VP-related tasks are available from the Storage section page within Unisphere for VMAX. The primary subsection pages used are:

- ◆ Storage Groups
- ◆ FAST
- ◆ Thin Pools
- ◆ Tiers



These subsection pages are also available by hovering the mouse pointer over the Storage item in the task bar.



The following sections detail the use of Unisphere for managing and controlling FAST VP.

Examining Symmetrix Virtual Provisioning thin pools

Prior to enabling and configuring the FAST controller, and configuring Symmetrix tiers and FAST policies, it is important to understand the existing configuration of the Symmetrix array. Unisphere can be used to take an inventory of the existing thin pools within the array.

Details on the existing available pools can be seen on the Thin Pools subsection page under Storage.

Unisphere for VMAX

Home System Storage Hosts Data Protection Support

000195700398 > Storage > Thin Pools

Thin Pools

Name	Technology	Disk Location	Configuration	Emulation	Allocated Capacity	Capacity (GB)
R0_FTS_Pool	N/A	External	Unprotected	FBA	0 %	195.31
R1_FC_Pool	FC	Internal	2-Way Mir	FBA	9 %	10722.66
R53_EFD_Pool	EFD	Internal	RAID-5 (3 + 1)	FBA	3 %	4394.51
R57_FC_Pool	FC	Internal	RAID-5 (7 + 1)	FBA	0 %	10722.66
R6_SATA_Pool	SATA	Internal	RAID-6 (14 + 2)	FBA	4 %	24538.88
small_pool	FC	Internal	2-Way Mir	FBA	0 %	268.07

Allocated Capacity Free Capacity

Create Expand View Details Delete

The display for each thin pool shows the technology type, location, RAID protection, and emulation. It also shows the usable capacity of the pool, as well as the percentage allocated for the pool.

Note: For use with FAST VP, the technology of a thin pool may not be Mixed.

To see more information on an individual thin pool, select the pool and click View Details.

Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Thin Pools

Thin Pools

Name	Technology	Disk Location	Configuration	Emulation	Allocated Capacity	Capacity (GB)
R0_FTS_Pool	N/A	External	Unprotected	FBA	0 %	195.31
R1_FC_Pool	FC	Internal	2-Way Mir	FBA	9 %	10722.66
R53_EFD_Pool	EFD	Internal	RAID-5 (3 + 1)	FBA	3 %	4394.51
R57_FC_Pool	FC	Internal	RAID-5 (7 + 1)	FBA	0 %	10722.66
R6_SATA_Pool	SATA	Internal	RAID-6 (14 + 2)	FBA	4 %	24538.88
small_pool	FC	Internal	2-Way Mir	FBA	0 %	268.07

Allocated Capacity ■ Free Capacity ■

Create Expand **View Details** Delete

Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Thin Pools > R1_FC_Pool

Details : Thin Pool : R1_FC_Pool

Properties

Name	R1_FC_Pool
RAID Protection	2-Way Mir
Type	Thin
Technology	FC
Emulation	FBA
Total Capacity (GB)	10722.66
Free Capacity (GB)	9757.92
Thin Volumes	10
Enabled Volumes	160

Related Objects

Contains : [DATA Volumes - 160](#)

Associated With : [Bound Volumes - 10](#)

Create Expand Bind >> Apply Cancel

Additional information on the data devices contained in the pool, or the thin devices bound to the pool, can be viewed by clicking the relevant related objects link to the right of the detailed pool information.

Configuring the Symmetrix FAST controller

There are multiple FAST VP settings and parameters that affect the behavior of the FAST controller. These include:

- ◆ FAST VP Data Movement Mode
- ◆ Pool Reserved Capacity (PRC)
- ◆ FAST VP Relocation Rate
- ◆ VP Allocation by FAST Policy
- ◆ Workload Analysis Period
- ◆ Initial Analysis Period
- ◆ Performance Time Window
- ◆ Data Movement Time Window

The following sections detail the Unisphere for VMAX interface to display and modify each of these settings.

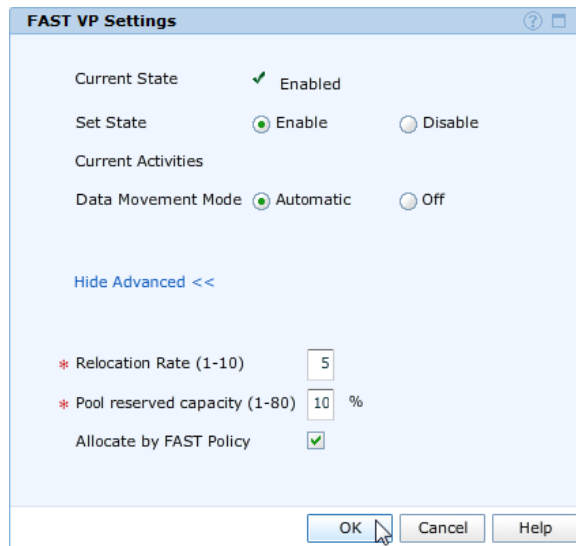
FAST controller settings list information

To view the existing FAST controller settings, go to the FAST subsection page under Storage and click Edit to the right of Settings.

The screenshot displays the Unisphere for VMAX interface. The top navigation bar includes Home, System, Storage, Hosts, Data Protection, and Support. The breadcrumb path is 000195700398 > Storage > FAST. The FAST Type is set to FAST VP. The main content area is divided into four panels:

- FAST Status Report:** A table showing settings for FAST VP. The State is Enabled (with a green checkmark and an Edit icon). Data Movement Mode is Automatic. Current Activities is Idle. Under Time Windows, Performance Time Window is Open (with a green checkmark and an Edit icon), and Move Time Window is Open (with a green checkmark and an Edit icon).
- FAST Policies:** A section titled "Manage Policies" with a table. The table has columns for Policy Name, Tier 1, Tier 2, Tier 3, and Storage Group. The table is currently empty.
- Tiers Demand Report:** A section with a table that is currently empty.
- Storage Groups under FAST control:** A section with a table. The table has columns for Storage Group Name, FAST Policy, Capacity Used Breach, and Compliant. The table is currently empty.

The resulting dialog box allows both viewing and editing of some of the FAST VP-related settings.



Setting the FASTVP Data Movement Mode

There are two possible values for the Data Movement Mode: Automatic and Off. The default value is Off.

To change the Data Movement Mode, click the appropriate radio button and then click OK.

Setting the system-wide FASTVP Pool Reserved Capacity

The FAST VP Pool Reserved Capacity (PRC) is specified as a percentage and can be set to be between 1 and 80. The default value is 10.

To set the system-wide PRC, enter a number between 1 and 80 in the text field to the right of Pool Reserved Capacity and click OK.

Setting the FASTVP Relocation Rate

The FAST VP Relocation Rate can be set to be between 1 and 10. The default value is 5.

To set the Relocation Rate, enter a number between 1 and 10 in the text field to the right of Relocation Rate and click OK.

Setting VP allocation by FAST policy

VP allocation by FAST policy can be set to enabled or disabled. The default value is disabled.

To enable VP allocation by FAST policy, select the checkbox to the right of Allocation by FAST policy and click OK.

To disable VP allocation by FAST policy, clear the checkbox to the right of Allocation by FAST policy and click OK.

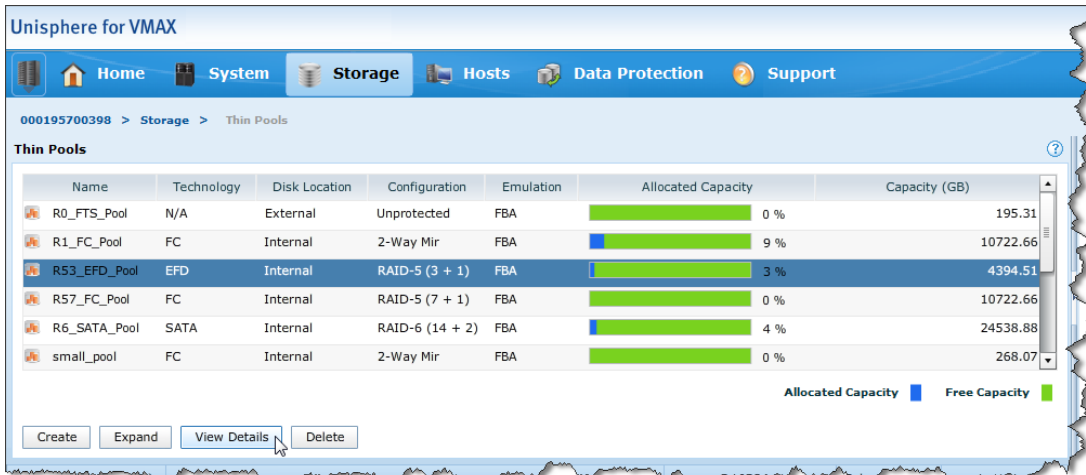
<p>Note: All of the above parameters may be edited prior to clicking the OK button.</p>
--

Setting the pool-level FASTVP Pool Reserved Capacity

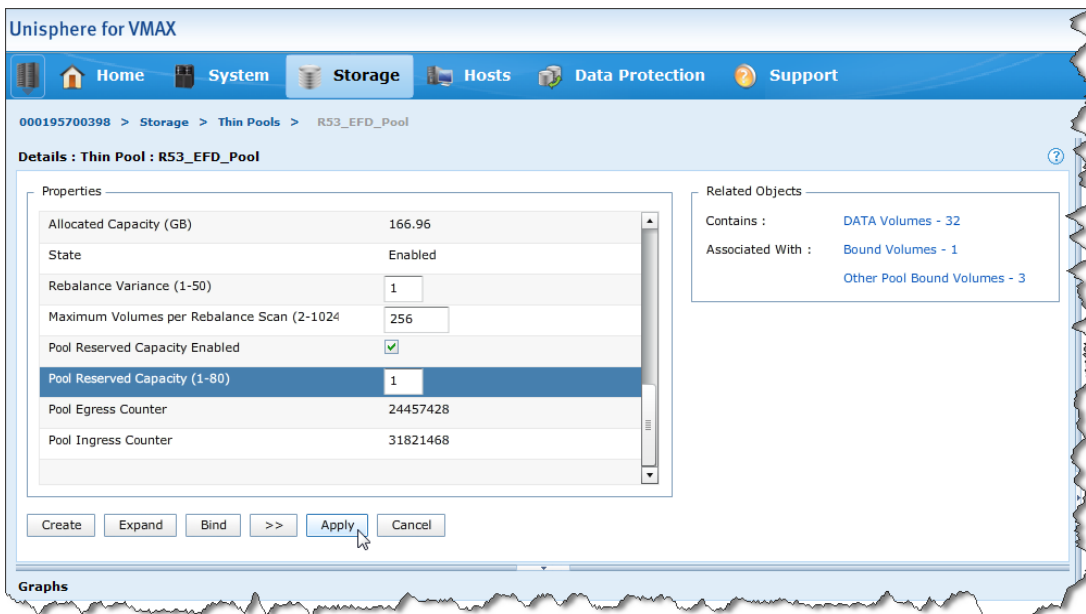
The pool-level FAST VP Pool Reserved Capacity (PRC) can be used to override the system-wide setting for each individual pool. At the pool-level, the PRC can be set between 1 and 80 percent, or NONE. The default value is NONE.

The value of NONE indicates that the system-wide setting should be used for the pool.

To set or change the PRC at the pool level, select the appropriate thin pool on the Thin Pools subsection page and click the View Details button.



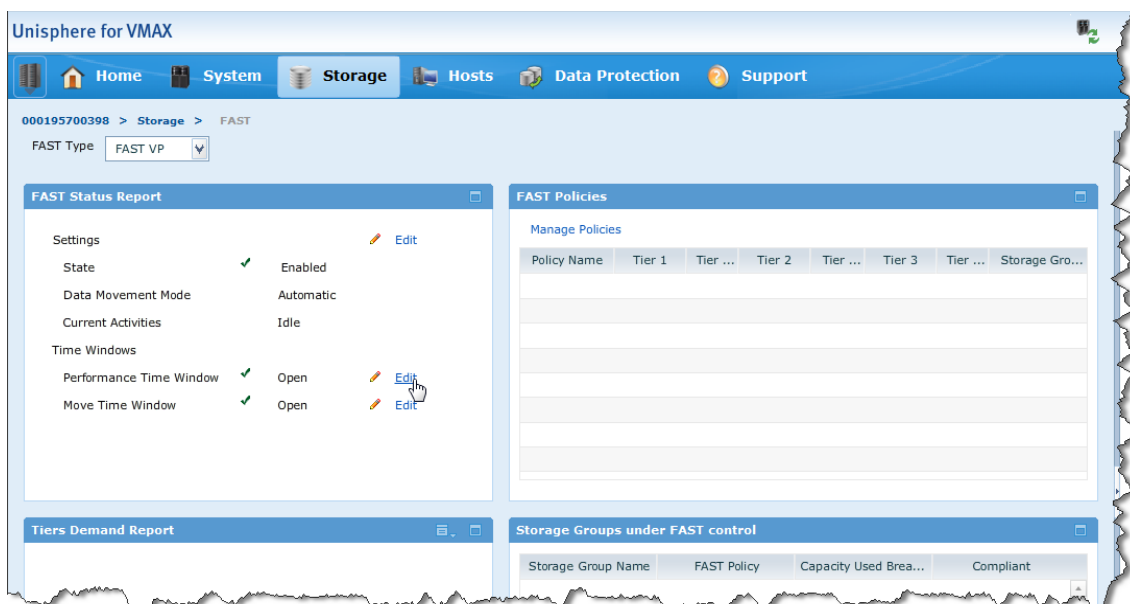
On the resulting object details page, enter the desired PRC value, between 1 and 80, in text field to the right of Pool Reserved Capacity and click Apply.



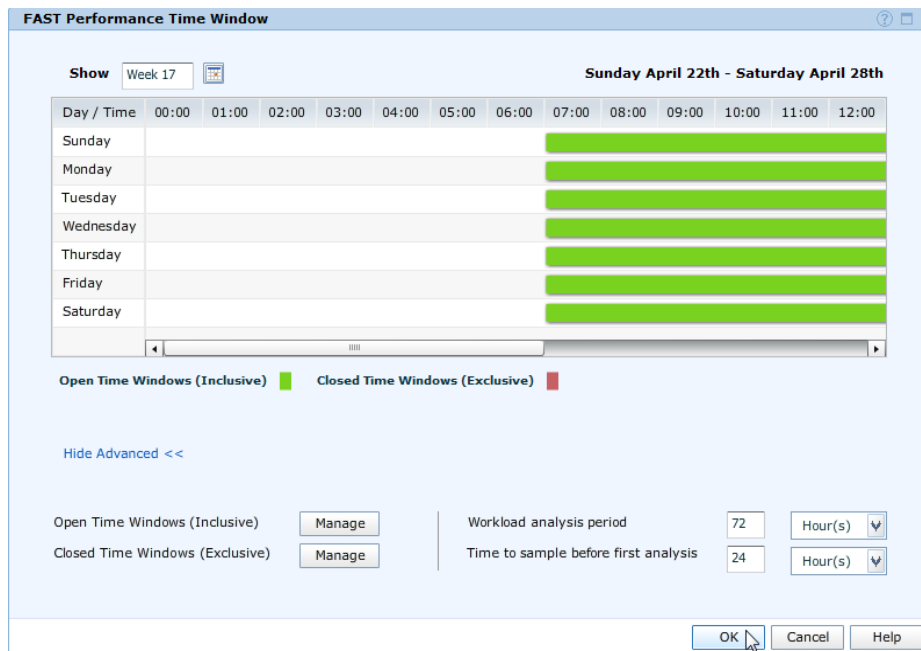
Managing Analysis Periods

Both the Workload Analysis Period and the Initial Analysis Period can be viewed and edited by either the Performance Time Window or data movement window management screens.

From the FAST subsection page under Storage, click the Edit link to the right of either Performance Time Window or Move Time Window.



In the resulting dialog box, click Show Advanced under the time window display. The analysis period parameters are available in the lower right-hand corner.



Setting the Workload Analysis Period

The Workload Analysis Period is shared with Optimizer and FAST. It can be set between 2 hours and 4 weeks.

To set the Workload Analysis Period, enter the desired value in the text field to the right of Workload Analysis Period, select the time unit (hours, days, or weeks), then click OK.

Setting the Initial Analysis Period

The Initial Analysis Period, also referred to as the Time to sample before first analysis, is a shared parameter with Optimizer and FAST. It can be set between 2 hours and 4 weeks. The default value is 8 hours.

To set the Initial Analysis Period, enter the desired value in the text field to the right of Time to sample before first analysis, select the time unit (hours, days, or weeks), then click OK.

Legacy time window management

Management of the legacy time windows is not available with Unisphere for VMAX; Unisphere for VMAX can only manage the enhanced format. Management of the legacy time windows is available using the SYMCLI `symoptmz` command.

For the management of FAST VP, it is recommended to use enhanced time window management.

Note: For information on managing legacy time windows, refer to the *Managing Legacy Time Windows* appendix in the *EMC Solutions Enabler Symmetrix Array Controls CLI* product guide available at <http://powerlink.emc.com>.

Converting time window type from legacy to enhanced

Legacy time windows cannot be converted to the enhanced format using Unisphere for VMAX. To perform the conversion, Solutions Enabler must be used.

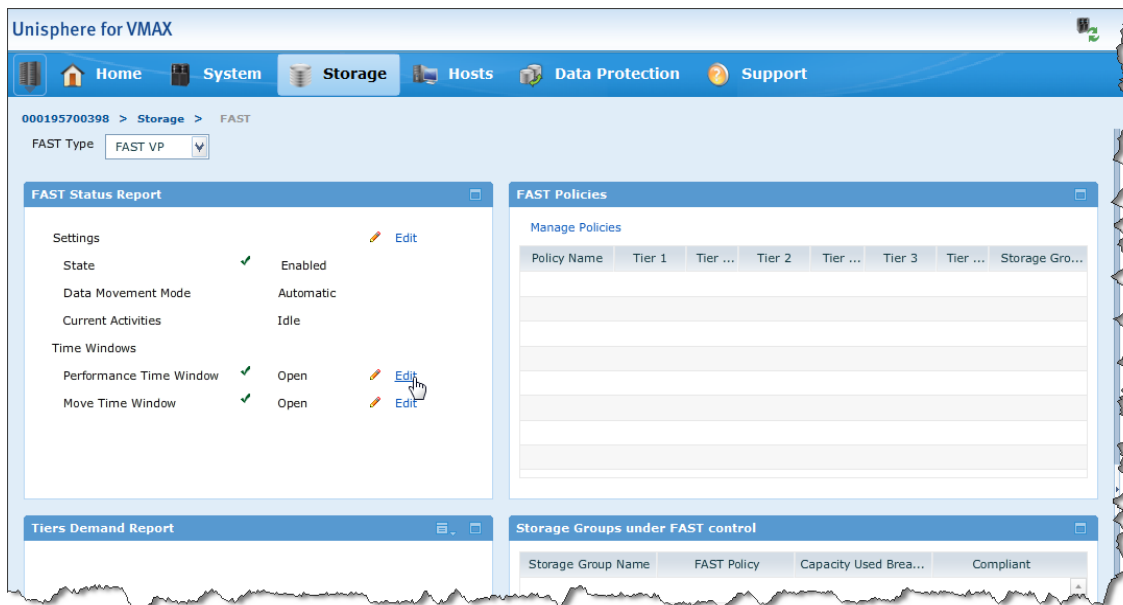
Note: For information on performing this conversion, see “*Converting time window type from legacy to enhanced*” on page 121 in the SYMCLI management interface section

Enhanced time window management

Management of the enhanced time windows is managed through the FAST subsection page of Unisphere.

Managing enhanced FASTVP performance time windows

To manage the performance time windows to be used by FAST VP, click the Edit link to the right of Performance Time Window in the FAST Status Report panel of the FAST subsection page.



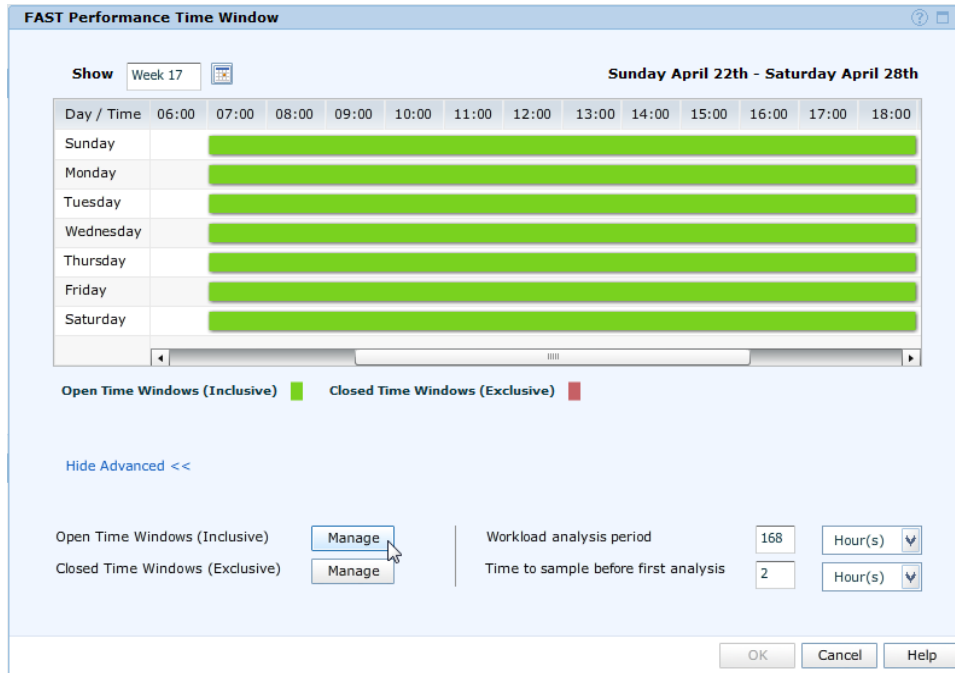
The resulting dialog box shows a graphical summary of the currently defined performance windows: A green bar indicates an inclusive time window has been defined; a red bar indicates when an exclusive time window has been defined; and white space where no time window exists.

Note: The time window is reported as closed for any exclusive time window definition and also where no explicit time window exists.

To manage the inclusive and exclusive time windows, click the Show Advanced link in the lower left-most corner of the dialog box.

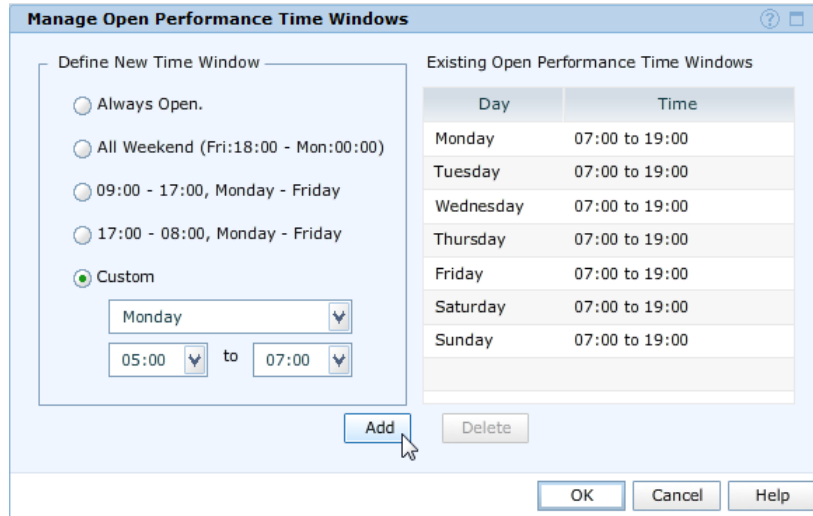
Editing inclusive performance time windows

To edit inclusive time windows, click the Manage button to the right of Open Time Windows (Inclusive).

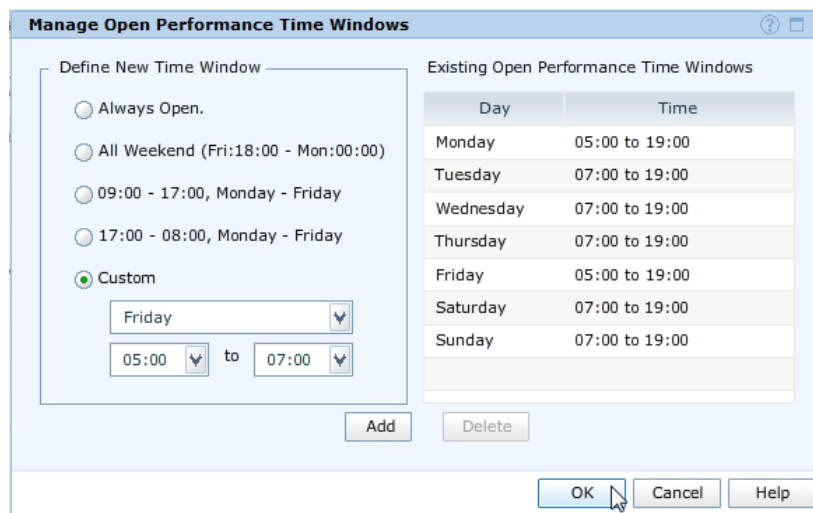


The resulting dialog box shows a text summary of the existing inclusive windows for each day of the week. It also provides several defaults for defining new windows and a custom option.

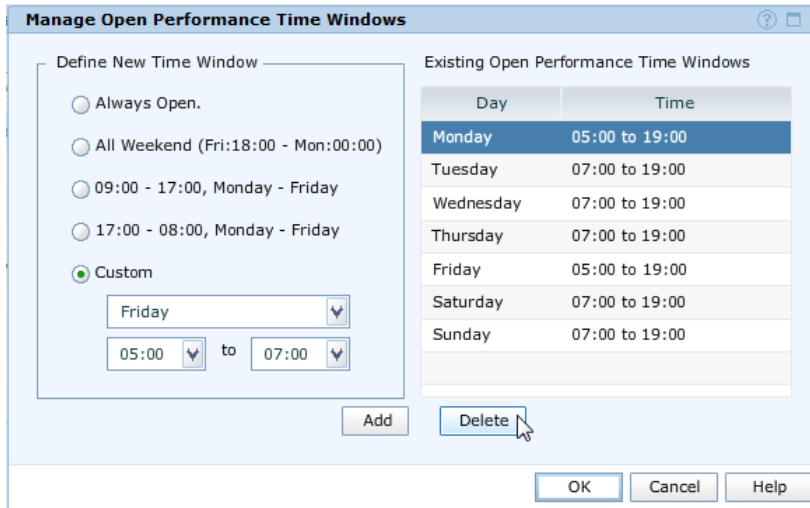
Select the desired option by choosing the appropriate radio button. If Custom is selected, choose the appropriate day of the week from the drop-down menu, and then the desired start time and end time. To create the window, click Add.



When all edits have been made, click OK to commit the changes.

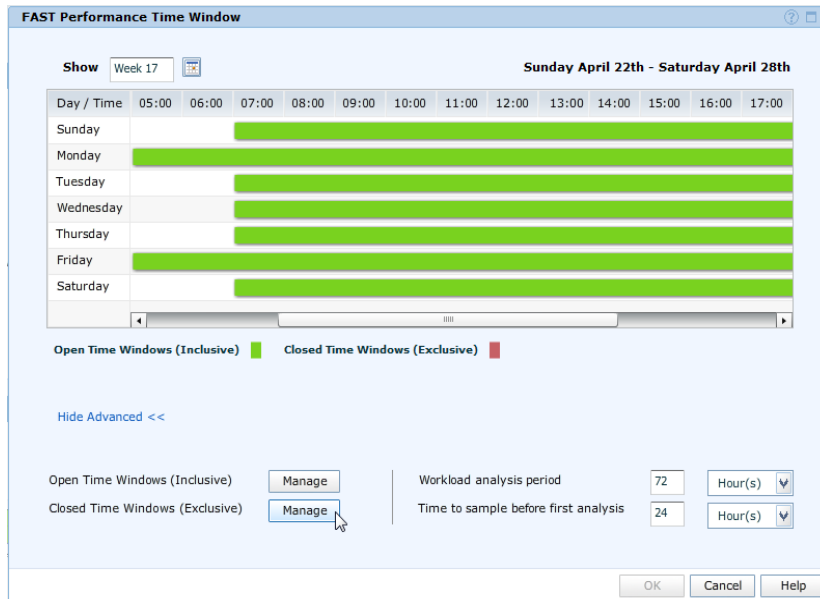


To delete an existing window, select the appropriate entry on the right-most side of the dialog box, and then click Delete.



Editing exclusive performance time windows

To edit exclusive time windows, click the Manage button to the right of Closed Time Windows (Exclusive).



The resulting dialog box shows a text summary of the existing exclusive windows.

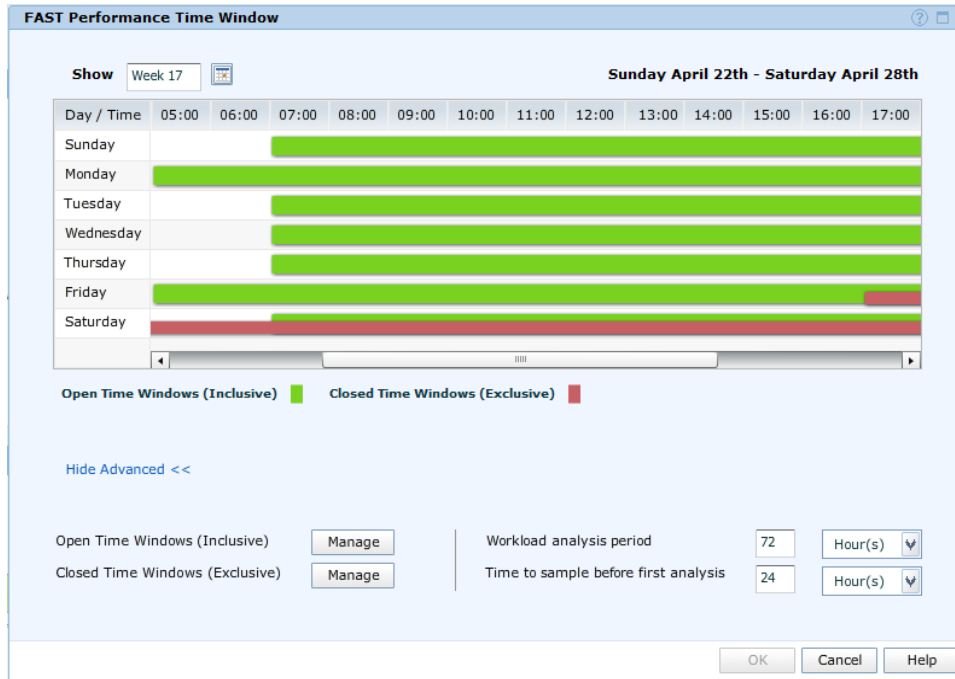
Select the desired start date and time, and end date and time, using the available drop-down menus and calendar icons. To create the window, click Add.

