

EMC[®] Symmetrix[®] Optimizer and Fully Automated Storage Tiering for Symmetrix DMX[™] Series Arrays

Technical Note

P/N 300-010-473

REV A01

April 27, 2010

This technical note contains information on these topics:

- ◆ Executive summary 2
- ◆ Introduction 2
- ◆ Symmetrix Optimizer 3
- ◆ How does Symmetrix Optimizer work? 4
- ◆ Symmetrix Optimizer configuration 9
- ◆ Rules-based analysis 14
- ◆ Optimizer device swaps 18
- ◆ Advanced data movement modes 26
- ◆ Replication interoperability 27
- ◆ Virtual LUN Technology 29
- ◆ Fully Automated Storage Tiering (FAST) for DMX 33
- ◆ Management Interface: SMC 39
- ◆ Management Interface: SYMCLI 74
- ◆ Logging 100
- ◆ Conclusion 111
- ◆ Appendix A: Affinity Groups 112
- ◆ References 115

Executive summary

Enterprises of all sizes require consistent and predictable access to mission-critical information as the fundamental criteria to achieving business success and competitive edge – this means the information infrastructure needs to be tuned and optimized at all times with minimal maintenance and intervention.

This technical note provides an in-depth operational description of EMC® Symmetrix® Optimizer, which provides both non-disruptive load rebalancing that reduces drive contention, and Symmetrix Virtual LUN technology that enables users to manually and non-disruptively relocate volumes between storage types to meet changing business requirements.

This document also examines Fully Automated Storage Tiering (FAST) for Symmetrix DMX, a product that automates the relocation of volumes between storage types, enabling customers to realize the benefits of tiered storage with less time and effort spent on performance tuning, management, and monitoring. FAST relies on Optimizer functionality.

Introduction

Symmetrix Optimizer improves array performance by continuously monitoring access patterns and migrating devices (Symmetrix logical volumes) to achieve balance across the drives within a physical disk group, and thereby reduce the risk of “hot spots.” This automated process is based on user-defined parameters and is completely transparent to end users, hosts, and applications in the environment.

Symmetrix Optimizer also allows for movement of devices between differing drive technologies to change the performance profile of the device. Such movement can be performed manually, using Virtual LUN technology, or automatically, based on performance workload, using Fully Automated Storage Tiering (FAST) for DMX.

Audience

This technical note provides an in-depth look at the Symmetrix Optimizer feature for Symmetrix DMX enterprise storage arrays, including Virtual LUN technology and FAST for DMX. The intended audience includes system and storage administrators, customers, EMC staff, and partners that need to understand how to implement and leverage FAST.

Symmetrix Optimizer

EMC Symmetrix Optimizer helps increase the performance of a Symmetrix system by spreading I/O activity evenly across the physical drives. When a particular drive is in high demand, there is excessive head movement on that drive that can slow down read and write activity. By balancing highly active and less active logical devices across drives, seek activity is balanced, and contention among drives is reduced. Throughput within the overall Symmetrix system is improved, and response times optimized.

Optimizer tunes of Symmetrix data configurations from the Symmetrix service processor by:

- ◆ Analyzing statistics about Symmetrix logical device activity.
- ◆ Determining which logical devices should have their physical locations swapped to enhance Symmetrix performance.
- ◆ Swapping logical devices and their data using internal Dynamic Reallocation Volumes (DRVs) to protect customer data while reconfiguring the system (on a device-to-device basis).

Symmetrix Optimizer:

- ◆ Automatically collects logical device activity data, based upon the devices and time window defined by the user.
- ◆ Identifies “hot” and “cold” logical devices, and determines on which physical drives they reside.
- ◆ Determines which logical device swaps would reduce physical drive contention and minimize average disk service times.
- ◆ Swaps logical devices to balance activity across the back end of the Symmetrix array.

Optimizer is designed to run automatically in the background, analyzing performance in the performance time windows specified and performing swaps in the swap time windows specified.

Management and operation of Symmetrix Optimizer is provided by Symmetrix Management Console (SMC), as well as the Solutions Enabler Command Line Interface (SYMCLI).

How does Symmetrix Optimizer work?

Symmetrix Optimizer operates within the confines of Symmetrix physical disk groups. Symmetrix disk groups are defined in a Symmetrix bin file, and separate physical drives based on drive technology, drive size, and drive rotational speed.

Note: Disk groups are formed when the bin file is initially created, or when additional physical drives are added to the system.

By default, the goal of the Symmetrix Optimizer is to minimize the *average* service time of the Symmetrix disks, within a single physical disk group. It achieves this by measuring and analyzing the performance of all devices contained within each disk group defined on the array. Once analyzed, Optimizer will relocate the devices across the physical disks within the disk group in the most optimal layout possible for the best backend performance.

Note: Differing service times for data on the same drive typically only applies to “spinning disks” and not to solid state devices. As such, the load balancing operations performed by Optimizer apply only to Fibre Channel and SATA disks and not to Enterprise Flash Drives.

Device relocation is performed by swapping the physical locations of two separate devices. During the swap the locations of each physical hyper of both devices are exchanged. A physical hyper is an individual member of the RAID group protecting the Symmetrix device.

Note: For more information on Symmetrix disk group and physical hyper configuration information, refer to *Appendix A: Affinity Groups*.

Optimizer Controller

The Optimizer controller runs as a service on the Symmetrix DMX service processor. When active, the Optimizer controller has four primary responsibilities:

1. Performance data collection
2. Performance data and capacity usage analysis
3. Configuration change plan generation
4. Device movement execution

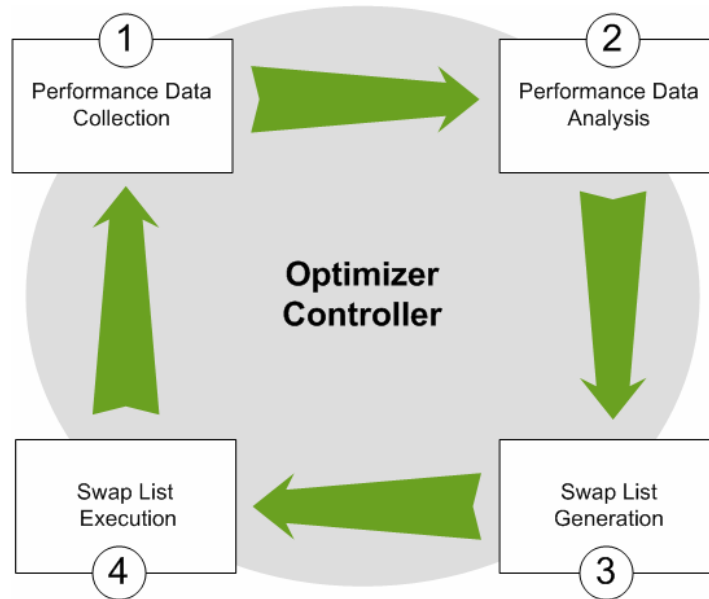


Figure 1 Symmetrix Optimizer controller operation

Performance statistics are collected at 10-minute intervals and are stored in a database file on the service processor.

At regular intervals, the data collected by the Optimizer controller is analyzed, and determinations made as to whether the physical locations of devices need to be swapped. The generated list of device swaps created by Optimizer is known as a swap list.

When a swap list exists, the Optimizer controller is responsible for executing the plan, committing the required changes to the Symmetrix back-end configuration. When created, a swap list can be executed automatically without user interaction, or execution can be delayed until user approval is granted manually.

Note: By default, Optimizer runs in “demo mode” on the service processor. In this mode, no software client connections can be made to Optimizer from Solutions Enabler or the Symmetrix Management Console. Also, no device swaps will be performed by Optimizer.

Prior to enabling Optimizer, or making any configuration changes to Optimizer, it must first be taken out of demo mode. This can only be done by an EMC customer service representative activating Optimizer on the service

processor.

Optimizer controller state

There are four possible states that the Optimizer controller can be reported in. These are:

- ◆ Started/Enabled—Optimizer will perform all of its functions: performance data collection, performance data analysis, swap list generation, and swap execution.
- ◆ Stopped/Disabled—Optimizer will only perform one of its functions: performance data collection. Data analysis and swap list generation or execution will not be performed.
- ◆ Stopping—Optimizer is transitioning from Started to Stopped.
- ◆ StoppedOnError—Optimizer has stopped operation due to an internal error. No Optimizer operations will be performed.

Started/Enabled state

When the state of the Optimizer controller is queried, and the state is Started/Enabled, the current activity being performed by the controller will also be displayed. Valid activities include:

- ◆ Idle—Optimizer is currently idle
- ◆ Fetching Statistics—Optimizer is collecting Symmetrix device performance statistics
- ◆ Analyzing Statistics—Optimizer is performing analysis on the collected device performance statistics
- ◆ PendingPlan—A swap list has been approved and its execution scheduled
- ◆ RunningPlan—A swap list is currently being executed

Optimizer performance analysis

When collecting device performance data for use in its analysis, Optimizer only collects statistics related to Symmetrix back-end activity.

The metrics collected are:

- ◆ DA logical volumes reads
- ◆ DA logical volumes writes
- ◆ Logical volume prefetch
- ◆ DA logical volumes blocks read
- ◆ DA logical volumes blocks written

With this information, Optimizer will model the disk service time for each Symmetrix device in the array.

The disk service time is defined as the sum of seek, latency, and transfer time.

- ◆ Seek time—The time it takes the disk arm to move and position the disk head on the correct track. That is, move from track X (serving the previous I/O) to track Y (serving the next I/O).
- ◆ Latency time—The delay for disk rotation. The latency time is a function of the disk rotational speed. Optimizer assumes one-third of a spin for each I/O assuming the internal disk optimization is on.
- ◆ Transfer time—The time that it takes the disk to transfer the data from/to the disk. Transfer time is a function of the data transfer rate, disk bandwidth, and data layout.

While modeling the disk service time, different weights are assigned to read, write, and prefetch activity. The assumption is that since writes are done as a background process, they are done in sequence and hence are more efficient.

Note: Optimizer does not try to follow DMSP policies; instead, it assumes that reads are equally spread among all the device mirrors, and all mirrors do all writes.

As accurate seek and latency times are impractical to get because they would require a complete trace of the I/O sequence, Optimizer uses a mathematical model to calculate these metrics. The Optimizer model makes the following assumptions in order to bridge the gap of missing data:

- ◆ Random Distribution—Optimizer assumes that the arrival rate and sequence of I/Os at the back end are randomly distributed.
- ◆ Independent I/Os—Optimizer does not assume any relationships between two consecutive I/Os.

Because of these assumptions seek and latency times calculated are, in a sense, a worst-case scenario. Optimizer calculated seek time is usually higher than the actual disk seek time because real-life workloads typically include sequences of I/Os and high percentage of locality of reference.

Optimizer swap strategies

When Symmetrix Optimizer evaluates performance statistics, it evaluates potential volume swaps based on how well they would

improve overall performance. This analysis is based on minimizing disk service time (rotational latency plus seek time plus transfer time).

Symmetrix Optimizer uses three strategies when determining which swaps to make:

- ◆ Load Balancing
- ◆ Minimize Seek Time
- ◆ Faster Media

Load balancing

The first strategy used is to swap highly active volumes to disks with lower activity. This evens out the load across the physical disks in the Symmetrix system. It also decreases the contention on the individual physical disks, thereby improving the performance of both highly active volumes and low-activity volumes.




Figure 2 Optimizer Strategy: Load Balancing

Minimize Seek Time

The second strategy performs swaps that relocate highly active volumes so that they are closer together on the physical disk. This is to decrease the seek distance for I/Os to these volumes. When Optimizer swaps volumes, it will try to achieve this type of configuration since it decreases the overall I/O service time.




Figure 3 Optimizer Strategy: Minimize Seek Time

Faster Media

The third strategy tries to relocate highly active volumes to the outer zones of the disk. Typically, volumes located on the outer zones of a physical disk have faster transfer speeds.




Figure 4 Optimizer Strategy: Faster Media

Finding the best swap

The Optimizer analysis consists of three high-level phases:

- ◆ Calculate service time—Model and sum the total service time of each disk, for every timestamp that was marked to be included in the analysis.
- ◆ Sort disks by activity—Sort all disks by their modeled total service time.
- ◆ Find best swap—starting with the busiest disk, check all potential swaps.

The analysis process models what-if scenarios using virtual swaps to estimate the impact of a swap on the service time of the affected disks. Optimizer will check as many swaps as possible in order to guarantee that one of the best swaps is indeed selected.

Symmetrix Optimizer configuration

Optimizer has multiple configuration parameters that control its behavior. These include time windows, which control what performance data to include in the data analysis and when device movements will take place. Other settings determine how much historical data is to be analyzed when creating swap lists, as well as quality of service.

Optimizer time windows

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

- ◆ Performance time window
- ◆ Swap time window

The *performance* time windows are used to specify which performance data samples to be included when Optimizer performs its data analysis.

The *swap* time windows define when to perform the configuration changes necessary to swap the physical locations of two volumes.

Time window characteristics

Both performance time windows and swap time 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 be performed outside the defined time window.

Multiple time windows of both types may be defined. Each type of window is hierarchical in nature and can be either periodic or non-periodic. If multiple time windows of the same type have time ranges that overlap one another, the most recently added time window will override the others.

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

A Symmetrix DMX storage array will support up to 128 defined time windows. Each time window name may be up to 32 alpha-numeric characters, hyphens (-), and underscores (_).

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 samples will be included in, or excluded from, Optimizer's performance data analysis. The intent of defining performance time windows is to distinguish periods of important system activity from periods when the Symmetrix is idle. And to only use performance data collected during these relevant periods to optimize system activity for the targeted IO patterns.

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

Swap time window

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

Optimizer will never start a swap session less than 30 minutes before the end of an inclusion swap time window; however, a swap that has started may continue beyond the specified time window.

A default swap time window defines an exclusive window, preventing swaps from occurring 24 hours a day, 7 days a week, 365 days a year.

Optimizer controller settings

There are multiple Optimizer settings that affect the behavior of the Optimizer controller. These include:

- ◆ Data Movement Mode
- ◆ Maximum Moves Per Day
- ◆ Maximum Simultaneous Moves
- ◆ Workload Analysis Period
- ◆ Initial Period

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

Data Movement Mode

Optimizer, when enabled, may operate in one of two modes—User Approved or Automatic.

In *User Approved* mode, Optimizer will generate swap lists on an hourly basis, but will not perform any swaps until such time that the list is approved and scheduled by a user. Swap lists approved by a user may be scheduled to run during the next open swap time window, or at a time specified by the user.

In *Automatic* mode, swap lists will be generated and executed at the beginning of each inclusion swap time window, without user involvement.

The default mode is Automatic.

Note: Even though the default mode is Automatic, the default swap time window excludes all device swaps.

Maximum Moves Per Day

The Maximum Moves Per Day setting controls the number of swaps or moves Optimizer will perform during a 24-hour period. As Optimizer swaps modify the back-end configuration of the Symmetrix, the number of swaps may be limited each day until the effect of such movements is understood within the array.

As two devices are involved in a single swap, each swap performed will count as two devices moved towards this total.

The maximum moves per day setting may be configured between 2 and 200. The default setting is 200.

Maximum Simultaneous Moves

The Maximum Simultaneous Moves setting controls the amount of swaps Optimizer will perform concurrently. As Optimizer device swaps consume Symmetrix resources, the amount of concurrent movements may be limited until their impact is understood.

Similarly to the Maximum Moves Per Day value, each swap performed will count as two devices moved towards this total.

The maximum simultaneous moves setting may be configured between 2 and 8. The default value is 8.

Note: For device swaps, the maximum simultaneous device moves value is dependent on the number of DRVs configured on the Symmetrix. If there are only four DRVs configured, Optimizer will only be able to swap a maximum of two device pairs, regardless of whether the maximum simultaneous moves is set higher.

Workload Analysis Period

Optimizer's Workload Analysis Period defines the amount of historical statistical information that will be maintained by the Optimizer controller. This will be the amount of data used for analysis in generating an Optimizer swap list.

Ideally, the workload analysis period will cover the time over which the Symmetrix experiences its normal workload cycle.

The workload analysis period can be configured between 1 hour and 4 weeks. The default is 1 week.

Initial Period

Optimizer's Initial Period defines the minimum amount of data Optimizer should wait to collect before generating its first swap list.

This value should be set to a long enough value so as to allow sufficient data samples for Optimizer to establish a good characterization of the typical workload in the Symmetrix. This value allows Optimizer to commence analysis and swap activities prior to the full workload analysis period elapsing, if so desired.

The initial analysis period can be configured to be between 1 hour and 4 weeks. It cannot exceed the value of the workload analysis period. The default is 1 week.

Advanced Optimizer settings

Several additional settings exist that can modify the default behavior of Optimizer. These can affect the type of data analysis performed by Optimizer, as well as the types of devices that Optimizer can automatically swap.

Hot Spot Analysis

By default, the goal of the Symmetrix Optimizer is to minimize the *average* service time of the Symmetrix disks. Optimizer's goal can be changed to lower the peak service times, rather than the average service time, by enabling hot spot analysis mode.

Hot spot analysis mode can be enabled via SMC or SYMCLI.

Allow SRDF device swapping

By default, Optimizer excludes SRDF devices that have no local RAID protection in generated swap lists. A setting can be enabled that will allow such SRDF devices to be swapped automatically by Optimizer.

The setting to allow SRDF device swapping can only be enabled by an EMC customer service representative via the Symmetrix service processor.

| |
|---|
| Note: Locally unprotected SRDF devices may be swapped manually without enabling this option. |
|---|

Allow unprotected device swapping

By default, Optimizer will not include unprotected devices in generated swap lists. A setting can be enabled that will allow unprotected devices to be swapped automatically by Optimizer.

The setting to allow unprotected device swapping can only be enabled by an EMC customer service representative, via the Symmetrix service processor.

Optimize DA

By default, Optimizer device swaps will perform load balancing across

disks within physical disk groups. An additional setting can be enabled within Optimizer to consider the existing load on the backend disk adapters (DAs) within the array. When enabled, swaps generated by Optimizer will attempt to balance the workload across both physical disks and DAs.

The setting to optimize workload across the DAs may only be enabled by an EMC customer service representative, via the Symmetrix service processor.

Enforce physical disk groups

By default, Optimizer will only perform device swaps within a single physical disk group. The setting to enforce Symmetrix physical disk groups can be disabled, thereby allowing Optimizer to suggest swaps between physical disk groups. When disabled, swaps will still be limited between devices that are configured on similar disks—disks must be of the same technology, size, and rotational speed.

The setting to enforce physical disk groups may only be disabled by an EMC customer service representative, via the Symmetrix service processor.

Note: this setting applies to the traditional load balancing feature of Optimizer, and does not affect the behavior of Virtual LUN technology, or FAST for DMX.

Maximum number of days for rollback

Swaps performed by Optimizer can be manually reversed, and the configuration of the array “rolled back” to a specific point-in-time. By default, Optimizer will allow rollbacks to be performed to any time within the previous 30 days. Valid values for the maximum number of days for rollback are between 1 and 180 days.

The maximum number of days for rollback may be modified using SYMCLI.

Start Up Mode

In the event of a service processor restart, Optimizer can be configured such that it restarts automatically, or that it must be manually restarted. By default, Optimizer will restart automatically. Setting the start up mode to manual means that user intervention will be required to start Optimizer should the service processor be rebooted.

Rules-based analysis

By default, Optimizer will treat all devices equally, and will perform swaps between all eligible devices, on all disks within each physical disk

group—Optimizer has no knowledge of the host environment relating to each device.

Rules-based analysis allows for finer control of the optimization process by allowing users of SMC, or Ionix ControlCenter, to define sets of rules that can be applied to sets of devices, or to sets of disks. These rules, when implemented, affect the composition of the swap lists generated by Optimizer.

Rules created can exclude certain devices, or disks, from Optimizer's analysis. They can also control the devices that can be swapped with each other, and the disks that devices can be swapped to.

There are two types of rules that can be created: disk rules and device rules.

Disk rules

Disk rules control the use of physical disks as part of the Optimizer data analysis and swap list creation processes. Disk rules can be used to define if certain disks should be included or excluded when generating a swap list.

There are three types of disk rules that can be created

- ◆ Disk pool
- ◆ Disk exclude
- ◆ Intra disk

Disk pool

By default, Optimizer considers all disks within a physical disk group when generating a swap list. The physical disks can be partitioned into smaller groupings of disks by creating a disk pool. When created, a disk pool rule can direct Optimizer to only optimize within a user-defined pool of disks, or to allocate a higher priority to that pool of disks.

Devices on the disks included in the disk pool rule will only be swapped with other devices on disks within the same disk pool.

Disk exclude

By creating a disk exclude rule, Optimizer will exclude a user-defined pool of disks from being included in Optimizer analysis.

Devices configured on the excluded disks will not be swapped by Optimizer.

Intra disk

When created, the intra-disk rule allows only intra-disk swaps for the

disks that the rule is applied to.

Device hypers configured on the disks included under this rule will only be swapped with other hypers on the same physical disk.

Note: Even though the benefits of load balancing cannot be achieved on disks included in an intra-disk rule, there can still be improvement in performance obtained by reducing seek time and transfer time.

Device rules

Device rules provide the ability to control the swapping of devices, and the possible destination of devices included under the rules.

There are three types of device rules:

- ◆ Device avoidance
- ◆ Migration
- ◆ Cross disk

Device avoidance

When implemented, the device avoidance rule prevents Optimizer from generating a swap list that would place two or more devices included under the rule on the same physical disk.

Migration

The migration rule specifies that Optimizer can move data among storage tiers within the same Symmetrix array with devices of the same RAID protection scheme.

This applies only to manual moves.

Note: For more information on manual moves between storage tiers, see the section on

Virtual LUN Technology.

Cross disk group

The cross disk group rule allows Optimizer to swap devices between physical disk groups that reside on different disk technology. With this rule, Optimizer will attempt to move devices with higher disk utilization to faster technologies, such as EFD or FC, while moving devices with lower disk utilization to more cost-effective technologies, such as SATA.

Note: For more information on the cross disk group device movements, see the section on *Fully Automated Storage Tiering (FAST) for DMX*.

Note: The management of device and disk rules is not available via the Solutions Enabler SYMCLI. It is only available via SMC or Ionix Control Center.

Device attributes

In addition to disk and device rules, individual device attributes can be set to assign priorities to each device. The priority of each device is taken into consideration during the creation of a swap list following Optimizer's data analysis.

Each device swappable by Optimizer may be assigned one of three priorities:

- ◆ Normal
- ◆ High
- ◆ No swap

Normal Priority

A device with normal priority is available to be swapped by Optimizer. This is the default attribute assigned to each device.

High Priority

Devices assigned the high priority may be relocated to achieve the best performance available for the device. However, attempts to improve performance will be made without sacrificing the performance of other devices that have also been assigned the high priority attribute.

No Swap

Giving a device the No Swap attribute indicates that the device should be excluded from data analysis performed by Optimizer, and any

resulting swap list generation. Such a device will never be moved from its existing location.

Optimizer device swaps

Only certain device types, and emulations, are eligible to be swapped. Also, there are certain configuration requirements on the Symmetrix in order for Optimizer swaps to be successful. The following section discusses the configuration requirements and swap eligible devices. Also, the process followed in swapping devices of the various RAID protections is presented.

Swap configuration requirements

There are three basic configuration requirements necessary for an Optimizer swap to complete. These are:

- ◆ Dynamic Reallocation Volumes (DRVs)
- ◆ Open mirror position
- ◆ Open Symmetrix configuration lock

Dynamic Reallocation Volumes (DRVs)

DRVs are non-user-addressable devices used to ensure that data is available and protected during the Optimizer swap process. Two DRVs are required for each device pair being swapped. As such, in order to swap multiple device pairs simultaneously, multiple DRVs should be configured.

At a minimum, two DRVs of each size and emulation of device to be swapped must be configured in the Symmetrix. DRV devices are configured as unprotected devices, and can be created as FBA or CKD. DRV devices can be created on EFD, FC, or SATA.

Note: RAID 6 protected devices, and actively protected R1 devices, do not require a DRV when being swapped.

Note: If no DRVs are configured, Optimizer will stop on error on the service processor and will require manual intervention from EMC to be restarted.

