

Washington Systems Center - Storage

**Accelerate with IBM
Storage:**

IBM Spectrum Virtualize Easy Tier Best Practices



-Andrew Martin

WW United Kingdom Team

Lead Spectrum Virtualize SVC Architect

-Brad Ashbaugh

Washington Systems Center

IBM Certified Storage SME

https://www.ibm.com/support/knowledgecenter/STSLR9_8.2.0/com.ibm.fs9100_820.doc/svc_easy_tier.html

Easy Tier
Quick
Reference
Link

Accelerate with IBM Storage Webinars

The Free IBM Storage Technical Webinar Series Continues in 2019...

Washington Systems Center – Storage experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

Located in the Accelerate with IBM Storage Blog:
<https://www.ibm.com/developerworks/mydeveloperworks/blogs/accelerate/?lang=en>

Also, check out the WSC YouTube Channel here:
https://www.youtube.com/channel/UCNuks0go01_ZrVVF1jgOD6Q

2019 Upcoming Webinars:

May 8 - Discover Latest Capabilities of Spectrum Virtualize Easy Tier Setup(s)

May 23 - How to use the IBM TS4500 and TS4300 Library Web GUI

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=edaae73a07257503198d0bdc04d18973e>

May 30 - New Features with IBM Cloud Object Storage including a File Access Demo

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=eed0398a8690f50d436c35ba50c0135ac>



IBM Spectrum Virtualize Software Easy Tier Function

Proven, Efficient, Dependable, Effective

- Easy Tier® eliminates manual intervention moving highly active data on volumes to faster responding storage on the fly, non-disruptively
- Five types of MDisks supported; only up to 3 in a pool:
 - Tier NVMe exists so the fastest storage media can communicate over NVMe protocol versus SCSI
 - Tier 0 flash tier exists when the pool contains high performance flash drives
 - Tier 1 flash tier exists when the pool contains tier 1 flash drives. Tier 1 flash drives typically offer larger capacities, but slightly lower performance and write endurance characteristics
 - Enterprise tier exists when the pool contains enterprise-class MDisk, which are disk drives that are optimized for performance
 - Nearline tier exists when the pool contains nearline-class MDisk, which are disk drives that are optimized for capacity



IBM
Spectrum
Virtualize

Enabling clients to do more with less

3x performance with only **5%** flash

Up to **100%** better storage utilization

IBM's Award-Winning Spectrum Virtualize Storage Portfolio, multiple years running all with Easy Tier® support