When all exclusive windows have been created, click OK.

Start Time	End Time
Apr-27-2012 17:00	Apr-30-2012 05:00

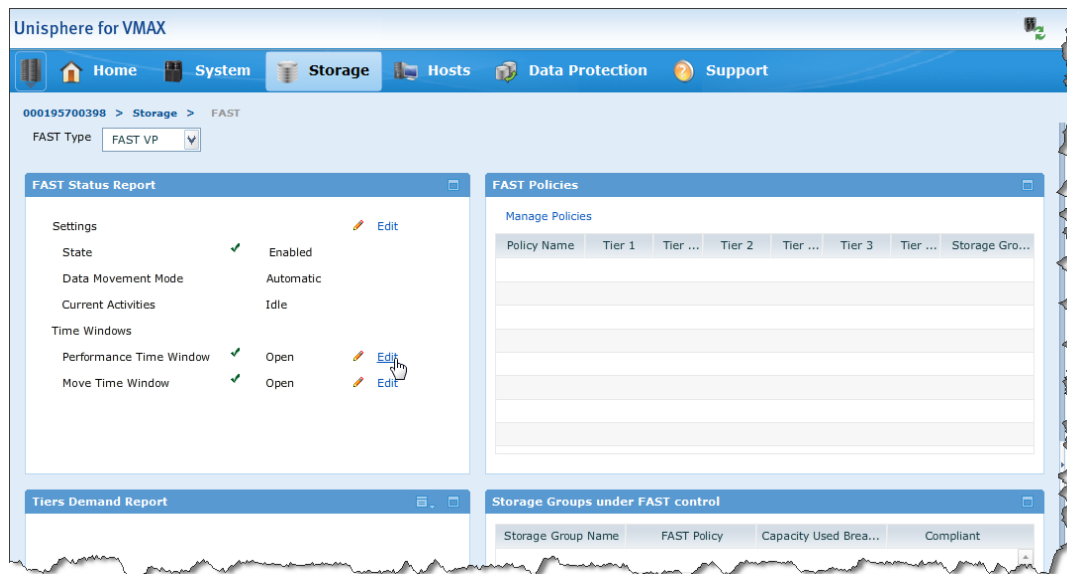
Performance Time Window summary

After all windows have been edited, the performance time window display looks like this.



Managing enhanced FASTVP data movement time windows

To manage the data movement time windows used by FAST VP, click the Edit link to the right of Move Time Window in the FAST Status Report panel of the FAST subsection page.



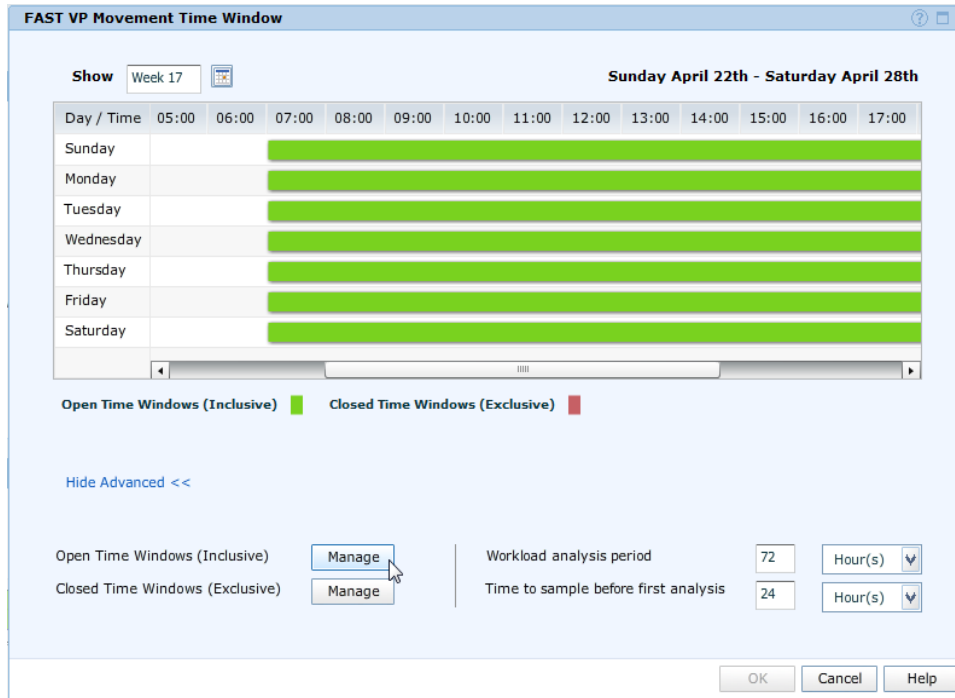
The resulting dialog box shows a graphical summary of the currently defined performance window: A green bar indicates an inclusive time window has been defined; a red bar indicates when an exclusive time window has been defined; and white space where no time window exists.

Note: The time window is reported as closed for any exclusive time window definition, and also where no explicit time window exists.

To manage the inclusive and exclusive time windows, click the Show Advanced link in the lower left-most corner of the dialog box.

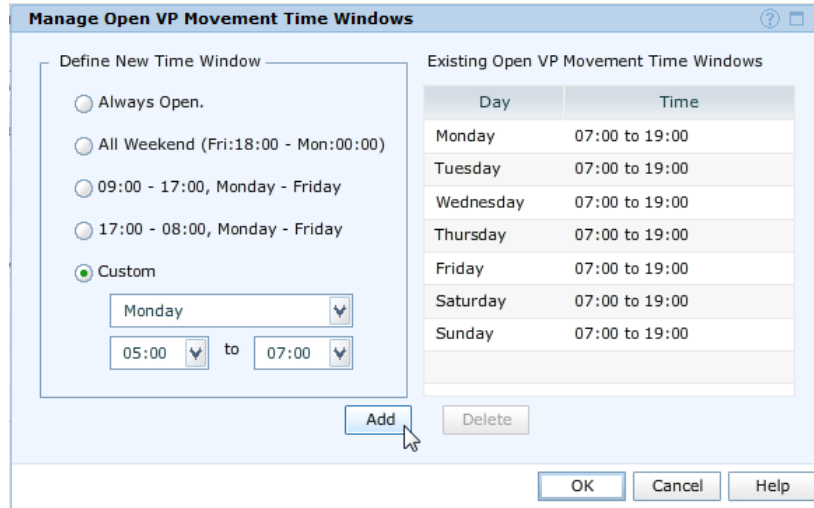
Editing inclusive move time windows

To edit inclusive move time windows, click the Manage button to the right of Open Time Windows (Inclusive).

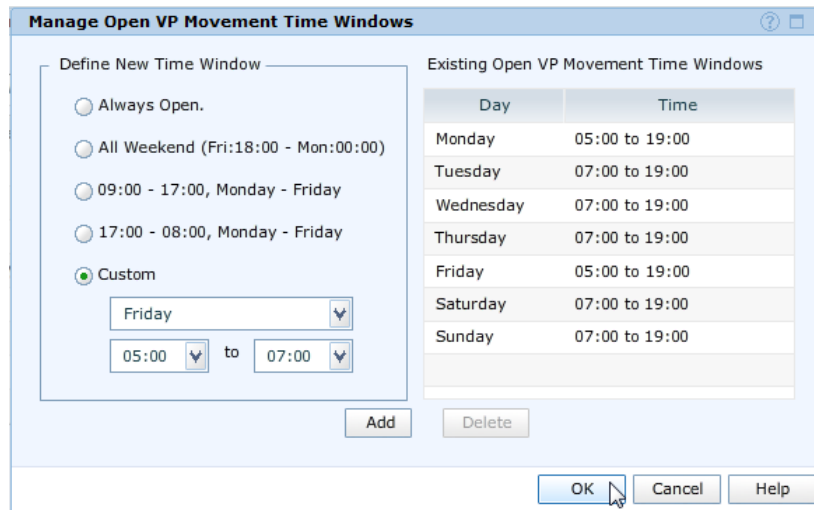


The resulting dialog box shows a text summary of the existing inclusive windows for each day of the week. It also provides several defaults for defining new windows, as well as a custom option.

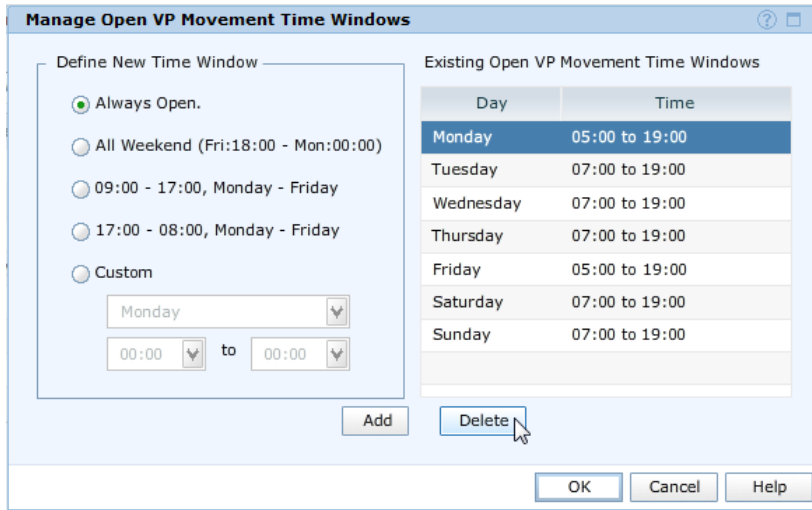
Choose the desired option by selecting the appropriate radio button. If Custom is selected, choose the appropriate day of the week from the drop-down menu, and then the desired start time and end time. To create the window, click Add.



When all edits have been made, click OK to commit the changes.

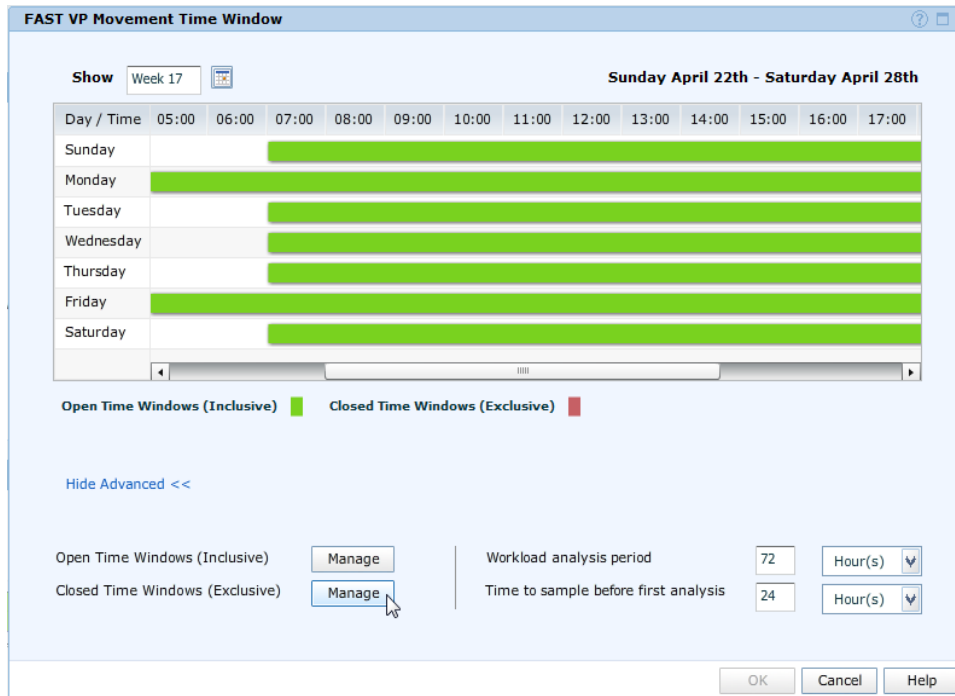


To delete an existing window, select the appropriate entry on the right-most side of the dialog box, and then click Delete.



Editing exclusive movement time windows

To edit exclusive time windows, click the Manage button to the right of Closed Time Windows (Exclusive).



The resulting dialog box shows a text summary of the existing exclusive windows.

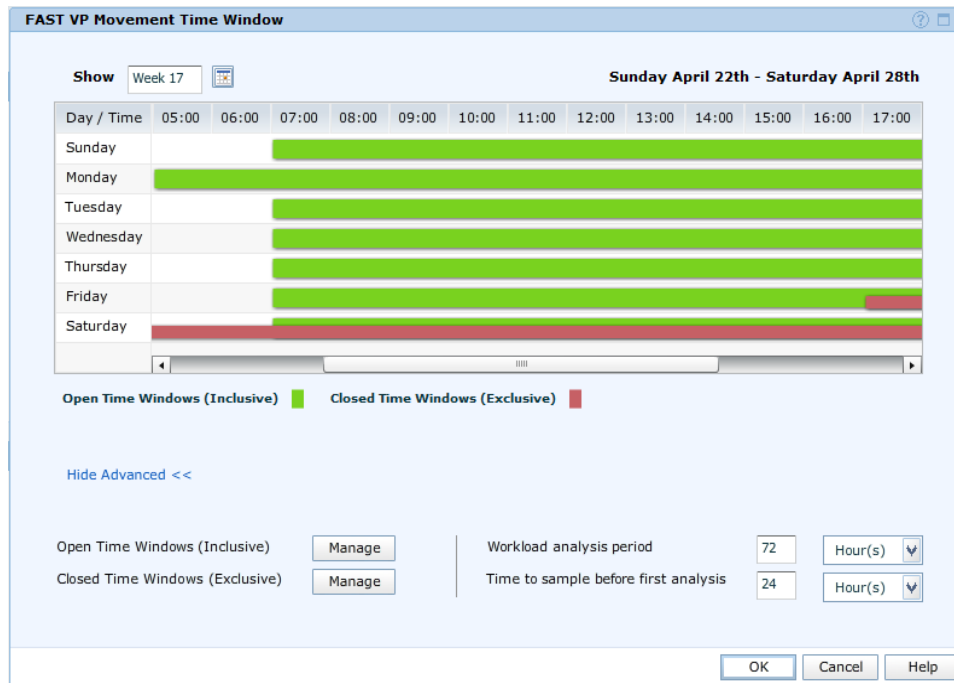
Select the desired start date and time, and end date and time, using the available drop-down menus and calendar icons. To create the window, click Add.

When all exclusive windows have been created, click OK.

Start Time	End Time
Apr-27-2012 17:00	Apr-30-2012 05:00

Move Time Window summary

After all windows have been edited, the move time window display looks like this.



Creating FASTmanaged objects

There are three managed objects related to the use of FAST VP in the Symmetrix VMAX Series array. These are:

- ◆ Symmetrix VP tiers
- ◆ FAST policies
- ◆ Storage groups

When created, storage groups can be associated with a FAST policy, which in turn associates the devices in the storage group with up to three VP tiers. The FAST policy also defines the upper usage limit for the storage group in each tier.

The following sections detail the Unisphere interfaces used to create each of the managed objects, and the methods for associating them. Information is also shown for removing these associations, and removing each of the objects.

Creating a Symmetrix VP tier

A Symmetrix VP tier may contain between one and four Virtual

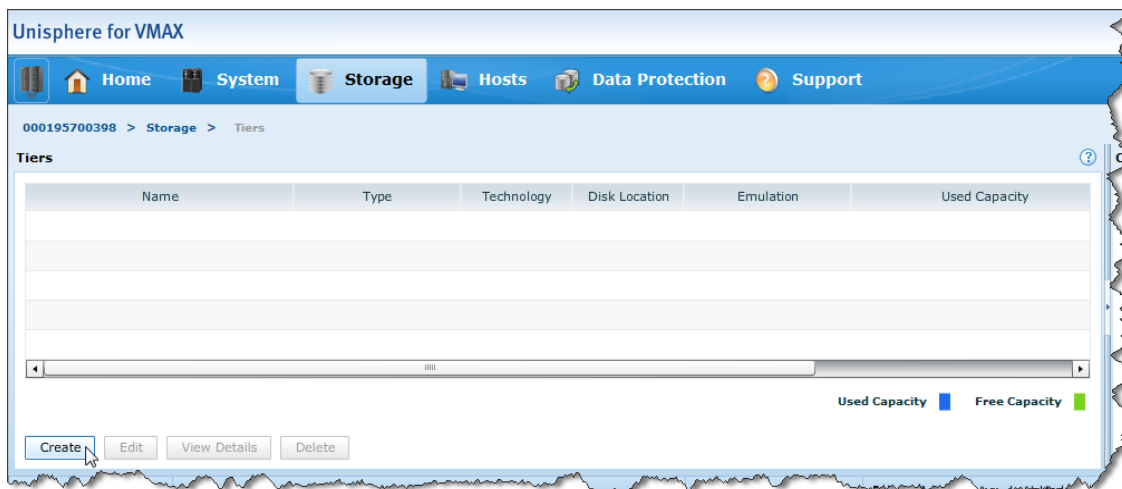
Provisioning thin pools. When creating a VP tier the following information must be known:

- ◆ The tier name
- ◆ The desired protection type of the tier
- ◆ The drive technology, or location, to be used for the tier
- ◆ The thin pool(s) to be added to the tier

Once this information has been decided, the tier can be created.

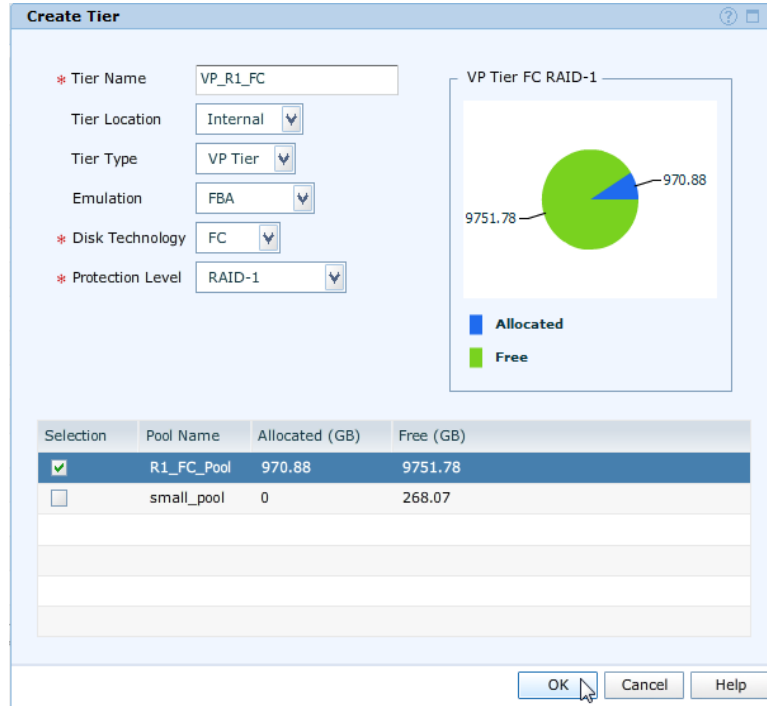
Creating an internal tier

To create an internal tier, go to the Tiers subsection page under Storage and click Create.




In the resulting dialog box, fill out the fields as desired, however, make sure to specify the location as internal and the tier type as VP Tier.

The list of available pools is filtered as the emulation, disk technology, and protection level is specified.

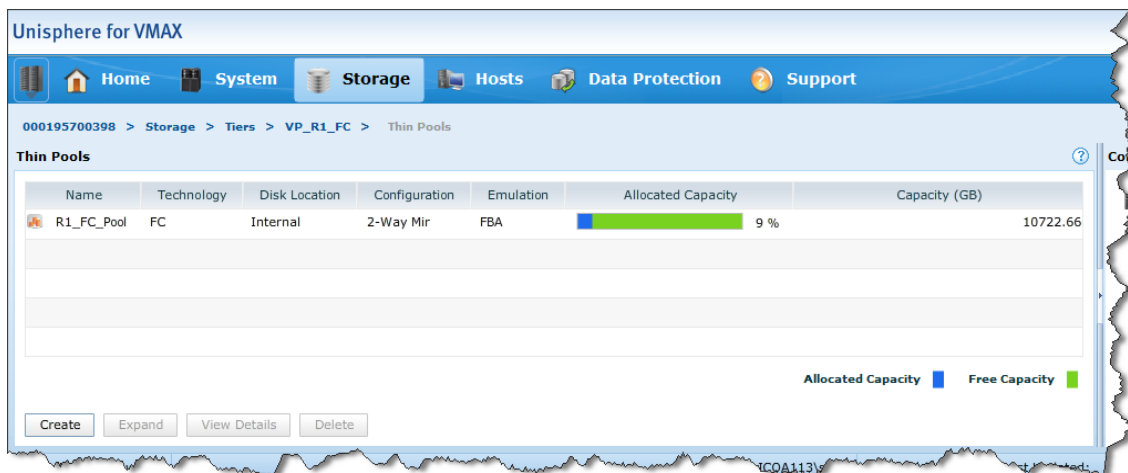


After selecting the pool, or pools, to be added to the tier, click OK to create the tier on the array. In this case, a tier containing the R1_FC_Pool pool, containing RAID 1 data devices configured on FC drives, is created. The Symmetrix tier name was chosen to indicate the RAID protection type (RAID 1), the drive type (FC), and the fact that it is a VP tier—VP_R1_FC.

Once created, the information on the Symmetrix tier can be seen on the Tiers subsection page.



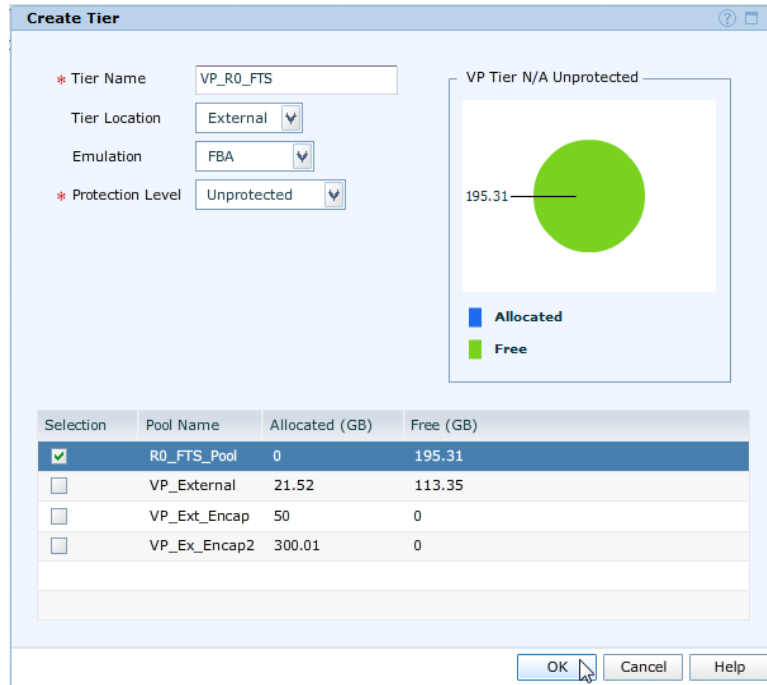
The thin pool information for the tier can be viewed by double-clicking the tier name, then clicking the Thin Pools related object link.



Creating an external tier

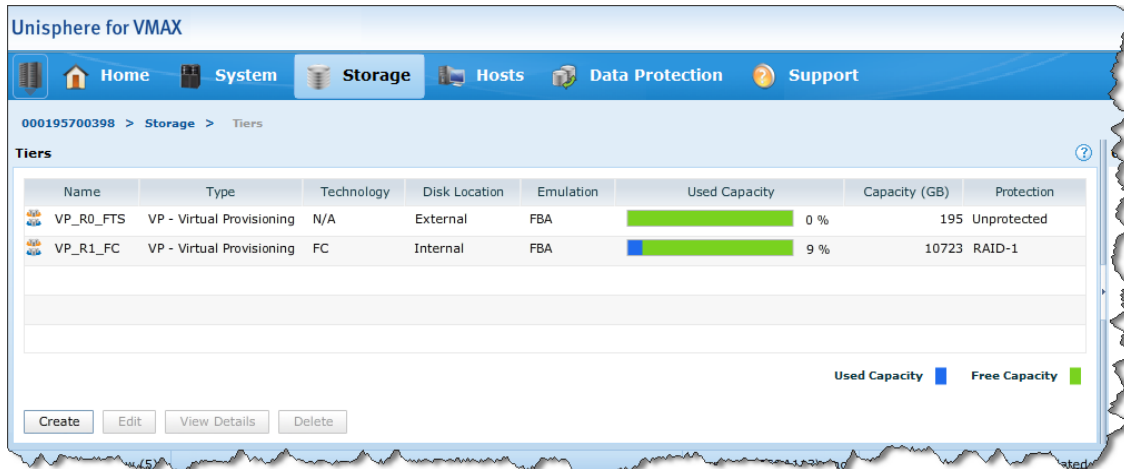
To create an external tier, click the Create button on the Tiers subsection page under Storage.

In the resulting dialog box, complete the fields as desired, ensuring to specify the location as external. The list of available external pools is displayed.

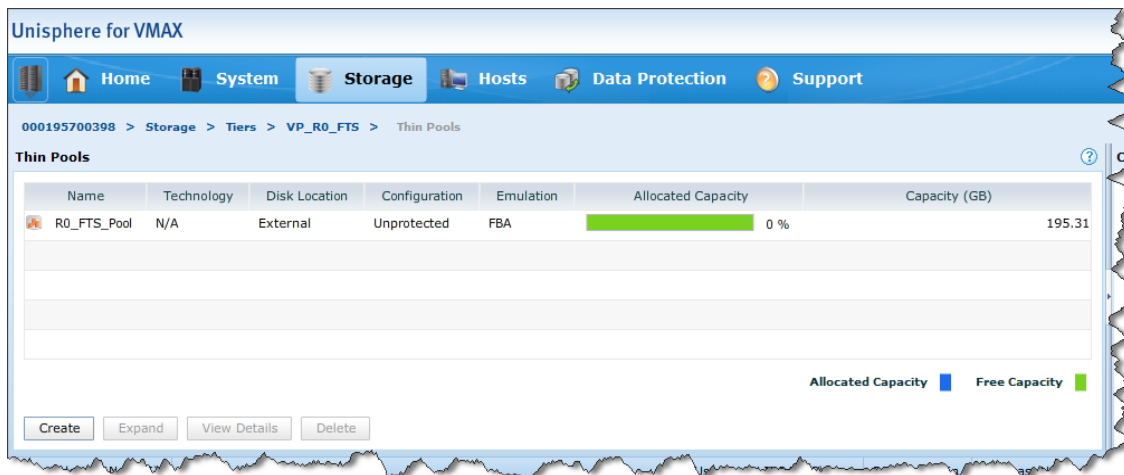


After selecting the pool, or pools, to be added to the tier, click OK to create the tier on the array. In this case, a tier containing the R0_FTS_Pool pool, containing locally unprotected data devices configured on an external array, is created. The Symmetrix tier name was chosen to indicate the RAID protection type (RAID 0), the location (FTS), and the fact that it is a VP tier—VP_R0_FTS.

Once created, the information on the Symmetrix tier can be seen on the Tiers subsection page.

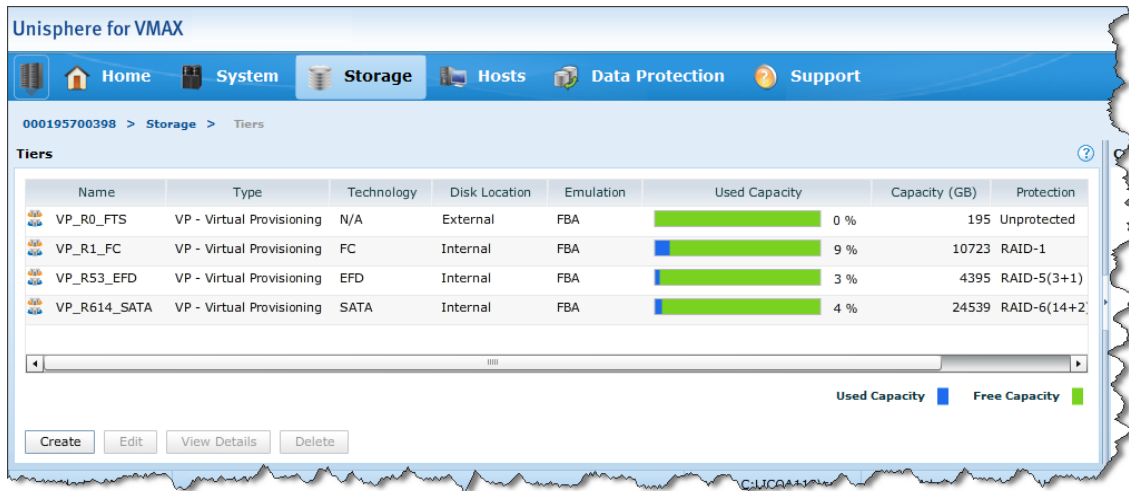


The thin pool information for the tier can be viewed by double-clicking the tier name, then clicking the Thin Pools related object link.



Symmetrix tier list information

After additional VP tiers have been created, information on all the VP tiers in the Symmetrix array can be viewed on the Tiers subsection page.