Open mirror position

During an Optimizer swap, DRVs behave similarly to traditional TimeFinder/Mirror BCV devices, in that they will occupy a mirror position of the device being protected. As such, for a device to be swapped, it must have an available mirror position for the DRV.

In automatic mode, if a mirror position is not available, Optimizer will stop the swap and rerun data analysis, and attempt to execute any subsequently generated swap lists.

In user approved mode, Optimizer will retry the swap a number of times and eventually stop on error if a mirror position does not become available.

Open Symmetrix configuration lock

Optimizer swaps are performed by completing a series of back-end configuration changes on the Symmetrix array, to relocate the individual physical hypervisors of the devices involved. Because of this, the Symmetrix external configuration lock (lock #15) is held for the entire duration of the swap process.

In automatic mode, if the configuration lock is held by another process, Optimizer will stop the swap and rerun data analysis, and attempt to execute any subsequently generated swap lists.

In user approved mode, Optimizer will retry the swap a number of times and eventually stop on error if the configuration lock does not become available.

Note: While Optimizer holds the configuration lock, other applications will not have the ability to commit configuration changes on the Symmetrix.

Swap eligible devices

The following devices are eligible to be swapped by Optimizer:

- ◆ RAID 1
- ◆ RAID 5 — 3+1 and 7+1
- ◆ RAID 6 — 6+2 and 14+2
- ◆ RDF devices with local RAID protection

The following devices may not be swapped by Optimizer, and as such will not be included in an Optimizer swap list:

- ◆ CKD Stripe (RAID 1/0)
- ◆ BCV
- ◆ DRV
- ◆ SFS
- ◆ WORM
- ◆ IBM i (AS400)
- ◆ TimeFinder/Snap devices — VDEVs, SAVE

- ◆ Virtual Provisioning devices — TDEVs, DATA
- ◆ Standard, unprotected volumes (default, can be modified by EMC)
- ◆ RDF devices with no local RAID protection (default, can be modified by EMC)

Swap configuration considerations

When Optimizer is determining device pairs to be swapped, the following considerations are made with regard to the Symmetrix configuration:

- ◆ Same disk — Two hypers of the same device cannot reside on the same disk
- ◆ Same DA—All mirrors of the same RAID 1 device cannot reside on the same disk adapter (DA) processor.
- ◆ Same port—Two hypers of the same device cannot reside on the same backend DA port.
- ◆ Same loop—Both mirrors of a RAID 1 device cannot reside on the same disk loop.
- ◆ Dual-Initiator—All mirrors of the same RAID 1 device cannot reside on the same dual-initiator pair.
- ◆ Affinity groups—All members of a RAID 1, RAID 5 or RAID 6 device must reside on the same affinity group.

Note: for more information on Symmetrix affinity groups, refer to *Appendix A: Affinity Groups*.

Metadevice considerations

When a metadevice is made up of members configured on either Fibre Channel or SATA disks, the physical hypers that make up each meta member may not be located on the same physical disk as another physical hyper of another member of the same metadevice.

If Optimizer determines that a metadevice member needs to be swapped, a swap candidate must be determined such that the swap will not result in hypers of that member being relocated to a physical disk that contains a hyper of the same metadevice.

Optimizer swap lists

The frequency at which Optimizer generates lists of devices to be swapped depends on the operating mode.

In *Automatic* mode, performance data analysis occurs when the swap time window first opens. Should the data analysis performed generate a swap list, it will be executed at that time.

In *User Approved* mode, performance data analysis occurs once every hour. Should the data analysis performed generate a swap list, it will be made available to the user for review, via SMC or SYMCLI. When reviewed, and approved, the execution of the swap list can be scheduled to occur during the next swap time window, or at a time of the users choosing.

When executed, swaps are performed in the order in which they appear in the swap list. The number of concurrent swaps will be determined by the Maximum Simultaneous Moves parameter setting, and the number of available DRVs in the array.

RAID 1 device swap

Performing a device swap of RAID 1 devices exchanges the physical locations of each physical hyper that constitutes the device. In order to preserve data on both devices involved in the swap, and to maintain full RAID protection on each device during the entire swap, two Dynamic Reallocation Volumes (DRVs) are used.

In the following example, Optimizer has determined that both physical hypers of device 100 will be swapped with the hypers of device 200.

Step 1

The two devices to be swapped are identified, along with the location of the physical hypers of each device




Figure 5 Optimizer RAID 1 device swap—Step 1

Optimizer submits the swap request to the Symmetrix. Upon submission, the external Symmetrix configuration lock (lock 15) is placed on the array in order for a configuration change to be performed.

Step 2

Once the lock has been taken, two DRVs are identified and established to the devices being swapped.




Figure 6 Optimizer RAID 1 device swap—Step 2

Step 3

Once established, the DRVs are synchronized from the hypers of the devices being swapped.




Figure 7 Optimizer RAID 1 device swap—Step 3

Step 4

Once the DRVs are synchronized, one hyper (mirror) of each device to be swapped is made not ready.




Figure 8 Optimizer RAID 1 device swap—Step 4

Step 5

Upon being made not ready, the physical locations of the hyperts being swapped are exchanged.




Figure 9 Optimizer RAID 1 device swap—Step 5

Step 6

The relocated hyperts are then made ready again in their new locations, and synchronized from the other hyperts (including the DRV).




Figure 10 Optimizer RAID 1 device swap—Step 6

Step 7

Steps 4, 5, and 6 are repeated for the other hyper of each device.

Step 8

Once the swap is complete, the DRVs are split from the devices and the configuration lock removed from the Symmetrix.




Figure 11 Optimizer RAID 1 device swap—Step 8**RAID 5 device swap**

As in the case with RAID 1, when two RAID 5 devices are swapped with one another, all physical hypers of the RAID groups are relocated. In the case of RAID 5 (3+1) all 4 hypers are relocated, and in the case of RAID 5 (7+1) all eight hypers are relocated.

To preserve data on both RAID 5 devices involved in the swap, and to maintain full RAID protection on each device during the entire swap, two Dynamic Reallocation Volumes (DRVs) are used.

In swapping two RAID 5 devices, the following steps are performed.

Step 1

Optimizer submits the swap request to the Symmetrix. Upon submission, the external Symmetrix configuration lock (lock 15) is placed on the array in order for a configuration change to be performed.

Step 2

Once the lock has been taken, two DRVs are identified and established to the devices being swapped.

Step 3

Once established, the DRVs are synchronized from the hypers of the devices being swapped.

Step 4

Once the DRVs are synchronized, one hyper of each device to be swapped is made not ready.

Step 5

Upon being made not ready, the physical locations of the hypers being swapped are exchanged, by means of a backend configuration change.

Step 6

The relocated hypers are then made ready again in their new locations, and rebuilt from the other hypers that constitute the RAID group.

Step 7

Steps 4, 5, and 6 are repeated for the remaining hypers of each device.

Step 8

Once all the RAID members have been relocated, the DRVs are split from the devices and the configuration lock removed from the Symmetrix.

RAID 6 device swap

When swapping two RAID 6 devices, all physical hyperrs of the RAID groups are relocated. In the case of RAID 6 (6+2) all 8 hyperrs are relocated, and in the case of RAID 6 (14+2) all 16 hyperrs are relocated.

As RAID 6 is designed to withstand the failure of up to two physical disks that contain hyperrs of the RAID group. As such, when a hyper is made not ready in order to be relocated, RAID protection is still afforded within the RAID group. Because of this capability, DRVs are not used in the swapping of RAID 6 devices.

In swapping two RAID 6 devices, the following steps are performed.

Step 1

Optimizer submits the swap request to the Symmetrix. Upon submission, the external Symmetrix configuration lock (lock 15) is placed on the array in order for a configuration change to be performed.

Step 2

Once the configuration lock has been taken, one hyper of each device to be swapped is made not ready.

Step 3

Upon being made not ready, the physical locations of the hyperrs being swapped are exchanged, by means of a backend configuration change.

Step 4

The relocated hyperrs are then made ready again in their new locations, and rebuilt from the other hyperrs that constitute the RAID group.

Step 5

Steps 2, 3, and 4 are repeated for the remaining hyperrs of each device.

Step 6

Once all the RAID members have been relocated, the configuration lock is removed from the Symmetrix.

Device swap control

Once an Optimizer swap begins, there is no interface which allows the swap to be stopped. Any swap that has started will need to be completed. This is true even if the swap will continue beyond the end of a defined swap time window.

Advanced data movement modes

While Optimizer will typically operate in one of two operating modes (automatic or user approved) there are two additional modes—manual and rollback. Manual mode allows users to manually swap a pair of devices, while rollback mode allows previously completed swaps to be reversed.

Manual device swap

Aside from recommending device swaps to improve overall system performance, Optimizer also allows the user to manually define device swaps to be performed. User defined swaps can be used to move heavier utilized devices to newly added disks, or to undo a specific device swap.

Manual device swaps can take place within a single physical disk group, but can also swap devices between two separate disk groups.

When defining a manual swap, the devices involved must be of the same size and RAID protection. Also, if swapping devices between disk groups, the disk groups must contain similar disk types, based on technology, capacity, and speed.

Due to affinity group rules, all physical hypervisors of the devices must be included when defining a manual swap. There are two exceptions to this; if the swap is exchanging the location of two hypervisors on the same physical disk, or if the manual swap is being used to fix a broken affinity group.

Note: for more information on affinity groups please refer to Appendix A: Affinity Groups.

When a manual swap is submitted, Optimizer will validate the swap to ensure that Symmetrix configuration rules are not violated. However, Optimizer will not examine the potential impact on performance of the completed swap. As such, it is the user's responsibility to define swaps that will not have a negative impact on the array.

Device swap rollback

In certain circumstances, the Symmetrix can provide the ability to undo or "rollback" device swaps to return the Symmetrix configuration to a point at which it was prior to an Optimizer swap or set of swaps.

The rollback of a device swap, or swaps, is performed by reversing the original swap of the device hypervisors. The rollback of a swap results in all device hypervisors being returned to their original locations.

The Symmetrix Optimizer rollback is an “all or nothing” feature. All swaps are reversed from the present to a selected prior point-in-time. For example, in the timeline shown in Figure 12, Optimizer has performed swaps each night for six straight days. If it was decided that the swaps that occurred on January 6th were to be rolled back, the swaps that occurred on January 7th and 8th would also have to be rolled back.



Figure 12 Optimizer swap timeline

Replication interoperability

Symmetrix Optimizer is fully interoperable with all Symmetrix replication technologies—EMC SRDF[®], TimeFinder[®]/Mirror, TimeFinder/Clone, TimeFinder/Snap, and Open Replicator. Any active replication on a Symmetrix device remains intact while the device is being swapped. Similarly, all incremental relationships are maintained for the swapped devices. However, what must be kept in mind is that Optimizer device swaps will consume array resources, so planning must be performed in determining when device movements should occur so as to minimize impact on other replication processes.

SRDF

By default, Optimizer will swap locally protected R1 and R2 devices. Swaps can take place while the device pairs are actively replicating and while the pairs are in any operating mode—synchronous, asynchronous, adaptive copy.

An advanced Optimizer configuration setting will allow SRDF devices with no local protection to also be swapped by Optimizer.

SRDF pair state

While Optimizer has the ability to swap both R1 and R2 devices, Optimizers behavior during a swap of an SRDF device may differ depending on the state of the SRDF pairing. If a locally protected R1 device is actively replicating, and so is also protected remotely,

Optimizer will not establish a DRV as the device would still be protected in the event of a disk failure during the swap.

Concurrent SRDF

Additional rules apply to devices configured as concurrent R1. If the concurrent R1 device is locally protected as RAID 1 or RAID 5, then there is no available mirror position to allow a DRV to be established. If both remote legs of the concurrent R1 are suspended or split, then Optimizer will be unable to swap the device.

TimeFinder

Optimizer can swap any TimeFinder source device—Mirror, Clone, or Snap. Standard devices that are the target of a Clone session can also be swapped.

Concurrent BCVs

When used with native TimeFinder/Mirror, unprotected or RAID 1 BCVs will occupy a mirror position of the source device when established. If two BCVs are established concurrently to a device then there will be no available position to allow a DRV to be established.

RAID 5 and RAID 6 protected BCVs can only operate in TimeFinder/Clone emulation mode and do not occupy a mirror position when established. As such, a device with two RAID5 or two RAID 6 BCVs established concurrently can be swapped by Optimizer.

SRDF and BCVs

A RAID 1 or RAID 5 protected R1 device, with an unprotected or RAID 1 BCV established in native TimeFinder/Mirror mode, may not be swapped by Optimizer as a DRV cannot be established due to the lack of an available mirror position.

Note: if the R1 device is actively replicating to the R2 device a DRV may not be required, and so Optimizer may still be able to swap the device while a BCV is established.

Open Replicator for Symmetrix

The control device in an Open Replicator session, push or pull, hot or cold, can be swapped by Optimizer.

Virtual LUN Technology

Virtual LUN technology enables transparent, non-disruptive data migration between storage types that have the same RAID protection scheme. This migration feature can be used to populate newly added disk drives or move devices between high performance and high capacity disks, delivering tiered storage capabilities within a single Symmetrix array.

The migration of a Symmetrix device involves moving the device from its current location on a set of physical disks, to a new location on a different set of disks, while preserving the RAID protection type of the device. The destination of the migration can be specified as a group of similar disks. Because of the use of physical disk groups on the Symmetrix, high-performance disks are grouped separately from high-capacity disks, making it easy to migrate devices to the appropriate tier of storage.

Virtual LUN migration contains the following features:

- ◆ One or more devices are moved to one or more target disks.
- ◆ All physical hypers of the devices are moved.
- ◆ The RAID protection type of the device will be preserved through the completion of the migration.

Note: The information contained in this document regarding Virtual LUN refers only to Symmetrix DMX series arrays. For information on Enhanced Virtual LUN for Symmetrix V-Max, please refer to the *Best Practices for Nondisruptive Tiering via EMC Symmetrix Virtual LUN Technical Note* available via <http://powerlink.emc.com>.

Virtual LUN migrations

Virtual LUN migrations can be managed with the Solutions Enabler SYMCLI `symoptmz` command, or from Symmetrix Management Console (SMC). The control host for the migration must be locally connected to the Symmetrix on which the migration is being performed.

Note: Virtual LUN technology requires a minimum of Enginuity 5772.83.75, and Solutions Enabler 6.4.2.

When performing Virtual LUN migrations, data can be migrated to either unconfigured or configured space. In each case, the devices to be migrated and the target physical disks must be provided to Optimizer.

Source criteria

From SYMCLI, the source devices for a Virtual LUN migration can be specified in one of three ways:

- ◆ A device group
- ◆ A range of devices
- ◆ A comma separated list of devices

From SMC, the source devices for a Virtual LUN migration can be specified as:

- ◆ A device group
- ◆ Manual selection

Target criteria

From both SYMCLI and SMC the target storage of a migration can be specified as:

- ◆ A physical disk group
- ◆ A list of physical disks

Note: Migrated devices are distributed across specified recipient disks by a sophisticated layout algorithm. Final placement of the volumes is not known in advance.

Migration to unconfigured space

In the case of the migration to unconfigured space, the device is migrated by moving each of the devices physical hypers, one at a time, to free space on the group of disks specified as the destination. As each hyper is moved, free space is then created on the physical disks from which the device is being migrated from.

The target disk group, or list of target disks, specified must contain sufficient physical disks to support the RAID protection of the device being migrated. Also, there must be enough available free space to migrate the device to.

Once the data transfer has been completed, the physical space previously associated with the migrated device will be returned to the free pool.

Migration to configured space

When performing a migration to configured space, a target device is selected by Engenuity on the target physical disks, and a swap is performed with the device being migrated.

The device selected as the target will be a configured, but unused device—either unmapped and/or unmasked. The target device will also be of the same size and RAID protection as the device being migrated.

Following the migration, the device selected as the target will occupy the storage originally associated with the device being migrated.

Migration configuration requirements

Virtual LUN technology is a feature of Symmetrix Optimizer and as such it has similar configuration requirements in order to perform a migration. These are:

- ◆ Dynamic Reallocation Volumes (DRVs)
- ◆ Open mirror position
- ◆ Open Symmetrix configuration lock

For more information on these requirements, please see *Swap configuration requirements* on page 18.

Virtual LUN DRV requirements

In order to migrate devices, DRVs are required to provide full data availability during the migration. The DRV requirements for Virtual LUN are the same as those for Optimizer. For more information, please see *Dynamic Reallocation Volumes (DRVs)* on page 18.

When performing a migration to configured space, the locations of two Symmetrix devices are swapped and, as such, two DRVs are required to complete each migration.

When a migration to unconfigured space is performed, only one Symmetrix device is affected by the relocation. Therefore, only one DRV is required to complete the migration.

Migration eligible devices

The following devices are eligible to be migrated by Virtual LUN:

- ◆ RAID 1
- ◆ RAID 5 — 3+1 and 7+1
- ◆ RAID 6 — 6+2 and 14+2
- ◆ SRDF devices (R1 and R2)

The following devices may not be migrated:

- ◆ CKD Stripe (RAID 1/0)
- ◆ BCV

- ◆ DRV
- ◆ SFS
- ◆ WORM
- ◆ IBM i (AS400)
- ◆ TimeFinder/Snap devices — VDEVs, SAVE
- ◆ Virtual Provisioning devices — TDEVs, DATA

Device migration control

Once a Virtual LUN migration begins, there is no interface which allows the migration to be stopped. Any migration that has started will need to be completed.

In the event that an executed migration is no longer wanted, the migration can be reversed using the Symmetrix Optimizer rollback feature.

Migrating metadevices

Virtual LUN allows for the migration of metadevices between disk technologies on Symmetrix DMX. Both SYMCLI and SMC allow for the entire metadevice to be migrated, or for individual metadevice members to be migrated.

Migrating the entire metadevice

Symmetrix configuration rules mandate that all members of a metadevice must be configured on similar disks, in terms of technology, size, and rotational speed. In addition to this rule, the metadevice members must be configured within the same physical disk group.

When performing a migration of a metadevice from one disk technology to another or from one disk group to another, all members of the metadevice must be migrated within the same swap list.

When migrating the entire metadevice, only the metadevice head device number should be specified.

Note: If a metadevice has more members than can be moved concurrently by Optimizer, then the Symmetrix will allow metadevice members to exist on different technologies, or within different disk groups, on a temporary basis—until the migration is completed.

Migrating individual metadevice members

Within both the SYMCLI and SMC it is possible to specify an individual metadevice member to be migrated to a new set of disks. However,

migrations of metadvice members can only take place within the physical diskgroup the device is currently configured. This is due to the configuration requirement that all metadvice members be configured within the same physical diskgroup.

While individual members can be migrated, it is not possible to specify just the head of the metadvice to be migrated. As previously discussed, specifying the metadvice head as the source of a migration will cause all members of the metadvice to be migrated.

Note: It is possible for an EMC representative to migrate only the metadvice head, using the Optimizer client on the Symmetrix service processor.

Virtual LUN replication interoperability

Virtual LUN is fully interoperable with all Symmetrix replication technologies—SRDF, TimeFinder/Mirror, TimeFinder/Clone, TimeFinder/Snap, and Open Replicator. Any active replication on a Symmetrix device remains intact while the device is being migrated. Similarly, all incremental relationships are maintained for the migrated devices.

Similar caveats apply, however, to that of traditional Optimizer operations. For full details, please refer to Replication interoperability on page 27.

Fully Automated Storage Tiering (FAST) for DMX

EMC Symmetrix FAST for DMX, based on Symmetrix Optimizer technology, enables customers to realize the benefits of tiered storage with less time and effort spent on performance tuning, management, and monitoring. FAST automatically relocates data across storage types as needed to optimize the overall performance of the system, making the best use of high-performance enterprise Flash drives and high-capacity SATA drives. FAST makes it easy and cost effective to provide the right levels of service to critical applications. With FAST, the workload on an array can be balanced across all drives.

The primary benefits of FAST for DMX include:

- ◆ Reduced management costs due to reduced time required to monitor and optimize IT infrastructures
- ◆ Improved application performance at the same cost, or the same application performance at lower cost. Costs include acquisition (both HW and SW), space/energy, and management time

Implementation

A new type of device movement rule, called a Cross Disk Group rule, has been added to Optimizer to support FAST for DMX. When created, the rule is associated with a list of Symmetrix devices, an Optimizer device list.

The Optimizer device list can contain devices of multiple sizes and RAID protection types, configured across multiple physical disk groups and drive types. However, when associated with a Cross Disk Group rule, Optimizer will only perform swaps among devices of the same size and protection type.

When created, the Cross Disk Group rule will serve as an implicit allocation policy for the Symmetrix device lists associated with the rule. If a device list containing ten devices of the same size and RAID protection is created—with, for example, two devices configured on EFD, four devices configured on FC, and four devices configured on SATA—FAST for DMX will attempt to optimize the workload on each of those ten devices by swapping devices with other devices on the same list. While the drive type the individual devices are configured on may change, the device list will always be configured with the same initial layout.

Multiple Cross Disk Group rules may be created by creating and associating additional device lists. However, FAST for DMX will only perform swaps of devices pairs contained within the same list, and that are of the same size and RAID protection type.

Management and operation of FAST for Symmetrix DMX is provided by the Symmetrix Management Console (SMC).

Note: Management of FAST for DMX requires a minimum of Engenuity 5773.155.107 and SMC 7.1.

FAST algorithms

FAST for DMX uses two distinct algorithms when determining the appropriate drive technology a device should be configured on. The algorithms, in order of probability, are:

- ◆ EFD promotion/demotion algorithm
- ◆ FC/SATA cross-tier algorithm

Enterprise Flash Drive promotion/demotion algorithm

The goal of the Enterprise Flash Drive (EFD) promotion/demotion algorithm is to maximize EFD utilization within the array. It only considers whether a device should, or should not, reside on EFDs.

When run, this algorithm models EFD performance for each device under FAST control. This modeling is based on each device's write percentage, average I/O size (in blocks) and total I/Os per second (IOPS). The I/O response time, or utilization, of a Flash drive depends largely on the write percentage, in addition to IOPS and I/O size.

In order to determine if a device is "busy," the number of reads per GB is used. With this, a weighted read density is calculated by combining the average read I/Os per GB (read density) and the average read I/O ratio.

When complete, the algorithm will list all the devices in the array in order of their EFD weighted read density. Those devices with a higher read density will be considered to be better served being configured on EFDs.

When generating the swap list, the algorithm will attempt to find all possible swap suggestions that will move busier devices to EFDs while demoting less busy devices to the lower tiers.

FC/SATA cross-tier algorithm

The goal of the FC/SATA cross-tier algorithm is to balance utilization across Fibre Channel and SATA technologies. Only devices configured on Fibre Channel and SATA drives are considered.

When run, this algorithm will calculate the disk service time for each device configured on FC and SATA disks. The disk service time is calculated from the following metrics:

- ◆ Seek time—The time it takes the disk arm to move and position the disk head on the correct track.
- ◆ Latency time—The delay for disk rotation.
- ◆ Transfer time—The time that it takes the disk to transfer the data from/to the disk.

Devices are sorted by disk service time, in descending order and the most utilized devices will be moved to the least utilized disks.

When generating the swap list, the algorithm will attempt to find all possible swap suggestions between higher and lower utilized devices.

Note: Even when a cross disk group rule exists, the traditional Optimizer algorithm will also be used to balance workloads within a physical disk group

for the devices associated with the rule.

FAST configuration requirements

FAST for DMX is a feature of Symmetrix Optimizer and as such it has similar configuration requirements in order to perform a migration. These are:

- ◆ Dynamic Reallocation Volumes (DRVs)
- ◆ Open mirror position
- ◆ Open Symmetrix configuration lock

For more information on these requirements, please see *Swap configuration requirements* on page 18.

FAST DRV requirements

In order to swap two devices on differing drive technologies, DRVs are required to provide full data availability on both devices during the swap. The DRV requirements for FAST for DMX are the same as those for Optimizer. For more information, please see *Dynamic Reallocation Volumes (DRVs)* on page 18.