With Easy Tier
Capabilities

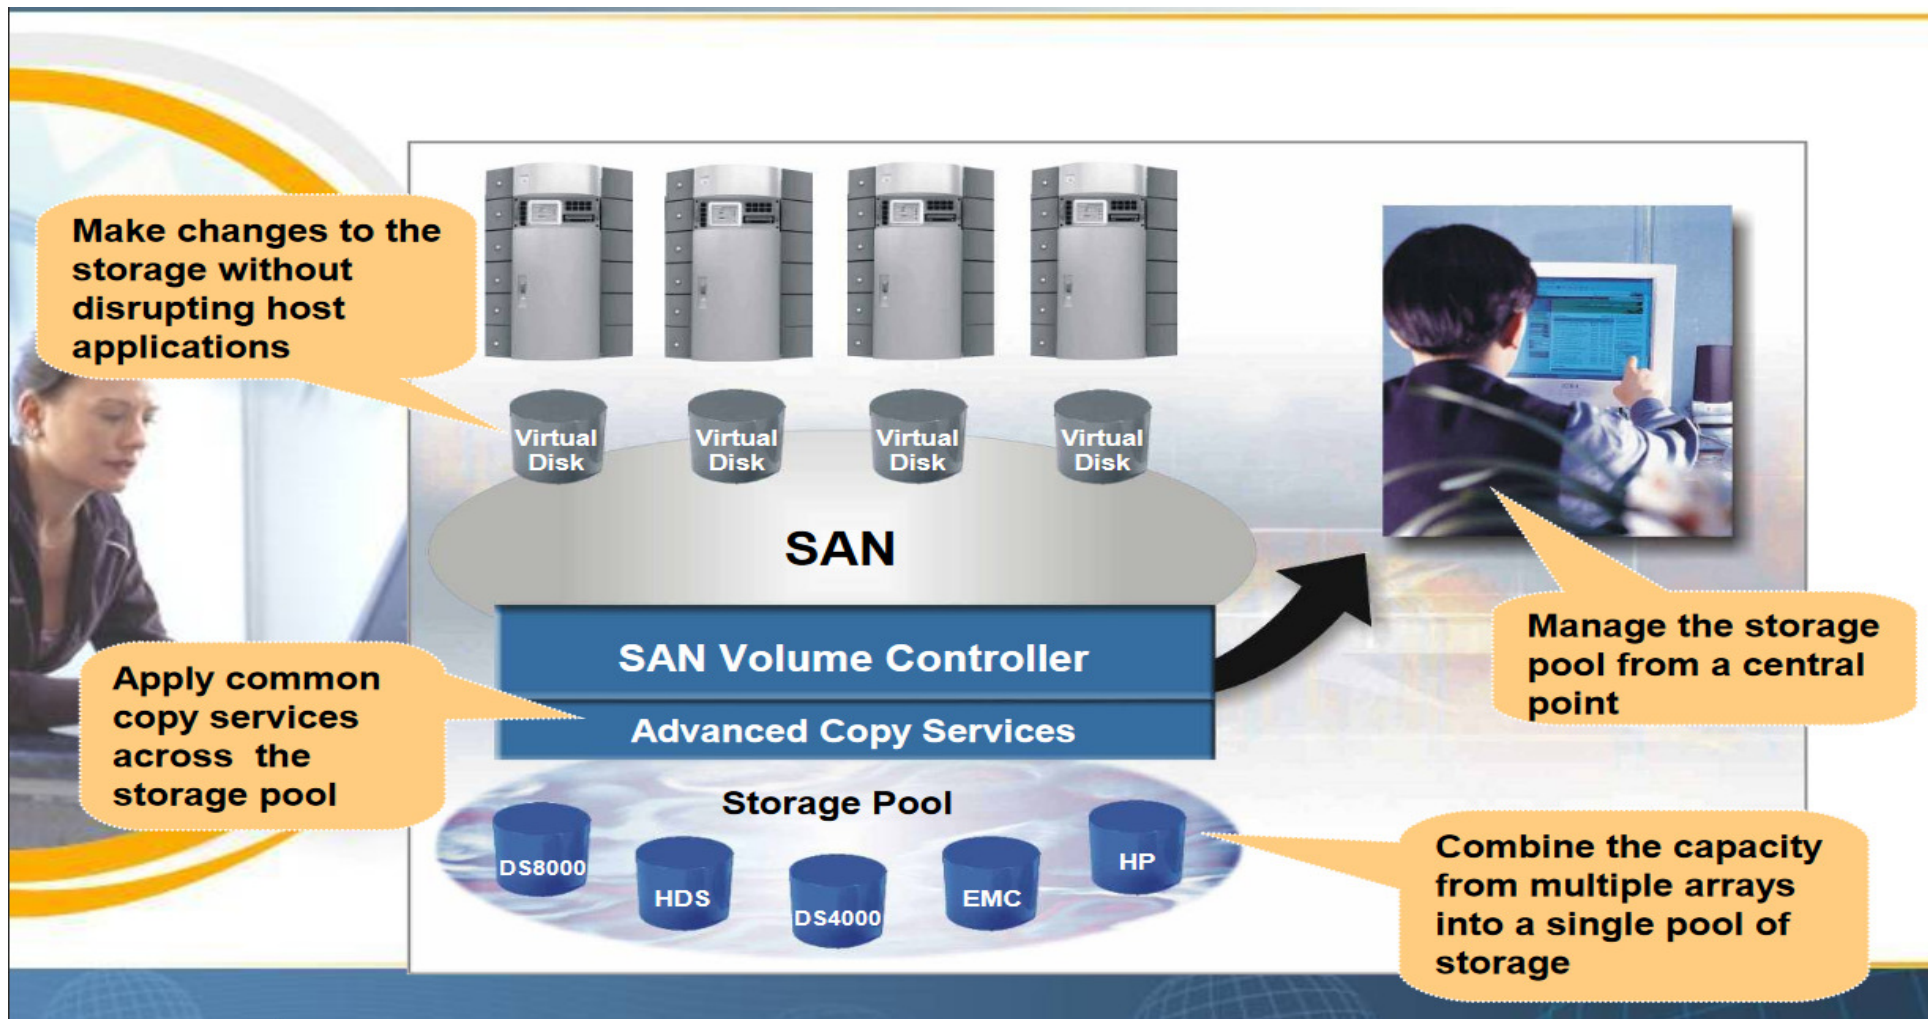
- Once you add multiple MDisks storage media types to a pool (up to 3) it automatically turns storage data sets optimization versus always on load balancing