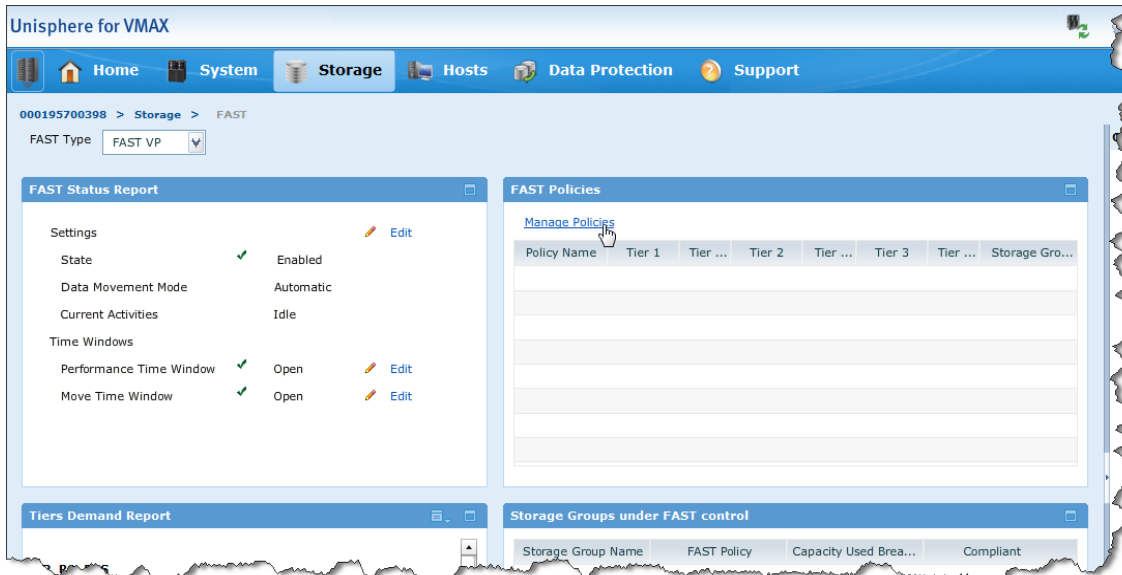
Creating a Symmetrix FAST policy

A FAST policy defines the upper usage limits for up to three tiers for any storage groups associated with the policy.

When creating a FAST policy, the following information must be known:

- ◆ The FAST policy name
- ◆ The VP tiers (maximum three) to be used in the policy
- ◆ The upper usage limits for each of the VP tiers being added

Once this information has been decided, the FAST policy can be created by first clicking the Manage Policies link on the FAST subsection page under Storage.



Then click the Create button on the Manage Policy related objects page.



In this example, a policy called System_Optimization is created. This policy allows up to 100 percent of the logical capacity of any associated storage groups to be moved to any of the tiers within the policy (in this

case, an EFD tier, a FC tier, and a SATA tier).

The 'Create FAST Policy' dialog box contains the following fields:

- Policy Name: System_Optimization
- Emulation: FBA
- Tier 1: VP_R53_EFD, 100%
- Tier 2: VP_R1_FC, 100%
- Tier 3: VP_R614_SATA, 100%

Buttons: OK, Cancel, Help

The policy is created when you click OK.

Once created, the information on the policy can be seen on the Manage Policies page.

Unisphere for VMAX

000195700398 > Storage > FAST > Manage Policies

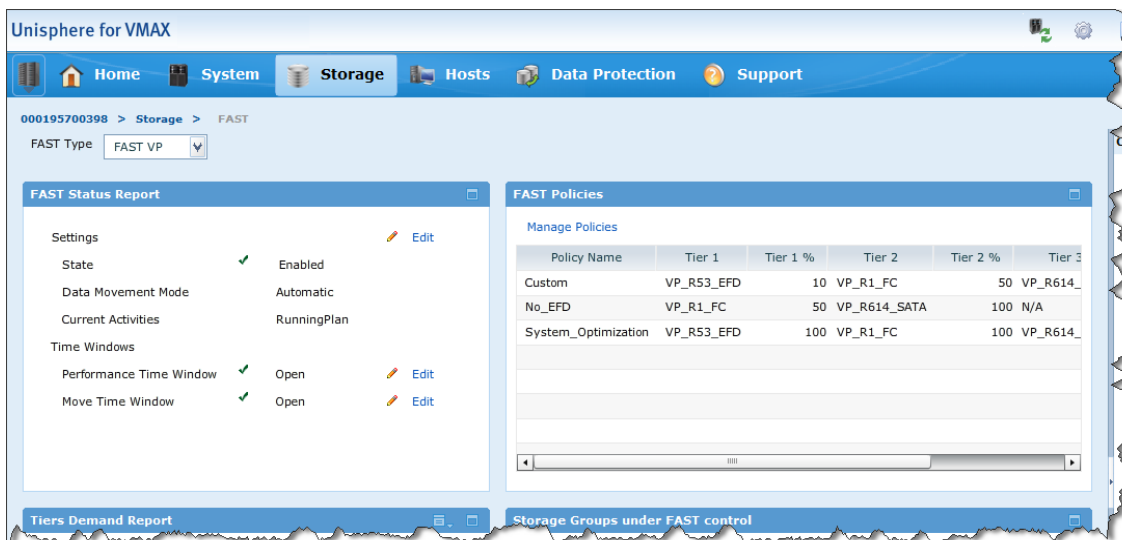
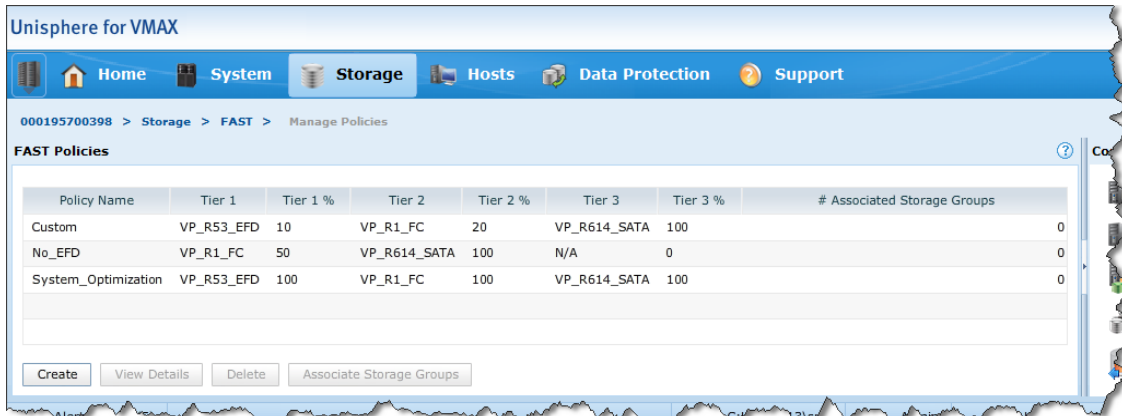
FAST Policies

Policy Name	Tier 1	Tier 1 %	Tier 2	Tier 2 %	Tier 3	Tier 3 %	# Associated Storage Groups
System_Optimization	VP_R53_EFD	100	VP_R1_FC	100	VP_R614_SATA	100	-1

Buttons: Create, View Details, Delete, Associate Storage Groups

FAST policy list information

After additional policies have been created, information on all FAST policies in the Symmetrix array can be viewed on either the Manage Policies page or the FAST subsection page.

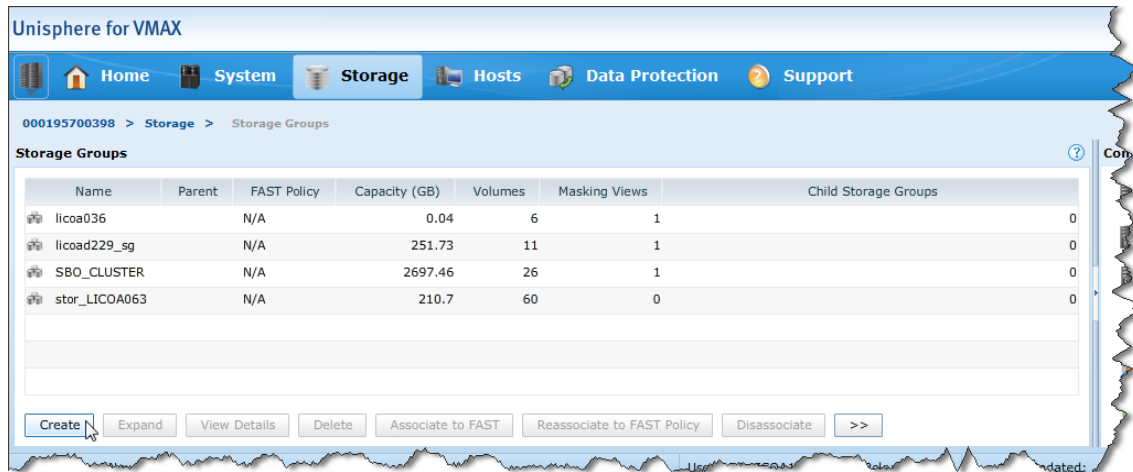


Creating a Symmetrix storage group

A storage group logically combines Symmetrix devices to be managed together. When creating a storage group, the following information must be known:

- ◆ The storage group name
- ◆ The Symmetrix devices to be added to the group

Once this information has been decided, the storage group is created by clicking the Create button on the Storage Groups subsection page under Storage.



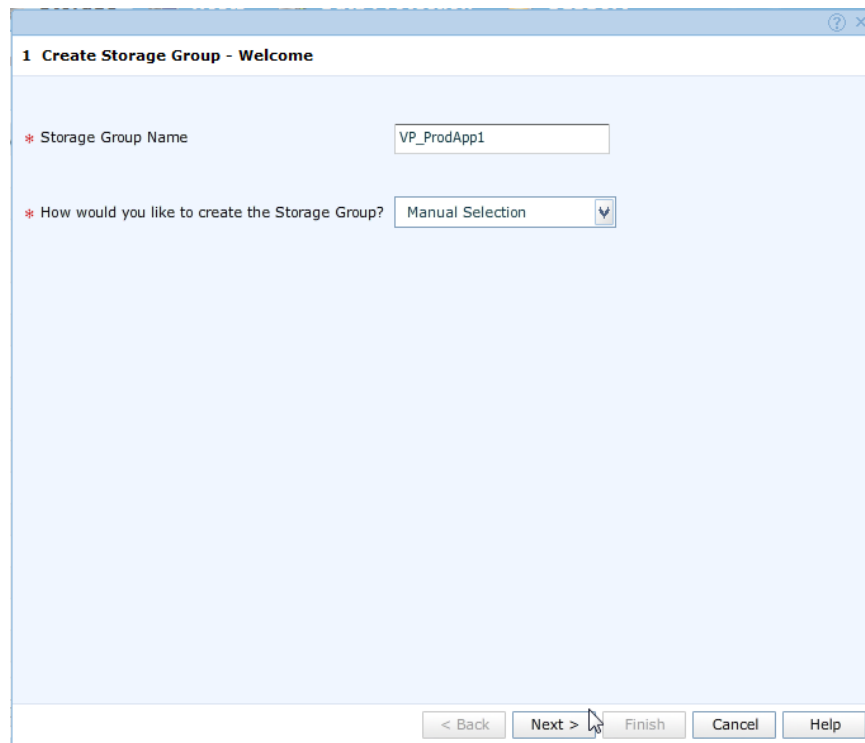
This launches the storage group creation wizard.

On the first page of the wizard, enter a storage group, and choose the method by which devices to be added are to be selected.

Choices for device selection include:

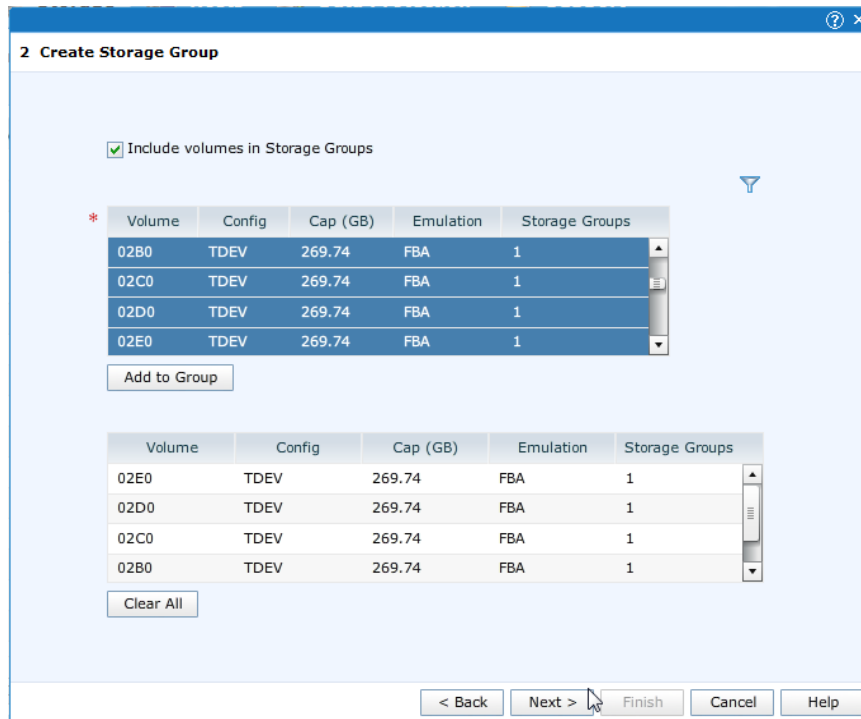
- ◆ Regular Volumes
- ◆ Virtual Volumes
- ◆ Manual Selection
- ◆ Cascaded Storage Group

The first two selections result in new devices being created and added to the group. The Manual Selection option allows already existing devices to be chosen and added to the group. Selecting Cascaded Storage Group allows the addition of child storage groups, containing devices, to the parent storage group.

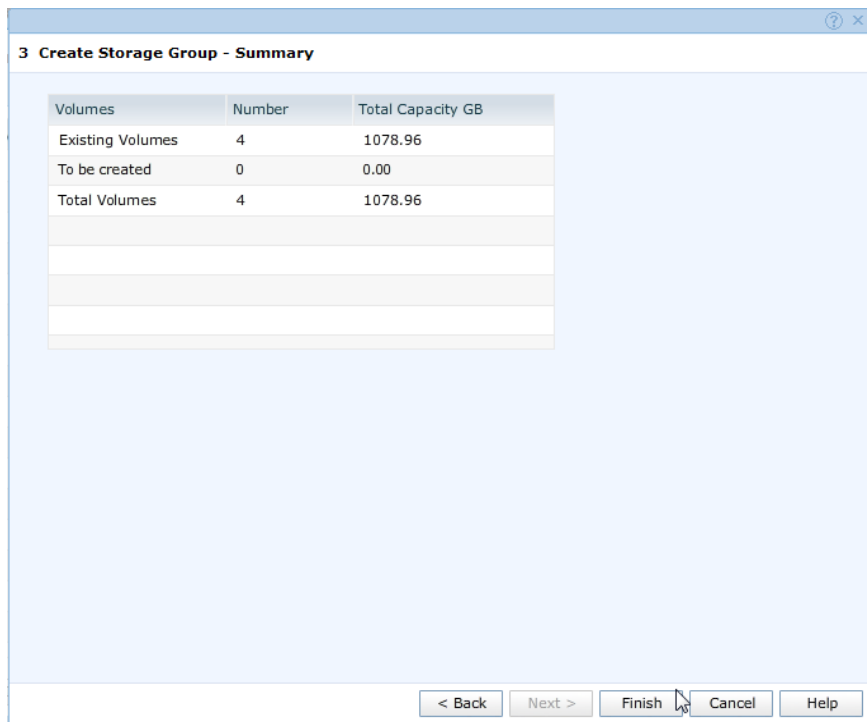


After choosing the selection method, in this case Manual Selection, click Next.

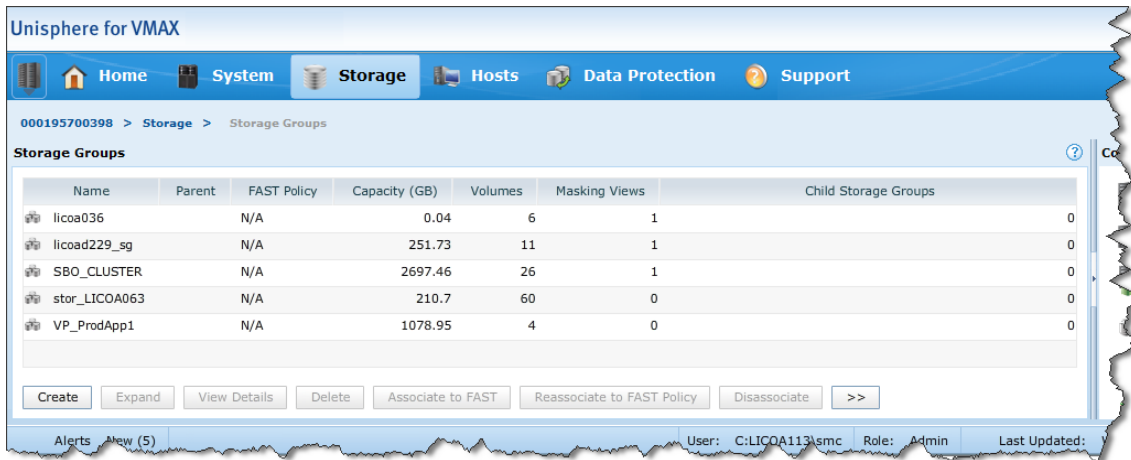
On the second screen of the wizard, after selecting the appropriate devices, click Add to Group, and then click Next.



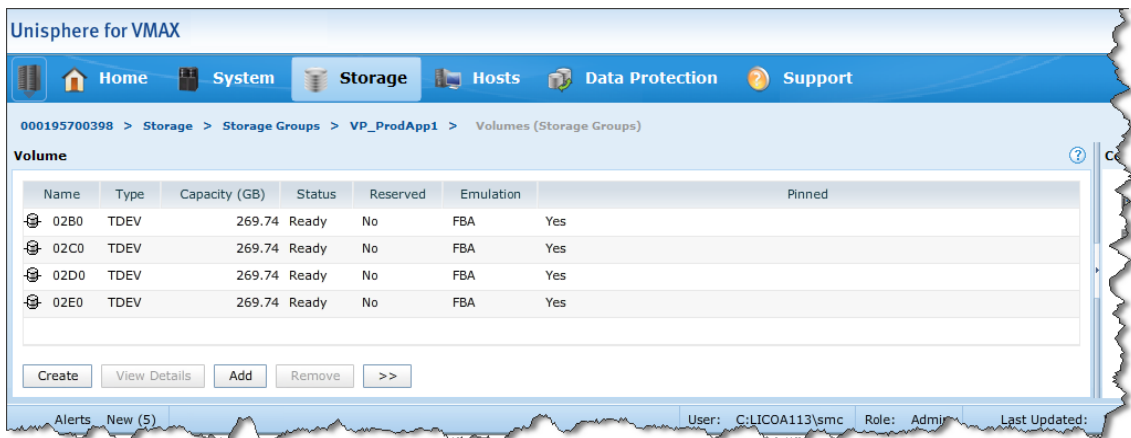
The final screen of the wizard provides a summary of the group being created and the devices being added. Click Finish to create the group.



Information on the created storage group can be displayed on the Storage Groups subsection page.



To verify the correct devices were added to the storage group, double-click the storage group name, and click the Volumes related object link.



Storage group list information

After additional storage groups have been created, information on all storage groups in the Symmetrix array can be viewed on the Storage Groups subsection page.

Information provided in this display includes:

- ◆ The names of all created storage groups
- ◆ The number of devices in each storage group

- ◆ The name of the FAST policy the group is associated with (if any)
- ◆ The number of masking views the storage group is included in

Unisphere for VMAX

Home System Storage Hosts Data Protection Support

000195700398 > Storage > Storage Groups

Storage Groups

Name	Parent	FAST Policy	Capacity (GB)	Volumes	Masking Views	Child Storage Groups
licoa036		N/A	0.04	6	1	0
licoad229_sg		N/A	251.73	11	1	0
SBO_CLUSTER		N/A	2697.46	26	1	0
stor_LICOA063		N/A	210.7	60	0	0
VP_Development		N/A	539.47	2	0	0
VP_ProdApp1		N/A	1078.95	4	0	0
VP_ProdApp2		N/A	1078.95	4	0	0

Create Expand View Details Delete Associate to FAST Reassociate to FAST Policy Disassociate >>

Alerts New (5) User: C:LICOA113\mc Role: Admin Last Updated: W

Associating a storage group with a FASTVP policy

Associating a storage group with a FAST VP policy brings the devices in the storage group under FAST VP control. All devices in the storage group are considered candidates to have data moved between the tiers included in the policy the storage group is associated with.

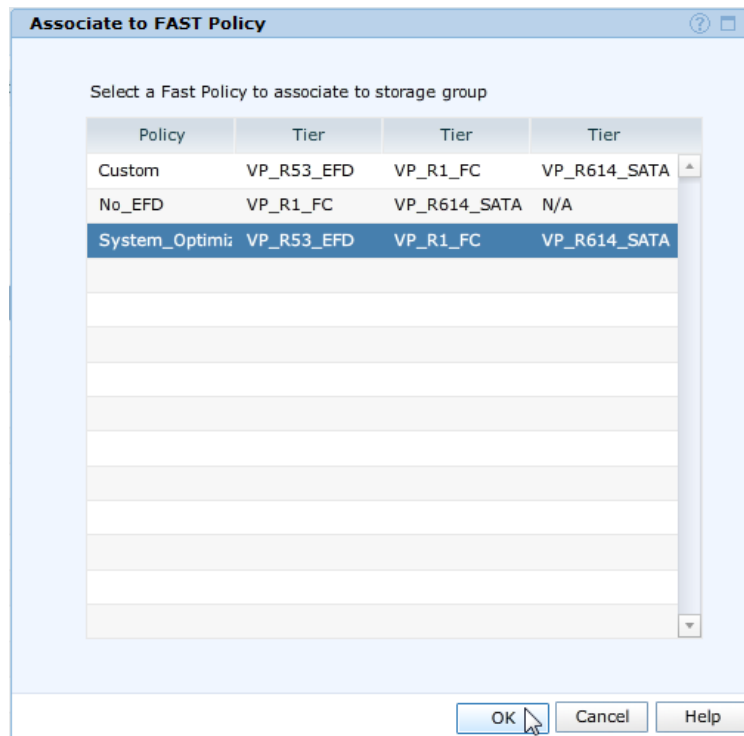
When creating a FAST policy association, the following information must be known:

- ◆ The FAST policy name
- ◆ The storage group name

Once you decide on the information, the association is performed by first selecting the appropriate storage group on the Storage Groups subsection page, then clicking the Associate to FAST.



In the resulting dialog box, choose the desired FAST policy and click OK.



Note: Storage groups are associated to the policy with a default value of 2. The following section describes how to modify a storage group priority within a FAST policy.

To verify the successful association of the storage group to the FAST policy, return to the Storage Groups subsection page.

Unisphere for VMAX

Home System Storage Hosts Data Protection Support

000195700398 > Storage > Storage Groups

Storage Groups

Name	Parent	FAST Policy	Capacity (GB)	Volumes	Masking Views	Child Storage Groups
licoa036		N/A	0.04	6	1	0
licoad229_sg		N/A	251.73	11	1	0
SBO_CLUSTER		N/A	2697.46	26	1	0
stor_LICOA063		N/A	210.7	60	0	0
VP_Development		N/A	539.47	2	0	0
VP_ProdApp1		System_Optimization	1078.95	4	0	0
VP_ProdApp2		N/A	1078.95	4	0	0

Create Expand View Details Delete Associate to FAST Reassociate to FAST Policy Disassociate >>

Alerts (5) User: C:\LICOA113\smc Role: Admin Last Updated:

When additional associations have been made, all storage group associations can be viewed on the Storage Groups subsection page.

Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Storage Groups

Storage Groups

Name	Parent	FAST Policy	Capacity (GB)	Volumes	Masking Views	Child Storage Groups
licoa036		N/A	0.04	6	1	0
licoad229_sg		N/A	251.73	11	1	0
SBO_CLUSTER		N/A	2697.46	26	1	0
stor_LICOA063		N/A	210.7	60	0	0
VP_Development		No_EFD	539.47	2	0	0
VP_ProdApp1		System_Optimization	1078.95	4	0	0
VP_ProdApp2		Custom	1078.95	4	0	0

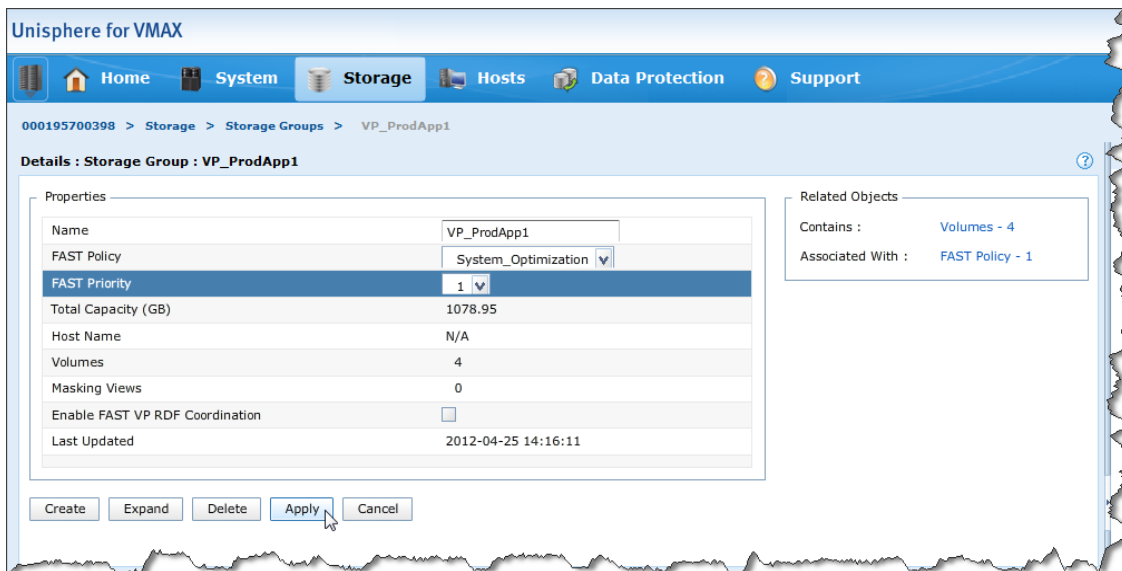
Create Expand View Details Delete Associate to FAST Reassociate to FAST Policy Disassociate >>

Alerts New (5) User: C:\LICOA113\smc Role: Administrator Last Updated: V

Modifying a storage group's priority in a FAST policy

When a storage group is associated to a FAST policy in Unisphere, it is added with a default priority of 2. After the association is complete, the priority can be raised to 1, or lowered to 3.

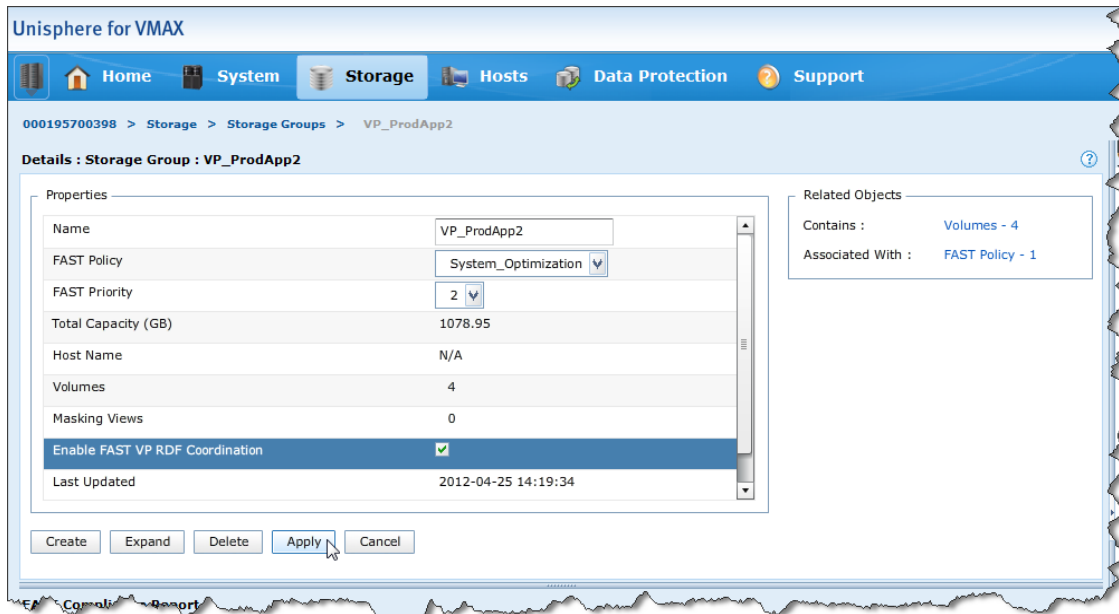
To modify the storage group's priority, double-click the storage group's name on the Storage Group subsection page. Change the value using the drop-down list to the right of FAST priority, then click Apply.