FAST eligible devices

The following devices are eligible to be managed by FAST, under a cross disk device rule:

- ◆ RAID 1
- ◆ RAID 5 — 3+1 and 7+1
- ◆ RAID 6 — 6+2 and 14+2
- ◆ SRDF devices (R1 and R2)

The following devices may not be placed in a device list managed by a cross disk device rule:

- ◆ CKD Stripe (RAID 1/0)
- ◆ BCV
- ◆ DRV
- ◆ SFS
- ◆ WORM
- ◆ IBM i (AS400)
- ◆ TimeFinder/Snap devices — VDEVs, SAVE
- ◆ Virtual Provisioning devices — TDEVs, DATA

FAST device movement

Device movements performed by FAST for DMX are accomplished by swapping the physical locations of two devices associated with the same Cross Disk Group rule.

The device swaps that take place follow the same process as performed during a normal Optimizer swap. For more details see *Optimizer device swaps* on page 18.

Device movement control

Once a FAST swap has begun, there is no interface which allows the swap to be stopped. Any swap that has started will need to be completed, even if the swap time window has expired.

In the event that a FAST swap needs to be reversed, the Symmetrix Optimizer rollback feature can be used.

FAST movement of metadevices

FAST for DMX supports the relocation of metadevices between disk technologies.

Symmetrix configuration rules mandate that all members of a metadevice must be configured on similar disks, in terms of technology, size, and rotational speed. In addition to this rule, the metadevice members must be configured within the same physical disk group.

When FAST determines that a metadevice needs to be relocated to another disk group within the array all members of the metadevice must be swapped within the same swap list.

Note: if a metadevice has more members than can be moved concurrently by Optimizer, then the Symmetrix will allow metadevice members to exist on different technologies, or within different disk groups, on a temporary basis—until the migration is completed.

FAST interoperability

FAST is fully interoperable with all Symmetrix replication technologies—SRDF, TimeFinder/Mirror, TimeFinder/Clone, TimeFinder/Snap, and Open Replicator. Any active replication on a Symmetrix device remains intact while the device is being migrated. Similarly, all incremental relationships are maintained for the migrated devices.

Similar caveats apply, however, to that of traditional Optimizer operations. For full details, please refer to Replication interoperability on

page 27.

FAST SRDF considerations

While there are some restrictions in the ability to swap SRDF devices with FAST, what must be considered is that device movements are restricted to the array upon which FAST is operating. That is to say that if an RDF1 device is moved between two tiers, FAST will not automatically perform a corresponding move of the respective RDF2 device on a remote array. This means that, in a SRDF failover scenario, the remote Symmetrix array will have different performance characteristics than the local, production array being failed over from.

Also, in an SRDF/Asynchronous environment, FAST swaps on the production R1 array could result in an unbalanced configuration between R1 and R2 (where the performance characteristics of the R2 device is lower than that of the paired R1 device).

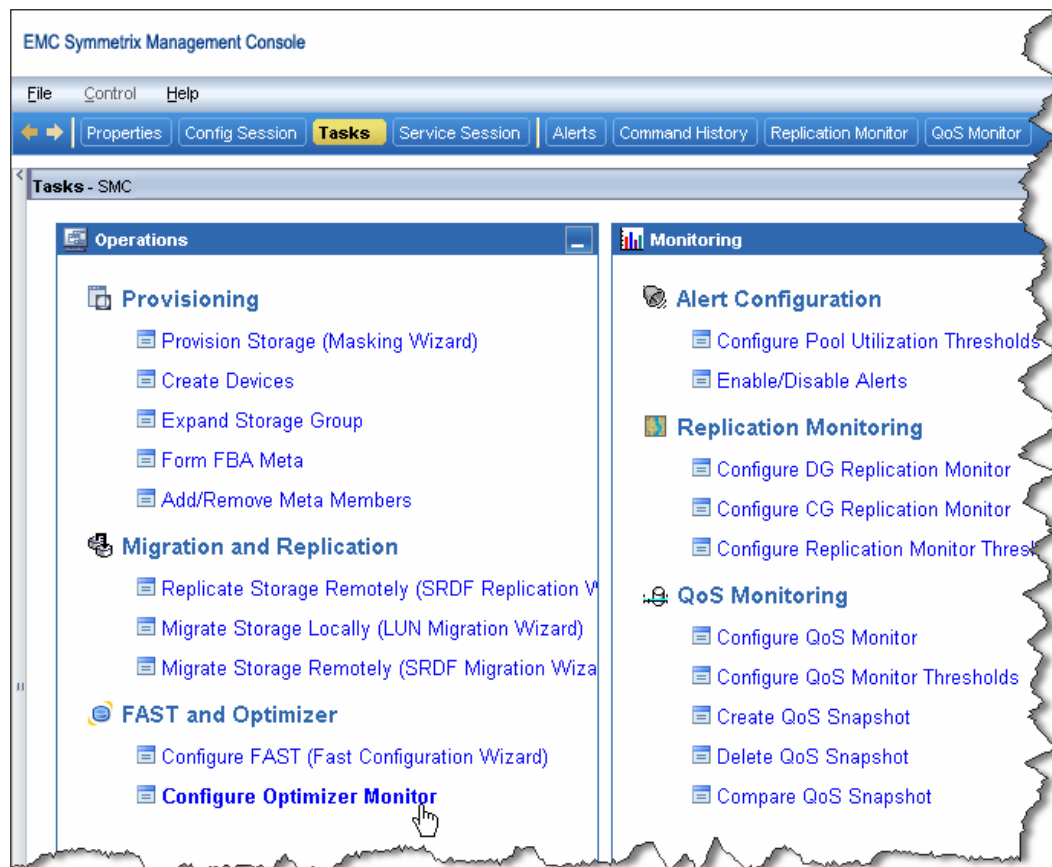
Management Interface: SMC

Symmetrix Management Console 7.1 adds features and functionality for managing FAST. All FAST managed objects and associations may be created and managed directly through menu options. Configuration settings and time windows may also be created through the menu driven interface.

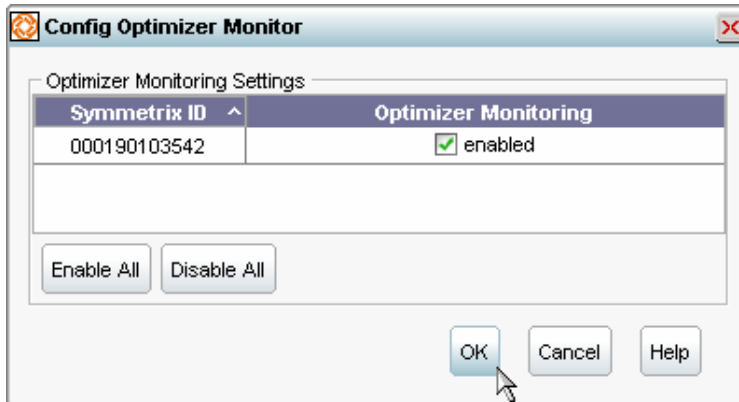
The following sections detail the use of the individual menu options within SMC for managing and controlling FAST.

Enabling Symmetrix Optimizer monitoring

Prior to managing Optimizer from SMC, monitoring of the Optimizer feature must first be enabled. To enable it, go to the Tasks view in SMC and click on Configure Optimizer Monitor in the Operations panel.



In the resulting dialog box, click in the checkbox to the left of enabled in the Optimizer Monitoring field for the appropriate Symmetrix storage arrays and click OK.

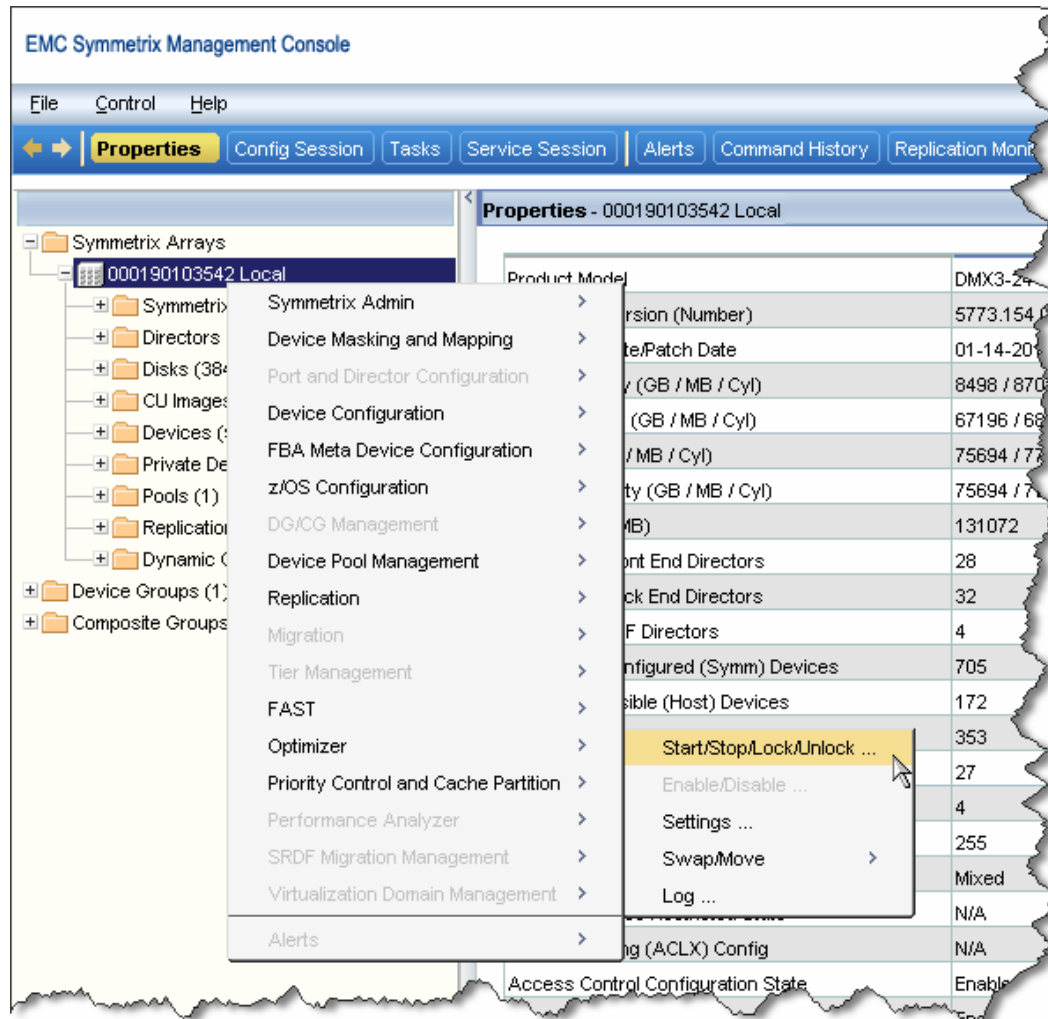


If more than one Symmetrix is being managed, the Enable All button can be clicked to place a check mark next to each Symmetrix serial number prior to clicking OK.

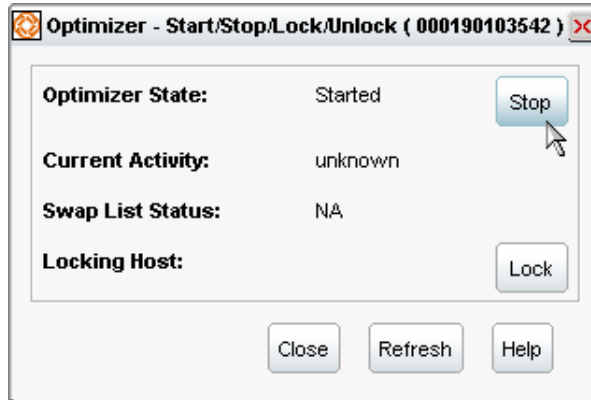

Stopping/Starting Symmetrix Optimizer

In order for Optimizer to perform data analysis and generate swap lists, the Optimizer controller must be started on the Symmetrix.

To check the current state of Optimizer, right-click on the appropriate Symmetrix serial number in the left-hand navigation pane, and select Optimizer, Start/Stop/Lock/Unlock.



The state of the Optimizer can be changed by clicking the Start or Stop button that is available.



Configuring the Symmetrix Optimizer controller

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

- ◆ Operating Mode
- ◆ Maximum Moves Per Day
- ◆ Maximum Simultaneous Moves
- ◆ Workload Analysis Period
- ◆ Initial Period
- ◆ Start Up Mode
- ◆ Hot Spot Analysis Mode
- ◆ Performance Time Window
- ◆ Device Movement Time Window
- ◆ Device Attributes

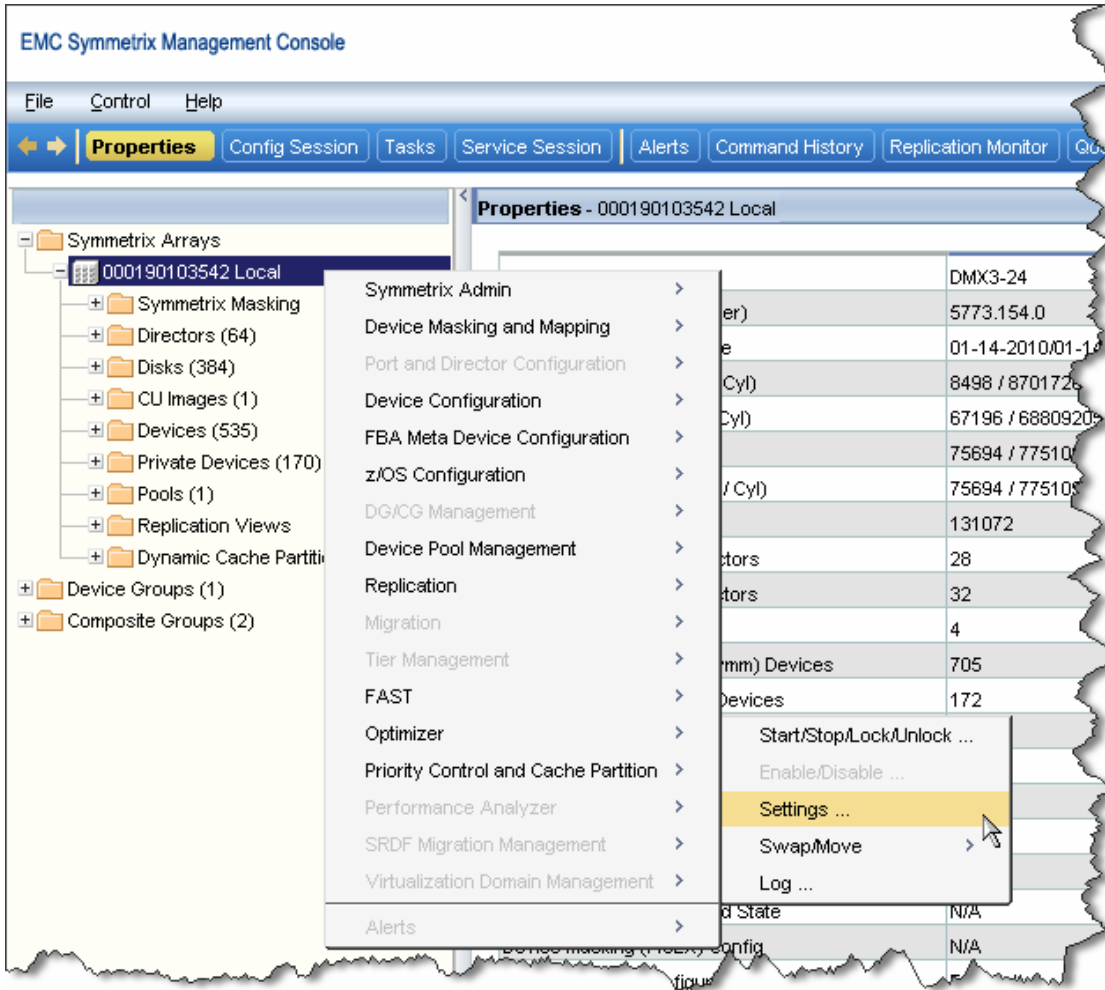
- ◆ Lists & rules

The following sections detail the SMC interface used to list modify each of these settings.

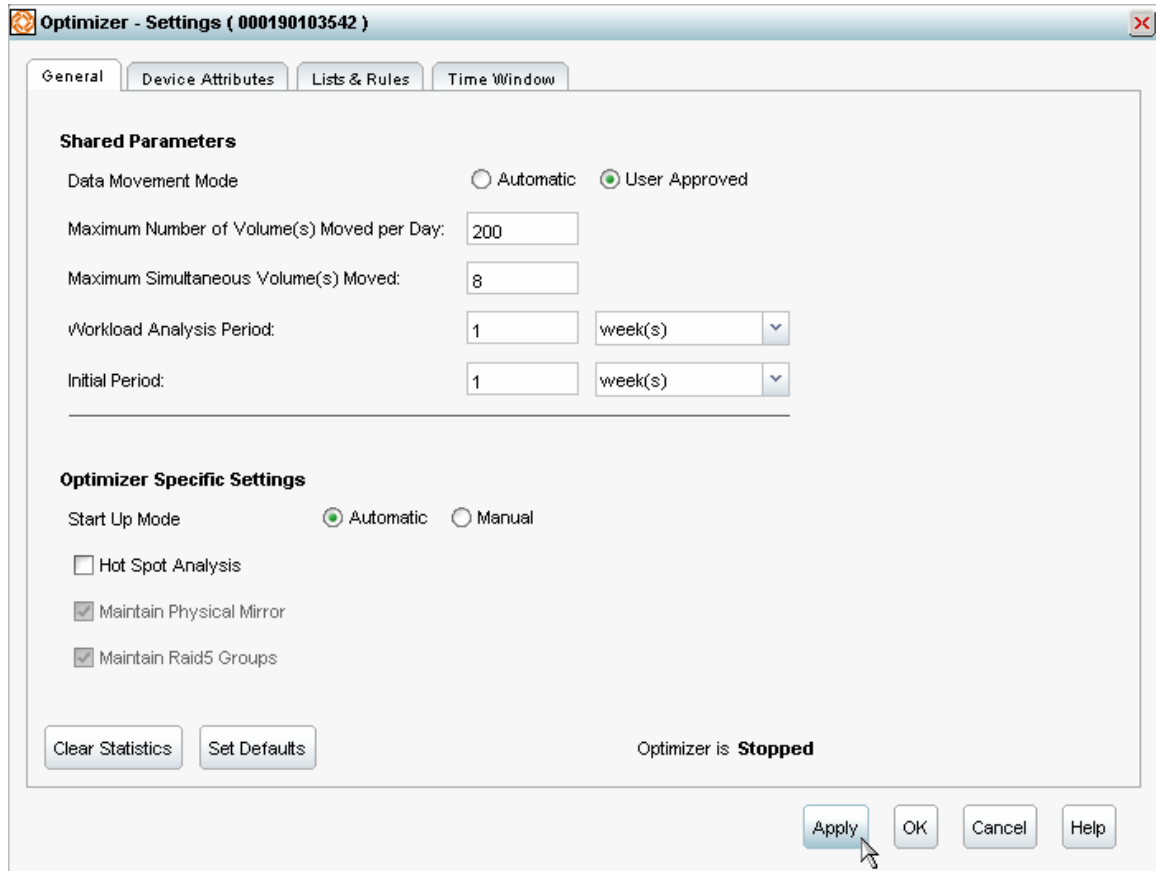
Note: prior to changing any Symmetrix Optimizer controller setting, Optimizer must first be Stopped.

Optimizer controller Settings

The same interface is used to display the existing Optimizer settings and to modify them. To launch the Optimizer settings display, right-click on the appropriate Symmetrix serial number in the left-hand navigation pane, then select Optimizer, and Settings.



The following dialog box will be displayed.



Setting the data movement mode

There are two possible values for the data movement mode: Automatic and User Approved. The default value is Automatic.

To set the mode to User Approved, click on the radio button to the left of User Approved, and click Apply.

Setting the Maximum Moves Per Day

Maximum Device Moves Per Day can be set between 2 and 200. The default value is 200.

To set the Maximum Device Moves Per Day, type the desired value in the field to the right of Maximum Number of Volume(s) Moved per Day, and click Apply.

Setting the Maximum Simultaneous Moves

Maximum Simultaneous Device Moves can be set between 2 and 32. The default value is 8.

To set the Maximum Simultaneous Device Moves, type the desired value in the field to the right of Maximum Simultaneous Volumes Moved, and click Apply.

Setting the Workload Analysis Period

The Workload Analysis Period can be set between 1 hour and 4 weeks. The default value is 1 week.

To set the Workload Analysis Period, type the desired value in the field to the right of Workload Analysis Period, select the desired unit of time, and click Apply.

Setting the Initial Period

The Initial Period can be set between 1 hour and 4 weeks. The default value is 1 week (168 hours).

To set the Initial Period, type the desired value in the field to the right of Initial Period, select the desired unit of time, and click Apply.

Start Up Mode

Start Up Mode can be automatic or manual. The default is automatic.

To set the start up mode to manual, click on the radio button to the left of Manual and click Apply.

Enabling Hot Spot Analysis

Hot spot analysis can be enabled or disabled. The default is disabled.

To enable hot spot analysis, click on the checkbox to the left of Hot Spot Analysis, and click Apply.

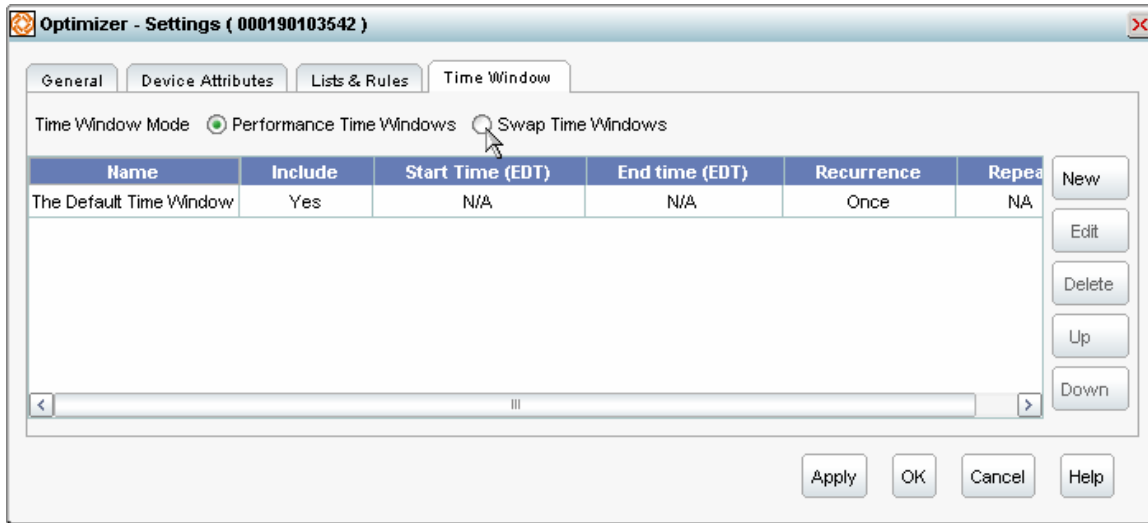
Modifying all Optimizer controller settings

Alternatively, all settings can be modified at the same time, by editing each value and clicking Apply, or OK.

Optimizer time window list information

To view, create, edit, or delete Optimizer time windows the same Optimizer Settings dialog box as shown in the previous section is used. However, in this case the Time Window tab will be selected.

To launch the Optimizer settings display, right-click on the appropriate Symmetrix serial number in the left-hand navigation pane, then select Optimizer, and Settings. After clicking the Time Window tab on the resulting dialog box, the following will be displayed.