With Automatic
Machine Learning

Infrastructure



Also Apply Easy Tier and other Management Functions to External Storage



Easy Tier Modes of Operation

- Easy Tier® has three modes of operation:

- Load Balancing Mode

Easy Tier Load Balancing mode collects usage statistics for all the volumes in the single pool at the extent level for each HDD tier.

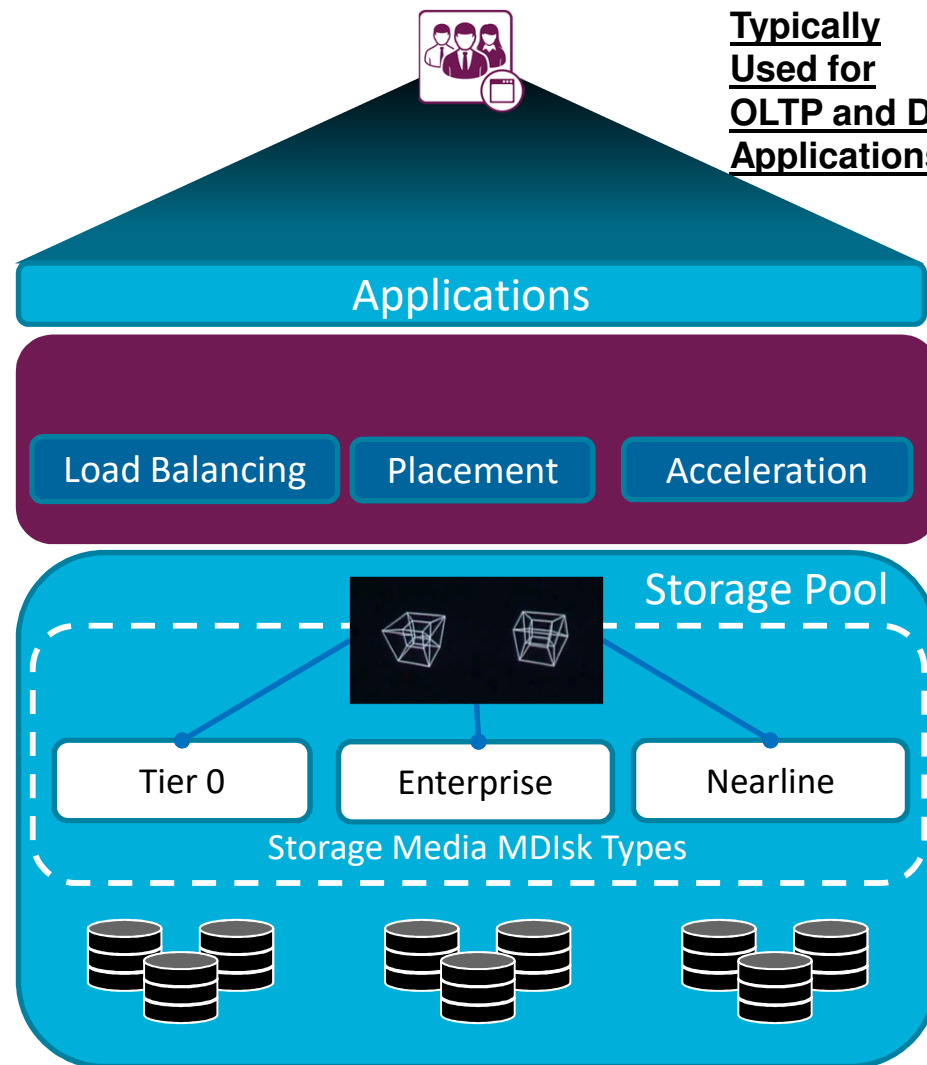
- Automatic Data Placement

Also measures the amount of data access. Easy Tier then acts on the measurements to automatically place the data into the appropriate tier of a storage pool that contains both MDisk tiers. System completes several functions to automatically place the dataset in the appropriate storage tier. Easy Tier measures the host access activity to the data on each storage extent. It also provides a mapping that identifies high activity extents, and then moves the high-activity data according to its relocation plan algorithms.

- Acceleration Mode

Normally, your system runs with the Easy Tier accelerated mode set to off. If turned on it allows you to change the mode rates of extent migrations through the IBM® Easy Tier® function and pool balancing functions from a default migration rate from 12GB per 5 minutes to 48GB per 5 minutes. To avoid the possibility of performance issues that are caused by overloading the managed disks, use the accelerated mode during periods of reduced system activity.

Typically
Used for
OLTP and DB
Applications



Easy Tier Modes of Operation – Load Balancing Con't

- Always on Load Balancing Mode
 - Easy Tier Load Balancing mode is enabled for a storage pool within each single tier of storage. Easy Tier collects usage statistics for all the volumes in the pool at the extent level.
 - Determining the amount of data activity in an extent and when to move the extent to the proper storage tier is usually too complex a task to manage manually.
 - Easy Tier Load Balancing mode collects usage statistics for each storage extent for a storage pool where the capability of moving data from one tier to the other tier is not possible or is disabled.
 - An example of such a storage pool is a pool of homogeneous MDisk, where all MDisk are the same size/type of Flash or HDDs.