Enabling/disabling SRDF coordination

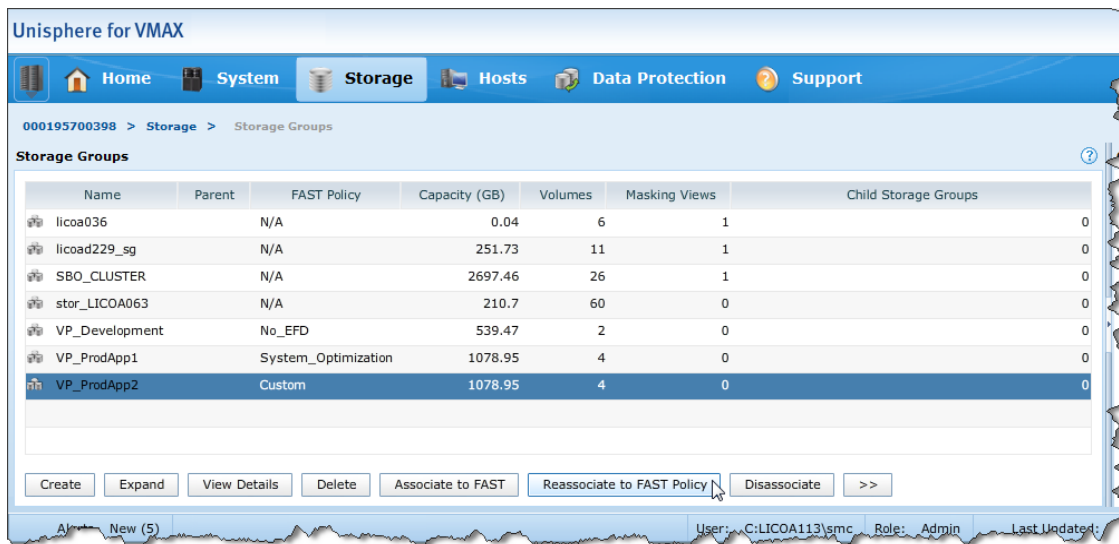
By default, SRDF coordination is disabled for any storage group associated with a FAST policy. SRDF coordination may be enabled, however, after the group has been associated.

To enable SRDF coordination during association, double-click the storage group's name on the Storage Group subsection page. Select the checkbox to the right of Enable FAST VP RDF Coordination, then click Apply.

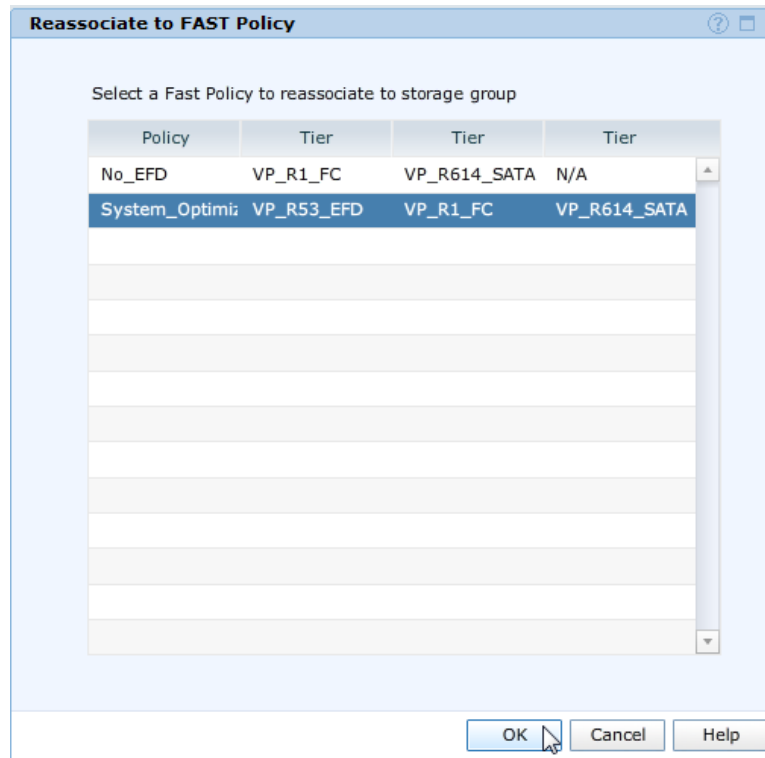


Reassociating a storage group to a different FASTVP policy

To move a storage group from one FAST policy to another, select the storage group on the Storage Groups subsection page, and then click Reassociate to FAST Policy.



In the resulting dialog box, choose the new policy the group is to be associated with and click OK.



To verify the successful reassociation of the storage group to the FAST policy, go to the Storage Groups subsection page under Storage.

Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Storage Groups

Storage Groups

Name	Parent	FAST Policy	Capacity (GB)	Volumes	Masking Views	Child Storage Groups
licoa036		N/A	0.04	6	1	0
licoad229_sg		N/A	251.73	11	1	0
SBO_CLUSTER		N/A	2697.46	26	1	0
stor_LICOA063		N/A	210.7	60	0	0
VP_Development		No_EFD	539.47	2	0	0
VP_ProdApp1		System_Optimization	1078.95	4	0	0
VP_ProdApp2		System_Optimization	1078.95	4	0	0

Create Expand View Details Delete Associate to FAST Reassociate to FAST Policy Disassociate >>

Alerts New (5) User: C:\LICOA113\smc Role: Admin Last Updated: W

By viewing the details of the reassociated storage group, you can confirm that the priority and SRDF coordination attributes remained the same during the reassociation.

Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Storage Groups > VP_ProdApp2

Details : Storage Group : VP_ProdApp2

Properties

Name	VP_ProdApp2
FAST Policy	System_Optimization
FAST Priority	2
Total Capacity (GB)	1078.95
Host Name	N/A
Volumes	4
Masking Views	0
Enable FAST VP RDF Coordination	<input checked="" type="checkbox"/>
Last Updated	2012-04-25 14:19:34

Related Objects

Contains : Volumes - 4

Associated With : FAST Policy - 1

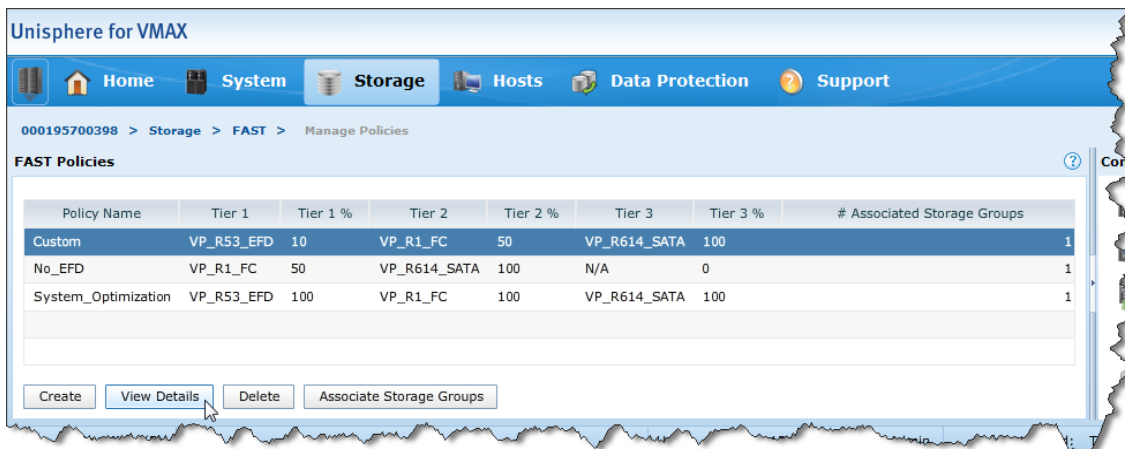
Create Expand Delete Apply Cancel

Modifying a Symmetrix tier in a FAST policy

After some time, it may be determined that the upper usage limit of a particular tier within a FAST policy needs to be adjusted. This can be done dynamically in Unisphere.

If any storage groups are associated with the policy being modified, the change in the usage limit cannot cause the sum of the usage limits for all tiers in the policy to fall below 100 percent. In this case, the Custom policy is modified.

To modify the usage limit for a Symmetrix tier within a policy, select the policy on the Manage Policies page, and click View Details.

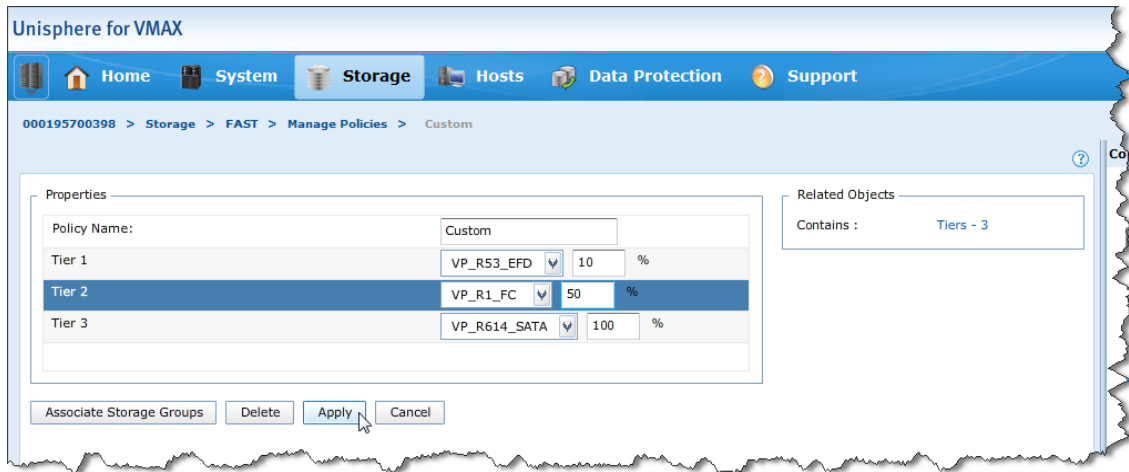


The screenshot shows the 'Manage Policies' page in Unisphere for VMAX. The breadcrumb trail is '000195700398 > Storage > FAST > Manage Policies'. The page title is 'FAST Policies'. Below the title is a table with the following data:

Policy Name	Tier 1	Tier 1 %	Tier 2	Tier 2 %	Tier 3	Tier 3 %	# Associated Storage Groups
Custom	VP_R53_EFD	10	VP_R1_FC	50	VP_R614_SATA	100	1
No_EFD	VP_R1_FC	50	VP_R614_SATA	100	N/A	0	1
System_Optimization	VP_R53_EFD	100	VP_R1_FC	100	VP_R614_SATA	100	1

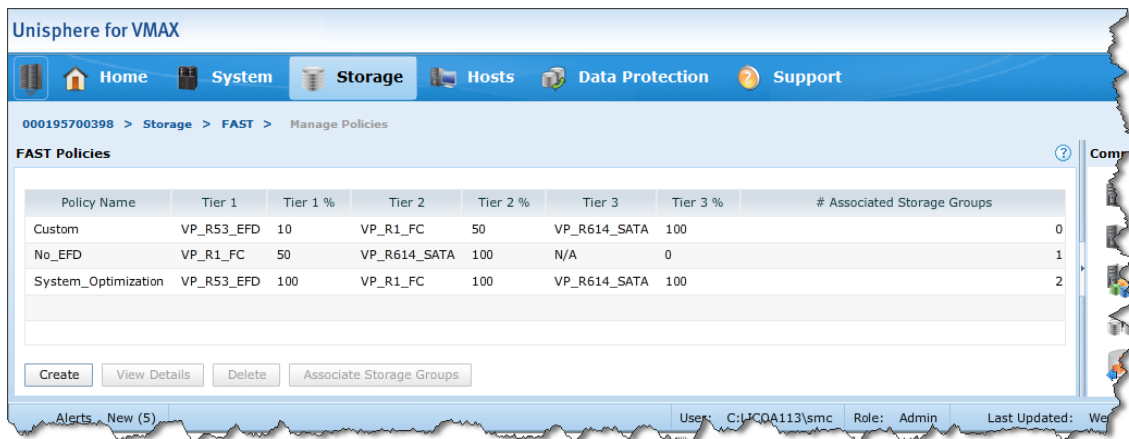
Below the table are four buttons: 'Create', 'View Details', 'Delete', and 'Associate Storage Groups'. A mouse cursor is pointing at the 'View Details' button.

On the resulting page, edit the percent value for the desired tier, and click Apply.



In this case the percentage of capacity allowed on the V_R1_FC tier was increased from 20 percent to 50 percent.

The change can be verified on the Manage Policies page.



Controlling FAST device movement

Aside from using the FAST controller device movement window, there are several other ways of controlling when device movements can take place. These include:

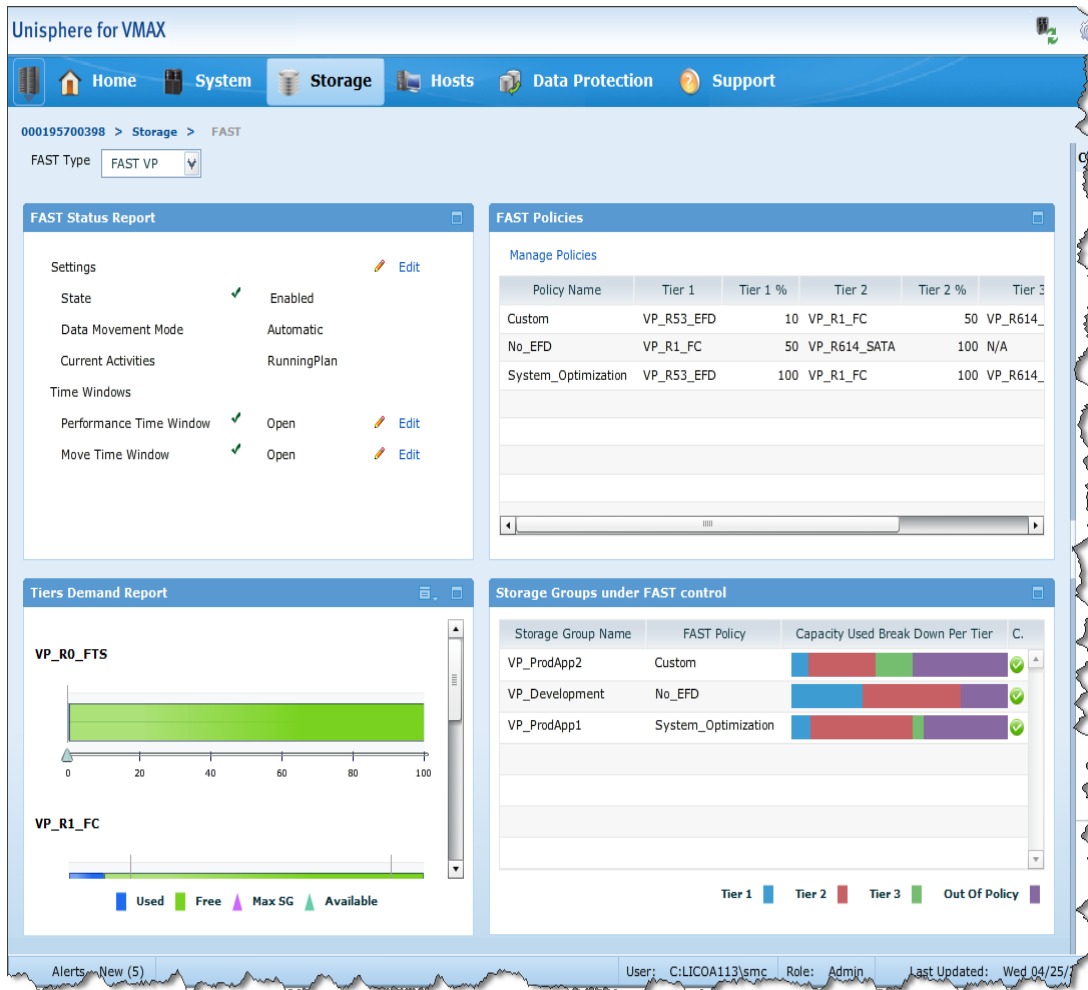
- ◆ Disabling the FAST controller
- ◆ Pinning devices under FAST VP control

- ◆ Changing the data movement mode
- ◆ Modifying data movement windows

Monitoring FASTVP status

The current status and current activity of FAST VP can be monitored in Unisphere by using the FAST status report on the FAST subsection page. Information provided by this report includes:

- ◆ The FAST VP state
- ◆ Degraded reason
- ◆ FAST VP Current Activities
- ◆ FAST VP Performance Time Window state
- ◆ FAST VP Move Time Window state



The FAST VP state is typically Enabled, Disabled, or Degraded. If the state is degraded, a reason is listed.

Note: For more information on other possible state values and degraded reason codes, see “Appendix A: FAST VP state” on page 180.

FAST VP current activities either display as Idle or Running Plan. An activity of Idle indicates that there are currently no active data movement tasks related to FAST VP in the array. Running Plan indicates

that data is actively being moved at the sub-LUN level for thin devices.

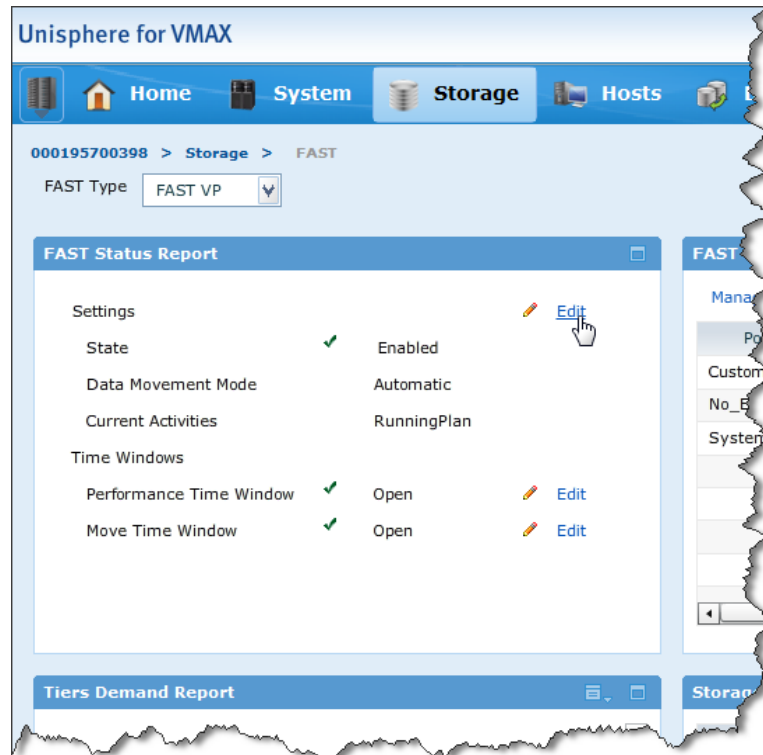
The state of the performance and move time windows are displayed as Open or Closed. Windows are considered open if the current time corresponds to a defined inclusion time window. Similarly, the windows are considered closed if the current time falls outside of any defined inclusion windows.

Note: The FAST VP move time window is displayed as closed if the data movement mode is off, regardless of any defined inclusion time windows.

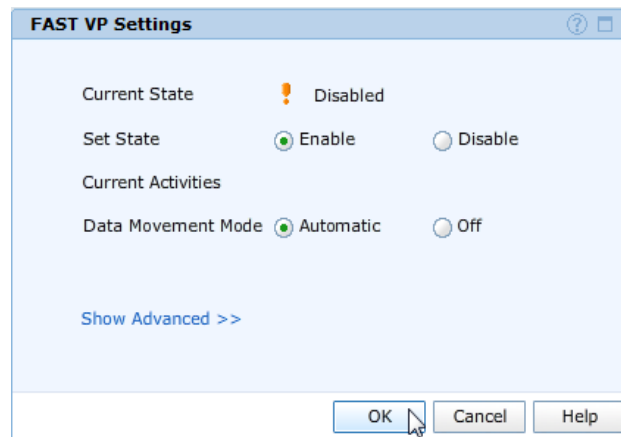
Enabling/disabling FASTVP

In order for FAST VP to perform device movements, FAST VP must first be enabled on the Symmetrix array.

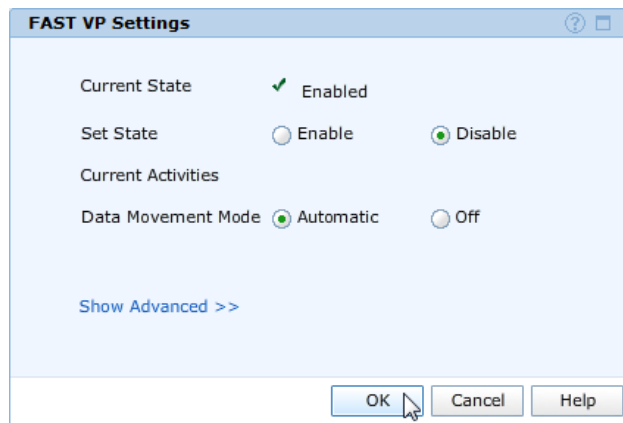
The state of the controller can be changed by clicking the Edit link to the right of Settings.



If disabled, FAST VP can be enabled by selecting the radio button to the left of Enable.



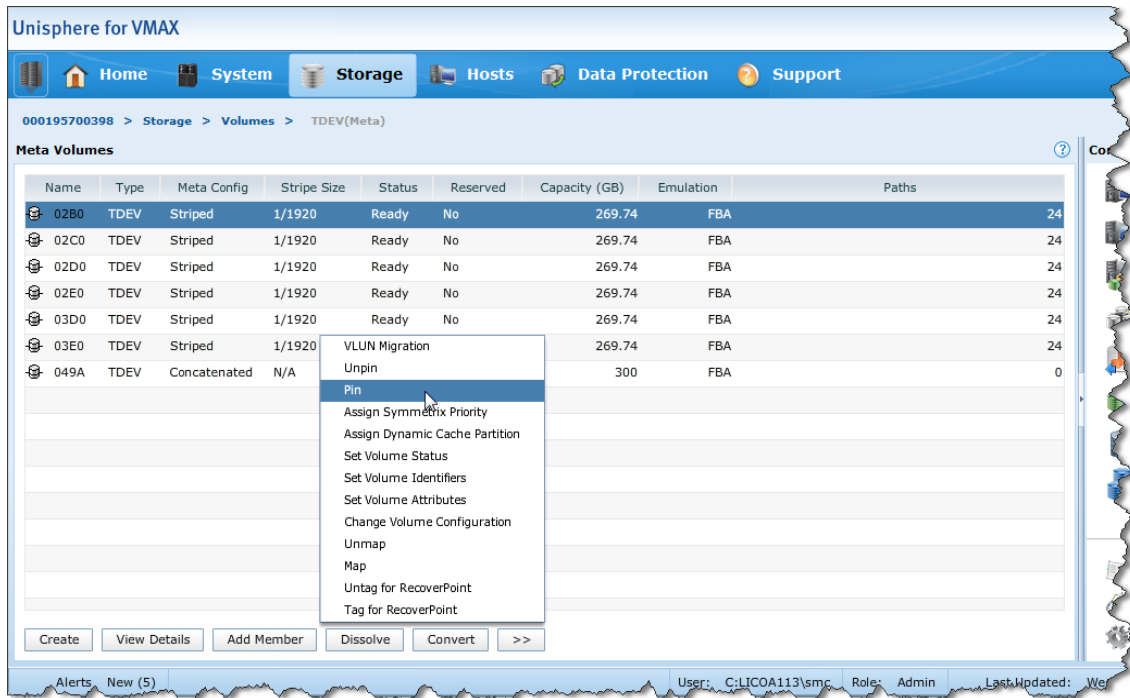
If enabled, FAST VP can be disabled by selecting the radio button to the left of Disable.



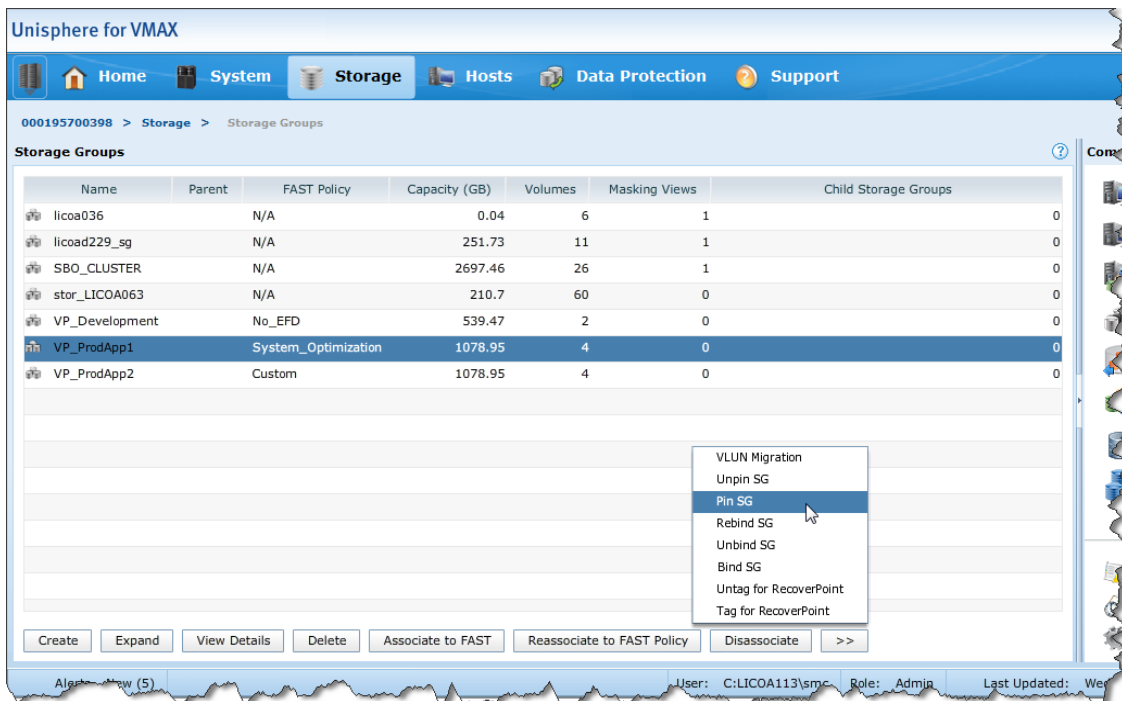
Pinning a thin device under FASTVP control

Pinning a device that is associated with a FAST VP policy prevents any data movement for that device. Devices can be pinned individually, or as a storage group, or device group.

To pin an individual device, select a device on any page where individual devices are listed. For example, on the TDEV(Meta) page under Volumes, click >>, and then select Pin.



To pin all devices in a storage group, select a storage group on the Storage Groups subsection page, click >>, and then select Pin SG.

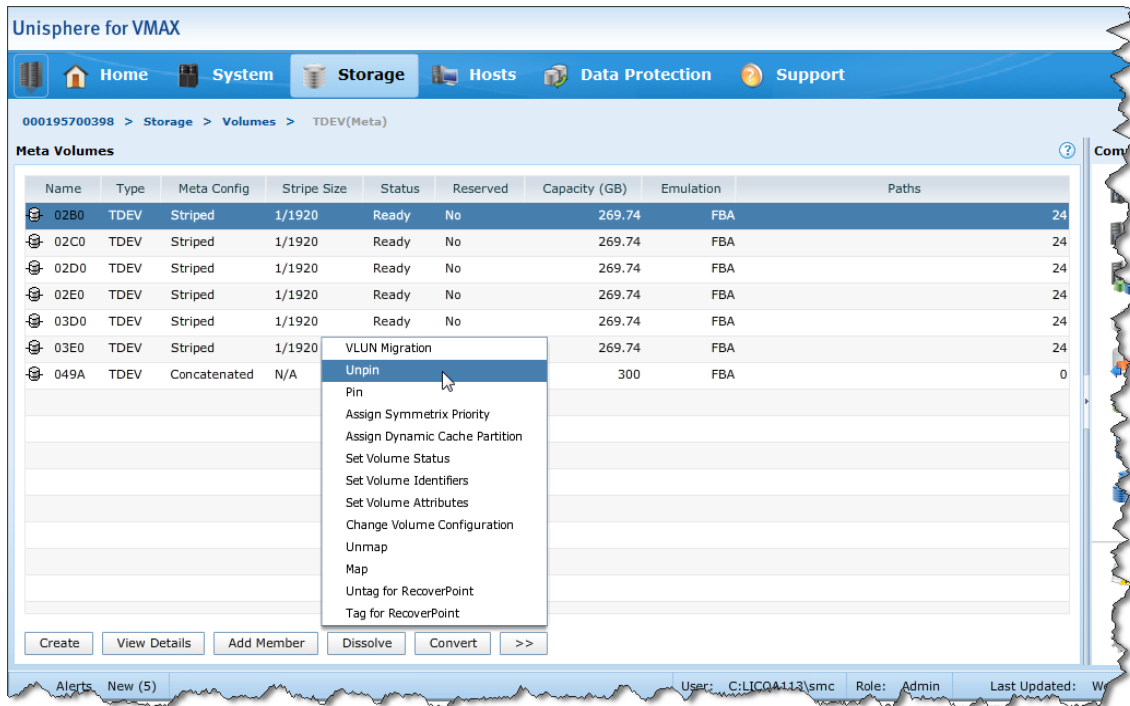


Unpinning a thin device under FASTVP control

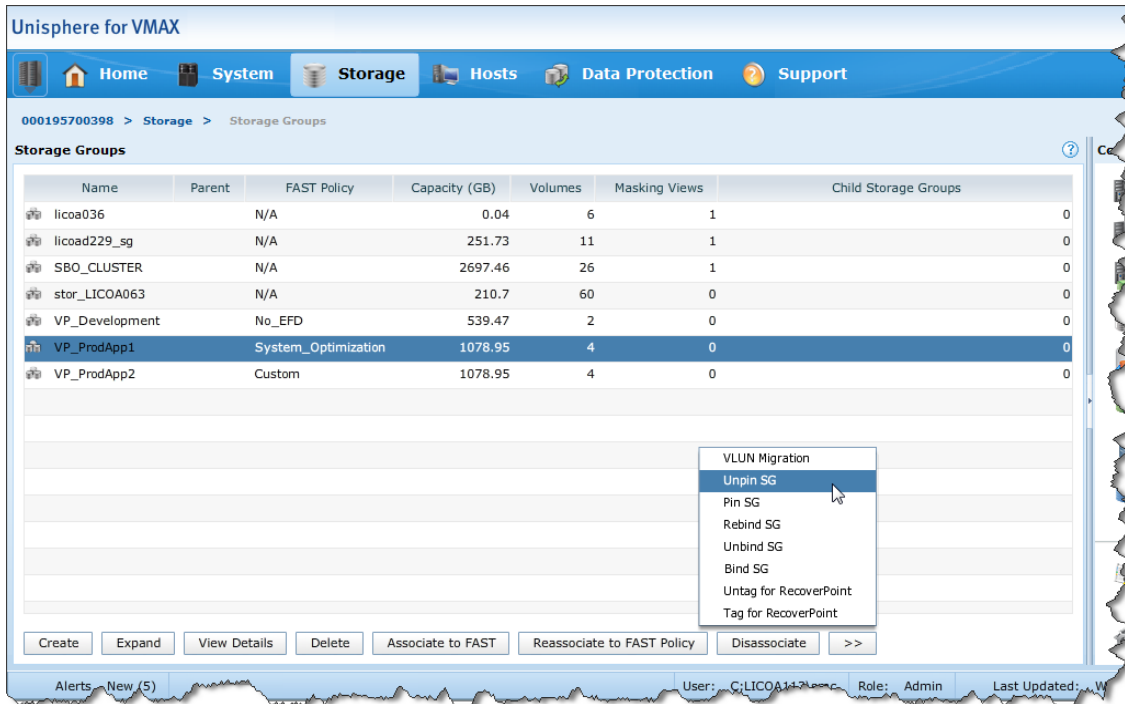
In order for FAST VP to resume data movements after a device has been pinned, it must be unpinned.