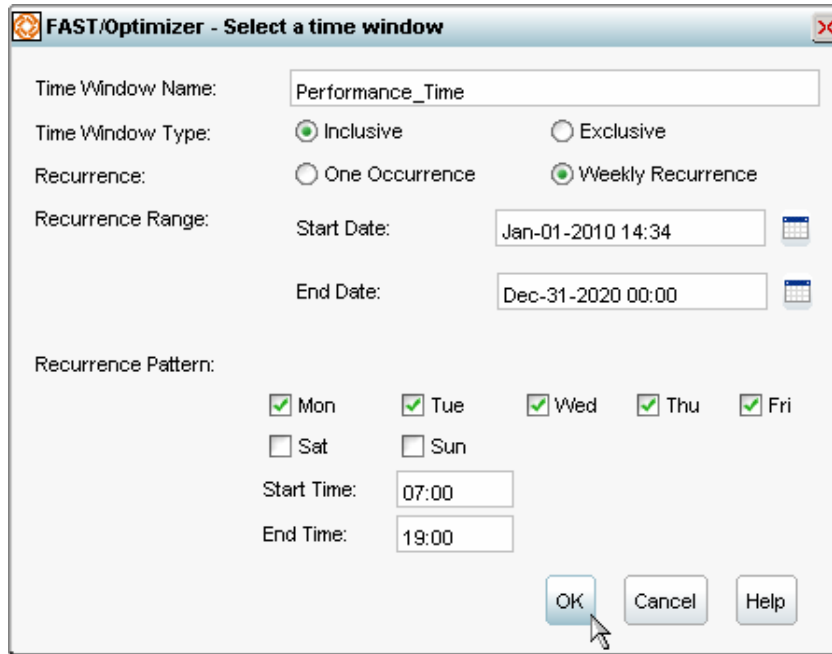


To view existing time windows of each type, the radio buttons next to Performance Time Windows or Swap Time Windows may be clicked. In the above output, only the default time performance time window has been defined.

Creating an Optimizer performance time window

To create a performance time window, ensure that the radio button to the left of Performance Time Windows has been clicked, and click on New.

As an example, to create a performance time window that will include for analysis data samples collected between 7 a.m. and 7 p.m. on each day, Monday through Friday, fill out the information as shown below, and click OK.



Creating an Optimizer swap time window

To create a swap time window, ensure that the radio button to the left of Swap Time Windows has been clicked, and click New.

As an example, to create a swap time window that will allow swaps to be performed between 1 a.m. and 5 a.m. on Saturdays and Sundays, fill out the information as shown below, and click OK.

FAST/Optimizer - Select a time window

Time Window Name:

Time Window Type: Inclusive Exclusive

Recurrence: One Occurrence Weekly Recurrence

Recurrence Range: Start Date:

End Date:

Recurrence Pattern:

Mon Tue Wed Thu Fri

Sat Sun

Start Time:

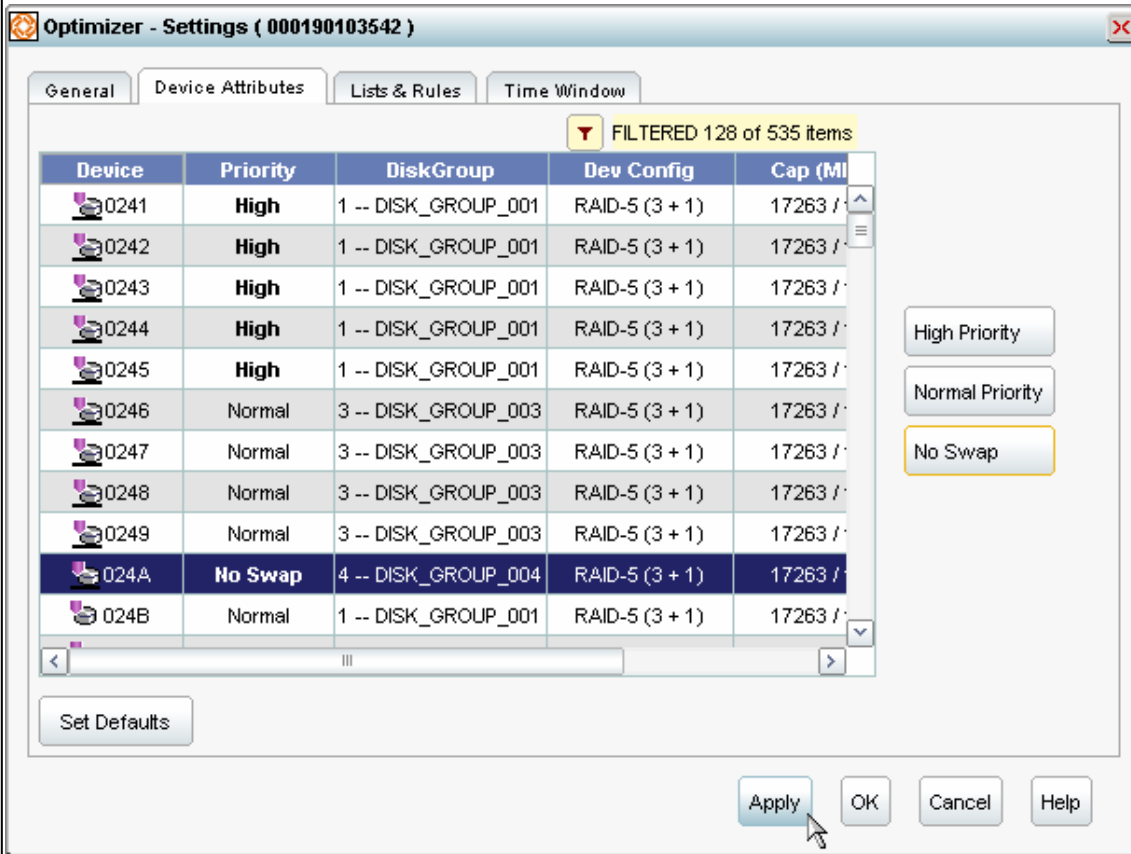
End Time:

Setting device attributes

To set device attributes on one or more devices, the same Optimizer Settings dialog box as shown in the previous sections is used. However, in this case the Device Attributes tab will be selected.

To change the priority on a device, the device should first be highlighted. Once selected, the appropriate button—high priority, normal priority, no swap – should be clicked. It is possible to select multiple devices, by holding down the CTRL or SHIFT key, and assign the same priority.

When the device attributes have been selected for each device, the priorities can be assigned by clicking on either the Apply or OK button.

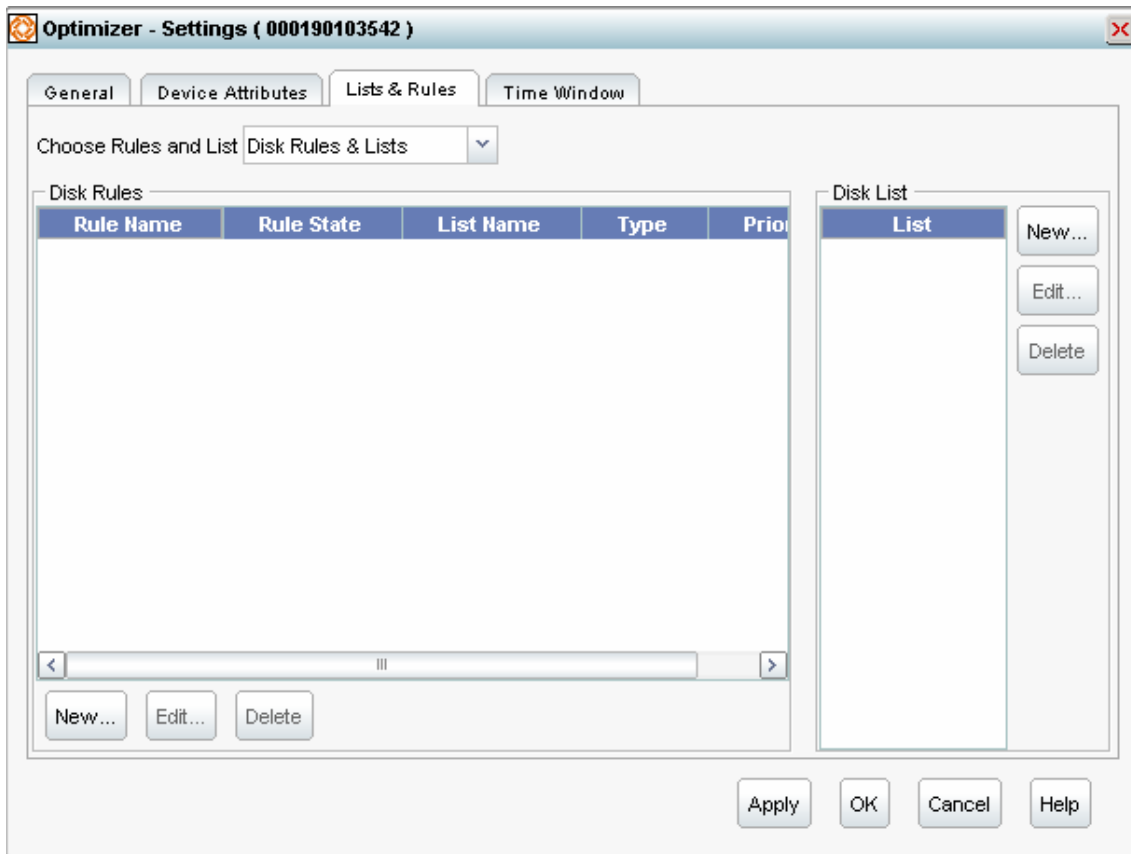


Creating Optimizer lists and rules

To create Optimizer lists and rules, the same Optimizer Settings dialog box as shown in the previous sections is used. However, in this case the List & Rules tab will be selected.

Disks rules and device rules will be created separately. A drop down menu located to the right of Choose Rules and List allows users to toggle between managing disk rules and lists or device rules and lists.

Any existing rules will be displayed in the left hand panel, while existing lists are displayed to the right.

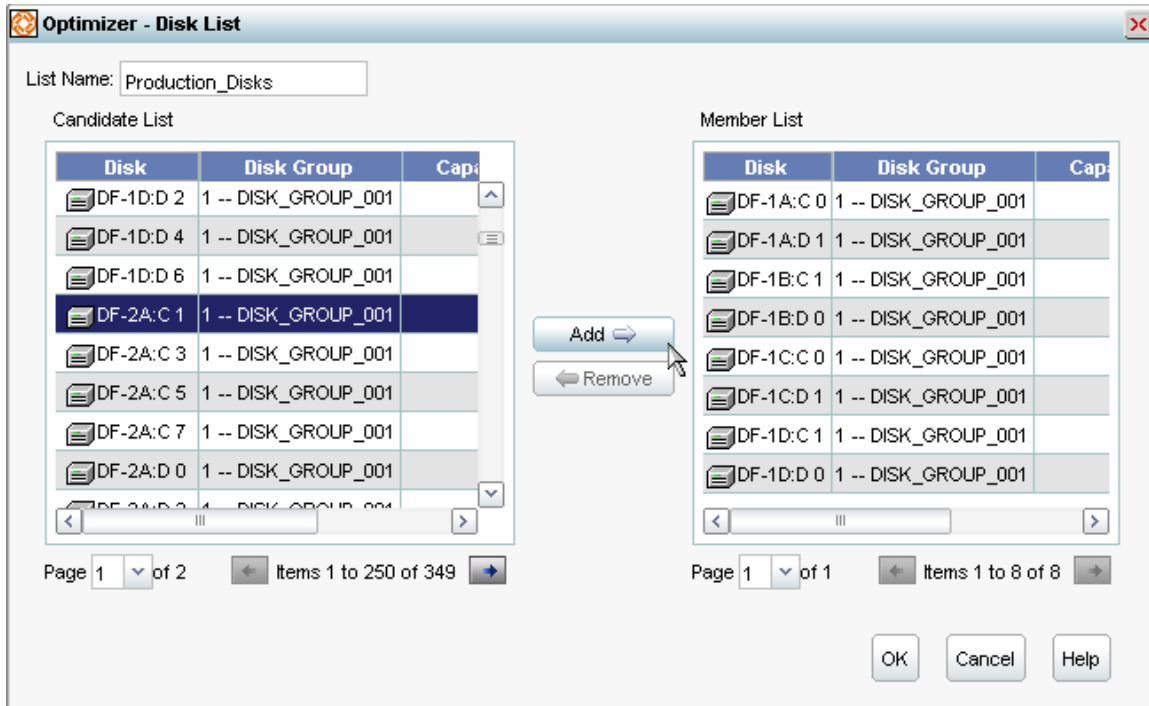


Creating an Optimizer disk list

To create an Optimizer disk list, the Choose Rules and List drop-down menu should be set to Disk Rules & List. With this selected, click on the New button located to the right of the Disk List list. When clicked the following dialog box will be displayed.

Once a list name has been provided, physical disks in the array can be added to the list by selecting them in the left-hand panel and clicking the Add button.

Disks can be removed from the list by selecting them in the right-hand panel and clicking Remove.



Once all the required disks have been added to the list the OK button should be clicked. Clicking OK returns the user to the main settings dialog.

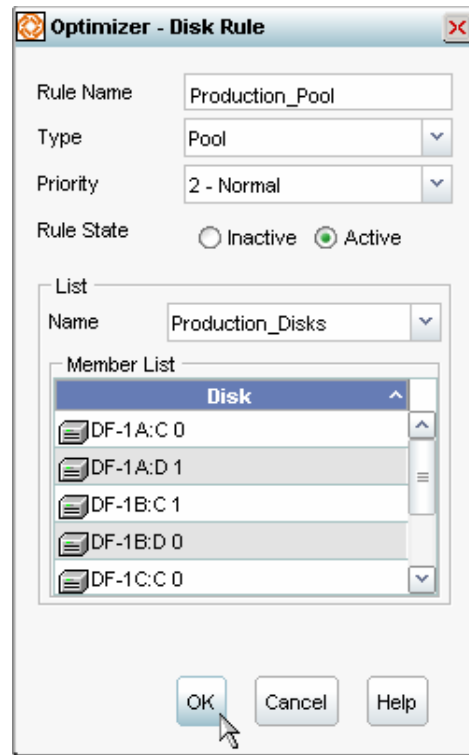
Creating an Optimizer disk rule

To create an Optimizer disk rule, click on the New button located below the Disk Rules list. When clicked the following dialog box will be displayed.

After providing a rule name, the rule type needs to be chosen—Pool, Exclude, or Intra.

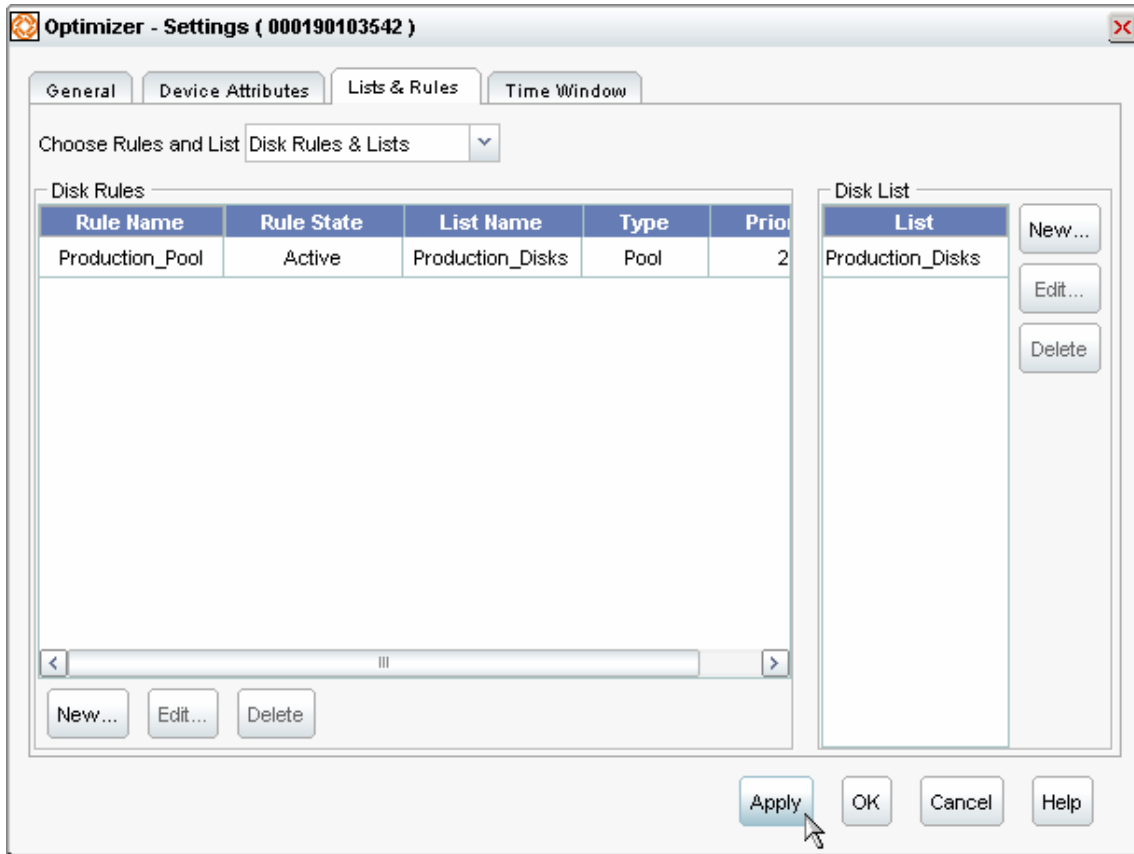
A priority can be assigned to the rule, and the rule can be created as active or inactive.

Finally, the list of disks the rule will be applied to will be chosen. All available disk lists will be displayed in a drop down list.



Clicking the OK button returns the user to the main Settings dialog.

To execute the creation of new disk rules and lists, the Apply button should be clicked on the Settings dialog.

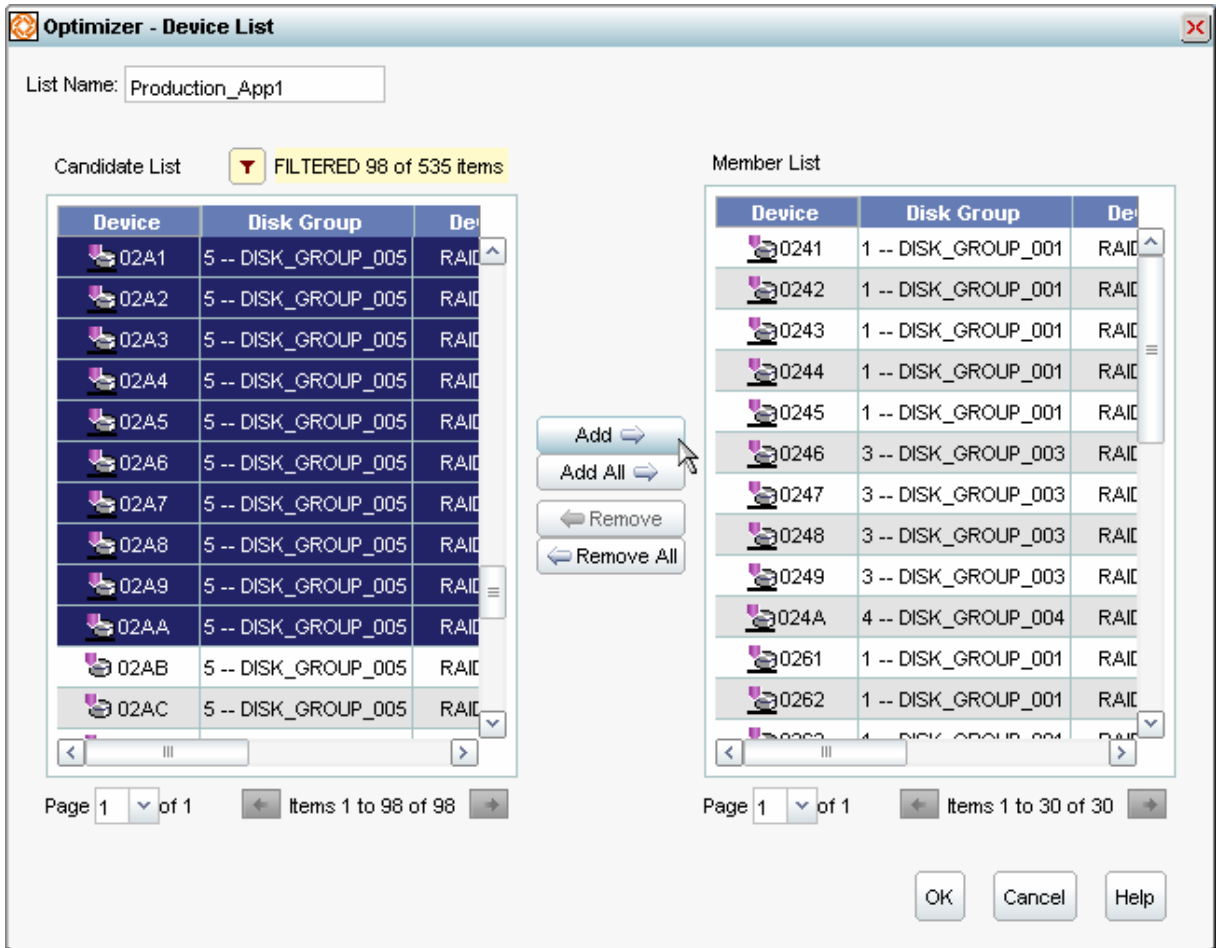


Creating an Optimizer device list

To create an Optimizer device list, the Choose Rules and List drop-down menu should be set to Device Rules & List. With this selected, click on the New button located to the right of the Device Lists list. When clicked the following dialog box will be displayed.

Once a list name has been provided, swappable devices in the array can be added to the list by selecting them in the left-hand panel and clicking the Add button.

Devices can be removed from the list by selecting them in the right-hand panel and clicking Remove.



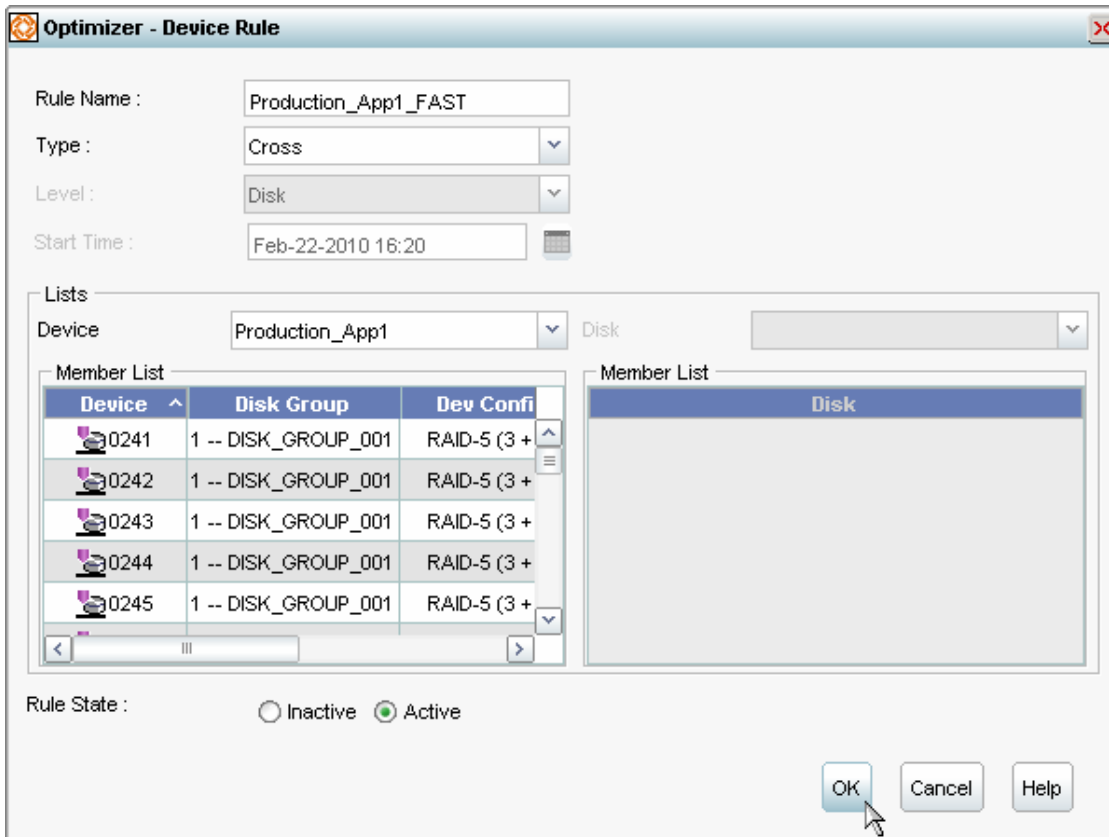
Once all the required devices have been added to the list the OK button should be clicked. Clicking OK returns the user to the main settings dialog.