 - A summary file is created in the /dumps directory on the configuration node (**dpa_heat.node_name.date.time.data**), which can be offloaded and viewed by using the IBM® Storage Tier Advisor Tool.

Easy Tier Modes of Operation – Evaluation Con't

- IBM® Storage Tier Advisor Tool
 - Information gather can be used to analyze workload statistics and evaluate which logical volumes might be candidates for Easy Tier management and reports can be generated.
 - You can use the usage statistics that are gathered by the monitoring process to help you better understand our environment.

System Summary Report

- The number of storage pools that are monitored
- The number of volumes that are monitored
- The total capacity that is monitored
- The estimated total hot data capacity
- List of storage pools that are monitored, including pool ID, capacity, tier configuration, tier status, and data management status.

Volume Heat Distribution Report


- The number of storage pools that are monitored
- The number of volumes that are monitored
- The total capacity that is monitored
- The estimated total hot data capacity
- List of storage pools that are monitored, including pool ID, capacity, tier configuration, tier status, and data management status.

System Recommendation Report

- A sorted list of suggested flash MDisk's to add to the storage pools (sorted by estimated performance improvement)
- A sorted list of suggested Enterprise MDisk's to add to the storage pools (sorted by estimated IOPS improvement)
- A sorted list of suggested Nearline MDisk's to add to the storage pools (sorted by estimated Bandwidth improvement)
- For each suggested MDisk: the estimated improvement

Storage Pool Recommendation Report

- For each storage pool: list MDisk's current IOPS utilization and projected IOPS utilization
- For each storage pool: show the workload distribution on each tier
- For each storage pool: A list of suggested flash MDisk's to add to the storage pools
- For each storage pool: A list of suggested Enterprise MDisk's to add to the storage pools
- For each storage pool: A list of suggested Nearline MDisk's to add to the storage pools
- For each suggested MDisk: the estimated improvement