Devices can be unpinned individually, or as a storage group, or device group.

To unpin an individual device, select a device on any page where individual devices are listed. For example, on the TDEV(Meta) page under Volumes, click >>, and then select Unpin.



To unpin all devices in a storage group, select a storage group on the Storage Groups subsection page, click >>, and then select Unpin SG.



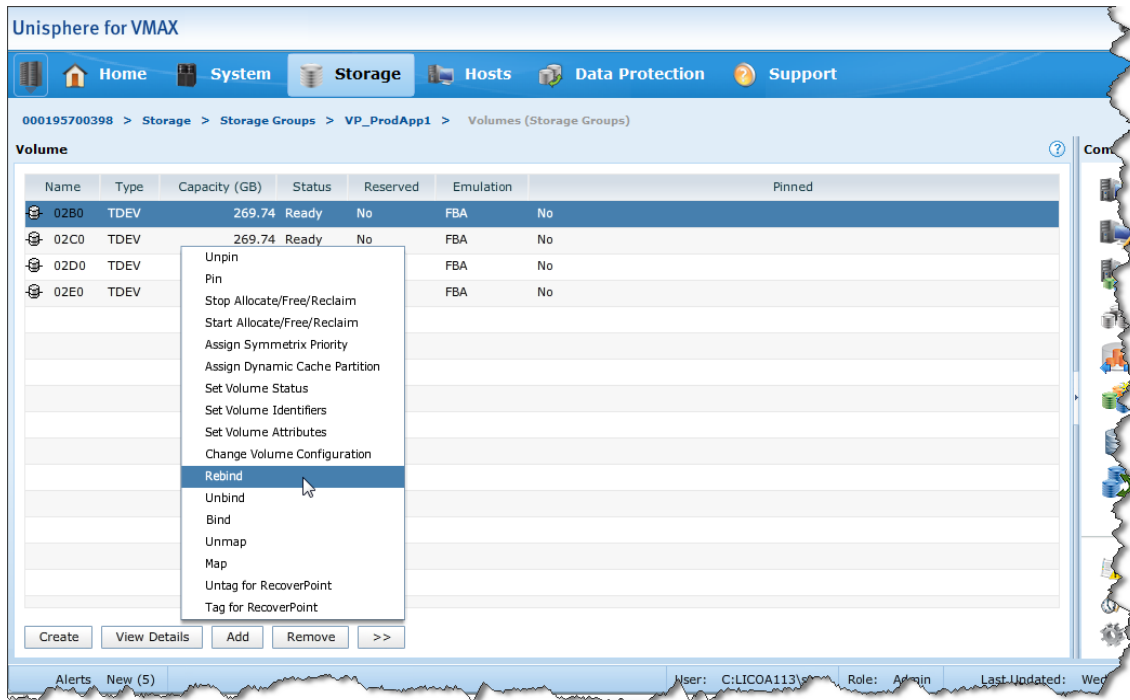
Rebinding a thin device

While FAST VP moves data at the sub-LUN level between device pools, the thin devices associated with a policy still remain bound to a single pool. Any new allocations that occur as the result of a host write come from the bound pool, unless VP allocation by FAST policy is enabled.

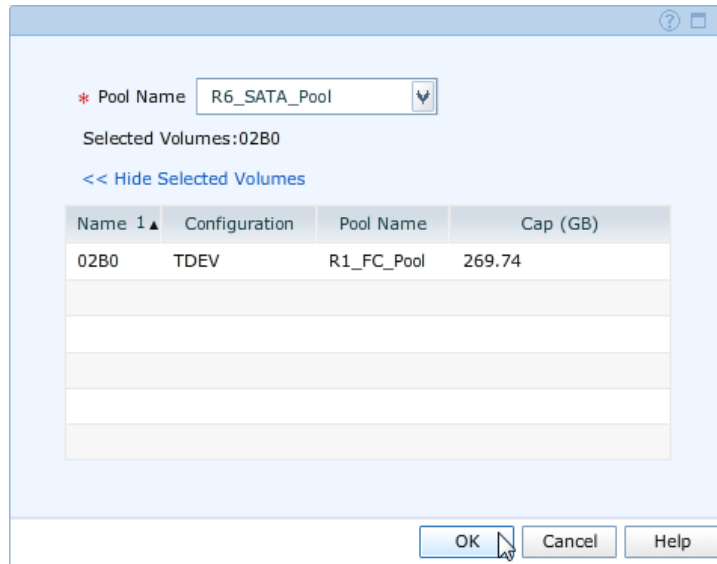
The thin device rebind feature allows the binding information for a thin device to be changed, without changing the current allocation of data across pools.

Note: If the devices being rebound are associated with a FAST VP Policy, only pools that are contained within the policy can be specified as the new bind pool.

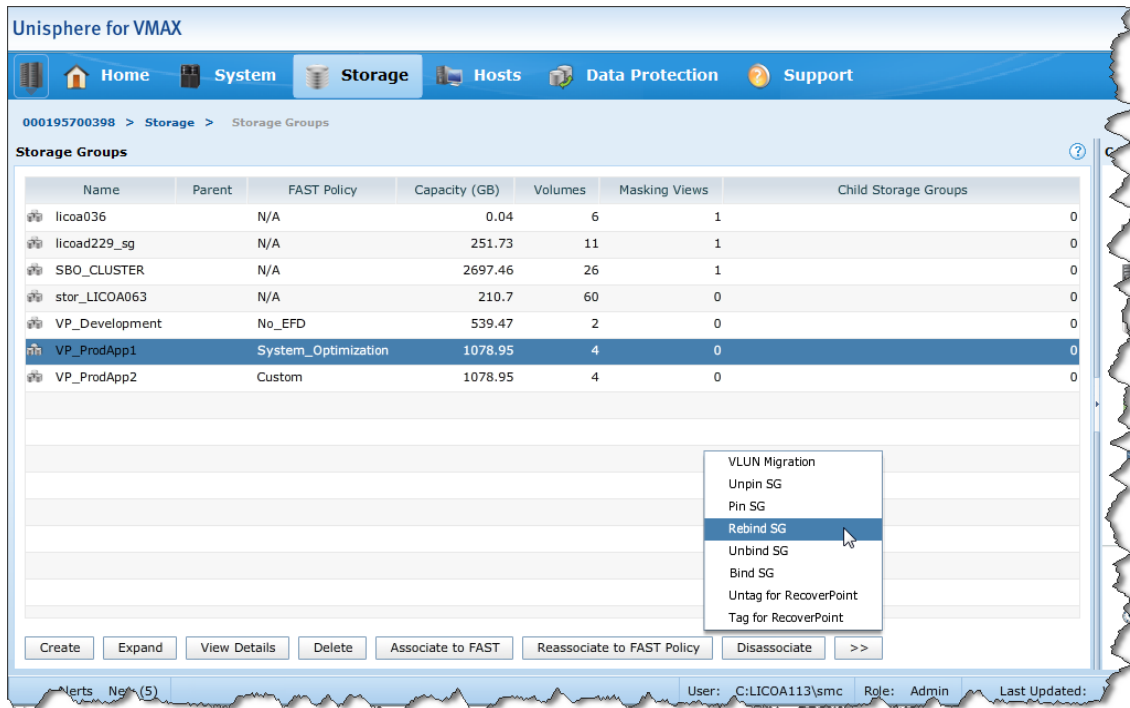
In this example, to rebind thin device 2B0 to the RAID 6 protected SATA thin pool, R6_SATA_Pool, select the device on any page that lists Symmetrix devices, click >>, then select Rebind.



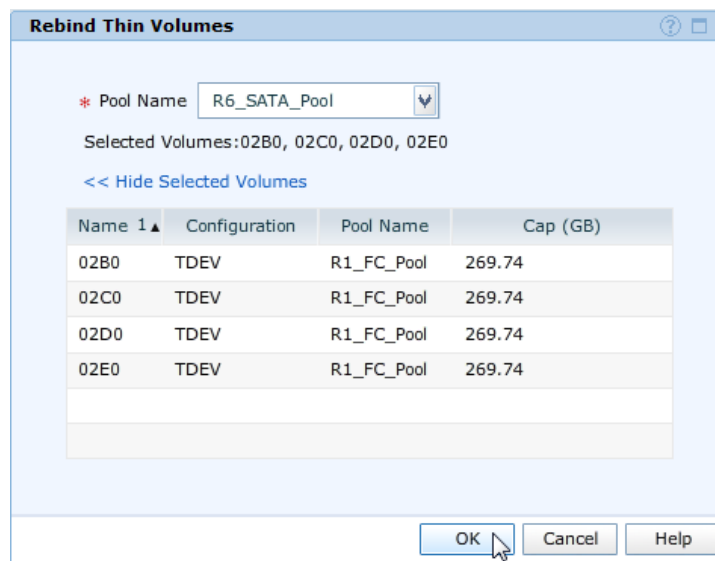
The resulting dialog box provides a drop-down list of available pools for the device to be rebound to. Select the pool, and then click OK.



In this example, to rebind all the thin devices in the VP_ProdApp1 storage group to the R6_SATA_Pool thin pool, select the group on the Storage Group subsection page, click >>, then select Rebind SG.



The resulting dialog box provides a drop-down list of available pools for the device to be rebound to. Select the pool, and then click OK.



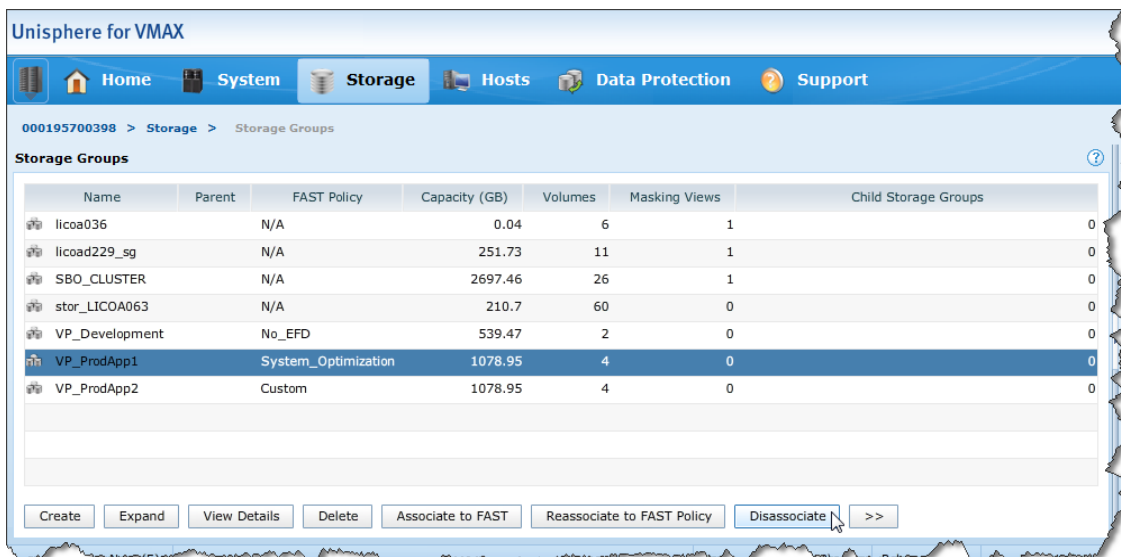
Removing FAST managed objects

The following sections detail the Unisphere interfaces used to break each of the associations, and to remove each of the managed objects.

Disassociating a storage group from a FAST policy

Disassociating a storage group from a FAST VP policy removes the devices in the storage group from FAST VP control. Devices no longer under FAST VP control will not have data automatically moved between tiers.

To disassociate a storage group from a FAST policy, select the appropriate storage group on the Storage Groups subsection page, then click the Disassociate button.



Unisphere for VMAX

Home System Storage Hosts Data Protection Support

000195700398 > Storage > Storage Groups

Storage Groups

Name	Parent	FAST Policy	Capacity (GB)	Volumes	Masking Views	Child Storage Groups
licoa036		N/A	0.04	6	1	0
licoad229_sg		N/A	251.73	11	1	0
SBO_CLUSTER		N/A	2697.46	26	1	0
stor_LICOA063		N/A	210.7	60	0	0
VP_Development		No_EFD	539.47	2	0	0
VP_ProdApp1		System_Optimization	1078.95	4	0	0
VP_ProdApp2		Custom	1078.95	4	0	0

Create Expand View Details Delete Associate to FAST Reassociate to FAST Policy Disassociate >>

Removing a Symmetrix tier from a FAST policy

As previously stated, the sum of the upper usage limits for each Symmetrix tier contained in a policy must total a minimum of 100

percent. If a policy has associated storage groups, the removal of a tier will fail if it causes the total tier usage limits to drop below 100 percent.

To successfully remove a Symmetrix tier from a FAST policy, the upper usage limits of the remaining tiers must be modified in order that they total more than 100 percent. Alternatively, the Symmetrix tier can be removed if all storage group associations for the policy are removed in advance.

To remove a Symmetrix tier from a FAST policy, click the Manage Policies link on the FAST subsection page.

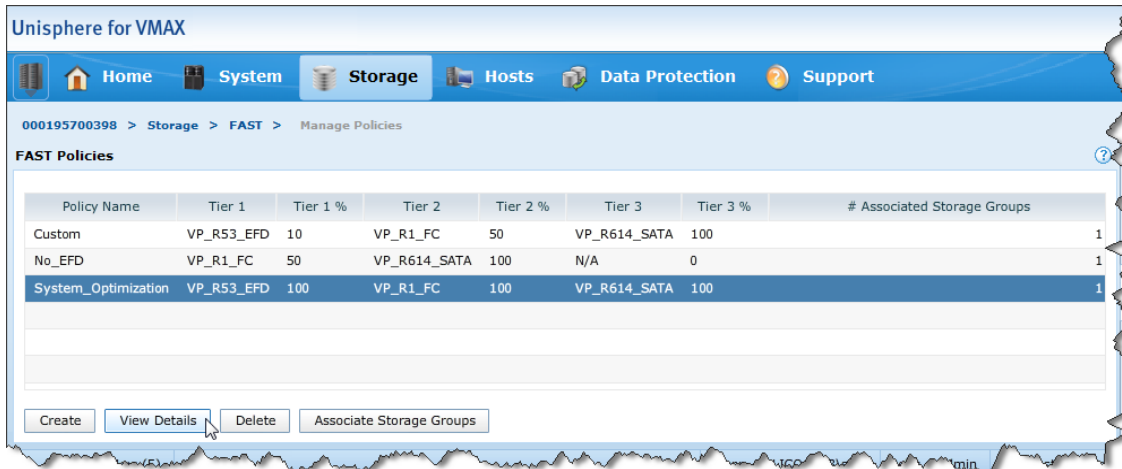
The screenshot displays the Unisphere for VMAX interface. The top navigation bar includes Home, System, Storage, Hosts, Data Protection, and Support. The breadcrumb path is 000195700398 > Storage > FAST. The FAST Type is set to FAST VP. The interface is divided into several sections:

- FAST Status Report:** Shows settings for State (Enabled), Data Movement Mode (Automatic), Current Activities (Idle), and Time Windows (Performance and Move Time Windows are Open).
- FAST Policies:** A table listing policies with their tiers and usage percentages. A "Manage Policies" link is visible above the table.
- Tiers Demand Report:** Shows a report for VP_R0_FTS.
- Storage Groups under FAST control:** A table showing storage groups and their associated FAST policies.

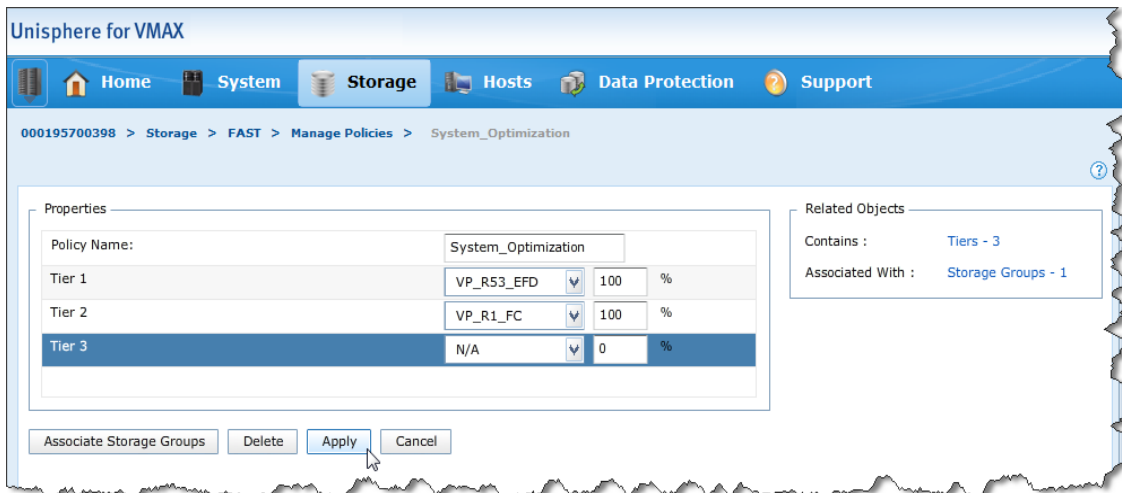
Policy Name	Tier 1	Tier 1 %	Tier 2	Tier 2 %	Tier 3
Custom	VP_R53_EFD	10	VP_R1_FC	50	VP_R614
No_EFD	VP_R1_FC	50	VP_R614_SATA	100	N/A
System_Optimization	VP_R53_EFD	100	VP_R1_FC	100	VP_R614

Storage Group Name	FAST Policy	Capacity Used Break Down Per Tier	C.
VP_ProdApp2	Custom		

Next, select the policy to be modified, and click View Details.



To remove a tier from the policy, select N/ A from the drop-down menu for the tier being removed, and click Apply.

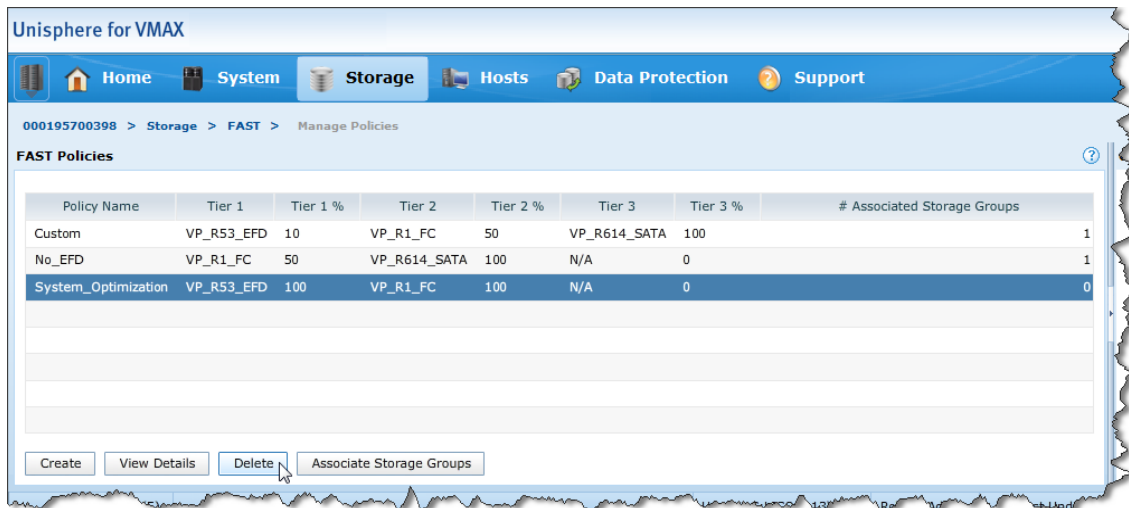


Deleting a FAST policy

A FAST policy may not be deleted if any storage groups are associated with it. Prior to deleting the policy, all storage group associations should

be removed. Also, prior to deleting a policy, all tiers contained in the policy should be removed.

To delete a FAST policy, go to the Manage Policies page. Select the policy to be deleted, then click Delete.



Deleting a Symmetrix VP tier

If you want to delete a VP tier, it cannot be included within any FAST policy on the Symmetrix array.

To delete a VP tier, go to the Tiers subsection page. Select the tier to be deleted, and then click Delete.

Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Tiers

Tiers

Name	Type	Technology	Disk Location	Emulation	Used Capacity	Capacity (GB)	Protection
VP_R0_FTS	VP - Virtual Provisioning	N/A	External	FBA	0 %	195	Unprotected
VP_R1_FC	VP - Virtual Provisioning	FC	Internal	FBA	9 %	10723	RAID-1
VP_R53_EFD	VP - Virtual Provisioning	EFD	Internal	FBA	3 %	4395	RAID-5(3+1)
VP_R614_SATA	VP - Virtual Provisioning	SATA	Internal	FBA	4 %	24539	RAID-6(14+2)

Used Capacity Free Capacity

Create Edit View Details Delete

Deleting a storage group

Before deleting a storage group, ensure that the group is not associated with any FAST VP policy.

To delete a storage group, go to the Storage Groups subsection page, select the storage group to be deleted, and then click Delete.



Note: Since storage groups may also be used for the purposes of FAST and Auto-provisioning, prior to deleting the storage group, you must delete any masking views that use the storage group.

Management interface: SYMCLI

Solutions Enabler provides features and functionality for managing FAST VP in both Open Systems and mainframe environments.

There are five Solutions Enabler SYMCLI commands that can be used for the purposes of managing FAST VP:

- ◆ `symtier`
- ◆ `symfast`
- ◆ `symoptmz`
- ◆ `symtw`
- ◆ `symsg`

The `symtier` command provides the ability to create and manage Symmetrix tiers to be used with FAST policies.

The `symfast` command allows for the creation and management of FAST policies, including their associations with storage groups and Symmetrix tiers. The command also provides management control of the FAST controller, including modifying settings, and enabling or disabling

the controller.

The `symoptmz` command is used to create and manage legacy time windows, for both performance and data movement.

The `symtw` command is used to create and manage enhanced time windows, for both performance and data movement.

The `symsg` command is used to create and manage storage groups on the Symmetrix array for the purpose of being used with FAST.

Note: The `symaccess` command introduced in Solutions Enabler V7.0 can also be used to create storage groups for use with FAST. The `symsg` command may also be used to create storage groups to be used in creating masking views as a part of Auto-provisioning Groups.

The following sections detail the use of all of these commands in building, managing, and enabling a FAST VP environment using the Solutions Enabler SYMCLI.

Examining Symmetrix Virtual Provisioning thin pools

Prior to configuration of the FAST controller, and configuring Symmetrix tiers and FAST policies, it is important to understand the existing configuration of the Symmetrix array. Solutions Enabler SYMCLI can inventory the existing thin pools within the array. This can be done with the following command:

```
symcfg -sid 0398 list -pools -thin -mb
```

```
Symmetrix ID: 000195700398
```

```
          S Y M M E T R I X   P O O L S
-----
Pool      Flags Dev           Usable    Free      Used Full
Name      PTESL Config           MBs       MBs       MBs  (%)
-----
R53_EFD_Pool TEFEI RAID-5 (3+1)  71999616  71987820   11796    0
R1_FC_Pool  TFFEI 2-Way Mir    175680000 172478388  3201612  1
R57_FC_Pool TFFEI RAID-5 (7+1) 175680000 175680000         0    0
```


R6_SATA_Pool	TSFEI RAID-6(14+2)	402044928	382260876	19784052	4
R0_FTS_Pool	T-FEX Unprotected	3200040	3200040	0	0
Total		-----	-----	-----	----
MBs		51787827	50351341	1436453	3

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

(S)tate:

E = Enabled, D = Disabled, B = Balancing

Disk (L)ocation:

I = Internal, X = External, M = Mixed, - = N/A

The output for each thin pool shows the technology type, emulation, and location of the pool. It also shows the usable capacity of the pool (the free and used capacity), as well as the percentage allocated for the pool.

Running the command without the `-mb` option shows similar information, but the capacity values are displayed in tracks, rather than megabytes (MBs).

Note: To use a thin pool for FAST VP, all data devices in the pool must be configured on the same drive technology.

To see more information on an individual thin pool, including the data devices in the pool and the thin devices bound to the pool, run:

```
symcfg -sid 0398 show -pool R1_FC_Pool -thin -detail -mb
```

Symmetrix ID: 000195700398

```
Symmetrix ID           : 000195700398
Pool Name              : R1_FC_Pool
Pool Type              : Thin
Disk Location          : Internal
Technology             : FC
Dev Emulation          : FBA
Dev Configuration     : 2-Way Mir
Pool State             : Enabled
```

```

# of Devices in Pool           : 160
# of Enabled Devices in Pool   : 160
# of Usable Tracks in Pool     : 175680000
# of Allocated Tracks in Pool  : 10148976
# of Shared Tracks in Pool     : 0
Pool Utilization (%)          : 0
Max. Subscription Percent      : None
Rebalance Variance            : 1%
Max devs per rebalance scan    : 256
Pool Reserved Capacity         : None

```

Enabled Devices(160):

```

{
-----
Sym      Usable      Alloc      Free Full FLG Device
Dev      MBs        MBs        MBs  (%)  S  State
-----
00F0     68625      8207      60419  11  .  Enabled
00F1     68625      8210      60416  11  .  Enabled
00F2     68625      8194      60431  11  .  Enabled
...
018D     68625      8198      60427  11  .  Enabled
018E     68625      8187      60438  11  .  Enabled
018F     68625      8222      60403  11  .  Enabled
-----
MBs      10980000   1315357   9664643  11
}

```

Pool Bound Thin Devices(10):

```

{
-----
Sym  FLG      Pool      Pool      Total
Dev  T      Total Sub  Allocated  Written
      T      MBs (%)   MBs (%)   MBs (%)
-----
02B0  B      276210  3      77650  28      81729  30
02C0  B      276210  3      81275  29      81921  30
02D0  B      276210  3      78875  29      81761  30
02E0  B      276210  3      82926  30      82688  30
02F0  B      276210  3      110594  40     110379  40
0300  B      276210  3      110594  40     110121  40

```

```

0310  B      276210  3      110505  40      110031  40
0320  B      276210  3      110506  40      110041  40
03D0  B      276210  3      276216 100      276207 100
03E0  B      276210  3      276216 100      276207 100
-----
MBs      2762100 25      1315357 48      1321084 48
}

```

Legend:

Enabled devices FLG:

(S)hared Tracks : X = Shared Tracks , . = No Shared Tracks

Bound Devices FLG:

S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
D = Deallocating, R = Reclaiming, . = Unbound

Running the command without the `-mb` option shows similar information, but the capacity values are displayed in tracks, rather than megabytes (MBs).

Configuring the Symmetrix FASTcontroller

There are multiple FAST VP settings and parameters that affect the behavior of the FAST controller. These include:

- ◆ FAST VP Data Movement Mode
- ◆ Workload Analysis Period
- ◆ Initial Analysis Period
- ◆ Pool Reserved Capacity (PRC)
- ◆ FAST VP Relocation Rate
- ◆ VP Allocation by FAST Policy
- ◆ Performance Time Window
- ◆ Data Movement Time Window

The following sections detail the SYMCLI commands used to list and modify each of these settings.

FAST controller settings list information

To view the existing FAST controller settings, the `symfast list` command can be used with the `-control_parms` option:

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

```
Symmetrix ID: 000195700398
```

```
Optimizer and FAST Control Parameters:
```

```
Data Movement Mode           : User_Approve
Max Simultaneous Device Moves : 8
Max Device Moves Per Day      : 200
```

```
Optimizer, FAST and FAST VP Control Parameters:
```

```
Min Initial Workload Period(hrs) : 8
Workload Analysis Period(hrs)    : 168
```

```
FAST Control Parameters:
```

```
Swap Not Visible Devices       : Disabled
Allow Only Swap                : Disabled
```

```
FAST VP Control Parameters:
```

```
FAST VP Data Movement Mode    : NONE
FAST VP Data Relocation Rate   : 5
Thin Pool Reserved Capacity(%) : 10
VP Allocation By FAST policy    : Disabled
```

Setting the FASTVP Data Movement Mode

There are two possible values for the Data Movement Mode: Automatic (AUTO) and Off (NONE). The default value is NONE.

To set the mode to Automatic, run:

```
symfast -sid 0398 set -control_parms -vp_data_move_mode AUTO
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

```
The FAST SET CONTROL PARAMETERS operation finished successfully
```

To set the mode to Off, run:

```
symfast -sid 0398 set -control_parms -vp_data_move_mode NONE
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

```
The FAST SET CONTROL PARAMETERS operation finished successfully
```

Setting the Workload Analysis Period

The Workload Analysis Period is shared with Optimizer and FAST. It can be set between 2 hours and 4 weeks, but its value is specified in hours. The default value is 168 hours (1 week).

To set the Workload Analysis Period for FAST, run:

```
symfast -sid 0398 set -control_parms -workload_period 72
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

```
The FAST SET CONTROL PARAMETERS operation finished successfully
```

Setting the Initial Analysis Period

The Initial Analysis Period is a shared parameter with Optimizer and FAST. It can be set between 2 hours and 4 weeks. The default value is 8 hours.

To set the Initial Analysis Period for FAST, run:

```
symfast -sid 0398 set -control_parms -min_perf_period 24
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

```
The FAST SET CONTROL PARAMETERS operation finished successfully
```

Setting the system-wide FASTVP Pool Reserved Capacity

The FAST VP Pool Reserved Capacity (PRC) is specified as a percentage and can be set to be between 1 and 80. The default value is 10.