Creating an Optimizer device rule

To create an Optimizer device rule, click on the New button located below the Device Rules list. When clicked the following dialog box will be displayed.

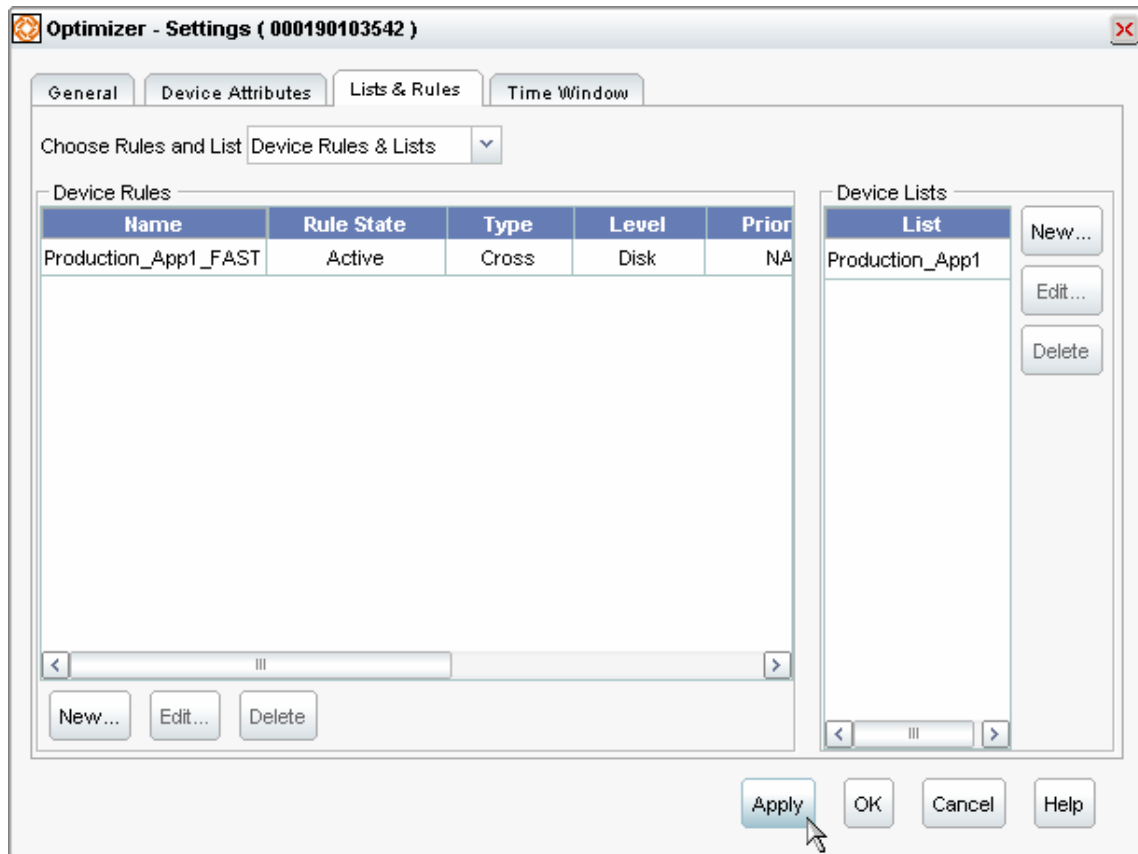
After providing a rule name, the rule type needs to be chosen—Avoid, Migration, or Cross.

Finally, the list of devices the rule will be applied to will be chosen. All available device lists will be displayed in a drop down list.



Clicking the OK button returns the user to the main Settings dialog.

To execute the creation of new disk rules and lists, the Apply button should be clicked on the Settings dialog.



Managing Optimizer swap lists

When running Optimizer in user approved mode, Optimizer will attempt to generate a swap list approximately once an hour. Generated swap lists are available to be reviewed and approved or declined by the user. The approval of a swap list also allows for the swap list execution to be scheduled.

Reviewing Optimizer swap list

When running in user approved mode, in order for a swap to take place, it must be manually reviewed and then approved or declined.

The following information is available when reviewing a swap list:

- ◆ The devices to be swapped
- ◆ The current location of each device hyper being swapped

- ◆ The destination of each device hyper being swapped
- ◆ The swap status

To review a generated swap list, right-click on the appropriate Symmetrix serial number in the left-hand navigation pane, and select Optimizer, Swap/Move, Swap/Move List.

| Type | Group | Device | Mirror / Member | From | To | Status |
|---------------------|-------|--------|-----------------|---------------|---------------|-------------|
| Optimizer Generated | 1 | 0289 | 3 | DF-11D:C B H2 | DF-1D:C 5 H8 | Not Started |
| Optimizer Generated | 1 | 0288 | 1 | DF-6D:C A H1 | DF-15A:D 7 H8 | Not Started |
| Optimizer Generated | 1 | 0263 | 4 | DF-16A:C 5 H8 | DF-6B:D 9 H2 | Not Started |
| Optimizer Generated | 1 | 0263 | 2 | DF-2A:C 5 H8 | DF-11A:C C H2 | Not Started |
| Optimizer Generated | 1 | 0288 | 2 | DF-11D:D E H1 | DF-2D:C 4 H8 | Not Started |
| Optimizer Generated | 1 | 0263 | 3 | DF-1D:C 5 H8 | DF-11D:C B H2 | Not Started |
| Optimizer Generated | 1 | 0263 | 1 | DF-15D:C 5 H8 | DF-6C:D C H2 | Not Started |
| Optimizer Generated | 1 | 0288 | 4 | DF-11B:D A H1 | DF-16D:C 4 H8 | Not Started |
| Optimizer Generated | 1 | 0289 | 1 | DF-6C:D C H2 | DF-15D:C 5 H8 | Not Started |
| Optimizer Generated | 1 | 0289 | 2 | DF-6B:D 9 H2 | DF-16A:C 5 H8 | Not Started |
| Optimizer Generated | 1 | 0288 | 3 | DF-11A:D D H1 | DF-1A:D 7 H8 | Not Started |
| Optimizer Generated | 1 | 0284 | 4 | DF-1A:D 7 H8 | DF-11A:D D H1 | Not Started |

Approving an Optimizer swap list

When a swap list has been reviewed, and it is going to be approved, the execution of the list needs to be scheduled. The swaps can be scheduled to occur during the next available swap time window, or at a user-specified time.

To approve a plan, right-click on the appropriate Symmetrix serial number in the left-hand navigation pane, and select Optimizer, Swap/Move, Approve Swap.


To schedule the swap to occur during the next swap time window, click on the radio button to the left of According to Optimizer policy, and then click OK.

To schedule the swap to occur at a specific time, click on the radio button to the left of On and specify a date and time, and then click OK.

Swaps - Approval Swap (000190103542)

| Group | Device | Mirror / Member | From | To |
|-------|--------|-----------------|---------------|---------------|
| 1 | 0289 | 3 | DF-11D:C B H2 | DF-1D:C 5 H8 |
| 1 | 0284 | 2 | DF-2D:C 4 H8 | DF-11D:D E H1 |
| 1 | 0284 | 1 | DF-15A:D 7 H8 | DF-6D:C A H1 |
| 1 | 0288 | 1 | DF-6D:C A H1 | DF-15A:D 7 H8 |
| 1 | 0263 | 4 | DF-16A:C 5 H8 | DF-6B:D 9 H2 |
| 1 | 0263 | 1 | DF-15D:C 5 H8 | DF-6C:D C H2 |
| 1 | 0289 | 1 | DF-6C:D C H2 | DF-15D:C 5 H8 |
| 1 | 0288 | 3 | DF-11A:D D H1 | DF-1A:D 7 H8 |
| 1 | 0284 | 4 | DF-1A:D 7 H8 | DF-11A:D D H1 |
| 1 | 0289 | 4 | DF-11A:C C H2 | DF-2A:C 5 H8 |

Schedule for execution

According to Optimizer policy
 On Feb-23-2010 06:18 PM 
 Decline Plan

OK Cancel Help

Once approved and scheduled, the execution time can be viewed by reviewing the swap list again.

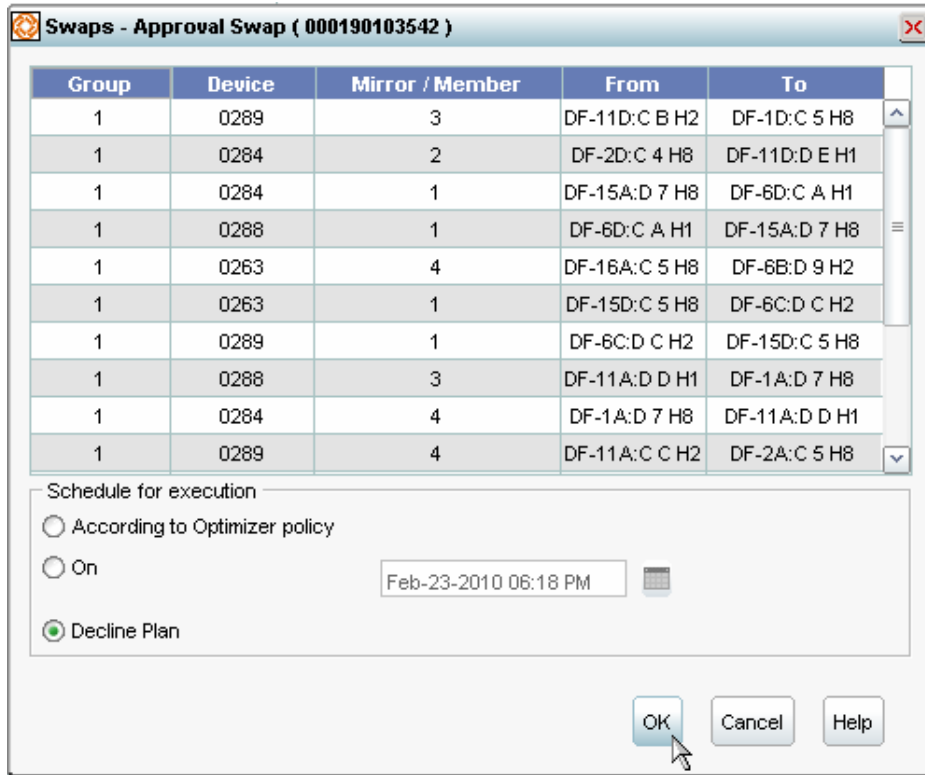
The swap can also be rescheduled anytime prior to the previously scheduled start time.

Declining an Optimizer swap list

When a swap list has been reviewed, and it is going to be declined, the decline operation will remove the plan from Optimizer, and cannot be subsequently approved.

To decline a swap list, right-click on the appropriate Symmetrix serial number in the left-hand navigation pane, and select Optimizer, Swap/Move, Approve Swap.

In the resulting dialog box, click on the radio button to the left of Decline Plan, and click OK.



When the plan has been declined, Optimizer will not display a swap list until a new one has been generated.

Monitoring an Optimizer swap

During the execution of the configuration change plan, the progress of the swap can be monitored by viewing the Optimizer Swap/Move list, and looking at the status column.

To monitor the swap list, right-click on the appropriate Symmetrix serial number in the left-hand navigation pane, and select Optimizer, Swap/Move, Swap/Move List.

Optimizer - Swap/Move List (000190103542)

| Type | Group | Device | Mirror / Member | From | To | Status |
|---------------------|-------|--------|-----------------|---------------|---------------|---------------|
| Optimizer Generated | 1 | 0289 | 2 | DF-6B:D 9 H2 | DF-16A:C 5 H8 | Step 41 of 82 |
| Optimizer Generated | 1 | 0289 | 4 | DF-11A:C C H2 | DF-2A:C 5 H8 | Step 41 of 82 |
| Optimizer Generated | 1 | 0289 | 1 | DF-6C:D C H2 | DF-15D:C 5 H8 | Step 41 of 82 |
| Optimizer Generated | 1 | 0288 | 1 | DF-6D:C A H1 | DF-15A:D 7 H8 | Step 41 of 82 |
| Optimizer Generated | 1 | 0263 | 3 | DF-1D:C 5 H8 | DF-11D:C B H2 | Step 41 of 82 |
| Optimizer Generated | 1 | 0288 | 4 | DF-11B:D A H1 | DF-16D:C 4 H8 | Step 41 of 82 |
| Optimizer Generated | 1 | 0263 | 1 | DF-15D:C 5 H8 | DF-6C:D C H2 | Step 41 of 82 |
| Optimizer Generated | 1 | 0284 | 3 | DF-16D:C 4 H8 | DF-11B:D A H1 | Step 41 of 82 |
| Optimizer Generated | 1 | 0288 | 2 | DF-11D:D E H1 | DF-2D:C 4 H8 | Step 41 of 82 |
| Optimizer Generated | 1 | 0263 | 2 | DF-2A:C 5 H8 | DF-11A:C C H2 | Step 41 of 82 |
| Optimizer Generated | 1 | 0284 | 1 | DF-15A:D 7 H8 | DF-6D:C A H1 | Step 41 of 82 |
| Optimizer Generated | 1 | 0289 | 3 | DF-11D:C B H2 | DF-1D:C 5 H8 | Step 41 of 82 |
| Optimizer Generated | 1 | 0263 | 4 | DF-16A:C 5 H8 | DF-6B:D 9 H2 | Step 41 of 82 |
| Optimizer Generated | 1 | 0288 | 3 | DF-11A:D D H1 | DF-1A:D 7 H8 | Step 41 of 82 |
| Optimizer Generated | 1 | 0284 | 4 | DF-1A:D 7 H8 | DF-11A:D D H1 | Step 41 of 82 |
| Optimizer Generated | 1 | 0284 | 2 | DF-2D:C 4 H8 | DF-11D:D E H1 | Step 41 of 82 |

Refresh Close Help

The Status field shows how many steps are complete in the Symmetrix-based script that is being run, performing the device swap.

To update the display, click on the Refresh button.

Note: when multiple hypervisors of the same device are being swapped they are swapped one at a time. As such the step number will reset as each hypervisor is being swapped.

When the swap is completed, the swap list will be empty.

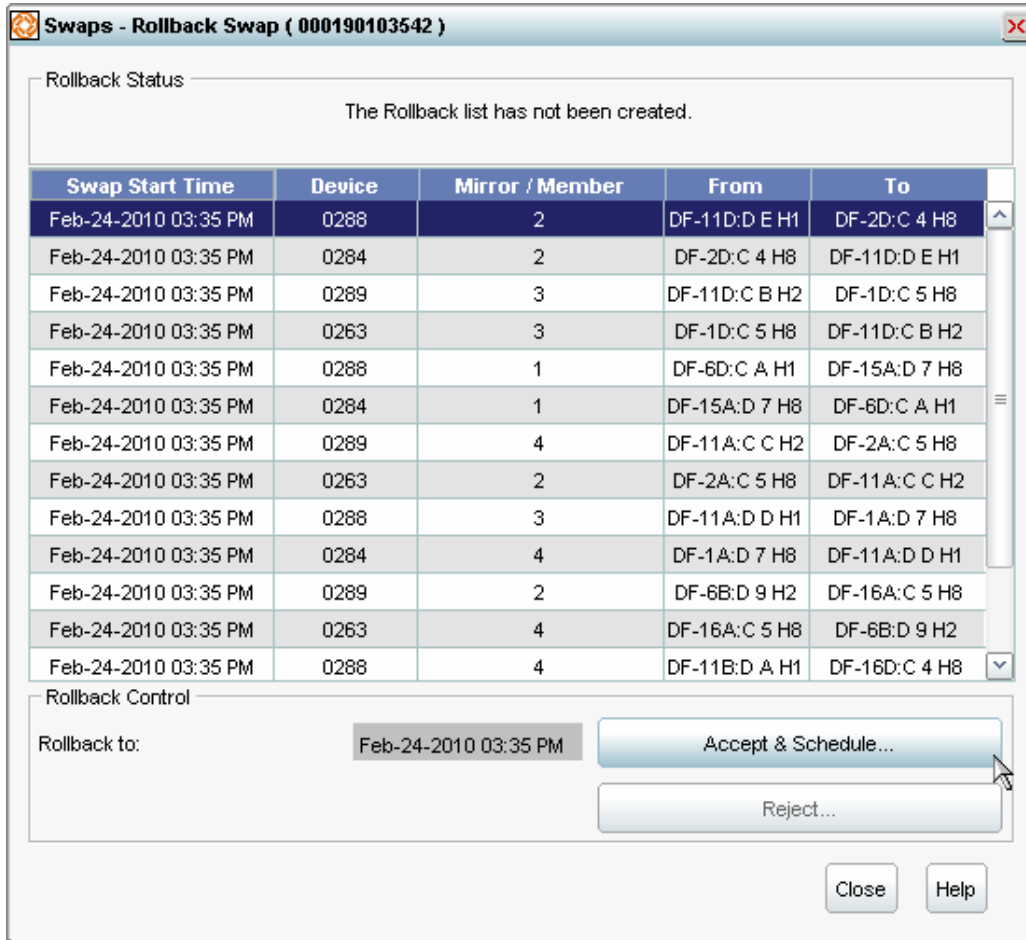
Performing an Optimizer swap rollback

Optimizer rollback allows previously performed swaps to be reversed. Swaps can be rolled back to a particular point-in-time.

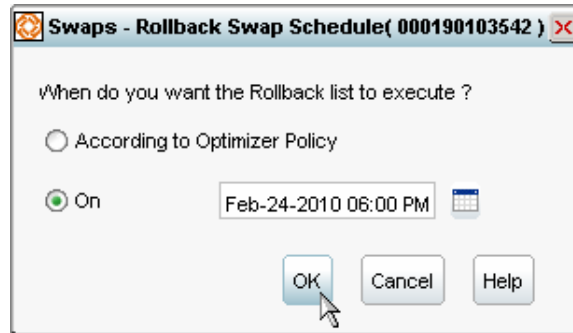
To perform the rollback, right-click on the appropriate Symmetrix serial number in the left-hand navigation pane, and select Optimizer, Swap/Move, Rollback.

Note: in order to run a rollback, Optimizer must be stopped/disabled.

In the resulting dialog box, select the point-in-time to rollback to—this is done by selecting any swap entry with the desired time to be rolled back to. Once selected, click on Accept & Schedule to specify when the rollback should be performed.



The rollback can be scheduled to occur at a specific time, or during the next available device movement window.



Performing a manual Optimizer swap

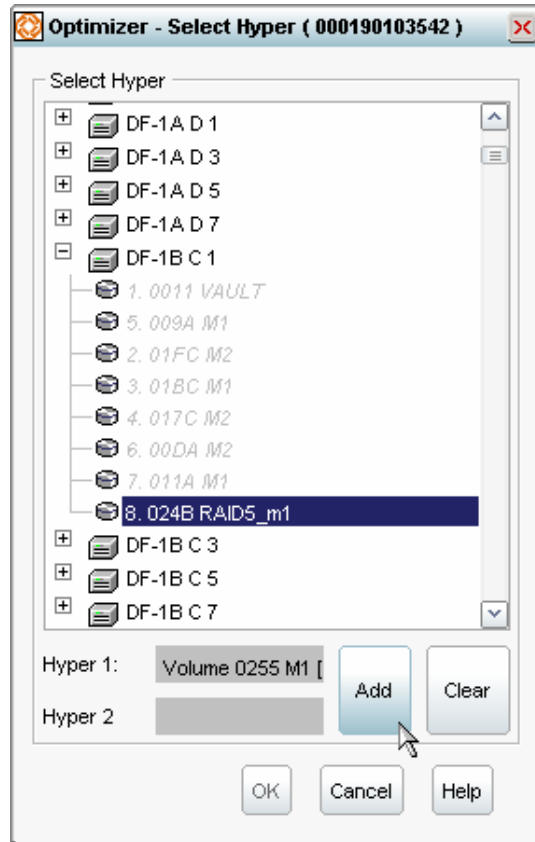
A manual swap can be performed where the user selects the individual physical hypervisors of two devices to be swapped.

To perform a manual swap, right-click on the Symmetrix serial number, then choose Optimizer, Swap/Move, Manual Swap. In the resulting dialog box, click the Add button to select the hypervisors to be swapped.

Note: prior to submitting a manual swap, Optimizer must first be stopped/disabled.

Expanding out the disks displays the hypervisors configured on each disk. Hypervisors belonging to swappable devices will be shown in black text (non-swappable hypervisors are shown in grey and italics).

Selecting a hypervisor and then clicking Add will place that hypervisor's configuration information in the Hyper 1 field. Click Add again after select the hypervisor to swap with this first hypervisor and it will be added to the Hyper 2 field.

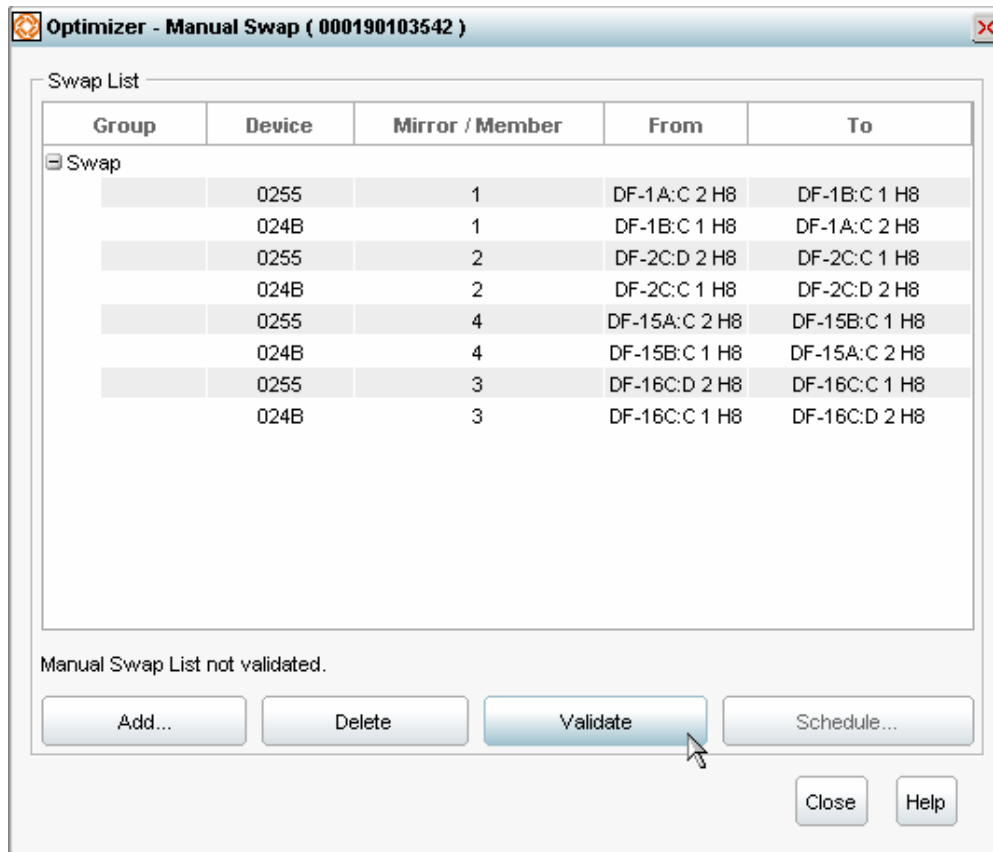


To complete the swap request click OK.

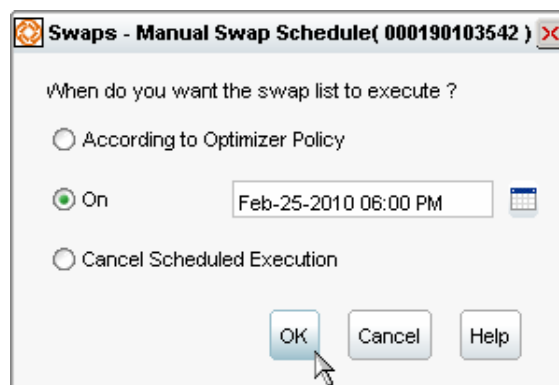
Details of the swap request will be displayed in the main Manual Swap dialog box. If the hypers being swapped belong to separate affinity groups, SMC will automatically add the remaining hypers of each device to the swap request.