Easy Tier Modes of Operation – ADP Con't

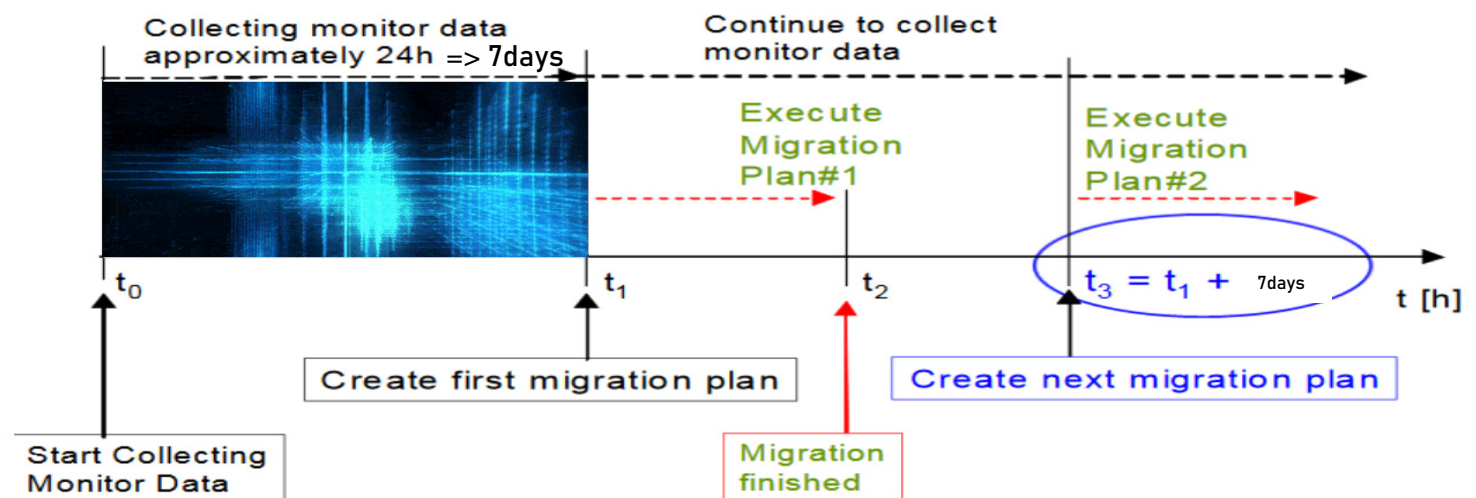
• Automatic Data Placement Mode

To automatically relocate the data, Easy Tier initiates the following processes:

- Monitors volumes for host access to collect average usage statistics for each extent over a rolling 24-hour period of I/O activity.
- Analyzes the amount of I/O activity for each extent to determine whether the extent is a candidate for migrating to or from the higher performing flash drive tier.
- Develops an extent relocation plan for each storage pool to determine exact data relocations within that storage pool. Easy Tier then automatically relocates the data according to the plan.


While it relocates volume extents, Easy Tier follows these actions:

- Attempts to migrate the most active volume extents first.
- Refreshes the task list as the plan changes. The previous plan and any queued extents that are not yet relocated are abandoned.
- Automatic data placement is enabled by default for storage pools with more than one tier of storage. When you enable automatic data placement, by default all striped volumes are candidates for automatic data placement. Image mode and sequential volumes are never candidates for automatic data placement.



Easy Tier Modes of Operation – Acceleration Con't

- Acceleration Mode
 - Enable you to change the mode rates of extent migration through the Easy Tier function and pool balancing functions.
 - Currently configured value, enter the lssystem command
 - Maximum migration rate if enabled is up to 48GB (from 12GB) per 5 minutes and cannot always be guaranteed.
 - use the chsystem command to turn on accelerated mode
 - `chsystem –easytieracceleration on`
 - `chsystem –easytieracceleration off`
 - Why use Acceleration Mode
 - You added more capacity to an existing storage pool either by adding to an existing tier or by adding a tier to the pool. When you enable the accelerated mode, the system can quickly spread the existing volumes into the new capacity.
 - You are migrating multiple volumes between pools, and the target pool has more tiers than the source pool. Volume migration between storage pools preserves the tier of the volume's extents. When you enable the accelerated mode, the system can quickly take advantage of the additional tiers in the target pool.