To set the PRC, run:

```
symfast -sid 0398 set -control_parms -pool_resv_cap 5
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

The FAST SET CONTROL PARAMETERS operation finished successfully

Setting the pool-level FASTVP Pool Reserved Capacity

The pool-level FAST VP Pool Reserved Capacity (PRC) can be used to override the system-wide setting for each individual pool. At the pool-level, the PRC can be set between 1 and 80 percent, or NONE. The default value is NONE.

The value of NONE indicates that the system-wide setting should be used for the pool.

To set or change the PRC at the pool level, run:

```
symconfigure -sid 0398 -cmd "set pool R53_EFD_Pool, type=thin, pool_resv_cap=1;" commit -v
```

Execute a symconfigure operation for symmetrix '000195700398' (y/[n]) ? y

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

```
Establishing a configuration change session.....Established.
{
  set pool R53_EFD_Pool type=thin, pool_resv_cap=1;
}
```

```
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Committing configuration changes.....Reordering.
Setting pool attributes .....Done.
Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

To set or change the PRC back to NONE, run:

```
symconfigure -sid 0398 -cmd "set pool R53_EFD_Pool, type=thin, pool_resv_cap=NONE;" commit
```

Execute a symconfigure operation for symmetrix '000195700398' (y/[n]) ? y

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

```
Establishing a configuration change session.....Established.
{
  set pool R53_EFD_Pool type=thin, pool_resv_cap=NONE;
}

Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Committing configuration changes.....Reordering.
Setting pool attributes .....Done.
Committing configuration changes.....Committed.
Terminating the configuration change session.....Done.
```

The configuration change session has successfully completed.

Setting the FASTVP Relocation Rate

The FAST VP Relocation Rate can be set to be between 1 and 10. The default value is 5.

To set the Relocation Rate, run:

```
symfast -sid 0398 set -control_parms -vp_reloc_rate 8
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

```
The FAST SET CONTROL PARAMETERS operation finished successfully
```

Setting VP allocation by FAST policy

VP allocation by FAST policy can be set to enabled or disabled. The default value is disabled.

To enable VP allocation by FAST policy, run:

```
symfast -sid 0398 set -control_parms -vp_allocation_by_fp ENABLE
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

```
The FAST SET CONTROL PARAMETERS operation finished successfully
```

To disable VP allocation by FAST policy, run:

```
symfast -sid 0398 set -control_parms -vp_allocation_by_fp DISABLE
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

```
The FAST SET CONTROL PARAMETERS operation finished successfully
```

Modifying all FAST controller settings

Alternatively, all settings can be modified at the same time with the following single command:

```
symfast -sid 0398 set -control_parms -vp_data_move_mode AUTO  
-workload_period 72 -min_perf_period 24 -vp_reloc_rate 8 -pool_resv_cap 5  
-vp_allocation_by_fp ENABLE
```

```
Execute SET CONTROL PARAMETERS operation for FAST (y/[n]) ? y
```

```
The FAST SET CONTROL PARAMETERS operation finished successfully
```

Verifying FAST controller settings

When all changes have been made to the FAST controller settings, the changes can be verified by listing all the control settings:

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

```
Symmetrix ID: 000195700398
```

```
Optimizer and FAST Control Parameters:
```

```
Data Movement Mode           : User_Approve  
Max Simultaneous Device Moves : 8  
Max Device Moves Per Day      : 200
```

```
Optimizer, FAST and FAST VP Control Parameters:
```

```
Min Initial Workload Period(hrs) : 2  
Workload Analysis Period(hrs)    : 72
```

```
FAST Control Parameters:
```



```
Swap Not Visible Devices      : Disabled
Allow Only Swap              : Disabled
```

FAST VP Control Parameters:

```
FAST VP Data Movement Mode   : AUTO
FAST VP Data Relocation Rate  : 8
Thin Pool Reserved Capacity(%) : 5
VP Allocation By FAST policy   : Enabled
```

Legacy time window management

Management of the legacy time windows is managed through the SYMCLI `symoptmz` command.

For the management of FAST VP, it is recommended to use enhanced time window management.

Note: For information on managing legacy time windows, refer to the *Managing Legacy Time Windows* appendix in the *EMC Solutions Enabler Symmetrix Array Controls CLI* product guide available at <http://powerlink.emc.com>.

Converting time window type from legacy to enhanced

On the Symmetrix VMAX Series array, in order to use the enhanced time window definitions, it is necessary to first convert the legacy time windows to the new format. This conversion is performed using the `symtw` command.

The following guidelines apply to convert to the enhanced time window format:

- ◆ Only the time windows defined during the 7-day period, beginning with the start, are converted. All inclusive time windows defined that do not include this period of time are deleted.
- ◆ If the enhanced time window format is defined as ONCE, exclusive time windows previously defined for a future date are converted to the enhanced format.
- ◆ The time windows' start and end time are rounded to the half hour. The start time is rounded down to the previous half hour, and the end time is rounded up to the next half hour.
- ◆ If the legacy time window definitions cannot be translated to the

enhanced format, an error returns indicating that a specific window is either non-convertible or partially convertible. In either case, the -force option can be specified to complete a best-effort conversion.

Note: Following a successful conversion, the symoptmz is no longer supported for the management of time windows.

To perform a conversion of the time windows to the enhanced format, run:

```
symtw -sid 0398 convert -date 04012012
```

```
Execute symtw 'CONVERT' operation (y/[n]) ? y
```

```
Execute conversion of legacy time windows for Symmetrix 000195700398
```

```
Evaluating Performance Time Window 'DEFAULT_PERF_TW': Expired  
Evaluating Performance Time Window 'Negate_Default': Convertible  
Evaluating Performance Time Window 'Production_Day': Convertible  
Evaluating Move DP Time Window 'DEFAULT_SWAP_TW': Expired  
Evaluating Move VP Time Window 'DEFAULT_THIN_SWAP_TW': Expired  
Evaluating Move VP Time Window 'Production_Day': Convertible
```

```
Converting Performance Time Window 'DEFAULT_PERF_TW': Skipped  
Converting Performance Time Window 'Negate_Default': Done  
Converting Performance Time Window 'Production_Day': Done  
Converting Move DP Time Window 'DEFAULT_SWAP_TW': Skipped  
Converting Move VP Time Window 'DEFAULT_THIN_SWAP_TW': Skipped  
Converting Move VP Time Window 'Production_Day': Done
```

```
Conversion of legacy time windows successfully completed
```

Enhanced time window list information

Management of the enhanced time windows is available through the SYMCLI `symtw` command.

To view existing enhanced time windows configured on the Symmetrix (both performance and data movement), the following command can be run:

symtw -sid 0398 list

Symmetrix ID: 000195700398

DP Data Movement Time Windows

Sunday : None
Monday : None
Tuesday : None
Wednesday : None
Thursday : None
Friday : None
Saturday : None

Exclusive Time Windows (0)

VP Data Movement Time Windows

Sunday : 07:00 - 19:00
Monday : 07:00 - 19:00
Tuesday : 07:00 - 19:00
Wednesday : 07:00 - 19:00
Thursday : 07:00 - 19:00
Friday : 07:00 - 19:00
Saturday : 07:00 - 19:00

Exclusive Time Windows (0)

Performance Time Windows

Sunday : 07:00 - 19:00
Monday : 07:00 - 19:00
Tuesday : 07:00 - 19:00
Wednesday : 07:00 - 19:00
Thursday : 07:00 - 19:00
Friday : 07:00 - 19:00
Saturday : 07:00 - 19:00

Exclusive Time Windows (0)

Adding enhanced FASTVP time windows

To add an enhanced time window to be used by FAST VP, the following information needs to be decided:

- ◆ Time window type (performance or data movement)

- ◆ Whether the time windows are inclusive or exclusive
- ◆ Days of the week the window will be used (inclusive only)
- ◆ Times of the day the window will be used (inclusive only)
- ◆ Time period for which the time window is valid (exclusive only)

In this example, to add an inclusive performance time window that collects performance metrics between 5 a.m. and 7 a.m., Monday through Friday, run:

```
symtw -sid 0398 -inclusive -type perf add -days MON,TUE,WED,THU,FRI
-start_time 05:00 -end_time 07:00
```

Execute symtw 'ADD' operation (y/[n]) ? y

The Time Window ADD operation finished successfully

In this example, to add an inclusive VP data movement time window that allows FAST VP data movement between 5 a.m. and 7 a.m., Monday through Friday, run:

```
symtw -sid 0398 -inclusive -type move_vp add -days MON,TUE,WED,THU,FRI
-start_time 05:00 -end_time 07:00
```

Execute symtw 'ADD' operation (y/[n]) ? y

The Time Window ADD operation finished successfully

Note: For inclusive time windows, the start and end time must be specified in 30-minute increments, from 00:00 to 24:00, based on the 24-hour clock format.

The start time and end time cannot be extended to the previous or next day, and the end time must be greater than the start time. Any time windows that extend to the next day need to be defined separately.

In this example, to add an exclusive VP performance time window that prevents FAST VP performance collection from 5 p.m. on April 20, 2012 until 5 a.m. on April 23, 2012, run:

```
symtw -sid 0398 -exclusive -type perf add -start_day 04202012:1700
-end_day 04232012:0500
```

Execute symtw 'ADD' operation (y/[n]) ? y

The Time Window ADD operation finished successfully

In this example, to add an exclusive VP data movement time window that prevents FAST VP data movement from 5 p.m. on April 20, 2012 until 5 a.m. on April 23, 2012, run:

```
symtw -sid 0398 -exclusive -type move_vp add -start_day 04202012:1700  
-end_day 04232012:0500
```

Execute symtw 'ADD' operation (y/[n]) ? y

The Time Window ADD operation finished successfully

Note: For exclusive time windows, the start day/ time and end day/ time are specified in the format MMDDYYYY:HHMM. The start and end times must be specified in 30-minute increments and are based on the 24-hour clock format.

To view the changes, run:

```
symtw -sid 0398 list
```

Symmetrix ID: 000195700398

DP Data Movement Time Windows

```
Sunday      : None  
Monday      : None  
Tuesday     : None  
Wednesday  : None  
Thursday    : None  
Friday      : None
```

```
Saturday    : None
```

Exclusive Time Windows (0)

VP Data Movement Time Windows

```
Sunday      : 07:00 - 19:00  
Monday      : 05:00 - 19:00  
Tuesday     : 05:00 - 19:00
```

```
Wednesday      : 05:00 - 19:00
Thursday       : 05:00 - 19:00
Friday        : 05:00 - 19:00
Saturday      : 07:00 - 19:00
```

Exclusive Time Windows (1)

```
{
  Fri Apr 20 17:00:00 2012 - Mon Apr 23 05:00:00 2012
}
```

Performance Time Windows

```
Sunday        : 07:00 - 19:00
Monday        : 05:00 - 19:00
Tuesday       : 05:00 - 19:00
Wednesday     : 05:00 - 19:00
Thursday      : 05:00 - 19:00
Friday        : 05:00 - 19:00
Saturday      : 07:00 - 19:00
```

Exclusive Time Windows (1)

```
{
  Fri Apr 20 17:00:00 2012 - Mon Apr 23 05:00:00 2012
}
```

Removing enhanced FASTVP time windows

To remove an enhanced time window used by FAST VP, the following information needs to be decided:

- ◆ Time window type (performance or data movement)
- ◆ Whether the time windows are inclusive or exclusive
- ◆ Days of the week the window will be removed from (inclusive only)
- ◆ Times of the day the window will be removed from (inclusive only)
- ◆ Time period for which the time window will be removed (exclusive only)

When removing a time window, the time periods specified may include all or part of existing time windows. Partially removing an existing time

window may generate additional time windows on either side of the removed time period.

In this example, to remove an inclusive performance time window that currently collects performance metrics between 11 a.m. and 1 p.m., Monday and Friday, run:

```
symtw -sid 0398 -inclusive -type perf remove -days MON,FRI -start_time 11:00 -  
end_time 13:00
```

```
Execute symtw 'REMOVE' operation (y/[n]) ? y
```

```
The Time Window REMOVE operation finished successfully
```

To remove an inclusive VP data movement time window that currently allows FAST VP data movement between 5 a.m. and 7 a.m., Monday and Friday, run:

```
symtw -sid 0398 -inclusive -type move_vp remove -days MON,FRI -start_time 11:00 -  
end_time 13:00
```

```
Execute symtw 'REMOVE' operation (y/[n]) ? y
```

```
The Time Window REMOVE operation finished successfully
```

Note: For inclusive time windows, the start and end time must be specified in 30-minute increments, from 00:00 to 24:00, based on the 24-hour clock format.

In this example, to remove an exclusive VP performance time window that is defined to prevent FAST VP performance collection from 5 p.m. on April 21, 2012 until 5 a.m. on April 22, 2012, run:

```
symtw -sid 0398 -exclusive -type perf remove -start_day 04212012:1700  
-end_day 04222012:0500
```

```
Execute symtw 'REMOVE' operation (y/[n]) ? y
```

```
The Time Window REMOVE operation finished successfully
```

In this example, to remove an exclusive VP data movement time window that is defined to prevent FAST VP data movement from 5 p.m. on April 21, 2012 until 5 a.m. on April 22, 2012, run:

```
symtw -sid 0398 -exclusive -type move_vp remove -start_day 04212012:1700
```

-end_day 04222012:0500

Execute symtw 'REMOVE' operation (y/[n]) ? y

The Time Window REMOVE operation finished successfully

Note: For exclusive time windows, the start day/ time and end day/ time are specified in the format MMDDYYYY:HHMM. The start and end times must be specified in 30-minute increments and are based on the 24-hour clock format.

To view the changes, run:

symtw -sid 0398 list

Symmetrix ID: 000195700398

DP Data Movement Time Windows

Sunday	: None
Monday	: None
Tuesday	: None
Wednesday	: None
Thursday	: None
Friday	: None
Saturday	: None

Exclusive Time Windows (0)

VP Data Movement Time Windows

Sunday	: 07:00 - 19:00
Monday	: 05:00 - 11:00 13:00 - 19:00
Tuesday	: 05:00 - 19:00
Wednesday	: 05:00 - 19:00
Thursday	: 05:00 - 19:00
Friday	: 05:00 - 11:00 13:00 - 19:00
Saturday	: 07:00 - 19:00


```
Exclusive Time Windows (2)
{
  Fri Apr 20 17:00:00 2012 - Sat Apr 21 17:00:00 2012
  Sun Apr 22 05:00:00 2012 - Mon Apr 23 05:00:00 2012
}
```

Performance Time Windows

```
Sunday      : 07:00 - 19:00
Monday      : 05:00 - 11:00
             13:00 - 19:00
Tuesday     : 05:00 - 19:00
Wednesday  : 05:00 - 19:00
Thursday   : 05:00 - 19:00
Friday     : 05:00 - 11:00
             13:00 - 19:00
Saturday   : 07:00 - 19:00
```

```
Exclusive Time Windows (2)
{
  Fri Apr 20 17:00:00 2012 - Sat Apr 21 17:00:00 2012
  Sun Apr 22 05:00:00 2012 - Mon Apr 23 05:00:00 2012
}
```

Viewing the FASTVP time window summary

A week time window summary can also be displayed. This summary displays the current week, starting on Sunday, and indicates the actions that are performed in increments of 30 minutes for that week.

The actions displayed indicate when performance data is collected, or not collected, and when data movement occurs for both FAST DP and FAST VP.

It also displays time periods when these actions are explicitly overridden by an exclusion window.

To view the time window summary, run:

```
symtw -sid 0398 list -summary
```

Symmetrix ID: 000195700398

```

                                Time Window Summary
|-----|
| SUN | MON | TUE | WED | THU | FRI | SAT |
|-----|
```

	-----	-----	-----	-----	-----	-----	-----
	D V P	D V P	D V P	D V P	D V P	D V P	D V P
00:00 - 00:30
00:30 - 01:00
01:00 - 01:30
01:30 - 02:00
02:00 - 02:30
02:30 - 03:00
03:00 - 03:30
03:30 - 04:00
04:00 - 04:30
04:30 - 05:00
05:00 - 05:30 V P	. V P	. V P	. V P	. V P	. . .
05:30 - 06:00 V P	. V P	. V P	. V P	. V P	. . .
06:00 - 06:30 V P	. V P	. V P	. V P	. V P	. . .
06:30 - 07:00 V P	. V P	. V P	. V P	. V P	. . .
07:00 - 07:30	. V P	. V P	. V P	. V P	. V P	. V P	. E E
07:30 - 08:00	. V P	. V P	. V P	. V P	. V P	. V P	. E E
08:00 - 08:30	. V P	. V P	. V P	. V P	. V P	. V P	. E E
08:30 - 09:00	. V P	. V P	. V P	. V P	. V P	. V P	. E E
09:00 - 09:30	. V P	. V P	. V P	. V P	. V P	. V P	. E E
09:30 - 10:00	. V P	. V P	. V P	. V P	. V P	. V P	. E E
10:00 - 10:30	. V P	. V P	. V P	. V P	. V P	. V P	. E E
10:30 - 11:00	. V P	. V P	. V P	. V P	. V P	. V P	. E E
11:00 - 11:30	. V P V P	. V P	. V P E E
11:30 - 12:00	. V P V P	. V P	. V P E E
12:00 - 12:30	. V P V P	. V P	. V P E E
12:30 - 13:00	. V P V P	. V P	. V P E E
13:00 - 13:30	. V P	. V P	. V P	. V P	. V P	. V P	. E E
13:30 - 14:00	. V P	. V P	. V P	. V P	. V P	. V P	. E E
14:00 - 14:30	. V P	. V P	. V P	. V P	. V P	. V P	. E E
14:30 - 15:00	. V P	. V P	. V P	. V P	. V P	. V P	. E E
15:00 - 15:30	. V P	. V P	. V P	. V P	. V P	. V P	. E E
15:30 - 16:00	. V P	. V P	. V P	. V P	. V P	. V P	. E E
16:00 - 16:30	. V P	. V P	. V P	. V P	. V P	. V P	. E E
16:30 - 17:00	. V P	. V P	. V P	. V P	. V P	. V P	. E E
17:00 - 17:30	. V P	. V P	. V P	. V P	. V P	. E E	. V P
17:30 - 18:00	. V P	. V P	. V P	->. V P<-	. V P	. E E	. V P
18:00 - 18:30	. V P	. V P	. V P	. V P	. V P	. E E	. V P
18:30 - 19:00	. V P	. V P	. V P	. V P	. V P	. E E	. V P

```

19:00 - 19:30 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
19:30 - 20:00 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
20:00 - 20:30 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
20:30 - 21:00 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
21:00 - 21:30 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
21:30 - 22:00 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
22:00 - 22:30 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
22:30 - 23:00 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
23:00 - 23:30 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
23:30 - 24:00 | . . . | . . . | . . . | . . . | . . . | . . . | . . . |
                |-----|-----|-----|-----|-----|-----|-----|

```

Legend: D = Disk Group Provisioning Movement Time Window
V = Virtual Provisioning Movement Time Window
P = Performance Time Window
E = Time Windows Overridden by the Exclusive Time Windows

Note: The time slot marked with the -> and <- symbols indicates the time slot during which the time window summary was displayed.

To see the summary for a specific week in the future, the `-date` option may also be specified.

Creating FASTmanaged objects

There are three managed objects related to the use of FAST VP in the Symmetrix VMAX Series array. They are:

- ◆ Symmetrix VP tiers
- ◆ FAST policies
- ◆ Storage groups

When created, storage groups can be associated with a FAST policy, which in turn associates the devices in the storage group with up to three VP tiers, while defining the upper usage limit for the storage group in each tier.

The following sections detail the SYMCLI commands used to create each of the managed objects, and the methods for associating them.

Information is also shown for removing these associations, and removing each of the objects.

Creating a Symmetrix VP tier

A Symmetrix VP tier may contain between one and four Virtual

Provisioning thin pools. When creating a VP tier, the following information must be known:

- ◆ The tier name
- ◆ The desired protection type of the tier
- ◆ The drive technology, or location, to be used for the tier
- ◆ The thin pool(s) to be added to the tier

Once this information has been decided, the tier can be created.

Creating an internal tier

To create an internal tier, run:

```
symtier -sid 0398 create -name VP_R1_FC -tgt_raid1 -technology FC -vp
```

This command created an empty Symmetrix VP tier, with target protection type of RAID 1, and disk technology of Fibre Channel (FC). The Symmetrix tier name was chosen to indicate the RAID protection type (RAID 1), the drive type (FC), and the fact that it is a VP tier, VP_R1_FC.

The `symtier list` command can be used to verify the successful creation of the Symmetrix VP tier.

```
symtier -sid 0398 list -vp
```

```
Symmetrix ID          : 000195700398
```

```
-----  
                L           I  
                O   Target   n  
Tier Name        Type C Tech Protection  Emul c  
-----  
VP_R1_FC         VP   I FC   RAID-1     FBA  S
```

Legend:

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

Inc Type : S = Static, D = Dynamic

Once the empty tier has been created, thin pools can be added to it, again using the `symtier` command.

To add a thin pool to the tier, run:

```
symtier -sid 0398 add -tier_name VP_R1_FC -pool R1_FC_Pool
```

To add more than one thin pool to the tier, the following can be run:

```
symtier -sid 0398 add -tier_name VP_R1_FC -pool R1_FC_Pool1,R1_FC_Pool2
```

Once the pool, or pools, has been added, details on the tier can be seen by using the `symtier show` command, as follows:

```
symtier -sid 0398 show -tier_name VP_R1_FC
```

```
Symmetrix ID : 000195700398
```

```
Tier Name : VP_R1_FC
Tier Type : VP
Disk Location : Internal
Technology : FC
Target Protection : RAID-1
Emulation : FBA
Include Type : Static
```

```
Thin Pools(1)
```

```
{
-----
                Logical Capacities (GB)
Pool Name      Dev  ----- Full
                Emul Enabled   Free   Used (%)
-----
R1_FC_Pool     FBA   10723   9438   1285  11
-----
Total                10723   9438   1285
}
```

Legend:

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

Creating an external tier

To create an external tier, run:

```
symtier -sid 0398 create -name VP_R0_FTS -tgt_unprotected -external -vp
```

This command created an empty Symmetrix VP tier, with target protection type of unprotected, with an external location.

Note: External tiers can only have a RAID level of unprotected. Also, drive technology is not specified, only location.

The Symmetrix tier name was chosen to indicate the RAID protection type (RAID 0), the drive type (FTS), and the fact that it is a VP tier, VP_R0_FTS.

The `symtier list` command can be used to verify the successful creation of the Symmetrix VP tier.

```
symtier -sid 0398 list -vp
```

```
Symmetrix ID          : 000195700398
```

```
-----  
                L           I  
                O   Target   n  
Tier Name      Type C Tech Protection Emul c  
-----  
VP_R0_FTS      VP   X N/A  Unprotected N/A S  
VP_R1_FC       VP   I FC   RAID-1      FBA S
```

Legend:

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

Once the empty tier has been created, thin pools can be added to it, again using the `symtier` command.

To add a thin pool to the external tier, run:

```
symtier -sid 0398 add -tier_name VP_R0_FTS -pool R0_FTS_Pool
```

To add more than one thin pool to the tier, the following can be run:

```
symtier -sid 0398 add -tier_name VP_R0_FTS -pool R0_FTS_Pool1,R0_FTS_Pool2
```

Once the pool, or pools, has been added, details on the tier can be seen by using the `symtier show` command, as follows:

```
symtier -sid 0398 show -tier_name VP_R0_FTS
```

```
Symmetrix ID      : 000195700398

Tier Name         : VP_R0_FTS
Tier Type         : VP
Disk Location     : External
Technology        : N/A
Target Protection : Unprotected
Emulation         : FBA
Include Type      : Static

Thin Pools(1)
{
-----
                Logical Capacities (GB)
Pool Name      Dev  ----- Full
                Emul Enabled   Free   Used  (%)
-----
R0_FTS_Pool    FBA      195    195    0    0
-----
Total          195    195    0
}

```

Legend:

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

Symmetrix tier list information

After additional VP tiers have been created, information on all the VP tiers in the Symmetrix can be viewed by running:

```
symtier -sid 0398 list -vp
```

```
Symmetrix ID      : 000195700398

-----
                L                               I
                O   Target                     n
Tier Name         Type C Tech Protection   Emul c

```

```

-----
VP_R0_FTS          VP  X N/A  Unprotected  FBA  S
VP_R1_FC           VP  I FC   RAID-1       FBA  S
VP_R53_EFD        VP  I EFD  RAID-5(3+1)  FBA  S
VP_R614_SATA      VP  I SATA RAID-6(14+2) FBA  S

```

Legend:

```

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

```

More verbose details, including the pools contained in each tier, can be viewed by running:

syntier -sid 0398 list -vp -v

Symmetrix ID : 000195700398

```

Tier Name          : VP_R0_FTS
Tier Type          : VP
Disk Location      : External
Technology         : N/A
Target Protection  : Unprotected
Emulation          : FBA
Include Type       : Static

```

Thin Pools(1)

```

{
-----
                Logical Capacities (GB)
Pool Name      Dev  ----- Full
                Emul  Enabled   Free    Used   (%)
-----
R0_FTS_Pool    FBA          195    195     0     0
-----
Total          195    195     0
}

```

```

Tier Name          : VP_R1_FC
Tier Type          : VP
Disk Location      : Internal

```



```

Technology          : FC
Target Protection   : RAID-1
Emulation           : FBA
Include Type        : Static

```

```
Thin Pools(1)
```

```

{
-----
                Logical Capacities (GB)
Pool Name      Dev  ----- Full
                Emul  Enabled   Free   Used  (%)
-----
R1_FC_Pool     FBA      10723   9438   1285   11
-----
Total          10723   9438   1285
}

```

```

Tier Name          : VP_R53_EFD
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  ----- Full
                Emul  Enabled   Free   Used  (%)
-----
R53_EFD_Pool   FBA      4395   4394     1     0
-----
Total          4395   4394     1
}

```

```

Tier Name          : VP_R614_SATA
Tier Type          : VP
Disk Location      : Internal
Technology         : SATA
Target Protection  : RAID-6(14+2)
Emulation          : FBA
Include Type       : Static

```

```
Thin Pools(1)
```

```

{
-----
          Logical Capacities (GB)
Pool Name   Dev  ----- Full
            Emul  Enabled   Free   Used  (%)
-----
R6_SATA_Pool FBA      24539   23877   662    2
-----
Total                24539   23877   662
}

```

Legend:

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

The Free column indicates the sum of the unallocated space of all the thin pools included in the tier.

The Used column displays the total pool-allocated capacity for all thin pools in the tier.

The Enabled column is total enabled capacity for all thin pools in the tier.

Note: For the Used values, all allocated capacity is counted, including capacity allocated on data devices that may be disabled. In this case, it is possible for the Used value to be greater than the Enabled value.

Creating a Symmetrix FAST policy

A FAST policy defines the upper usage limits for up to three tiers for any storage groups associated with the policy.

When creating a FAST policy, the following information must be known:

- ◆ The FAST policy name
- ◆ The VP tiers (maximum three) to be used in the policy
- ◆ The upper usage limits for each of the VP tiers being added

Once this information has been decided, the FAST policy can be created by running:

```
symfast -sid 0398 -fp create -name System_Optimization
```

Symmetrix tiers can then be added, one at a time, by running:

```
symfast -sid 0398 -fp -fp_name System_Optimization add -tier_name VP_R1_FC
symfast -sid 0398 -fp -fp_name System_Optimization add -tier_name VP_R53_EFD
-max_sg_percent 100
symfast -sid 0398 -fp -fp_name System_Optimization add -tier_name
VP_R614_SATA -max_sg_percent 100
```

Note: The `-max_sg_percentage` parameter is optional. If not included in the command, the percentage value defaults to 100 percent.

The creation of the FAST policy can be verified by running:

```
symfast -sid 0398 show -fp_name System_Optimization
```

```
Symmetrix ID          : 000195700398
```

```
Policy Name           : System_Optimization
```

```
Emulation              : FBA
```

```
Tiers(3)
```

```
{
```

```
-----
                                     L
                               Max SG O
Tier Name                        Type Percent C Tech Target
-----
VP_R53_EFD                       VP      100 I EFD  RAID-5(3+1)
VP_R1_FC                          VP      100 I FC   RAID-1
VP_R614_SATA                      VP      100 I SATA RAID-6(14+2)
}
```

```
No Storage Groups associated with Policy System_Optimization
```

Legend:

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

Disk (Location) : I = Internal, X = External

FAST policy list information

After additional policies have been created, information on all FAST

policies in the Symmetrix array can be viewed by running:

```
symfast -sid 0398 list -fp
```

```
Symmetrix ID          : 000195700398
```

```
-----  
Policy Name           Tiers Assocs  
-----  
Custom                3         0  
No_EFD                2         0  
System_Optimization  3         0
```

Creating a Symmetrix storage group

A storage group logically combines Symmetrix devices to be managed together. The creation and management of storage groups are performed using the `symsg` command.

When creating a storage group, the following information must be known:

- ◆ The storage group name
- ◆ The Symmetrix devices to be added to the group

Once this information has been decided, the storage group can be created by running:

```
symsg -sid 0398 create VP_ProdApp1
```

Devices can be added to the storage group in a number of ways:

- ◆ One device at a time

```
symsg -sid 0398 -sg VP_ProdApp1 add dev 2B0
```

- ◆ A contiguous range of devices

```
symsg -sid 0398 -sg VP_ProdApp1 addall devs -range 2B0:2E0
```

- ◆ A comma-separated list of devices

```
symmsg -sid 0398 -sg VP_ProdApp1 addall devs -devs 2B0,2C0,2D0,2E0
```

- ◆ A comma-separated list of devices and device ranges

```
symmsg -sid 0398 -sg VP_ProdApp1 addall devs -devs 2B0,2C0:2E0
```

- ◆ A device file