Additional hypers/devices can be added to the manual swap request by again clicking on the Add button and following the previous process.

When the swap request is complete, the next step is to validate the request. This will ensure that the swap will not violate any configuration rules of the Symmetrix. To do this, click on the Validate button.



When successfully validated, the Schedule button will become available. Click this to schedule when the manual swap should be executed.



Monitoring of a manual swap is performed in a similar fashion as

discussed earlier in this section.

Performing a Virtual LUN migration (manual move)

Symmetrix devices can be manually moved between physical disk groups, of different disk technologies, using the Virtual LUN migration feature of Optimizer.

The source devices for a migration may be specified individually, as a device list, or collectively by device group. The target destination for the device migration can be specified as configured space, or unconfigured space.

The following sections detail examples of performing migrations to both configured and unconfigured space.

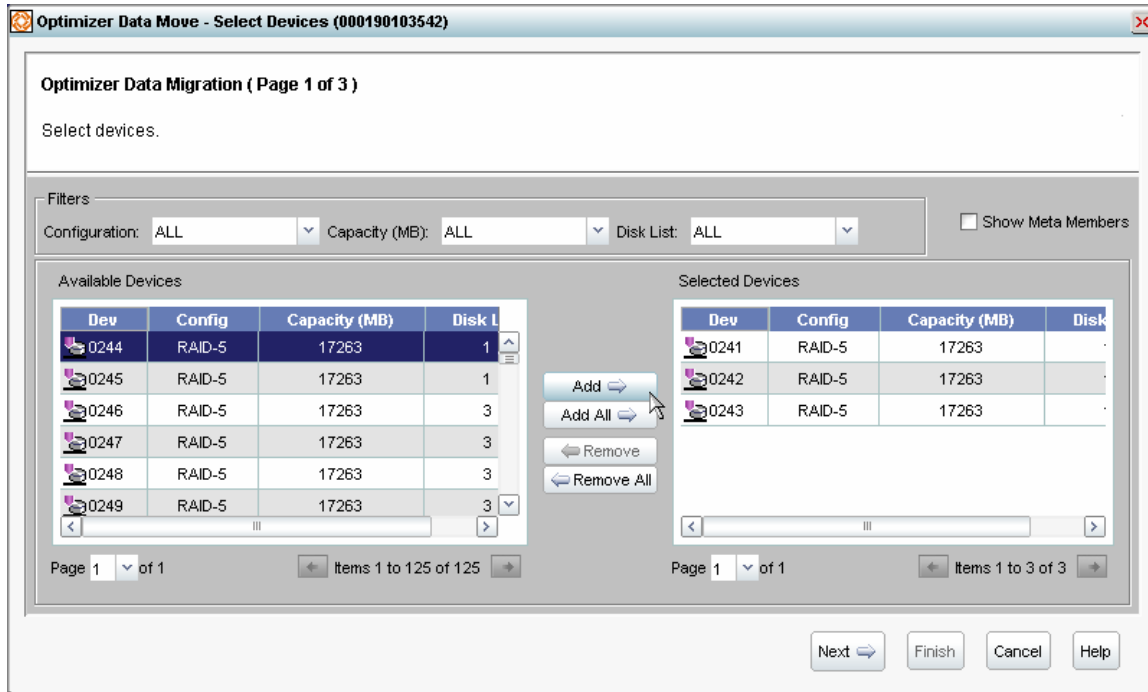
Note: prior to performing a manual migration Optimizer must be stopped/disabled.

Performing a migration to unconfigured space

To perform a migration to unconfigured space, and to specify individual devices to be migrated, right-click on the Symmetrix serial number, then choose Optimizer, Swap/Move, Manual Move.

The resulting dialog box will display all devices in the array that it is possible to migrate. Selecting a device in the left-hand panel, under Available Devices, and clicking Add will move the device to the Selected Devices panel in preparation to be migrated.

When all the devices to be migrated have been selected the Next button is clicked.

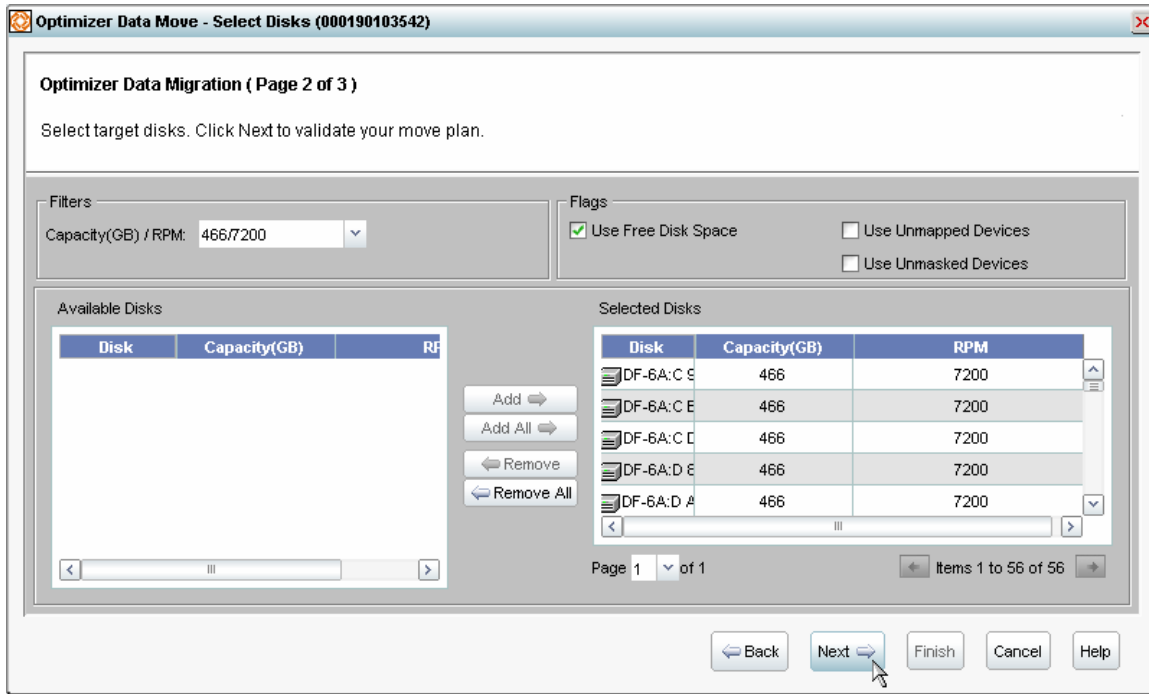


After selecting the devices to be migrated, the next step is to specify the target storage to which the devices will be migrated. The target storage is specified by selecting multiple physical disks. Available disks in the array can be filtered by choosing a capacity and speed in the Filters panel.

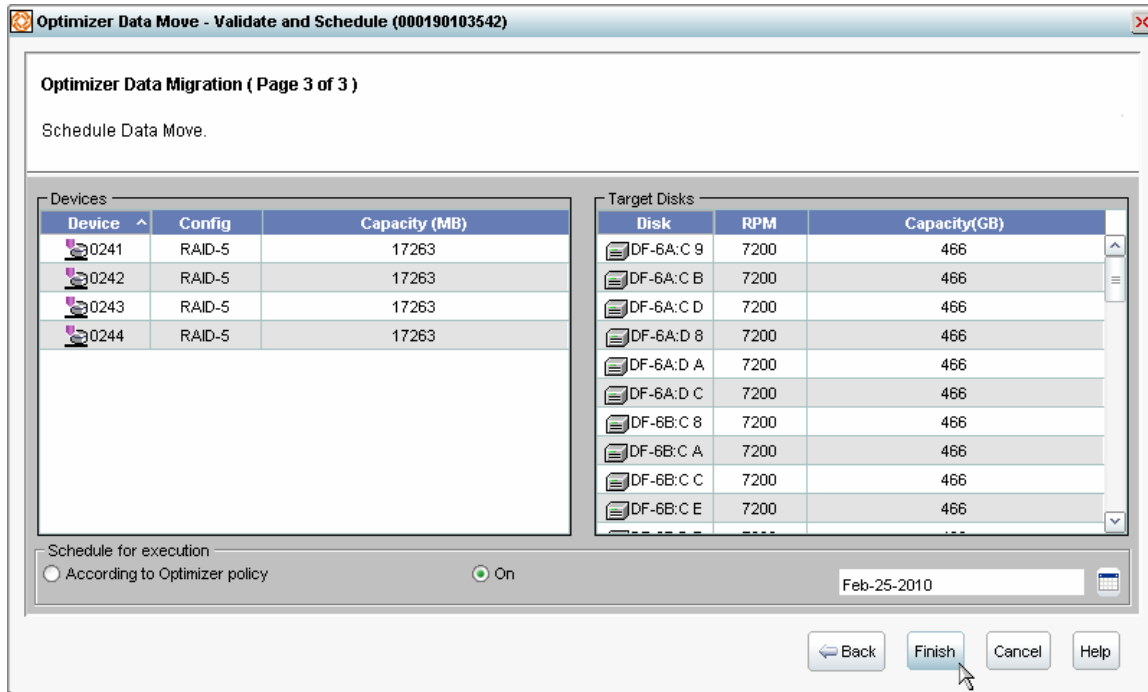
Once all disks of the same type have been displayed they can be added to the Selected Disks display by clicking on Add All.

By clicking on the checkbox to the left of Freespace in the Flags panel, specifies a migration to unconfigured space.

Note: while subsets of disks of the same type can be selected, it is recommended that all disks of the same type be specified as the target storage. In this way Symmetrix Engineuity can determine the most appropriate layout for the device once migrated and ensures that configuration rules will not be violated.



When the target storage has been specified, the Next button is clicked. The resulting dialog box shows a list of the source devices and target storage. From this dialog the execution of the migration can also be scheduled. Options for scheduling include allowing the migration to take place according to the Optimizer swap time window, or at a specific time of the users choosing.



Clicking on Finish will submit the migration to Optimizer for execution.

Virtual LUN migrations can be viewed and monitored in the same way as normal Optimizer swaps, as discussed earlier in this section.

Performing a migration to configured space

To perform a migration to configured space, and to specify a device group to be migrated, right-click on the device group containing the devices to be migrated, then choose Optimizer, Swap/Move, Manual Move.

Management Interface: SMC

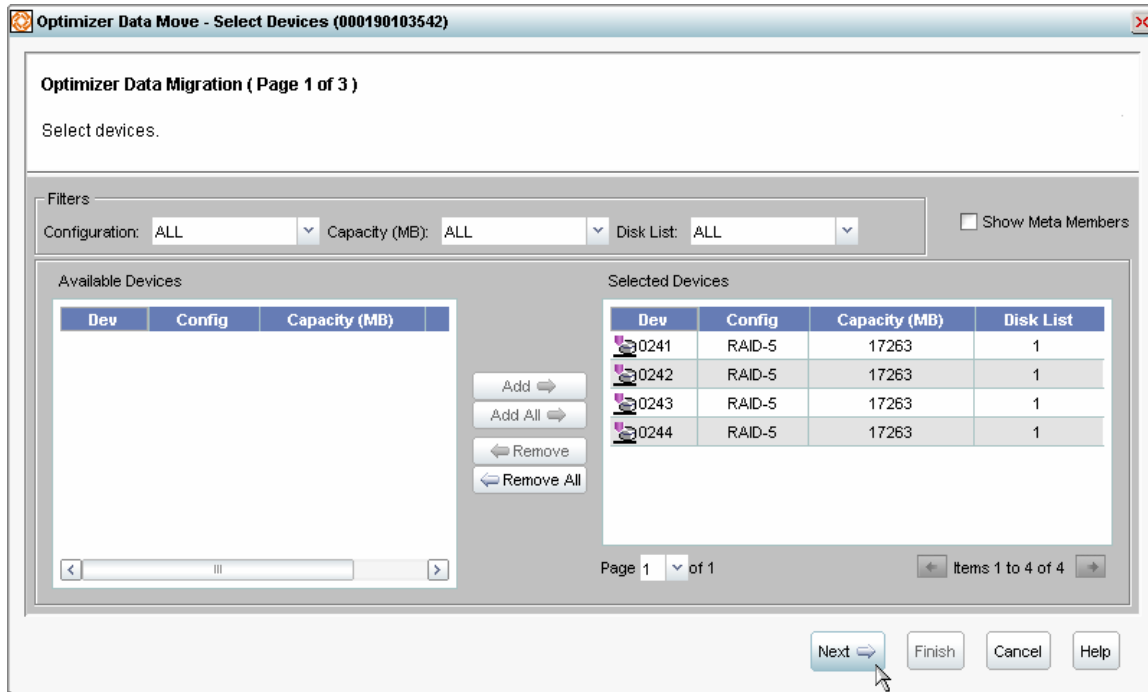
The screenshot displays the Symmetrix Management Console (SMC) interface. On the left, a tree view shows the hierarchy: Replication Views, Dynamic Cache Partitions, Device Groups (1), Production_App1 [103542], Standard Devices (4), and Composite Groups (2). The 'Production_App1 [103542]' group is selected, and a context menu is open over it. The menu items are: Symmetrix Admin, Device Masking and Mapping, Port and Director Configuration, Device Configuration, FBA Meta Device Configuration, z/OS Configuration, DG/CG Management, Device Pool Management, Replication, Migration, Tier Management, FAST, Optimizer, Priority Control and Cache Partition, Performance Analyzer, SRDF Migration Management, Virtualization Domain Management, and Alerts. The 'Migration' item is expanded, showing a sub-menu with: Start/Stop/Lock/Unlock ..., Enable/Disable ..., Settings ..., Swap/Move, and Log ... The 'Swap/Move' item is further expanded, showing a sub-menu with: Approve Swap ..., Manual Swap ..., Manual Move ... (highlighted), Rollback ..., and Swap/Move List ...

| Device Group Name | Creation Time | Modify Time | Symmetrix ID |
|--------------------------|---------------------|---------------------|--------------|
| Production_App1 [103542] | 2010-02-25 17:12:12 | 2010-02-25 17:12:12 | 000190103542 |

| Device Type | Count |
|---------------------|-------|
| peers | 0 |
| up | 4 |
| BCVs | 0 |
| VDEVs | 0 |
| TGTs | 0 |
| ed BCVs (STD RDF) | 0 |
| ed BCVs (BCV RDF) | 0 |
| ed RBCVs (RBCV RDF) | 0 |
| ed VDEVs | 0 |
| ed TGTs | 0 |

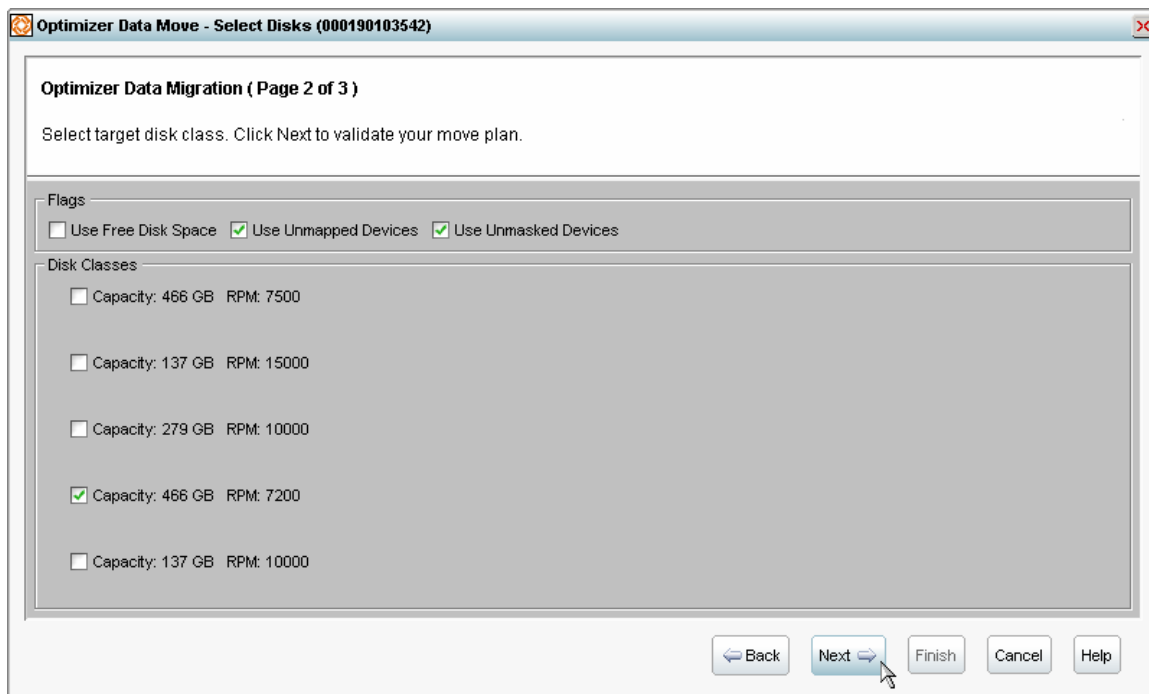
The resulting dialog box will display all the devices in the device group that it is possible to migrate. Selecting a device in the left-hand panel, under Available Devices, and clicking Add will move the device to the Selected Devices panel in preparation to be migrated.

When all the devices to be migrated have been selected the Next button is clicked.

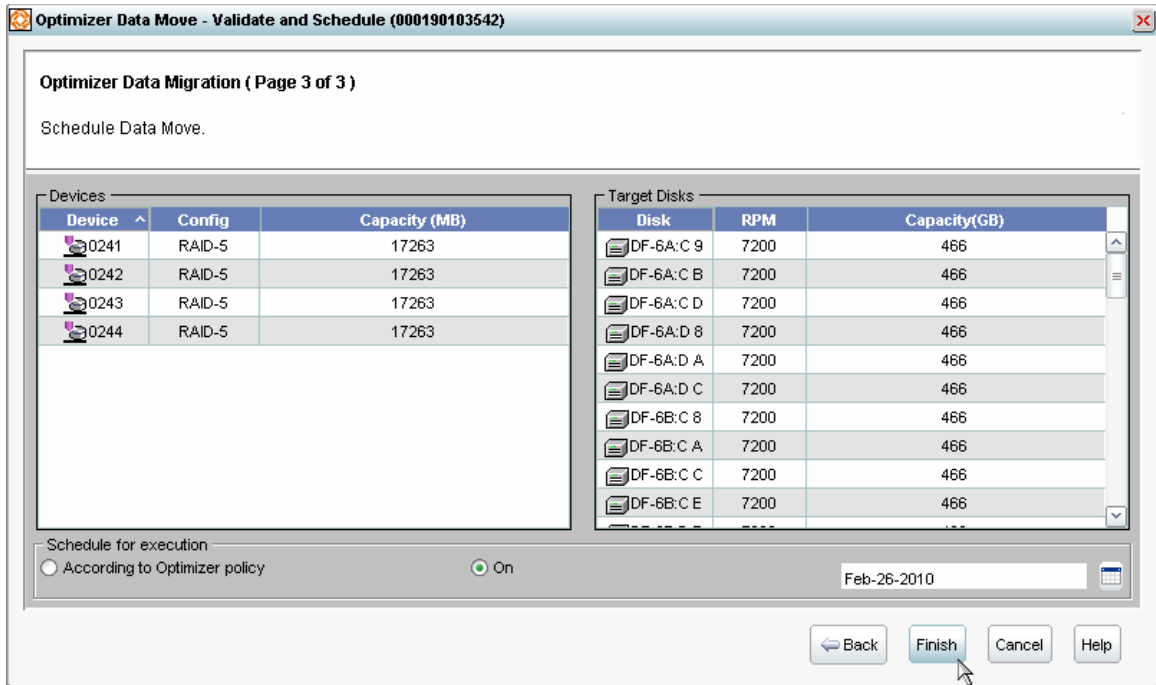


After selecting the devices to be migrated, the next step is to specify the target storage to which the devices will be migrated. The target storage is specified by selecting a category of disk, based on capacity and speed. This is done by clicking on the checkbox to the left of the appropriate disk category.

When migrating to configured space, the target storage must contain devices of the same size and RAID protection as the devices being migrated. These devices must also not be in use on a host. As such, a migration to configured space is specified by clicking on the checkboxes to the left of Use Unmapped Devices or Use Unmasked Devices, or both.



When the target storage has been specified, the Next button is clicked. The resulting dialog box shows a list of the source devices and target storage. From this dialog the execution of the migration can also be scheduled. Options for scheduling include allowing the migration to take place according to the Optimizer swap time window, or at a specific time of the users choosing.



Clicking on Finish will submit the migration to Optimizer for execution. Virtual LUN migrations can be viewed and monitored in the same way as normal Optimizer swaps, as discussed earlier in this section.

Management Interface: SYMCLI

Management of Optimizer from Solutions Enabler is performed using the SYMCLI `symoptmz` command

The following sections detail the use of the `symoptmz` command in configuring and managing Symmetrix Optimizer using the Solutions Enabler SYMCLI.

Stopping/Starting Symmetrix Optimizer

In order for Optimizer to perform data analysis and generate swap lists, the Optimizer controller must be enabled on the Symmetrix.

To check the current state of Optimizer, the `symoptmz` command can be used as follows:

```
symoptmz -sid 3542 query
```

This operation may take up to a few minutes. Please be patient...

```
Optimizer State   Current Activities
-----
Enabled           Idle
```

The state of the Optimizer can be changed by running the following commands as appropriate:

To stop, or disable, Optimizer run:

```
symoptmz -sid 3542 disable
```

This operation may take up to a few minutes. Please be patient...

```
Optimizer is Disabled for Symmetrix : 000190103542
```

To start, or enable, Optimizer run:

```
symoptmz -sid 3542 enable
```

This operation may take up to a few minutes. Please be patient...

```
Optimizer is Enabled for Symmetrix : 000190103542
```

Configuring the Symmetrix Optimizer controller

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

- ◆ Operating Mode
- ◆ Maximum Moves Per Day
- ◆ Maximum Simultaneous Moves
- ◆ Workload Analysis Period
- ◆ Initial Period
- ◆ Start Up Mode
- ◆ Hot Spot Analysis Mode
- ◆ Maximum Days for Rollback
- ◆ Performance Time Window
- ◆ Device Movement Time Window

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

Note: prior to changing any Symmetrix Optimizer controller setting, Optimizer must first be stopped/disabled.

Optimizer controller settings list information

To view the existing Optimizer controller settings and time windows, the `symoptmz list` command can be used with the `-parms` option:

```
symoptmz -sid 3542 show -parms
```

This operation may take up to a few minutes. Please be patient...

Optimizer Control Parameters

```
Start Mode           : Auto
Swap Mode            : Auto
Min Perf Period      : 168
Workload Period      : 168
Max Simult Swaps     : 8
Max Swaps Per Day    : 200
```

Advanced Parameter Settings

```
Max Days to keep Rollbacks : 30
Hot Spot Analysis           : False
Maintain Physical Mirrors   : True
Adv Parm Flag               : 1

Number of Swap Time Windows : 1

Time Window ID              : The Default Time Window
Type                        : Swap
```

```
Flags : Exclusive
Periodicity : Once
Start Date : Fri Dec 31 19:00:00 1999
Stop Date : Mon Dec 30 19:00:00 2030

Performance Time Windows : 1

Time Window ID : The Default Time Window
Type : Performance
Flags : Inclusive
Periodicity : Once
Start Date : Fri Dec 31 19:00:00 1999
Stop Date : Mon Dec 30 19:00:00 2030
```

Setting the operating mode

There are two possible values for the data movement mode: Automatic (AUTO) and User Approved (USER_OK). The default value is AUTO.

To set the mode to User Approved, first create a text file containing:

```
set control_parms swap_mode=USER_OK;
```

Then run:

```
symoptmz -sid 3542 -f optmz_parms.txt commit -v
```

```
Processing Command file : op_mode.txt
PREVIEW .....Started.
Set Control Params Command
{
  set control_parms swap_mode = USER_OK;
}

PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Setting the Maximum Moves Per Day

Maximum Device Moves Per Day (swap_rate) can be set between 2 and 200. The default value is 200.

To set the Maximum Device Moves Per Day to 64, first create a text file containing:

```
set control_parms swap_rate=64;
```

Then run:

```
symoptmz -sid 3542 -f optmz_parms.txt commit -v
```

```
Processing Command file : optmz_parms.txt
PREVIEW .....Started.
Set Control Params Command
{
  set control_parms swap_rate = 64;
}

PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Setting the Maximum Simultaneous Moves

Maximum Simultaneous Device Moves (`max_simult_swaps`) can be set between 2 and 8. The default value is 8.