Easy Tier Benchmark With and Without Easy Tier

Easy Tier and compression is up to three times faster in terms of application response time by having 5% of SSDs in the configuration. Figure below shows the test results of an IBM benchmark on compressed volume with Easy Tier enabled and disabled on V7000. The benchmark was used with an Oracle database and represents a realistic online transaction processing (OLTP) workload.

Storwize V7000 system used in the benchmark was running software version 7.1.0.1 and was using the disk configuration without and with Easy Tier.

Without Easy Tier: 72 x 300 GB SAS HDD With Easy Tier: 68 x 300 GB SAS HDD & 4 x 300 GB SAS SSD



Spectrum Virtualize Easy Tier (ET) Best Practices

For a storage pool and volume to be automatically managed by Easy Tier, ensure that the following conditions are met:

- The volume must be striped.
- When you create a pool in the management GUI that uses Easy Tier, ensure that you add an HDD-based array to the pool first and then add flash drive based arrays.
- Remember that EasyTier only does data migration based on a cost benefit analysis. Demoting something from the flash tier when there isn't anything hotter to replace it with has a cost (moving the data takes time and IOPs) and a negative value (the performance will get worse). Therefore it's a "lose lose" situation that ET will not normally do.
- Note that demoting to Nearline has value even though the performance might be decreased because you are making space for new vdisk allocations.
- The result of this is that the top tier of EasyTier is designed to run almost full. Now full can either mean "capacity is 98% filled up" or it can mean "the total IOPs or MB/s workload currently being sent to this managed disk has reached the predicted capabilities of the mdisk". I refer this as "capacity full" and "performance full"
- To ensure optimal performance, all MDisks in a storage pool tier must have the same technology and performance characteristics.
- Remember Easy Tier is supported on compressed volumes, but it migrates data based only on read activity.
- The Easy Tier function does support ADP or Evaluation mode for image mode or sequential volumes.
- If, for any reason, Easy Tier must be disabled on a certain volume, disable it by using the following command:`svctask chvdisk – easytier off volume_name`

Spectrum Virtualize Easy Tier Best Practices Con't

You probably already knew this - but this fact is becoming much more important in modern systems with overallocated backend storage controllers such as FS9100 with FCM compression in a tiered storage pool.

The lesson is this:

Even if you keep your storage pool 60% full - EasyTier will be trying to keep your top tier about 100% full. So if you have over-allocated your top tier with 2:1 and your data is only achieving 1.8:1 then there is a good chance that your overprovisioned backend controller could run out of space.

Tips to help in this area:

- I. If your top tier is overallocated - be very conservative about the overallocation ratio (at least to start with)
- II. Monitor your overallocated backend and set up alerts if it is running low on space
- III. Run 8.1 or higher on the ET system so you get UNMAP support to help avoid needless out of space
 - A. Note - in 8.2.1 there are 2 unmap settings. Host unmap and storage unmap. **Add content about when and why you should/shouldn't disable the different unmappings**
- IV. Have a backup plan for how to stop the top tier running out of space:
 - A. At 8.2.1 or earlier you can create a temporary volume per Mdisk to prevent ET using any more of the free capacity in that mdisk. At some point later you should reduce the over-provisioning to SVC by removing one or more managed disk from the pool. The properties of that temporary vdisk should be:
 1. EasyTier disabled
 2. format disabled
 3. specify -mdisk to make sure you fill up the free space in the mdisk
 4. Isfreeextents command can tell you how much free space there is on that mdisk. Use that to size the volume
 5. Give it a meaningful name that will remind you not to delete it.
 6. If you are seeing performance issues you may turn of host unmap, however IBM never recommends you disable storage unmap
 - B. NOT YET SHIPPED: At 8.3 (and probably 8.2.1.x ptf) there is a new command that will let you tell SVC to assume that the managed disk is at critical space threshold. This will prevent any new data being migrated to the mdisk and will cause ET to start (slowly) demoting data to free up space from that mdisk

Other notes:

All out of space monitoring for FCMs physically attached to the box running ET is automated and will attempt to prevent out of space using the critical space thresholds mentioned above. So the manual intervention should not be required in a standalone FS9100/V7000 G3/V5100