```
symmsg -sid 0398 -sg VP_ProdApp1 addall devs -file mydevices.txt
```

To verify the successful creation of the storage group, and the addition of the devices, run:

```
symmsg -sid 0398 show VP_ProdApp1
```

```
Name: VP_ProdApp1
```

```
Symmetrix ID           : 000195700398
Last updated at       : Mon Apr 16 15:53:40 2012
Masking Views         : No
FAST Policy           : No
Number of Storage Groups : 0
Storage Group Names   : N/A
```

```
Devices (4):
```

```
{
-----
Sym          Device          Cap
Dev   Pdev Name   Config      Sts   (MB)
-----
02B0   N/A          TDEV              RW  276210
02C0   N/A          TDEV              RW  276210
02D0   N/A          TDEV              RW  276210
02E0   N/A          TDEV              RW  276210
}
```

Storage group list information

The `symmsg list` command is available with several options to provide information for storage groups, including:

- ◆ The names of all created storage groups

- ◆ The number of devices in each storage group
- ◆ Whether the storage group is associated with a FAST policy
- ◆ Whether the storage group is used in a masking view

symsg -sid 0398 list

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

Symmetrix ID: 000195700398

Storage Group Name	Number Devices	Child SGs	Flags FMS
licoa036	6	0	.X.
SBO_CLUSTER	26	0	.X.
VP_Development	2	0	.X.
VP_ProdApp1	4	0	.X.
VP_ProdApp2	4	0	.X.

Legend:

Flags:

(F)ast Policy, X = Associated with Fast Policy, . = N/A
 (M)asking View, X = Contained in Mask View(s), . = N/A
 Cascade (S)tatus, P = Parent SG, C = Child SG, . = N/A

Associating a storage group with a FASTVP policy

Associating a storage group with a FAST VP policy brings the devices in the storage group under FAST VP control. All devices in the storage group are considered candidates to have data moved between the tiers included in the policy the storage group is associated with.

When creating a FAST policy association, the following information must be known:

- ◆ The FAST policy name
- ◆ The storage group name
- ◆ The priority of the storage group within the policy

Once this information has been decided, the association is performed by running:

```
symfast -sid 0398 -fp_name System_Optimization associate -sg VP_ProdApp1 -priority 2
```

To verify the successful association of the storage group to the FAST policy, run:

```
symfast -sid 0398 show -association -sg VP_ProdApp1
```

Symmetrix ID : 000195700398

Storage Group : VP_ProdApp1

Thin Devices (4)

```
{
-----
      Flgs Dev      Total Bound      Allocated
Sym   P   Emul      Tracks Pool Name      Tracks
-----
02B0 N   FBA      4419360 R1_FC_Pool      1326816
02C0 N   FBA      4419360 R1_FC_Pool      1326816
02D0 N   FBA      4419360 R1_FC_Pool      1326816
02E0 N   FBA      4419360 R1_FC_Pool      1326816

Total          -----
Tracks          17677440          5307264
      GBs          1079          324
}
```

Policy Name : System_Optimization

Priority : 2

RDF Coordination : Disabled

Tiers (3)

```
{
-----
                                          L
                                          Max SG O      Target
Tier Name          Type  Percent C Tech  Protection
-----
VP_R53_EFD          VP      100 I EFD  RAID-5(3+1)
VP_R1_FC            VP      100 I FC   RAID-1
VP_R614_SATA        VP      100 I SATA RAID-6(14+2)
```

```
}
```

Legend:

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

Device Flags:

(P)inned : Y = Device is Pinned, N = Device is not Pinned

Tier Flags:

Disk (Loc)ation : I = Internal, X = External

When additional associations have been made, all storage groups associations can be viewed by running:

```
symfast -sid 0398 list -association
```

```
Symmetrix ID : 000195700398
```

```
-----  
Storage Group Name          Policy Name                Pri Flgs  
                             R  
-----  
VP_Development              No_EFD                     2 .  
VP_ProdApp1                  System_Optimization        2 .  
VP_ProdApp2                  Custom                      2 .
```

Legend:

Flgs:

(R)DF Coordination : X = Enabled, . = Disabled

Modifying a storage group's priority in a FAST policy

After a storage group has been associated with a FAST policy, it is possible to modify the priority of the storage group within the policy using the `symfast` command.

To modify the storage group's priority in the FAST policy, run the following command:

```
symfast -sid 0398 -fp_name System_Optimization modify -sg VP_ProdApp1 -  
priority 1
```


To verify the priority of the storage group was changed successfully,
run:

```
symfast -sid 0398 list -association
```

```
Symmetrix ID          : 000195700398
```

```
-----  
Storage Group Name      Policy Name              Pri Flgs  
                        R  
-----  
VP_Development         No_EFD                  2 .  
VP_ProdApp1            System_Optimization     1 .  
VP_ProdApp2            Custom                  2 .
```

Legend:

Flgs:

(R)DF Coordination : X = Enabled, . =

Enabling/disabling SRDF coordination

By default, SRDF coordination is disabled for any storage group associated with a FAST policy. SRDF coordination may be enabled on a storage group while it is being associated with a policy, or it may be enabled after the group has been associated.

To enable SRDF coordination during association, run:

```
symfast -sid 0398 -fp_name Custom associate -sg VP_ProdApp2  
-priority 2 -rdf_coordination ENABLE
```

To enable SRDF coordination after the association, run:

```
symfast -sid 0398 -fp_name Custom modify -sg VP_ProdApp2  
-rdf_coordination ENABLE
```

To verify that SRDF coordination has been enabled, run:

```
symfast -sid 0398 list -association
```

```
Symmetrix ID          : 000195700398
```

```
-----  
Storage Group Name      Policy Name              Pri Flgs  
                        R  
-----
```

```

VP_Development          No_EFD          2      .
VP_ProdApp1            System_Optimization  1      .
VP_ProdApp2            Custom          2      X

```

Legend:

Flgs:

(R)DF Coordination : X = Enabled, . =

To disable SRDF coordination on a storage group associated with a policy, run:

```

symfast -sid 0398 -fp_name Custom modify -sg VP_ProdApp2
-rdf_coordination DISABLE

```

Reassociating a storage group to a different FASTVP policy

To move a storage group from one FAST policy to another, run:

```

symfast -sid 0398 -sg VP_ProdApp2 reassociate -fp_name System_Optimization

```

To verify the successful reassociation of the storage group to the FAST policy, run:

```

symfast -sid 0398 show -association -sg VP_ProdApp2

```

```

Symmetrix ID          : 000195700398

```

```

Storage Group         : VP_ProdApp2

```

Thin Devices(4)

{

```

-----
      Flgs Dev      Total Bound      Allocated
Sym   P   Emul      Tracks Pool Name      Tracks
-----
02F0  N   FBA      4419360 R1_FC_Pool      1769508
0300  N   FBA      4419360 R1_FC_Pool      1769508

```

```

0310 N FBA 4419360 R1_FC_Pool 1768080
0320 N FBA 4419360 R1_FC_Pool 1768092

Total -----
Tracks 17677440 7075188
GBs 1079 432
}

```

```

Policy Name : System_Optimization
Priority : 2
RDF Coordination : Enabled

```

Tiers(3)

```

{
-----
L
Max SG O Target
Tier Name Type Percent C Tech Protection
-----
VP_R53_EFD VP 100 I EFD RAID-5(3+1)
VP_R1_FC VP 100 I FC RAID-1
VP_R614_SATA VP 100 I SATA RAID-6(14+2)
}

```

Legend:

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

Device Flags:

(P)inned : Y = Device is Pinned, N = Device is not Pinned

Tier Flags:

Disk (Loc)ation : I = Internal, X = External

Note: During the reassociation, any attributes of the association, such as priority or SRDF coordination, that have been changed from the default are maintained with the new policy.

Modifying a Symmetrix tier in a FAST policy

After some time, it may be determined that the upper usage limit of a particular tier within a FAST policy needs to be adjusted. This can be done dynamically through the `symfast` command.

If any storage groups are associated with the policy being modified, the change in the usage limit cannot cause the sum of the usage limits for all tiers in the policy to fall below 100 percent. In this case, the Custom

policy is modified. The current tier percentages can be viewed by running:

```
symfast -sid 0398 show -fp_name Custom
```

```
Symmetrix ID      : 000195700398
```

```
Policy Name      : Custom
```

```
Emulation        : FBA
```

```
Tiers (3)
```

```
{
```

```
-----  
                                     L  
                               Max SG O      Target  
Tier Name                        Type  Percent C Tech  Protection  
-----  
VP_R53_EFD                        VP      10 I EFD  RAID-5 (3+1)  
VP_R1_FC                           VP      20 I FC   RAID-1  
VP_R614_SATA                       VP      100 I SATA RAID-6 (14+2)  
}
```

```
Storage Groups (1)
```

```
{
```

```
-----  
Storage Group Name                Pri  
-----  
VP_ProdApp2                        2  
}
```

```
Legend:
```

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

```
Disk (Loc)ation : I = Internal, X = External
```

To modify the usage limit for a Symmetrix tier within a policy, run:

```
symfast -sid 0398 -fp -fp_name Custom modify -tier_name VP_R1_FC  
-max_sg_percent 50
```

In this case, the percentage of capacity allowed on the V_R1_FC tier was increased from 20 percent to 50 percent.

To verify the change was successful, run:

```
symfast -sid 0398 show -fp_name Custom
```

```
Symmetrix ID      : 000195700398
```

```
Policy Name      : Custom
```

```
Emulation        : FBA
```

```
Tiers(3)
```

```
{
-----
                                L
                                Max SG O
Tier Name                      Type Percent C Tech Target
-----
VP_R53_EFD                     VP      10 I EFD  RAID-5(3+1)
VP_R1_FC                       VP      50 I FC   RAID-1
VP_R614_SATA                   VP     100 I SATA RAID-6(14+2)
}
```

```
Storage Groups(1)
```

```
{
-----
Storage Group Name             Pri
-----
VP_ProdApp2                   2
}
```

Legend:

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

Disk (Loc)ation : I = Internal, X = External

Controlling FAST device movement

Aside from using the FAST controller device movement window, there are several other ways of controlling when device movements can take place. These include:

- ◆ Disabling the FAST controller
- ◆ Pinning devices under FAST VP control
- ◆ Changing the data movement mode

- ◆ Modifying data movement windows

Monitoring FASTVP status

The current status and current activity of FAST VP can be monitored using SYMCLI by means of the FAST status report. Information provided by this report includes:

- ◆ The FAST VP state
- ◆ Degraded reason
- ◆ FAST VP Current Activities
- ◆ FAST VP Performance Time Window state
- ◆ FAST VP Move Time Window state

To display the FAST status report, run:

```
symfast -sid 0398 list -state -vp
```

```
Symmetrix ID: 000195700398
```

```
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
```

The FAST VP state is typically Enabled, Disabled, or Degraded. If the state is degraded, a reason is listed.

Note: For more information on possible state values and degraded reason codes, see “Appendix A: FAST VP state” on page 180.

FAST VP current activities are displayed as either Idle or Running Plan. An activity of Idle indicates that there are currently no active data movement tasks related to FAST VP in the array. Running Plan indicates that data is actively being moved at the sub-LUN level for thin devices.

The state of the performance and move time windows are displayed as Open or Closed. Windows are considered open if the current time corresponds to a defined inclusion time window. Similarly, the windows are considered closed if the current time falls outside of any defined inclusion windows.

Note: The FAST VP move time window is displayed as closed if the data movement mode is off, regardless of any defined inclusion time windows.

Enabling/disabling FASTVP

In order for FAST VP to perform device movements, it must first be enabled on the Symmetrix array.

The state of the controller can be changed using the `symfast` command.

If disabled, FAST VP can be enabled by running:

```
symfast -sid 0398 enable -vp
```

```
Execute ENABLE operation for FAST (y/[n]) ? y
```

```
The FAST ENABLE operation finished successfully
```

If enabled, the FAST controller can be disabled by running:

```
symfast -sid 0398 disable -vp
```

```
Execute DISABLE operation for FAST (y/[n]) ? y
```

```
The FAST DISABLE operation finished successfully
```

Pinning a thin device under FASTVP control

Pinning a device that is associated with a FAST VP policy prevents any data movement for that device. Devices can be pinned individually, by device range, or device file, using the `symdev` command. Devices may also be pinned by storage group or device group using the `symsg` and `symsg` commands, respectively.

To pin an individual device, run:

```
symdev -sid 0398 pin 2B0
```

Execute a 'Pin' Device operation for device '2B0' (y/[n]) ? y

'Pin' Device operation successfully completed for the device.

To pin a range of devices, run:

```
symdev -sid 0398 pin -range 2B0:2E0
```

Execute a 'Pin' operation for devices in range '2b0' to '2e0' (y/[n]) ? y

'Pin' operation succeeded for devices in RANGE 2B0:2E0.

To pin all the devices in a device file, run:

```
symdev -sid 0398 pin -file my_devices.txt
```

Execute a 'Pin' operation for devices in file 'my_devices.txt' (y/[n]) ? y

'Pin' operation succeeded for devices in file 'my_devices.txt'.

To pin all devices in a storage group, run:

```
symsg -sid 0398 -sg VP_ProdApp1 pin
```

Execute a 'Pin' Device operation for all devices
in storage group 'VP_ProdApp1' (y/[n]) ? y

'Pin' Device operation successfully completed for the storage group.

To pin all devices in a device group, run:

```
symdg -g VP_ProdApp1 pin
```

Execute a 'Pin' Device operation for all devices
in device group 'VP_ProdApp1' (y/[n]) ? y

'Pin' Device operation successfully completed for the device group.

Unpinning a thin device under FASTVP control

After a device has been pinned, in order for FAST VP to resume data movements on the device it must be unpinned.

Devices can be unpinned individually, by device range, or device file using the `symdev` command. Devices may also be unpinned by storage group or device group using the `symsg` and `symsg` commands, respectively.

To unpin an individual device, run:

```
symdev -sid 0398 unpin 2B0
```

```
Execute a 'Unpin' Device operation for device '2B0' (y/[n]) ? y
```

```
'Unpin' Device operation successfully completed for the device.
```

To unpin a range of devices, run:

```
symdev -sid 0398 unpin -range 2B0:2E0
```

```
Execute a 'Unpin' operation for devices in range '2b0' to '2e0' (y/[n]) ? y
```

```
'Unpin' operation succeeded for devices in RANGE 2B0:2E0.
```

To unpin all the devices in a device file, run:

```
symdev -sid 0398 unpin -file my_devices.txt
```

```
Execute a 'Unpin' operation for devices in file 'my_devices.txt' (y/[n]) ? y
```

```
'Unpin' operation succeeded for devices in file 'my_devices.txt'.
```

To unpin all devices in a storage group, run:

```
symsg -sid 0398 unpin -sg VP_ProdApp1
```

```
Execute a 'Unpin' Device operation for all devices  
in storage group 'VP_ProdApp1' (y/[n]) ? y
```

```
'Unpin' Device operation successfully completed for the storage group.
```

To unpin all devices in a device group, run:

```
symdmg -g VP_ProdApp1 unpin
```

```
Execute a 'Unpin' Device operation for all devices  
in device group 'VP_ProdApp1' (y/[n]) ? y
```

```
'Unpin' Device operation successfully completed for the device group.
```

Modifying the copy pace of a FAST device movement

As previously stated, Symmetrix Quality of Service (QoS) tools can be used to control the pace at which data is copied for individual devices during FAST VP data movements. This is in addition to the system-wide relocation rate setting that applies to all thin devices associated with a FAST VP policy. The QoS setting that needs to be changed is the VLUN copy pace.

To change this setting, the `symqos` command can be run against a device group or a storage group.

To change the VLUN QoS setting on all devices in a device group, run:

```
symqos -g VP_ProdApp1 set VLUN pace 8
```

To verify the change was successful, run:

```
symqos -g VP_ProdApp1 list
```

```
Device Group (DG) Name: VP_ProdApp1  
DG's Type           : REGULAR  
DG's Symmetrix ID   : 000195700398
```

Device Name		Copy Pace					
Sym	Physical	Config	BCV	RDF	MIR	CLN	VLN
2B0	Not Visible	TDEV	0	0	0	0	8
2B1	Not Visible	TDEV	0	0	0	0	8
2B2	Not Visible	TDEV	0	0	0	0	8
...							

2ED Not Visible	TDEV	0	0	0	0	8
2EE Not Visible	TDEV	0	0	0	0	8
2EF Not Visible	TDEV	0	0	0	0	8

Rebinding a thin device

While FAST VP moves data at the sub-LUN level between device pools, the thin devices associated with a policy still remain bound to a single pool. Any new allocations that occur as the result of a host write come from the bound pool.

The thin device rebind feature allows the binding information for a thin device to be changed, without changing the current allocation of data across pools.

Note: If the devices being rebound are associated with a FAST VP Policy, only pools that are contained within the policy can be specified as the new bind pool.

In this example, to rebind thin device 2B0 to the RAID 6 protected SATA thin pool, R6_SATA_Pool, run:

```
symconfigure -sid 0398 -cmd "rebind tdev 90D to pool R6_SATA_Pool;" commit
```

In this example, to rebind a range of thin devices, 2B0 through 2E0, to the R6_SATA_Pool thin pool, run:

```
symconfigure -sid 0398 -cmd "rebind tdev 90D:925 to pool R6_SATA_Pool;" commit
```

In this example, to rebind all the thin devices in the VP_ProdApp1 storage group to the R6_SATA_Pool thin pool, run:

```
symconfigure -sid 0398 -cmd "rebind tdev in SG VP_ProdApp1 to pool R6_SATA_Pool;" commit
```

In this example, to rebind all the thin devices in the VP_ProdApp1 device group to the R6_SATA_Pool thin pool, run:

```
symconfigure -sid 0398 -cmd "rebind tdev in DG VP_ProdApp1 to pool R6_SATA_Pool;" commit
```

Removing FAST managed objects

The following sections detail the SYMCLI commands used to break each

of the associations, and to remove each of the managed objects.

Disassociating a storage group from a FAST policy

Disassociating a storage group from a FAST VP policy removes the devices in the storage group from FAST VP control. Devices no longer under FAST VP control will not have data automatically moved between tiers.

To disassociate a storage group from a FAST policy, run:

```
symfast -sid 0398 -fp_name System_Optimization disassociate -sg VP_ProdApp1
```

Removing a Symmetrix tier from a FAST policy

As previously stated, the sum of the upper usage limits for each Symmetrix tier contained in a policy must total a minimum of 100 percent. If a policy has associated storage groups, the removal of a tier will fail if in doing so causes the total tier usage limits to drop below 100 percent.

To successfully remove a Symmetrix tier from a FAST policy, the upper usage limits of the remaining tiers must be modified in order that they total more than 100 percent. Alternatively, the Symmetrix tier can be removed if all storage group associations for the policy are removed in advance.

To remove a Symmetrix tier from a FAST policy, run:

```
symfast -sid 0398 -fp -fp_name System_Optimization remove -tier_name VP_R1_FC
```

Deleting a FAST policy

A FAST policy may not be deleted if any storage groups are associated with the policy. Prior to deleting the policy, all storage group associations should be removed. Also, prior to deleting a policy, all tiers contained in the policy should be removed.

To delete a FAST policy, run:

```
symfast -sid 0398 delete -fp -fp_name System_Optimization
```

Deleting a Symmetrix VP tier

If you want to delete a VP tier, it cannot be included within any FAST policy on the Symmetrix array.

To delete a VP tier, run:

```
symtier -sid 0398 delete -tier_name VP_R53_EFD
```

Deleting a storage group

Before deleting a storage group, the group should not be associated with any FAST VP policy.

To delete a storage group, run:

```
symmsg -sid 0398 delete VP_ProdApp1
```

A storage group that contains devices may be deleted by adding the `-force` option to the `symmsg delete` command.

Note: Since storage groups may also be used for the purposes of FAST and Auto-provisioning, prior to deleting the storage group, you must delete any masking views that use the storage group.

Reporting

Both Solutions Enabler and Unisphere for VMAX provide two types of capacity utilization reports to assist in monitoring the FAST VP environment: Compliance reports and technology demand reports.

The compliance report displays tier utilization for all storage groups associated with FAST policies, DP, or VP. The report displays compliance for each storage group individually. The compliance report does not account for competition among storage groups sharing a policy, or capacity consumed by devices not under FAST control.

The technology demand report provides detail on thin capacity utilization and demand from the point of view of the technology type and tier definitions.

There are also reports that show the allocation of thin devices across multiple pools, due to FAST VP data movements, as well as the utilization of capacity within a thin pool for devices not bound to that pool.

The following sections discuss each of these reports.

Compliance reporting

As previously stated, a storage group is considered to be compliant with the FAST policy it is associated with when all the devices in the storage group are fully configured within the bounds of the upper usage limits for each tier contained with the policy.

The information contained in the compliance report includes:

- ◆ FAST policy name
- ◆ Associated storage group name
- ◆ Associated Symmetrix tiers
- ◆ Storage group usage of associated Symmetrix tiers

Each compliance report output contains the following values for each tier included in the policy:

- ◆ **Type:** Indicates whether the tiers included in the policy are disk group provisioning (DP) or Virtual Provisioning (VP). All tiers in a FAST policy can only be of a single type.
- ◆ **Target Prot:** Displays the RAID protection associated with the tier, based on the thin pools that are included in the tier.
- ◆ **Max SG Percent:** Indicates the maximum usage limit, as a percentage, of the storage group's configured capacity per tier. This upper usage limit is as defined in the FAST policy.
- ◆ **Max SG Demand:** Shows the calculated upper limit, in GB, for the storage group on the tier, based on the configured capacity of the storage group.
- ◆ **FAST SG Usage:** Shows the current occupancy of the storage group's allocated capacity in a Symmetrix tier.
- ◆ **Growth:** Indicates how much additional capacity of the storage group can be added to that tier, as determined by the FAST policy.

Note: If the growth value is negative, the storage group has exceeded the capacity limit for the tier, and is considered to be non-compliant.

SYMCLI

The `symfast` command can be run with certain options to determine current capacity usage by storage groups of the Symmetrix tiers contained within their associated FAST policies. The report indicates the compliance of the storage group within the FAST policy, or any non-compliance, including storage that is out of policy.

To list the association, and storage demand, of all storage groups under FAST VP control, the following command can be run:

```
symfast -sid 0398 list -association -demand
```

```
Symmetrix Id      : 000195700398
```

```
Policy Name       : Custom
Storage Group     : VP_ProdApp2
Priority          : 2
RDF Coordination  : Enabled
Tiers (3)
```

```
{
-----
                                     Logical Capacities (GB)
                                     -----
                                     Ty      Max
                                     p Target  SG
                                     e Prot   ()
Name                               -----
-----
VP_R53_EFD                        VP R5 (3+1)   10      108      73      +35
VP_R1_FC                           VP R1         50      540     201     +339
VP_R614_SATA                       VP R6 (14+2) 100     1079    300     +779

      Total                               -----
                                     1727      574
-----
}
```

```
Policy Name       : No_EFD
Storage Group     : VP_Development
Priority          : 2
RDF Coordination  : Disabled
Tiers (2)
```

```
{
-----
                                     Logical Capacities (GB)
                                     -----
```

Name	Type	Target	Max SG	Max SG Demand	FAST SG Usage	Growth
VP_R1_FC	VP R1		50	270	230	+40
VP_R614_SATA	VP R6 (14+2)		100	539	169	+370
Total				809	399	

Policy Name : System_Optimization
Storage Group : VP_ProdApp1
Priority : 2
RDF Coordination : Disabled
Tiers (3)

{

Logical Capacities (GB)						
Name	Type	Target	Max SG	Max SG Demand	FAST SG Usage	Growth
VP_R53_EFD	VP R5 (3+1)		100	1079	108	+971
VP_R1_FC	VP R1		100	1079	191	+888
VP_R614_SATA	VP R6 (14+2)		100	1079	328	+751
Total				3237	627	

}

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

The compliance report can also be run for an individual storage group by running:

symfast -sid 0398 list -association -demand -sg VP_ProdApp1

Symmetrix Id : 000195700398


```

Policy Name      : System_Optimization
Storage Group   : VP_ProdApp1
Priority        : 2
RDF Coordination : Disabled
Tiers (3)
{

```

Name	Type Target Prot	Max SG ()	Logical Capacities (GB)			Growth
			Max SG Demand	FAST SG Usage		
VP_R53_EFD	VP R5 (3+1)	100	1079	108	+971	
VP_R1_FC	VP R1	100	1079	192	+887	
VP_R614_SATA	VP R6 (14+2)	100	1079	327	+752	
Total			3237	627		

Legend:

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

To see the compliance report for all storage groups associated with a particular FAST policy, run:

```
symfast -sid 0398 list -association -demand -fp_name Custom
```

```
Symmetrix Id : 000195700398
```

```

Policy Name      : Custom
Storage Group   : VP_ProdApp2
Priority        : 2
RDF Coordination : Enabled
Tiers (3)
{

```

Name	Type Target Prot	Max SG ()	Logical Capacities (GB)			Growth
			Max SG Demand	FAST SG Usage		
VP_R53_EFD	VP R5 (3+1)	10	108	61	+47	
VP_R1_FC	VP R1	50	540	193	+347	
VP_R614_SATA	VP R6 (14+2)	100	1079	320	+759	

```

Total -----
                1727      574
}

```

Legend:

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

Unisphere for VMAX

Similar reports can be viewed in Unisphere for each individual storage group. When looking at the properties for a storage group, the FAST Compliance report is shown in the lower half of the window.

FAST Compliance Report

Tier	Protection	Technology	Max SG Demand (%)	Limit (GB)	Fast SG Used (GB)	Growth (GB)
VP_R53_EFD	RAID-5 (3+1)	EFD	10	+107.89	+60.62	+47.28
VP_R1_FC	RAID-1	FC	50	+539.47	+192.39	+347.08
VP_R614_SATA	RAID-6 (14+2)	SATA	100	+1078.95	+320.68	+758.26

Technology demand reporting

Technology demand reporting exists to display capacity usage from the perspective of the disk technologies and defined Symmetrix tiers.

The technology demand report can be used to detect possible contention for tier resources among the thin devices under FAST VP control.

The information contained in the technology demand report for FAST VP, includes:

- ◆ Tier: Shows the names of all Symmetrix tiers that have been created on the technology type.
- ◆ Attr: Shows the status of the Symmetrix tier on the technology type. Possible values are:
 - F: In a FAST policy associated with a storage group
 - P: In a FAST policy, or policies, but none of the FAST policies are associated with a storage group
 - N: Not in any FAST policy
- ◆ Target Prot: Indicates the target protection of the tier.
- ◆ Logical Tier Enabled: Indicates the total capacity for all enabled data devices in the thin pools included in the tier.
- ◆ Logical Tier Free: Displays the amount of unallocated space available in the tier. This is calculated as the difference between the logical tier enabled value and the logical tier used value, with a minimum value of 0 (zero).
- ◆ Logical Tier Used: Shows the total capacity allocated on all data devices in the thin pools included in the tier.

Note: The allocated capacity in a thin pool may include tracks allocated on data devices that are not currently enabled (inactive or draining). As a result, the logical tier used capacity may actually be greater than the logical tier enabled capacity.

- ◆ Logical FAST SG Usage — This is the sum of the allocated capacity residing on the VP tier from the thin devices in all storage groups associated with FAST policies containing that tier.
- ◆ Logical FAST Available — Indicates the maximum logical capacity available in the tier to FAST VP. This is calculated as the sum of the logical SG usage and the logical tier free, minus the capacity reserved by the PRC value.
- ◆ Logical Max SG Demand — Displays the total amount of capacity

required for all thin devices in storage groups associated with policies containing the tier to occupy the maximum allotted quota, based on the limit defined on the FAST policy.

Note: If `-allocated` is specified, the Max SG Demand is calculated using allocated instead of configured capacity.

- ◆ Logical Excess: Displays the difference between the available value and the max SG demand value. A positive value indicates that there is sufficient capacity to satisfy the maximum demand. A negative value indicates that there is contention among storage groups for capacity in that tier.

SYMCLI

The `symfast` command can be run with certain options to determine the current FAST VP technology demands within a Symmetrix array.

To list the thin demand for all technologies, run:

Note: The `-vp` option is required to see information pertinent to FAST VP. Leaving out `-vp` option shows information related to FAST DP.