To set the Maximum Simultaneous Device Moves to 4, create a file containing:

```
set control_parms max_simult_swaps=4;
```

Then run:

```
symoptmz -sid 3542 -f optmz_parms.txt commit -v
```

```
Processing Command file : optmz_parms.txt
PREVIEW .....Started.
Set Control Params Command
{
  set control_parms max_simult_swaps = 4;
}

PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Setting the Workload Analysis Period

The Workload Analysis Period (`workload_period`) can be set between 1 hour and 4 weeks, but its value is specified in hours. The default value is 1 week (168 hours).

To set the Workload Analysis Period to three days (72 hours), create a text file containing:

```
set control_parms workload_period=72;
```

Then run:

```
symoptmz -sid 3542 -f optmz_parms.txt commit -v
```

```
Processing Command file : optmz_parms.txt
PREVIEW .....Started.
Set Control Params Command
{
  set control_parms workload_period = 72;
}

PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Setting the Initial Period

The Initial Period (min_perf_period) can be set between 1 hour and 4 weeks, but its value is specified in hours. The default value is 1 week (168 hours).

To set the Initial Period to three days (72 hours), create a text file containing:

```
set control_parms min_perf_period=72;
```

Then run:

```
symoptmz -sid 3542 -f optmz_parms.txt commit -v
```

```
Processing Command file : optmz_parms.txt
PREVIEW .....Started.
Set Control Params Command
{
  set control_parms min_perf_period = 72;
}

PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Setting the Start Up Mode

There are two possible values for the start up mode: Automatic (AUTO) and Manual (MANUAL). The default value is AUTO.

To set the start up mode to manual, create a text file containing:

```
set control_parms start_mode=MANUAL;
```

Then run:

```
symoptmz -sid 3542 -f optmz_parms.txt commit -v
```

```
Processing Command file : optmz_parms.txt
PREVIEW .....Started.
Set Control Params Command
{
  set control_parms start_mode = MANUAL;
}

PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Setting the Hot Spot Analysis Mode

There are two possible values for the hot spot analysis mode: True or False. The default value is False.

To enable hot spot analysis mode, create a text file containing:

```
set advanced_parms hot_spot=TRUE;
```

Then run:

```
symoptmz -sid 3542 -f optmz_parms.txt commit -v
```

```
Processing Command file : optmz_parms.txt
PREVIEW .....Started.
Set Control Params Command
{
  set advanced_parms hot_spot = TRUE;
  adv_parms_flag = 0;
}

PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Setting the Maximum Days for Rollback

The Maximum Days for Rollback (max_rollback) can be set between 1 and 180 days. The default value is 30 days.

To set the Maximum Days for Rollback to 15 days, create a text file containing:

```
set advanced_parms max_rollback=15;
```

Then run:

```
symoptmz -sid 3542 -f optmz_parms.txt commit -v
```

```
Processing Command file : optmz_parms.txt
PREVIEW .....Started.
Set Control Params Command
{
  set advanced_parms max_rollback = 15;
                        adv_parms_flag = 0;
}
PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Creating Optimizer time windows

To create time windows to be used by Optimizer a text file needs to be created containing syntax defining the desired windows. Prior to creating the text file, the following information will need to be decided:

- ◆ Time window name
- ◆ Time window type—performance of device movement
- ◆ Whether the time window is inclusive or exclusive
- ◆ The frequency of occurrence
- ◆ Days of the week the window will be used
- ◆ Times of the day the window will be used
- ◆ Time period for which the time window is valid

Note: Using Solutions Enabler, it is necessary to create all desired performance and device movement windows at the same time.

As an example, to create a performance time window that will collect data samples between 7 a.m. and 7 p.m. on each day, Monday through Friday, the following syntax would need to be placed in a text file:

```
set time_window id=Performance_Window,
type=PERF,
flag=INCLUDE,
period=WEEKLY_BY_DAY,
starting=09012009:080000,
ending=09012019:080000,
days=MON,TUE,WED,THU,FRI,
start_time=07:00,
end_time=19:00;
```

To create a swap time window that will perform Optimizer swaps between 1 a.m. and 5 a.m. on Saturdays and Sundays, the following syntax would need to be added to the previously created text file:

```
set time_window id=Swap_Window,
type=SWAP,
flag=INCLUDE,
period=WEEKLY_BY_DAY,
starting=08012009:080000,
ending=08012019:080000,
days=SAT,SUN,
start_time=01:00,
end_time=05:00;
```

When the text file has been created, to commit the changes, run:

```
symoptmz -sid 3542 -f time_windows.txt commit
```

```
Processing Command file : time_windows.txt
PREVIEW .....Started.
PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
```

Verifying Optimizer controller settings

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

```
symoptmz -sid 3542 show -parms
```

This operation may take up to a few minutes. Please be patient...

Optimizer Control Parameters

Start Mode : Manual
Swap Mode : UserOK
Min Perf Period : 72
Workload Period : 72
Max Simult Swaps : 4
Max Swaps Per Day : 64

Advanced Parameter Settings

Max Days to keep Rollbacks : 15
Hot Spot Analysis : True
Maintain Physical Mirrors : True
Adv Parm Flag : 1

Number of Swap Time Windows : 2

Time Window ID : The Default Time Window
Type : Swap
Flags : Exclusive
Periodicity : Once
Start Date : Fri Dec 31 19:00:00 1999
Stop Date : Mon Dec 30 19:00:00 2030

Time Window ID : Swap_Window
Type : Swap
Flags : Inclusive
Periodicity : WeeklyByDay
Start Date : Thu Dec 31 19:00:00 2009
Stop Date : Wed Dec 30 19:00:00 2020
Days of the week : Sunday Saturday
Start Time (hh:mm) : 01:00
Stop Time (hh:mm) : 05:00

Performance Time Windows : 2

Time Window ID : The Default Time Window
Type : Performance
Flags : Inclusive
Periodicity : Once
Start Date : Fri Dec 31 19:00:00 1999
Stop Date : Mon Dec 30 19:00:00 2030

Time Window ID : Performance_Window
Type : Performance
Flags : Inclusive
Periodicity : WeeklyByDay
Start Date : Thu Dec 31 19:00:00 2009
Stop Date : Wed Dec 30 19:00:00 2020

```

Days of the week      : Monday Tuesday Wednesday Thursday Friday
Start Time (hh:mm)   : 07:00
Stop Time (hh:mm)    : 19:00

```

Setting device attributes

By setting device attributes, a swap priority can be assigned to a single device or a range of devices. Valid values for the swap priority are

- ◆ High
- ◆ Normal (default)
- ◆ No_Swap

To assign a swap priority of high to a single device, a text file needs to be created as follows:

```
set swap_priority to HIGH for dev 200;
```

When the text file has been created, to commit the changes, run:

```
symoptmz -sid 3542 -f device_attribs.txt commit
```

```

Processing Command file : device_attribs.txt
PREVIEW .....Started.
Set Swap Priority Command
{
  set swap_priority to HIGH for
    dev 0200;
}
PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.

```

To assign a swap priority of no_swap to a range of devices, a text file needs to be created as follows:

```
set swap_priority to NO_SWAP for dev 200:203;
```

When the text file has been created, to commit the changes, run:

```
symoptmz -sid 3542 -f device_attribs.txt commit
```

```
Processing Command file : device_attribs.txt
```

```
PREVIEW .....Started.
Set Swap Priority Command
{
  set swap_priority to NO_SWAP for
                        dev 0200:0203;
}
PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

Managing Optimizer swap lists

When running Optimizer in user approved mode, Optimizer will attempt to generate a swap list approximately once an hour. Generated swap lists are available to be reviewed and approved or declined by the user. The approval of a swap list also allows for the swap list execution to be scheduled.

Reviewing Optimizer swap list

When running in user approved mode, in order for a swap to take place, it must be manually reviewed and then approved or declined.

The following information is available when reviewing a swap list:

- ◆ The devices to be swapped
- ◆ The current location of each device hyper being swapped
- ◆ The destination of each device hyper being swapped
- ◆ The swap status

To review a generated swap list, run:

```
symoptmz -sid 3542 show -swap_list
```

This operation may take up to a few minutes. Please be patient...

```
Number of Optimizer Swap Lists : 1
Swap list state                  : NotApproved
```

```
Swap #1:
{
  Time Defined                   : 02232010:174710
  Delayed Start                  : None
  Number of Swap groups         : 1
```

```
Swap Group #1:
```

```

{
Number of Moves/Swaps : 16
Swap Attribs          : SWAP OPTIMIZER-GENERATED
State                 : SwapNotStarted
Percent Complete      : Unknown

Move Number          #1:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
Hyper1    11D      D          E          1          0288    2
Hyper2    02D      C          4          8          0288    2
}

Move Number          #2:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
Hyper1    02D      C          4          8          0284    2
Hyper2    11D      D          E          1          0284    2
}

Move Number          #3:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
Hyper1    06D      C          A          1          0288    1
Hyper2    15A      D          7          8          0288    1
}

Move Number          #4:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
Hyper1    15A      D          7          8          0284    1
Hyper2    06D      C          A          1          0284    1
}

Move Number          #5:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
Hyper1    11A      D          D          1          0288    3
Hyper2    01A      D          7          8          0288    3
}

Move Number          #6:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
Hyper1    01A      D          7          8          0284    4
Hyper2    11A      D          D          1          0284    4
}

```

```

Move Number      #7:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
  Hyper1   11B      D          A          1         0288    4
  Hyper2   16D      C          4          8         0288    4
}

```

```

Move Number      #8:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
  Hyper1   16D      C          4          8         0284    3
  Hyper2   11B      D          A          1         0284    3
}

```

```

Move Number      #9:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
  Hyper1   11D      C          B          2         0289    3
  Hyper2   01D      C          5          8         0289    3
}

```

```

Move Number      #10:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
  Hyper1   01D      C          5          8         0263    3
  Hyper2   11D      C          B          2         0263    3
}

```

```

Move Number      #11:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
  Hyper1   11A      C          C          2         0289    4
  Hyper2   02A      C          5          8         0289    4
}

```

```

Move Number      #12:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----
  Hyper1   02A      C          5          8         0263    2
  Hyper2   11A      C          C          2         0263    2
}

```

```

Move Number      #13:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -----  -----  -----  -----  -----  -----

```

```

Hyper1      06B      D          9          2      0289      2
Hyper2      16A      C          5          8      0289      2
}

Move Number      #14:
{
Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
-----  -
Hyper1   16A     C          5          8      0263     4
Hyper2   06B     D          9          2      0263     4
}

Move Number      #15:
{
Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
-----  -
Hyper1   06C     D          C          2      0289     1
Hyper2   15D     C          5          8      0289     1
}

Move Number      #16:
{
Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
-----  -
Hyper1   15D     C          5          8      0263     1
Hyper2   06C     D          C          2      0263     1
}
Number of migrations : 0
}
}

```

Approving an Optimizer swap list

When a swap list has been reviewed, and it is going to be approved, the execution of the list needs to be scheduled. The swaps can be scheduled to occur during the next available swap time window, or at a user-specified time.

To schedule the swap to occur during the next swap time window, create a text file containing:

```
set swap APPROVE, TIMESTAMP=02232010:174710;
```

Then run:

```
symoptmz -sid 3542 -f approve_swap.txt commit
```

```
Processing Command file : approve_swap.txt
PREVIEW .....Started.
Set Swap Command
{
  set swap APPROVE
}

```

```
                TIMESTAMP = Tue Feb 23 17:47:10 2010
                ;
    }
    PREVIEW .....Done.
    PREPARE .....Started.
    PREPARE .....Done.
    COMMIT .....Started.
    COMMIT .....Done.
    Closing Optimizer Session .....Done.
```

To schedule the swap to occur at a specific time, create a text file containing:

```
                set swap APPROVE, begin_at=02242010:180000,
                TIMESTAMP=02232010:174710;
```

Then run:

```
symoptmz -sid 3542 -f approve_swap.txt commit
```

```
Processing Command file : approve_swap.txt
    PREVIEW .....Started.
    Set Swap Command
    {
        set swap APPROVE
            begin_at = Wed Feb 24 18:00:00 2010
            TIMESTAMP = Tue Feb 23 17:47:10 2010
            ;
    }
    PREVIEW .....Done.
    PREPARE .....Started.
    PREPARE .....Done.
    COMMIT .....Started.
    COMMIT .....Done.
    Closing Optimizer Session .....Done.
```

Note: the TIMESTAMP value should be specified as the Time Defined value shown in the swap list. This value uniquely identifies the swap list being approved.

Once approved and scheduled, the execution time can be viewed by reviewing the swap list again.

The swap can also be rescheduled anytime prior to the previously scheduled start time.

Declining an Optimizer swap list

When a swap list has been reviewed, and it is going to be declined, the decline operation will remove the plan from Optimizer, and cannot be subsequently approved.

To decline a swap list, create a text file containing:

```
set swap DECLINE, TIMESTAMP=02232010:174710;
```

Then run:

```
symoptmz -sid 3542 -f decline_swap.txt commit
```

```
Processing Command file : decline_swap.txt
PREVIEW .....Started.
Set Swap Command
{
  set swap DECLINE
    TIMESTAMP = Tue Feb 23 17:47:10 2010
  ;
}
PREVIEW .....Done.
PREPARE .....Started.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

When the plan has been declined, Optimizer will not display a swap list until a new one has been generated.

Monitoring an Optimizer swap

During the execution of the configuration change plan, the progress of the swap can be monitored by viewing the Optimizer Swap/Move list, and looking at the status column.

To monitor the swap, run:

```
symoptmz -sid 3542 show -swap_list
```

This operation may take up to a few minutes. Please be patient...

```
Number of Optimizer Swap Lists : 1
Swap list state                 : ApprovedWithDelay

Swap #1:
{
  Time Defined                   : 02232010:174710
  Delayed Start                  : Time
  Start Time                     : Wed Feb 24 15:35:00 2010
  Number of Swap groups         : 1
```

```

Swap Group #1:
{
  Number of Moves/Swaps   : 16
  Swap Attribs            : SWAP OPTIMIZER-GENERATED
  State                   : SwapInProgress
  Percent Complete        : 54%

  Move Number             #1:
  {
    Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
    -----  -----  -----  -----  -----  -----  -----
    Hyper1   11D     D         E             1         0288    2
    Hyper2   02D     C         4             8         0288    2
  }

  Move Number             #2:
  {
    Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
    -----  -----  -----  -----  -----  -----  -----
    Hyper1   02D     C         4             8         0284    2
    Hyper2   11D     D         E             1         0284    2
  }

  ...

  Move Number             #15:
  {
    Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
    -----  -----  -----  -----  -----  -----  -----
    Hyper1   06C     D         C             2         0289    1
    Hyper2   15D     C         5             8         0289    1
  }

  Move Number             #16:
  {
    Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
    -----  -----  -----  -----  -----  -----  -----
    Hyper1   15D     C         5             8         0263    1
    Hyper2   06C     D         C             2         0263    1
  }
  Number of migrations    : 0
}
}

```

The State field shows the percent complete of the swap process, based on the number of steps to be performed and the current step being executed.

Note: when multiple hypers of the same device are being swapped they are swapped one at a time. As such the step number will reset as each hyper is being swapped.

When the swap is completed, the swap list will be empty.

Performing an Optimizer swap rollback

Optimizer rollback allows previously performed swaps to be reversed. Swaps can be rolled back to a particular point-in-time.

The list of swaps which are available to be rolled back can be seen by running:

```
symoptmz -sid 3542 show -rollback_list
```

This operation may take up to a few minutes. Please be patient...

```
Number of Optimizer Swap History Entries : 1
```

```
Swap Group #1:
```

```
{
Time Started      : Wed Feb 24 15:35:56 2010
Time Completed    : Wed Feb 24 15:57:37 2010
Swap Attribs     : SWAP OPTIMIZER-GENERATED
Number of Moves/Swaps: 16
```

```
Move Number      #1:
{
Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
-----  -
Hyper1   11D      D          E          1        0288    2
Hyper2   02D      C          4          8        0288    2
}
```

```
Move Number      #2:
{
Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
-----  -
Hyper1   02D      C          4          8        0284    2
Hyper2   11D      D          E          1        0284    2
}
```

...

```
Move Number      #15:
{
Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
-----  -
Hyper1   06C      D          C          2        0289    1
Hyper2   15D      C          5          8        0289    1
}
```

```
Move Number      #16:
{
Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
```

```

-----
Hyper1      15D      C          5          8      0263      1
Hyper2      06C      D          C          2      0263      1
}
}

```

To perform the rollback, a text file with the following syntax should be created:

```

set swap approve,
TIMESTAMP=02242010:153536,
ROLLBACK;

```

The TIMESTAMP value corresponds to the point-in-time to rollback to, and is available in the Time Started field of the rollback list output.

Once created, the rollback can be executed by running:

```
symoptmz -sid 3542 -f rollback.txt commit -v
```

```

Processing Command file : rollback.txt
PREVIEW .....Started.
Set Swap Command
{
  set swap APPROVE
    TIMESTAMP = Wed Feb 24 15:35:36 2010
    ROLLBACK;
}
PREVIEW .....Done.
PREPARE .....Started.
Optimizer has Validated the Rollback List.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.

```

By default, the rollback will be executed during the next available device movement time window.

Optionally, a specific time can be given for the rollback to be performed:

```

set swap approve,
begin_at=02252010:180000,
timestamp=02242010:153536,
ROLLBACK;

```

```
symoptmz -sid 3542 -f rollback.txt commit -v
```

```

Processing Command file : rollback.txt
PREVIEW .....Started.
Set Swap Command
{

```

```

set swap APPROVE
    begin_at = Thu Feb 25 18:00:00 2010
    TIMESTAMP = Wed Feb 24 15:35:36 2010
    ROLLBACK;
}
PREVIEW .....Done.
PREPARE .....Started.
Optimizer has Validated the Rollback List.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.

```

Note: The timestamp fields in the rollback command file are of the format MMDDYYYY:HHMMSS.

Performing a manual Optimizer swap

A manual swap can be performed where the user selects the individual physical hypervisors of two devices to be swapped.

Note: prior to submitting a manual swap, Optimizer must first be stopped/disabled.

To perform a manual swap, a text file should be created containing the list of hypervisors to be swapped. Each hypervisor in the text file will be denoted by the director (DA) identifier, the director interface, target id, and hypervisor number.

The hypervisors location information can be found in the output of the `symdev show` command.

```
symdev -sid 3542 show 24B
```

```

...
Device Symmetrix Name      : 024B
Device Serial ID          : N/A
Symmetrix ID              : 000190103542
...
Device Configuration      : RAID-5                (Non-Exclusive Access)
...
RAID-5 Device Information
{

```

```

Number of Tracks in a Stripe          : 4
Overall Ready State of RAID-5 Device  : ReadyNoOtherMirror
Overall WriteProtect State of RAID-5 Device : EnabledNoOtherMirror
Member Number of the Failing Device    : None
Member Number that Invoked the Spare   : None
Disk Director (DA) that Owns the Spare : None
Copy Direction                        : N/A
RAID-5 Hyper Devices (3+1):
{
  Device : 024B
  {
    -----
    Disk      DA      Hyper      Member      Spare      Disk
    DA :IT    Vol#    Num Cap(MB)  Num Status  Status  Grp#  Cap(MB)
    -----
    01B:C1    8      8      5755      1 RW      N/A    1    140014
    02C:C1    8      8      5755      2 RW      N/A    1    140014
    16C:C1    8      8      5755      3 RW      N/A    1    140014
    15B:C1    8      8      5755      4 RW      N/A    1    140014
  }
}

```

In the above example, the location of the first hyper in the RAID 5 raid group representing device 24B is represented as follows:

```

Director (DA) Identifier : 01B
Director Interface       : C
Target ID                : 1
Hyper number             : 8

```

Using this information, the text file necessary to swap all 4 hypers of two RAID 5 (3+1) devices would look like this:

```

set swap_list
  {01B,C,1,8} with {01A,C,2,8},
  {02C,C,1,8} with {02C,D,2,8},
  {16C,C,1,8} with {16C,D,2,8},
  {15B,C,1,8} with {15A,C,2,8};

```

To execute the manual swap during the next swap time window run:

```
symoptmz -sid 3542 -f manual_swap.txt commit -v
```

```

Processing Command file : manual_swap.txt
PREVIEW .....Started.
Set Swap Command
{
  set swap_list {15B,C,1,8} with {15A,C,2,8},
                {15A,C,2,8} with {15B,C,1,8},
                {16C,C,1,8} with {16C,D,2,8},
                {16C,D,2,8} with {16C,C,1,8},
}

```

```

        {02C,C,1,8} with {02C,D,2,8},
        {02C,D,2,8} with {02C,C,1,8},
        {01B,C,1,8} with {01A,C,2,8},
        {01A,C,2,8} with {01B,C,1,8};
    }
PREVIEW .....Done.
PREPARE .....Started.
Optimizer has Validated the Manual Swap List.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.

```

To execute the manual swap at a specific time, a `begin_at` value can be added to the text file.

```

set swap_list
    {01B,C,1,8} with {01A,C,2,8},
    {02C,C,1,8} with {02C,D,2,8},
    {16C,C,1,8} with {16C,D,2,8},
    {15B,C,1,8} with {15A,C,2,8},
begin_at=02252010:180000;

```

```
symoptmz -sid 3542 -f manual_swap.txt commit -v
```

```

Processing Command file : manual_swap.txt
PREVIEW .....Started.
Set Swap Command
{
  set swap_list {15B,C,1,8} with {15A,C,2,8},
                {15A,C,2,8} with {15B,C,1,8},
                {16C,C,1,8} with {16C,D,2,8},
                {16C,D,2,8} with {16C,C,1,8},
                {02C,C,1,8} with {02C,D,2,8},
                {02C,D,2,8} with {02C,C,1,8},
                {01B,C,1,8} with {01A,C,2,8},
                {01A,C,2,8} with {01B,C,1,8}
  begin_at=Thu Feb 25 18:00:00 2010
}
PREVIEW .....Done.
PREPARE .....Started.
Optimizer has Validated the Manual Swap List.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.

```

Monitoring of a manual swap is performed in a similar fashion as discussed earlier in this section.

Performing a Virtual LUN migration (manual move)

Symmetrix devices can be manually moved between physical disk groups, of different disk technologies, using the Virtual LUN migration feature of Optimizer.

The source devices for a migration may be specified individually, as a device list, or collectively by device group. The target destination for the device migration can be specified as configured space, or unconfigured space.

The following sections detail examples of performing migrations to both configured and unconfigured space.

Note: prior to performing a manual migration Optimizer must be stopped/disabled.

Performing a VLUN migration to unconfigured space

To perform a migration to unconfigured space, and to specify individual devices to be migrated, a text file similar to the following should be created:

```
migrate devs 241,242,243,244
      TO disk_group_num 3;
```

The disk group number specified relates to the physical storage to which the devices are to be migrated to.

Note: While a list of specific disks of the same type can be listed after TO, it is recommended that a disk group number be specified. In this way Symmetrix Engenuity can determine the most appropriate layout for the device once migrated and ensures that configuration rules will not be violated.

Once the device file has been created, the migration can be submitted to Optimizer by running:

```
symoptmz -sid 3542 -f migrate_to_unconfig.txt commit -v
```

```
Processing Command file : migrate_to_unconfig.txt
PREVIEW .....Started.
Migration Command
{
  migrate dev[s] 0241, 0242, 0243, 0244
    to disk_group_num 3;
}
PREVIEW .....Done.
PREPARE .....Started.
Optimizer has Validated the Manual Swap List.
PREPARE .....Done.
```



```

COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.

```

To schedule the migration to occur at a specific time, the text file to perform the migration can be edited to add a `begin_at` value:

```

migrate devs 241,242,243,244
TO disk_group_num 3,
begin_at 02262010:180000;

```

And then running:

```
symoptmz -sid 3542 -f migrate_to_unconfig.txt commit -v
```

```

Processing Command file : migrate_to_unconfig.txt
PREVIEW .....Started.
Migration Command
{
  migrate dev[s] 0241, 0242, 0243, 0244
  to disk_group_num 3
  begin_at=Fri Feb 26 18:00:00 2010
  ;
}
PREVIEW .....Done.
PREPARE .....Started.
Optimizer has Validated the Manual Swap List.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.

```

Virtual LUN migrations can be viewed and monitored in the same way as normal Optimizer swaps, as discussed earlier in this section.

Performing a VLUN migration to unconfigured space

Performing a migration to configured space, and using a device group as the source, is achieved by specifying the device group name and that unmapped and/or unmasked devices in the target storage should be used for the migration. This is done by creating a text file as follows:

```

migrate device_group Production_App1
TO disk_group_num 3,
unmapped=TRUE,
unmasked=TRUE;

```

And then running:

```
symoptmz -sid 3542 -f migrate_config.txt commit -v
```

```
Processing Command file : migrate_config.txt
PREVIEW .....Started.
Migration Command
{
  migrate device_group Production_App1
    to disk_group_num 3
    unmapped=TRUE unmasked=TRUE
  ;
}
PREVIEW .....Done.
PREPARE .....Started.
Optimizer has Validated the Manual Swap List.
PREPARE .....Done.
COMMIT .....Started.
COMMIT .....Done.
Closing Optimizer Session .....Done.
```

To schedule the migration to occur at a specific time, the text file to perform the migration can be edited to add a `begin_at` value:

```
migrate device_group Production_App1
  TO disk_group_num 3,
  unmapped=TRUE,
  unmasked=TRUE,
  begin_at=02262010:180000;
```

And running:

```
symoptmz -sid 3542 -f migrate_config.txt commit -v
```

```
Processing Command file : migrate_config.txt
PREVIEW .....Started.
Migration Command
{
  migrate device_group Production_App1
    to disk_group_num 3
    unmapped=TRUE unmasked=TRUE
    begin_at=Fri Feb 26 18:00:00 2010
  ;
}
PREVIEW .....Done.
PREPARE .....Started.
Optimizer has Validated the Manual Swap List.
PREPARE .....Done.
COMMIT .....Started.
```

```
COMMIT .....Done.  
Closing Optimizer Session .....Done.
```

Virtual LUN migrations can be viewed and monitored in the same way as normal Optimizer swaps, as discussed earlier in this section.

Logging

As Optimizer will modify the back end configuration of the Symmetrix array it is running on, it is highly important from the perspective of the storage administrator to be able to track swap activity. Several methods exist within the Symmetrix and Solutions Enabler to allow an administrator to determine both what changes were made, and when the changes occurred, both from the Optimizer configuration and Symmetrix back end configuration point of view. These include:

- ◆ Optimizer history
- ◆ The Symmetrix audit log
- ◆ Symmetrix-based Optimizer logs
- ◆ The Solutions Enabler event daemon
- ◆ SMC Alerts

The following sections detail the methods for tracking FAST activity, and what information is available with each of the methods.

Optimizer history

The Optimizer controller maintains a history of all device swaps for up to 1 year on the Symmetrix service processor. The history includes all Optimizer activity—automated swaps, manual swaps, VLUN migrations, FAST swaps, and rollbacks.

The information included in the Optimizer history is:

- ◆ Devices involved in the swap
- ◆ Old and new location of device hypers swapped
- ◆ The swap start time
- ◆ The swap completion time

To view the Optimizer swap history from SMC, right-click on the Symmetrix serial number, then choose Optimizer, Swap/Move, Swap/Move History.

By default, the previous thirty days worth of activity will be displayed. The start date and end date can be modified and the results filtered to display activity other than the previous thirty days.

Optimizer History(000190103542)

From Start Time: Jan-25-2010 16:30

To End Time: Feb-24-2010 16:30

Filter Start Times

| Group | Device | Mirror / Member | Start Time (EDT) | End time (EDT) | From | To |
|-------|--------|-----------------|-------------------|-------------------|---------------|---------------|
| 2 | 0249 | 3 | Jan-29-2010 10:16 | Jan-29-2010 10:38 | DF-1B:D 0 H8 | DF-6A:C 3 H2 |
| 2 | 0262 | 4 | Jan-29-2010 10:16 | Jan-29-2010 10:38 | DF-6B:D 1 H2 | DF-15D:D 0 H8 |
| 3 | 0284 | 2 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-2D:C 4 H8 | DF-11D:D E H1 |
| 3 | 0263 | 4 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-16A:C 5 H8 | DF-6B:D 9 H2 |
| 3 | 0288 | 4 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-11B:D A H1 | DF-16D:C 4 H8 |
| 3 | 0284 | 3 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-16D:C 4 H8 | DF-11B:D A H1 |
| 3 | 0289 | 1 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-6C:D C H2 | DF-15D:C 5 H8 |
| 3 | 0263 | 1 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-15D:C 5 H8 | DF-6C:D C H2 |
| 3 | 0289 | 3 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-11D:C B H2 | DF-1D:C 5 H8 |
| 3 | 0263 | 3 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-1D:C 5 H8 | DF-11D:C B H2 |
| 3 | 0288 | 1 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-6D:C A H1 | DF-15A:D 7 H8 |
| 3 | 0284 | 1 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-15A:D 7 H8 | DF-6D:C A H1 |
| 3 | 0288 | 2 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-11D:D E H1 | DF-2D:C 4 H8 |
| 3 | 0289 | 4 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-11A:C C H2 | DF-2A:C 5 H8 |
| 3 | 0263 | 2 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-2A:C 5 H8 | DF-11A:C C H2 |
| 3 | 0288 | 3 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-11A:D D H1 | DF-1A:D 7 H8 |
| 3 | 0284 | 4 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-1A:D 7 H8 | DF-11A:D D H1 |
| 3 | 0289 | 2 | Feb-24-2010 15:35 | Feb-24-2010 15:57 | DF-6B:D 9 H2 | DF-16A:C 5 H8 |

Close Help

From the SYMCLI, to see a complete history of all FAST activity, the following command is run:

```
symoptmz -sid 3542 show -swap_hist
```

This operation may take up to a few minutes. Please be patient...

Number of Optimizer Swap History Entries : 3

Swap Group #1:

```
{
  Time Started      : Thu Jan 28 16:01:12 2010
  Time Completed   : Thu Jan 28 16:27:07 2010
  Swap Attribs    : SWAP OPTIMIZER-GENERATED
  Number of Moves/ Swaps: 32
```

Move Number #1:

```
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev Mirror
  -----  -----  -----  -----  -----  -----
```

Logging

```
Hyper1    11C    C          8          2    0284    1
Hyper2    15A    D          7          8    0284    1
}
```

...

```
Move Number    #32:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -
Hyper1    02C    C          7          8    0247    1
Hyper2    06D    D          5          2    0247    1
}
}
```

Swap Group #2:

```
{
Time Started      : Fri Jan 29 10:16:08 2010
Time Completed    : Fri Jan 29 10:38:24 2010
Swap Attribs     : SWAP OPTIMIZER-GENERATED
Number of Moves/Swaps: 16
```

```
Move Number    #1:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -
Hyper1    11C    D          3          2    0265    3
Hyper2    01C    C          2          8    0265    3
}
}
```

...

```
Move Number    #16:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -
Hyper1    01D    D          0          8    0264    3
Hyper2    11C    D          1          2    0264    3
}
}
```

Swap Group #3:

```
{
Time Started      : Wed Feb 24 15:35:56 2010
Time Completed    : Wed Feb 24 15:57:37 2010
Swap Attribs     : SWAP OPTIMIZER-GENERATED
Number of Moves/Swaps: 16
```

```
Move Number    #1:
{
  Device  DA-Num  DA-Int  Disk SCSI ID  HyperNum  SymDev  Mirror
  -----  -
Hyper1    11D    D          E          1    0288    2
}
}
```

```

Hyper2      02D      C          4          8      0288      2
}

...

Move Number #16:
{
Device      DA-Num   DA-Int   Disk SCSI ID   HyperNum   SymDev   Mirror
-----
Hyper1      15D      C          5          8      0263      1
Hyper2      06C      D          C          2      0263      1
}
}

```

Symmetrix audit log

The Symmetrix audit log is a Symmetrix-based log used to correlate application-based control operations run against the array. Entries to the audit log are categorized into applications, function classes, and action codes, with each product feature having its own application and function class.

Optimizer-related activities logged to the Symmetrix audit log include:

- ◆ Optimizer controller state changes
- ◆ Optimizer settings modifications
- ◆ Swap list approval/decline
- ◆ Swap/migration execution

Contents of this audit log can be viewed and filtered using the SYMCLI **symaudit** command. Optimizer activities are logged under the function class “Optmzr”.

The audit log entries contain a hostname indicating from where the control operation was initiated. In the case where the hostname is a Symmetrix serial number, this means the Optimizer application submitted a swap list to be executed.

Changes to the Optimizer configuration settings will be logged from the host upon which the control (**symoptmz**) commands were run.

To view Optimizer activity in the audit log via SMC, right-click on the Symmetrix serial number, choose Symmetrix Admin, View Symmetrix Audit Log.

In the Range area, specify a start and end date for the time period of interest. Under Filter, select “Optmzr” in the Function Class field. To retrieve the log, click the Get Log button.

Logging

The screenshot displays the Symmetrix Admin interface for SymAudit Records (000190103542). The interface includes a 'Log Controls' section with 'Range' and 'Filter' options. The 'Range' section allows filtering by date (selected) or record number, with start and end dates and record counts. The 'Filter' section includes fields for Function Class, Host, Vendor ID, Activity ID, Start Device, Action Code, Process ID, Application ID, Username, and End Device. A 'Get Log' button is highlighted. Below the controls is a table of audit records with columns: Record Number, Date, Application ID, Host, OS Name, Function Class, Action Code, Activity ID, and Username. The table contains 15 records, with the last one (Record 2325) highlighted. At the bottom, there are 'Print', 'Export', 'Details ...', 'Done', and 'Help' buttons.

| Record Number | Date | Application ID | Host | OS Name | Function Class | Action Code | Activity ID | Username |
|---------------|---------------------|----------------|----------|---------|----------------|-------------|--------------|-----------------|
| 1487 | 2010-02-04 14:49:55 | SMC | licod207 | SunOS | Optmzr | Enable | SEe64f2a017c | H:\licod207\smc |
| 1488 | 2010-02-04 14:49:55 | SMC | licod207 | SunOS | Optmzr | Enable | SEe64f2a017c | H:\licod207\smc |
| 2135 | 2010-02-18 15:45:38 | SMC | licod207 | SunOS | Optmzr | Disable | SE996d114c23 | H:\licod207\smc |
| 2136 | 2010-02-18 15:45:38 | SMC | licod207 | SunOS | Optmzr | Disable | SE996d114c23 | H:\licod207\smc |
| 2137 | 2010-02-18 15:46:34 | SMC | licod207 | SunOS | Optmzr | Enable | SE996d114c23 | H:\licod207\smc |
| 2138 | 2010-02-18 15:46:34 | SMC | licod207 | SunOS | Optmzr | Enable | SE996d114c23 | H:\licod207\smc |
| 2139 | 2010-02-18 15:51:43 | SMC | licod207 | SunOS | Optmzr | Enable | SE996d114c23 | H:\licod207\smc |
| 2140 | 2010-02-18 15:51:43 | SMC | licod207 | SunOS | Optmzr | Enable | SE996d114c23 | H:\licod207\smc |
| 2323 | 2010-02-22 11:54:07 | SMC | licod207 | SunOS | Optmzr | Enable | SE2f53dfe43c | H:\licod207\smc |
| 2324 | 2010-02-22 11:54:07 | SMC | licod207 | SunOS | Optmzr | Enable | SE2f53dfe43c | H:\licod207\smc |
| 2325 | 2010-02-22 12:08:25 | SMC | licod207 | SunOS | Optmzr | Disable | SE2f53dfe43c | H:\licod207\smc |

More information on a particular entry in the audit log can be retrieved by highlighting the entry and clicking on the Details button.

| Attribute | Value |
|---------------------|---|
| Record Number | 2448 |
| Records in Seq | 1 |
| Offset in Seq | 1 |
| Date | 2010-02-23 18:17:15.0 |
| Vendor ID | EMC Corp |
| Application ID | OPTIMIZER |
| Application Version | 7.1.0.36 |
| API Library | SEK |
| API Version | V7.0-956 |
| Host | HK190103542 |
| OS Name | WinNT |
| OS Revision | 5.1.2600Se |
| Client Host | |
| Process ID | 1984 |
| Task ID | 1992 |
| Function Class | Optmzr |
| Action Code | OptmzrSwap |
| Activity ID | SE7fe456b82f |
| Username | H:HK190103542\SYSTEM |
| Text | Hyper swap has been initiated by an optimizer client to be executed on 02/24/2010 11:00 |

Buttons: Print, Export, Back, Forward, Done, Help

To view a list of all Optimizer-related activities from SYMCLI, the following command is run:

```
symaudit -sid 3542 list -function_class Optmzr
```

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

```
Symmetrix ID           : 000190103542
```

```
Record Number  Date      Time      Application      Host      Function Class  Action Code
-----
```

Logging

```
    20 01/05/10 14:26:18 SMC          licod207  Optmzr  Disable
    21 01/05/10 14:26:18 SMC          licod207  Optmzr  Disable
...
 2439 02/23/10 12:16:11 SYMPTMZ     licod207  Optmzr  Commit
 2440 02/23/10 12:16:11 SYMPTMZ     licod207  Optmzr  Commit
 2448 02/23/10 18:17:15 OPTIMIZER   HK190103542 Optmzr  Swap
 2493 02/24/10 15:25:04 OPTIMIZER   HK190103542 Optmzr  Swap
 2494 02/24/10 15:31:14 OPTIMIZER   HK190103542 Optmzr  Swap
 2495 02/24/10 15:31:09 OPTIMIZER   HK190103542 Optmzr  Swap
 2496 02/24/10 15:31:09 OPTIMIZER   HK190103542 Optmzr  Swap
 2531 02/24/10 16:28:41 OPTIMIZER   HK190103542 Optmzr  Swap
 2532 02/24/10 16:28:41 OPTIMIZER   HK190103542 Optmzr  Swap
```

To see a particular audit log entry in more detail, the **symaudit** command can be run with the **-v** option and specifying the entry's record number:

```
symaudit -sid 3542 list -record_num 2448 -n 1 -v
```

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

Symmetrix ID          : 000190103542

Record Number         :      2448
Records in Seq        :          1
Offset in Seq         :          1
Time                  : 02/23/10 18:17:15
Vendor ID             : EMC Corp
Application ID         : OPTIMIZER
Application Version    : 7.1.0.36
API Library           : SEK
API Version           : V7.0.1.0 (Edit Level: 956)
Host Name             : HK190103542
OS Name               : WinNT
OS Revision           : 5.1.2600Se
Client Host           :
Process ID            : 00001984
Task ID               : 00001992
Function Class        : Optmzr
Action Code           : OptmzrSwap
Text                  : Hyper swap has been initiated by an optimizer client to
be executed on 02/24/2010 11:00:00 PM.
Username              : H:HK190103542\SYSTEM
Activity ID           : SE7fe456b82f
```

Symmetrix-based Optimizer logs

The Optimizer controller maintains runtime and error logs on the Symmetrix service processor. Retrieval of these log files is performed using the **symoptmz** command.

The runtime log file contains the following information:

- ◆ Device performance statistics being collected
- ◆ Configuration change plan generation
- ◆ Configuration change plan execution
- ◆ Changes to Optimizer controller settings and time windows
- ◆ Client connections to the Optimizer controller

To view the runtime log, run:

```
symoptmz -sid 3542 read -log_type RUNTIME -start 02242010:0000
```

This operation may take up to a few minutes. Please be patient...

Size of Log file data read: 21310 characters

Optimizer File Data:

Symmetrix Id: 000190103542

File Type : LOG

Log Type : RUNTIME

Start date: 02242010:000000

Stop date: NONE

Data :

02/24/2010 05:09:10 AM (05:09:10 AM GMT)

(I00500) Successfully fetched volume statistics

...

02/24/2010 08:59:15 AM (08:59:15 AM GMT)

(I00500) Successfully fetched volume statistics

02/24/2010 09:09:15 AM (09:09:15 AM GMT)

(I00500) Successfully fetched volume statistics

...

02/24/2010 09:55:02 PM (09:55:02 PM GMT)

(I07504) No swap suggestions found!

02/24/2010 09:55:02 PM (09:55:02 PM GMT)

(I00503) Optimizer does not recommend any configuration changes at this time

02/24/2010 09:55:02 PM (09:55:02 PM GMT)

(I10500) New configuration change plan: Empty

...

The error log file contains the following information:

- ◆ Errors encountered by the Optimizer controller

To view the error log, run:

```
symoptmz -sid 3542 read -log_type ERROR -start 01012010:000000
```

This operation may take up to a few minutes. Please be patient...

Size of Log file data read: 662 characters

Optimizer File Data:

Symmetrix Id: 000190103542

File Type : LOG

Log Type : ERROR

Start date: 01012010:0000

Stop date: NONE

Data :

01/04/2010 03:51:28 PM (03:51:28 PM GMT)

(E00008) Optimizer is stopping on error!

01/04/2010 03:51:28 PM (03:51:28 PM GMT)

(E10016) Contacting EMC support personnel:

There are not enough DRVs to create swaps.

Note: Retrieval of runtime or error log activity can be filtered by start and end times. The format of the date and time for each option is mmddyyyy:hhmm.

Note: The date and timestamps displayed in the Optimizer runtime and error logs are GMT, and do not adjust for daylight savings time.

Event Daemon

In UNIX, Linux, and Windows environments, the event daemon—**storevntd**—enables monitoring of Symmetrix operations by detecting and reporting events as they happen. The event daemon continually collects Symmetrix event information in real-time, filters the events by severity and type, and responds by logging events to specified targets. When configuring the daemon to log events, you can specify to log the events to the UNIX Syslog, the Windows Event log, SNMP, or a file on disk.

Several event codes have been added to Symmetrix operations to track Optimizer activities.

Optimizer event codes

Event codes tracking administrative changes to Optimizer controller configuration settings are shown in Table 1.

Table 1. FAST administration event codes

| Event code | Event description |
|------------|---|
| 1292 | A FAST/Optimizer time window has been changed |
| 1293 | A FAST/Optimizer configuration setting has been changed |

Optimizer activity event codes

Event codes tracking Optimizer activities such as device movements and the availability of swap lists for approval are shown in Table 2.

Table 2. FAST activity event codes

| Event code | Event description |
|------------|--|
| 1506 | A FAST/Optimizer rollback has been performed |
| 1507 | A FAST/Optimizer configuration change plan has been generated, and user approval is required |
| 1508 | The state of the FAST Controller has changed |

The *EMC Solutions Enabler Installation Guide* provides more information on configuring and using the Solutions Enabler event daemon.

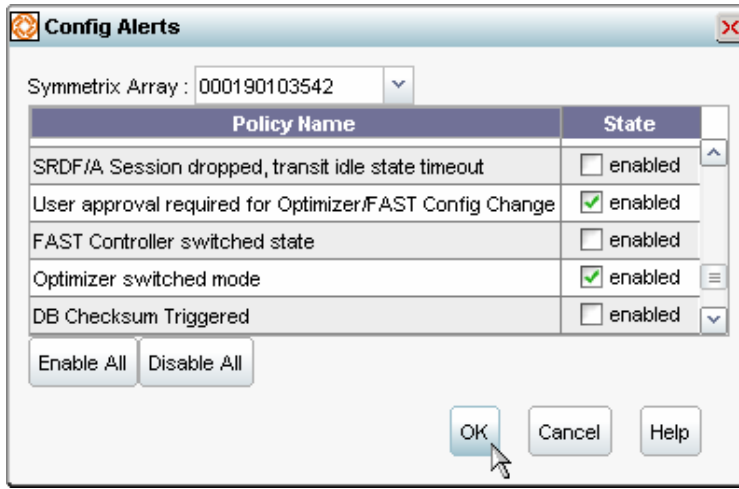
Symmetrix Management Console alerts

In addition to the events that can be monitored through the Solutions Enabler event daemon, two alerts can be configured in SMC to also track Optimizer activity. The two events that can be tracked are:

- ◆ An Optimizer swap list has been generated, and user approval is required
- ◆ Optimizer switched mode

To configure these alerts, the user must select the Tasks view in SMC and click on the Enable/Disable Alerts under Alert Configuration in the Monitoring pane.

From the resulting dialog box, the Optimizer-related alerts can be selected.



Once configured, alerts can be viewed through the Alerts view for a particular Symmetrix.

Conclusion

Symmetrix Optimizer improves array performance by continuously monitoring access patterns and migrating devices (Symmetrix logical volumes) to achieve balance across the drives within a physical disk group, and thereby reduce the risk of “hot spots.” This automated process is based on user-defined parameters and is completely transparent to end users, hosts, and applications in the environment.

Symmetrix Optimizer also allows for movement of devices between differing disk technologies to change the performance profile of the device. Such movement can be performed manually, using Virtual LUN technology, or automatically, based on performance workload, using FAST for DMX.

Appendix A: Affinity Groups

Symmetrix Engenuity 5771 introduced a new configuration rule for all RAID protection types, known as affinity groups. An affinity group binds together all physical drives containing members of a particular RAID type. Subsequent RAID devices, of the same type, created on any one of the drives contained in the RAID group must stay within the bounds of that affinity group.

To illustrate the use of affinity groups within the Symmetrix, the creation and configuration of RAID 5 (3+1) devices will be used as an example.

RAID 5 Overview

RAID 5 is architected around block-based multi-spindle parity protection. RAID 5 implements data striping and rotating parity across all physical hypervisors of a RAID 5 device (logical volume). Data blocks are striped (or interleaved) horizontally across the members of a RAID 5 group, similar to a striped metadata device. Each member owns some data tracks and some parity tracks.


The entire set of RAID 5 hypervisors is presented to the host as a single RAID 5 device. The physical hypervisors must reside on different physical disks.

RAID 5 (3+1)


A RAID 5 (3+1) configuration consists of four Symmetrix physical hypervisors with parity and data striped across each. As such, each Symmetrix device configured as RAID 5 (3+1) will consume 133% of the configured capacity—33% of the device capacity is configured as parity.

RAID 5 (3+1) affinity group

When a RAID 5 (3+1) device is created, physical hypervisors are created on four separate physical disks. These hypervisors constitute the RAID 5 group. The four disks the RAID group members are configured on constitute the affinity group.



As additional RAID 5 (3+1) devices are created on the same physicals, all hypers will be configured on the same physical disks contained within the affinity group.



Other RAID type affinity groups

As each RAID protection contains a different number of RAID members (hypers), the affinity group sizes will vary based on RAID type.

RAID 1 devices contain two physical hypers, and the affinity group size

is two physical disks.

RAID 5 (7+1) devices contain eight physical hypervisors, and the affinity group size is eight physical disks.

RAID 6 (6+2) and RAID 6 (14+2) contain eight and sixteen physical hypervisors, respectively. As such, the affinity group size for each are eight and sixteen.

Physical disk groups

A disk group is a group of physical disks, configured in the Symmetrix bin file, formed when the bin file is initially written or when disks are added to the array. Disk groups are used to separate physical disks by technology, capacity, and rotational speed (if applicable).

The number of affinity groups of each RAID protection type available within a disk group will depend on the number of disks within the disk group. To determine how many affinity groups of each RAID type are available, divide the number of hypervisors in the RAID group into the number of disks in the disk group.

For example, a disk group containing 128 physical disks can support the following number of affinity groups:

- ◆ 64 — RAID 1
- ◆ 32 — RAID 5 (3+1)
- ◆ 16 — RAID 5 (7+1)
- ◆ 16 — RAID 6 (6+2)
- ◆ 8 — RAID 6 (14+2)

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 V-Max Product Guide*
- ◆ *FAST Theory and Best Practices for Planning and Performance*
- ◆ *Managing your FAST environment with Ionix ControlCenter and Symmetrix Performance Analyzer*
- ◆ *Best Practices for Nondisruptive Tiering via EMC Symmetrix Virtual LUN*

Copyright © 2010 EMC Corporation. All Rights Reserved.

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

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

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

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

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