```
symfast -sid 0398 list -demand -tech ALL -vp
```

```
Symmetrix ID : 000195700398
```

```
Technology      : N/A
Disk Location   : External
```

```
VP Tiers (1)
```

```
{
-----
          A                      Logical Capacities (GB)
          T
Tier      T Target      Tier   Tier   Tier FAST SG   FAST Max SG Excess
          R Prot      Enabled Free   Used   Usage Avail Demand
-----
VP_R0_FTS N Unprot      195   195   0     0     0     -     -
Total
          195   195   0     0     0     0     0
}
```

Technology : EFD
 Disk Location : Internal

VP Tiers (1)

```
{
-----
          A                               Logical Capacities (GB)
          T                               -----
Tier      T Target      Tier      Tier      Tier FAST SG   FAST Max SG Excess
          R Prot        Enabled   Free     Used   Usage  Avail Demand
-----
VP_R53_EFD F R5(3+1)    4395    4094    301    300    4350    1187 +3163
Total
          4395    4094    301    300    4350    1187 +3163
}
```

Technology : FC
 Disk Location : Internal

VP Tiers (1)

```
{
-----
          A                               Logical Capacities (GB)
          T                               -----
Tier      T Target      Tier      Tier      Tier FAST SG   FAST Max SG Excess
          R Prot        Enabled   Free     Used   Usage  Avail Demand
-----
VP_R1_FC   F R1          10723   10110   613    613    9650    1888 +7762
Total
          10723   10110   613    613    9650    1888 +7762
}
```

Technology : SATA
 Disk Location : Internal

VP Tiers (1)

```
{
-----
          A                               Logical Capacities (GB)
          T                               -----
Tier      T Target      Tier      Tier      Tier FAST SG   FAST Max SG Excess
          R Prot        Enabled   Free     Used   Usage  Avail Demand
-----
VP_R614_SATA F R6(14+2) 24539   23201   1338    686   21433    2697 +18736
Total
          24539   23201   1338    686   21433    2697 +18736
}
```

Legend:

ATTR : F = Tier in a FAST policy associated with SG(s)
 : P = Tier in a FAST policy unassociated with SG(s)

: N = Tier not in any FAST policy

The thin demand report can also be run for a single technology type (EFD, FC, or SATA) by running:

```
symfast -sid 0398 list -demand -tech EFD -vp
```

Symmetrix ID : 000195700398

Technology : EFD
Disk Location : Internal

VP Tiers (1)

```
{
-----
          A                               Logical Capacities (GB)
          T                               -----
Tier      T Target      Tier   Tier   Tier FAST SG   FAST Max SG Excess
          R Prot        Enabled  Free   Used  Usage  Avail Demand
-----
VP_R53_EFD F R5(3+1)      4395  4189  206   205   4350   1187  +3163
Total                                     -----
          4395  4189  206   205   4350   1187  +3163
}
```

Legend:

ATTR : F = Tier in a FAST policy associated with SG(s)
: P = Tier in a FAST policy unassociated with SG(s)
: N = Tier not in any FAST policy

The `-v` option can be added to the previous command to provide more detailed information on the storage groups that are associated with each of the tiers that contain the disk groups of a particular technology. The `-v` option also provides summary information for each technology and each tier:

```
symfast -sid 0398 list -tech EFD -demand -vp -v
```

Symmetrix ID : 000195700398

Technology : EFD

```

Logical Tier Enabled Total (GB)      :      4395
Logical Tier Free Total (GB)        :      4185
Logical Tier Used Total (GB)        :        210
Logical FAST SG Usage Total (GB)    :        209
Logical FAST Available Total (GB)   :      4350
Logical Max SG Demand Total (GB)    :      1187
Logical Excess (GB)                 :     +3163

```

VP Tiers (1)

```

{
Tier Name           : VP_R53_EFD
Target Prot        : R5(3+1)
Logical Tier Enabled (GB) :      4395
Logical Tier Free (GB)   :      4185
Logical PRC Total (GB)  :         44
Logical Tier Used (GB)  :        210
Logical FAST SG Usage Total (GB) :        209
Logical FAST Available (GB) :      4350
Logical Max SG Demand Total (GB) :      1187
Logical Excess (GB)    :     +3163
Tier Status : Tier in a FAST policy associated with SG(s)

```

Storage Groups (2)

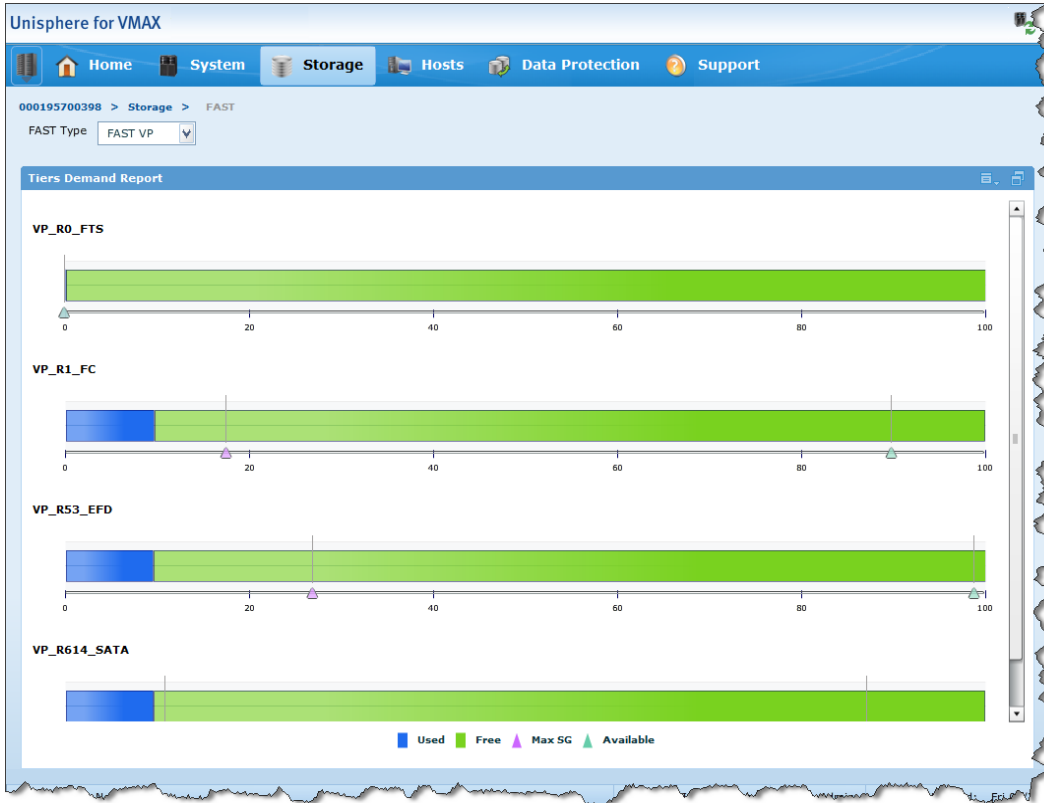
```

{
-----
                P  FAST SG   Max SG
                r   Usage   Demand
SG Name        Policy    i Log (GB) Log (GB)
-----
VP_ProdApp2    Custom      2     89    108
VP_ProdApp1    System_Optim* 2    120   1079
-----
                Total                209    1187
}

```

Unisphere for VMAX

In Unisphere, the technology demand reports can be accessed by the FAST view, located under Storage. The demand report can be viewed graphically in a chart or as a table.



Unisphere for VMAX

000195700398 > Storage > FAST

FAST Type: FAST VP

Tiers Demand Report

Name	Tech + Prot	Used (GB)	Free (GB)	Maximum SG Demand (GB)	Available (GB)	Excess (GB)
VP_R0_FTS	N/A Unprotected	0	195	0	0	0
VP_R1_FC	FC RAID-1	667	10056	1888	9650	+7762
VP_R53_EFD	EFD RAID-5(3+1)	216	4179	1187	4350	+3163
VP_R614_SATA	SATA RAID-6(14+2)	1369	23170	2697	21433	+18736

Thin pool utilization

As FAST VP migrates data between thin pools at the sub-LUN level, the thin devices that are under FAST VP control remain bound to a single pool. Each thin pool in the Symmetrix array has the possibility of having data allocated not only from the devices bound to it, but also from thin devices bound to other pools.

The thin pool utilization report can be used to determine how much capacity of the pool is currently allocated. The report provides a breakdown of the capacity allocated to devices bound to the thin pool. It also details information on the allocation to thin devices bound to other pools (capacity allocated as a result of FAST VP data movements).

The information contained in the thin pool utilization report includes:

- ◆ Pool Name: The name of pool for which the report was generated.
- ◆ Pool Type: The type of pool being viewed (in this case, it is thin).
- ◆ Dev Emulation: The device emulation of the data devices contained in the pool.
- ◆ Dev Configuration: The RAID protection type of the data devices contained in the pool.
- ◆ Pool State: The overall state of the pool, Enabled or Disabled.
- ◆ # of Devices in Pool: The total number of data devices contained in the pool.
- ◆ # of Enabled Devices in Pool: The total number of enabled data devices contained in the pool.
- ◆ Max. Subscription Percent: Indicates the maximum capacity of thin devices that can be bound to the pool. It is expressed in terms of a percentage of the total enabled capacity of the pool.
- ◆ Rebalance Variance: A target percentage for the variance in utilization of any one data device as measured against the thin pool's utilization.
- ◆ Max devs per rebalance scan: The maximum number of data devices in a thin pool that the rebalancing algorithm can work on concurrently.
- ◆ Enabled Devices: A listing of the enabled data devices contained in the pool.

List includes:

- SymDev: The Symmetrix device number for the data device.
- Total Tracks: The total capacity of the data device, expressed

- in 64 KB tracks.
- Alloc Tracks: The number of 64 KB tracks currently allocated from the data device to thin devices.
- Free Tracks: The number of 64 KB tracks not currently allocated from the data device.
- Full (%): The percent capacity currently allocated from the data device.
- Device State: The current data device state.
- ♦ Pool Bound Thin Devices: A listing of the thin devices bound to the pool.
List includes:
 - SymDev: The Symmetrix device number for the thin device.
 - Total Tracks: The total capacity of the thin device, expressed in 64KB tracks.
 - Pool Subs (%): Indicates the subscribed percentage of the thin device to the pool. This is the ratio of the capacity of the thin device to the total enabled capacity of the pool.
 - Pool Allocated Tracks: The number of tracks allocated for the thin device in the thin pool.
 - Pool Allocated (%): The percentage of the thin devices tracks that are allocated in the thin pool.
 - Total Written Tracks: The total number of written tracks for the thin device.
 - Total Written (%): The percentage of the thin device tracks that have been written to.
- ♦ Other Pool Bound Thin Devices: A listing of the thin devices that are bound to other thin pools, but have data allocated in the thin pool being viewed.
List includes:
 - SymDev: The Symmetrix device number for the thin device.
 - Bound Pool Name: The name of the pool the thin device is actually bound to.
 - Total Tracks: The total capacity of the thin device, expressed in 64 KB tracks.
 - Pool Allocated Tracks: The number of tracks allocated for the thin device in the thin pool being viewed.

- Pool Allocated (%): The percentage of the thin devices tracks that are allocated in the thin pool being viewed.

SYMCLI

To view the thin pool utilization report for a particular pool using SYMCLI, run:

```
symcfg -sid 0398 show -pool R53_EFD_Pool -thin -detail
```

```
Symmetrix ID: 000195700398
```

```
Symmetrix ID           : 000195700398
Pool Name              : R53_EFD_Pool
Pool Type              : Thin
Disk Location          : Internal
Technology             : EFD
Dev Emulation          : FBA
Dev Configuration     : RAID-5 (3+1)
Pool State             : Enabled
# of Devices in Pool   : 32
# of Enabled Devices in Pool : 32
# of Usable Tracks in Pool : 71999616
# of Allocated Tracks in Pool : 4094784
# of Shared Tracks in Pool : 0
Pool Utilization (%)   : 0
Max. Subscription Percent : None
Rebalance Variance     : 1%
Max devs per rebalance scan : 256
Pool Reserved Capacity : 1
```

```
Enabled Devices(32):
```

```
{
-----
Sym      Usable      Alloc      Free Full FLG Device
Dev      Tracks      Tracks      Tracks (%) S  State
-----
00D0     2249988     127152     2122836  5  .  Enabled
00D1     2249988     127536     2122452  5  .  Enabled
...
00EE     2249988     126936     2123052  5  .  Enabled
00EF     2249988     127584     2122404  5  .  Enabled
-----
Tracks   71999616     4094784     67904832  5
}
```

```
Pool Bound Thin Devices(1):
```

```

{
-----
Sym   FLG      Pool          Pool          Total
Dev   T        Total Sub    Allocated     Written
                Tracks (%)    Tracks (%)    Tracks (%)
-----
0430  B          276210  0      11784  4      11315  4
-----
Tracks          276210  0      11784  4      11315  4
}

```

Other-Pool Bound Thin Devices(6):

```

{
-----
                Pool
          Bound      Total      Allocated
Sym  Pool Name      Tracks      Tracks (%)
-----
02B0 R1_FC_Pool      4419360      783768  18
02C0 R1_FC_Pool      4419360       46140   1
02D0 R1_FC_Pool      4419360     144948   3
02E0 R1_FC_Pool      4419360    1366404  31
02F0 R1_FC_Pool      4419360    1105272  25
0300 R1_FC_Pool      4419360     636456  14
-----
Tracks          26516160    4082988  15
}

```

Legend:

Enabled devices FLG:

(S)hared Tracks : X = Shared Tracks , . = No Shared Tracks

Bound Devices FLG:

S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
D = Deallocating, R = Reclaiming, . = Unbound

Unisphere for VMAX

In Unisphere, the pool utilization report can be accessed through the Thin Pools view, located under Storage.

Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Thin Pools

Thin Pools

Name	Technology	Configuration	Emulation	Allocated Capacity	Capacity (GB)
R0_FTS_Pool	N/A	Unprotected	FBA	0 %	195.31
R1_FC_Pool	FC	2-Way Mir	FBA	5 %	10722.66
R53_EFD_Pool	EFD	RAID-5 (3 + 1)	FBA	5 %	4394.51
R57_FC_Pool	FC	RAID-5 (7 + 1)	FBA	0 %	10722.66
R6_SATA_Pool	SATA	RAID-6 (14 + 2)	FBA	5 %	24538.88
small_pool	FC	2-Way Mir	FBA	0 %	268.07

More detail on individual pools can be seen by double-clicking a pool name and selecting one of the related objects.

Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Thin Pools > R53_EFD_Pool

Details : Thin Pool : R53_EFD_Pool

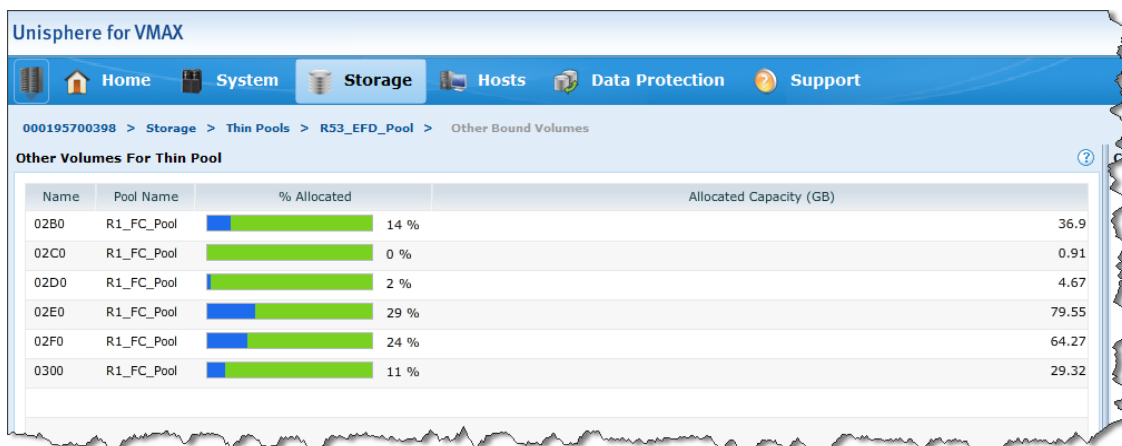
Properties

Name	R53_EFD_Pool
RAID Protection	RAID-5 (3 + 1)
Type	Thin
Technology	EFD
Emulation	FBA
Total Capacity (GB)	4394.51
Free Capacity (GB)	4140.63
Thin Volumes	1
Enabled Volumes	32

Related Objects

- Contains : DATA Volumes - 32
- Associated With : Bound Volumes - 1
- Other Pool Bound Volumes - 6

Create Expand Bind >> Apply Cancel



Thin device distribution

As mentioned previously, a thin device associated with a FAST VP policy may have track extents allocated across multiple pools. However, it always remains bound to a single pool.

The thin device distribution report can be used to determine both the thin pool a thin device is bound to and any other pools that have tracks allocated for that device.

The information contained in the thin device distribution report is:

- ◆ Sym: The Symmetrix device number for the thin device.
- ◆ Pool Name: The name of pool to which the thin device is bound, or the pool the device has tracks allocated in.
- ◆ Flags EM: Indicates the emulation of the thin device (E), and if there are multi-pool allocations for the thin device (M).
- ◆ Total Tracks: The total capacity of the thin device, expressed in 64 KB tracks.
- ◆ Pool Subs (%): Indicates the subscribed percentage of the thin device to the pool. This is the ratio of the capacity of the thin device to the total enabled capacity of the pool.
- ◆ Pool Allocated Tracks: The number of tracks allocated for the thin device in the thin pool.

- ◆ Pool Allocated (%): The percentage of the thin devices tracks that are allocated in the thin pool.
- ◆ Total Written Tracks: The total number of written tracks for the thin device. It is only shown for the pool the thin device is bound to.
- ◆ Total Written (%): The percentage of the thin device tracks that have been written to. It is only shown for the pool the thin device is bound to.
- ◆ Status: Indicates whether the device is bound or unbound. If bound, the pool the thin device is bound to is listed first.

SYMCLI

To view the thin device allocation report for all thin devices using SYMCLI, run:

```
symcfg -sid 0398 list -tdev -detail
```

```
Symmetrix ID: 000195700398
```

```
Enabled Capacity (Tracks) : 838235112
Bound Capacity (Tracks) : 73826310
```

S Y M M E T R I X T H I N D E V I C E S							
Sym	Bound Pool Name	Flags ESPT	Total Tracks	Pool Subs (%)	Pool Allocated Tracks (%)	Total Written Tracks (%)	
02B0	R1_FC_Pool	F..B	4419360	3	437256 10	2037532	46
	R53_EFD_Pool	-.--	-	-	926292 21	-	-
	R6_SATA_Pool	-.--	-	-	1164024 26	-	-
02C0	R1_FC_Pool	F..B	4419360	3	1171332 27	2030317	46
	R53_EFD_Pool	-.--	-	-	83160 2	-	-
	R6_SATA_Pool	-.--	-	-	1163964 26	-	-
02D0	R1_FC_Pool	F..B	4419360	3	1016688 23	2027570	46
	R53_EFD_Pool	-.--	-	-	217836 5	-	-
	R6_SATA_Pool	-.--	-	-	1184208 27	-	-
...							
04B4	R6_SATA_Pool	F..B	1638405	0	619032 38	222044	14
04B5	R6_SATA_Pool	F..B	1638405	0	624696 38	223593	14
04B6	R6_SATA_Pool	F..B	1638405	0	630252 38	224578	14
Total Tracks			113600550	14	38061804 5	24464742	3

Legend:

Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
(S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
(P)ersistent Allocs : A = All, S = Some, . = None
S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
D = Deallocating, R = Reclaiming, . = Unbound

To view the thin device allocation report for a single thin device, or a range of devices, run:

symcfg -sid 0398 list -tdev -RANGE 2B0:2C0 -detail

Symmetrix ID: 000195700398

Enabled Capacity (Tracks) : 838235112

Bound Capacity (Tracks) : 8838720

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

	Bound	Flags	Pool	Pool	Total		
Sym	Pool Name	ESPT	Total Tracks	Subs (%)	Allocated Tracks (%)	Written Tracks (%)	

02B0	R1_FC_Pool	F..B	4419360	3	273348 6	2037613	46
	R53_EFD_Pool	-.--	-	-	1217532 28	-	-
	R6_SATA_Pool	-.--	-	-	1036692 23	-	-
02C0	R1_FC_Pool	F..B	4419360	3	1050540 24	2030457	46
	R53_EFD_Pool	-.--	-	-	209496 5	-	-
	R6_SATA_Pool	-.--	-	-	1158420 26	-	-
Total							
Tracks			8838720	1	4946028 1	4068070	0

Legend:

Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
(S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
(P)ersistent Allocs : A = All, S = Some, . = None
S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
D = Deallocating, R = Reclaiming, . = Unbound

Note: To display information for a single device, specify the same device number for both the start and end device in the range. An example is 2B0:2B0.

To view the thin device allocation report for all the thin devices in a storage group, run:

```
symcfg -sid 0398 list -tdev -sg VP_ProdApp1 -detail
```

Symmetrix ID: 000195700398

Enabled Capacity (Tracks) : 838235112
Bound Capacity (Tracks) : 17677440

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

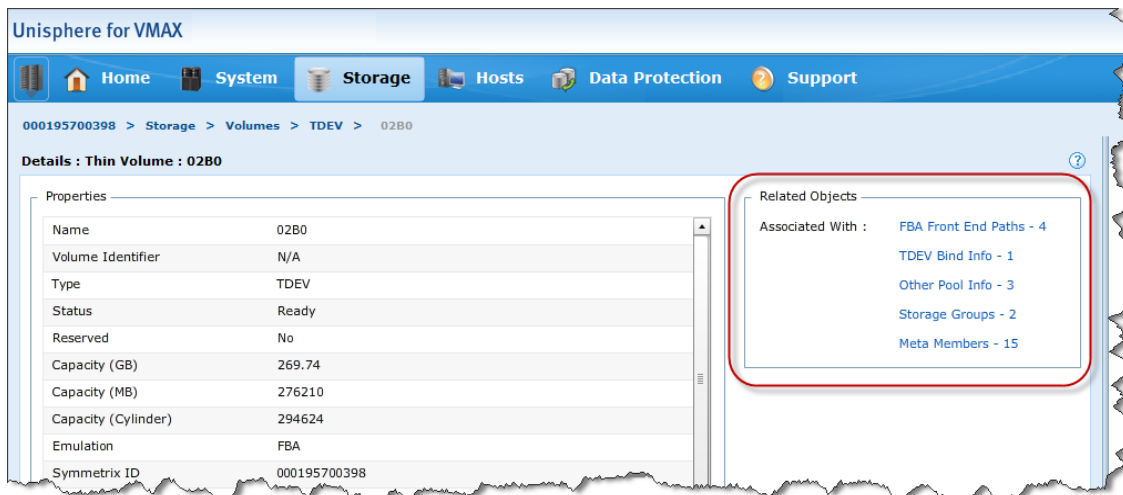
Sym	Bound Pool Name	Flags ESPT	Total Tracks	Pool Subs (%)	Pool Allocated Tracks (%)	Total Written Tracks (%)
02B0	R1_FC_Pool	F..B	4419360	3	401676 9	2037550 46
	R53_EFD_Pool	-.--	-	-	992472 22	- -
	R6_SATA_Pool	-.--	-	-	1133424 26	- -
02C0	R1_FC_Pool	F..B	4419360	3	1152072 26	2030344 46
	R53_EFD_Pool	-.--	-	-	102420 2	- -
	R6_SATA_Pool	-.--	-	-	1163964 26	- -
02D0	R1_FC_Pool	F..B	4419360	3	990036 22	2027606 46
	R53_EFD_Pool	-.--	-	-	255108 6	- -
	R6_SATA_Pool	-.--	-	-	1173588 27	- -
02E0	R1_FC_Pool	F..B	4419360	3	5760 0	2038445 46
	R53_EFD_Pool	-.--	-	-	1419552 32	- -
	R6_SATA_Pool	-.--	-	-	1482036 34	- -
Total			17677440	2	10272108 1	8133945 1

Legend:

Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
(S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
(P)ersistent Allocs : A = All, S = Some, . = None
S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
D = Deallocating, R = Reclaiming, . = Unbound

Unisphere for VMAX

In Unisphere, the pool utilization report can be accessed by the Other Pool Info related object, which is available when viewing the properties of a specific thin device.



Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Volumes > TDEV > 02B0

Details : Thin Volume : 02B0

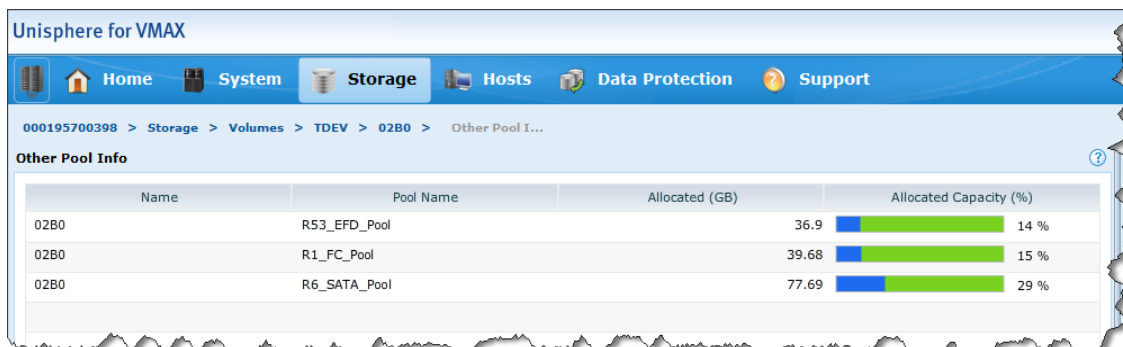
Properties

Name	02B0
Volume Identifier	N/A
Type	TDEV
Status	Ready
Reserved	No
Capacity (GB)	269.74
Capacity (MB)	276210
Capacity (Cylinder)	294624
Emulation	FBA
Symmetrix ID	000195700398

Related Objects

Associated With :

- FBA Front End Paths - 4
- TDEV Bind Info - 1
- Other Pool Info - 3**
- Storage Groups - 2
- Meta Members - 15



Unisphere for VMAX

Home System **Storage** Hosts Data Protection Support

000195700398 > Storage > Volumes > TDEV > 02B0 > Other Pool I...

Other Pool Info

Name	Pool Name	Allocated (GB)	Allocated Capacity (%)
02B0	R53_EFD_Pool	36.9	14 %
02B0	R1_FC_Pool	39.68	15 %
02B0	R6_SATA_Pool	77.69	29 %

Conclusion

EMC Symmetrix VMAX FAST VP for Virtual Provisioning environments automates the identification of active or inactive application data for the purposes of reallocating that data across different performance/ capacity tiers within an array. FAST VP proactively monitors workloads at both the LUN and sub-LUN level in order to identify busy data that would benefit from being moved to higher-performing drives. FAST VP also identifies less-busy data that could be moved to higher-capacity drives, without existing performance being affected. This promotion/ demotion activity is based on policies that associate a storage group to multiple drive technologies, or RAID protection schemes, by way of thin storage pools, as well as the performance requirements of the application contained within the storage group. Data movement executed during this activity is performed non-disruptively, without affecting business continuity and data availability.

Appendix A: FASTVP state

There are five possible states that the FAST controller can be reported in. These are:

- ◆ **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 movement is not executed.
- ◆ **Disabling:** The transition of the FAST controller from Enabled to Disabled.
- ◆ **DisabledwithError:** 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.

Enabled state

When the state of the FAST controller is queried, and the state is Enabled, the current activity being performed by the controller is also displayed. Valid activities include:

- ◆ **Idle:** The FAST controller is currently idle.
- ◆ **RunningPlan:** There are currently active data movement tasks running, moving thin device data between tiers.

Degraded state

When the state of the FAST controller is Degraded, a reason code is displayed when the FAST state is queried, and it indicates the cause of the degraded state.

These reason codes include:

- ◆ **Invalid Swap/ Performance time windows:** At least one of the defined time windows is invalid. To correct, each time window

should be checked, and any invalid time windows should be deleted or modified.

- ◆ Invalid device attributes: One or more storage groups have an invalid priority in a FAST policy. To correct, each storage group's priority should be checked in the FAST policy they are associated with. Any invalid priority should be modified to a valid value.
- ◆ Invalid FAST parameters: One or more of the FAST controller configuration settings are invalid. To correct, each configuration setting should be checked and set to a valid value.
- ◆ Performance time window is not present or does not extend into the future: No performance time window, default or user-defined, exists, or any that do exist have expired. To correct, a valid, inclusion performance time window should be created.
- ◆ FAST thin move time window is not present or does not extend into the future: No thin data movement time window, default or user-defined, exists, or any that do exist have expired. To correct, a valid, inclusion thin data movement time window should be created.
- ◆ FAST VP compliance movement failed: The most recent attempt to perform a FAST VP compliance movement was not successful. EMC customer service should be contacted to investigate the reason for the failure. If a subsequent attempt to perform a compliance movement is successful, the degraded state is cleared.
- ◆ FAST VP performance movement policy update failed: The most recent attempt to generate a data movement policy failed. EMC customer service should be contacted to investigate. If a subsequent attempt to generate a movement policy is successful, the degraded state is cleared.
- ◆ FAST VP is not licensed: An entitlement file including FAST VP has not been loaded to the Symmetrix array. To correct, the appropriate entitlement file should be obtained from EMC and loaded to the Symmetrix array.
- ◆ Statistics collection is failing for thin devices - No Performance movement will happen: Performance statistics are not being collected for thin devices under FAST VP control. EMC Customer Service should be contacted to investigate. If a subsequent attempt to collect statistics is successful, the degraded state is cleared.
- ◆ Timed out attempting to communicate with the FAST controller: Either the FAST controller running on the service processor is unavailable, or the service processor itself is unavailable. EMC Customer Service should be contacted to investigate.

Appendix B: Feature support

The following table describes the minimum Enginuity and management interface levels needed to support various FAST VP features.

Feature	Enginuity	Management Interface
FAST VP (Base)	5875.135.91	Solutions Enabler 7.3 SMC 7.3 Unisphere for VMAX 1.0
Setting PRC per pool	5875.198.38	Solutions Enabler 7.3.1 SMC 7.3.1 Unisphere for VMAX 1.0
VP allocation by FAST policy	5876.82.57	Solutions Enabler 7.4 Unisphere for VMAX 1.0
FAST VP SRDF coordination	5876.82.57	Solutions Enabler 7.4 Unisphere for VMAX 1.0
External tier (FTS)	5876.82.57	Solutions Enabler 7.4 Unisphere for VMAX 1.0
Storage group reassociation	5876.82.57	Solutions Enabler 7.4 Unisphere for VMAX 1.0
FAST VP for CKD	5876.82.57	Solutions Enabler 7.4 Unisphere for VMAX 1.0
FAST VP for IBM i	5876.82.57	Solutions Enabler 7.4 Unisphere for VMAX 1.0

References

- ◆ *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide*
- ◆ *EMC Solutions Enabler Symmetrix Array Management CLI Product Guide*
- ◆ *EMC Solutions Enabler Symmetrix CLI Command Reference HTML Help*
- ◆ *EMC Solutions Enabler Installation Guide*
- ◆ *EMC Symmetrix VMAX Series Product Guide*
- ◆ *FAST VP for EMC® Symmetrix® VMAX™ Theory and Best Practices for Planning and Performance*
- ◆ *Best Practices for Fast, Simple Capacity Allocation with EMC Symmetrix Virtual Provisioning Technical Note*
- ◆ *z/OS and Virtual Provisioning Best Practices*
- ◆ *Design and Implementation Best Practices for EMC Symmetrix Federated Tiered Storage (FTS) Technical Note*
- ◆ *Best Practices for Nondisruptive Tiering via EMC Symmetrix Virtual LUN Technical Note*

Copyright © 2012 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.

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