

EMC Symmetrix Virtual Provisioning

Applied Technology

Abstract

This white paper provides a detailed overview of business requirements as well as technology concepts, choices, and considerations related to Symmetrix[®] Virtual Provisioning for EMC[®] Symmetrix DMX-3 and DMX-4 environments.

June 2008

Copyright © 2008 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.

Part Number H4139.1

Table of Contents

Executive summary	4
Introduction	4
Audience	4
Terminology	5
Business requirements and considerations	5
Benefits	6
Application capacity growth: Virtual Provisioning vs. standard provisioning	6
Use cases	7
Virtual Provisioning on Symmetrix: Implementation overview	8
Product features	10
Configuration	10
Replication	10
Monitoring	11
Tiered storage	11
Flash drives	11
Performance	11
Implementation considerations	11
Conclusion	12

Executive summary

Virtual Provisioning enables organizations to improve speed and ease of use, enhance performance, and increase capacity utilization for certain applications and workloads. EMC® Symmetrix® Virtual Provisioning integrates with existing device management, replication, and management tools, enabling customers to easily build Virtual Provisioning into their existing storage management processes. Virtual Provisioning, which marks a significant advancement over technologies commonly known in the industry as “thin provisioning,” will add a new dimension to tiered storage in the array, without disrupting organizational processes.

Introduction

Organizations continually search for ways to both simplify storage management processes and improve storage capacity utilization. When provisioning storage for a new application, administrators must consider that application’s future capacity requirements rather than simply its current requirements. In order to reduce the risk that storage capacity will be exhausted, disrupting the application and business processes, and to reduce the number of times new storage must be provisioned for that application, organizations often have allocated more physical storage to an application than is needed for a significant amount of time. This allocated but unused storage introduces acquisition and operational costs. Even with the most careful planning, it often is necessary to provision additional storage in the future, which could potentially require an application outage.

EMC Virtual Provisioning addresses these challenges. It builds on the base “thin provisioning” functionality, which is the ability to have a large “thin” device (that is, volume) configured and presented to the host while consuming physical storage from a shared pool only as needed. Symmetrix Virtual Provisioning can improve storage capacity utilization and simplify storage management by presenting the application with sufficient capacity for an extended period of time, reducing the need to provision new storage frequently and avoiding costly allocated but unused storage.




Figure 1. Symmetrix Virtual Provisioning

Audience

This white paper is intended for IT and storage decision makers with organizations that use Symmetrix DMX-3 and DMX-4.

Terminology

Bind	Refers to the act of associating one or more “thin devices” with a “thin pool”
Data device	An “internal device” that provides storage capacity to be used by “thin devices”
Device	A logical unit of storage defined within an array
Device extent	Specifies the size of the smallest contiguous region of a “device” for which a extent mapping can occur
Metadevice	An aggregation of “host accessible devices,” seen from the host as a single “device”
Storage pool	A collection of “internal devices” for some specific purpose
Thin device	A “host accessible device” that has no storage directly associated with it
Thin device allocated capacity	The capacity that has been withdrawn from the “thin pool enabled capacity” for the exclusive use of a “thin device”
Thin pool	A collection of “data devices” that provides storage capacity for “thin devices”

Business requirements and considerations

Organizations large and small seek to reduce the cost of managing their storage infrastructure, while still meeting rigorous service level requirements and accommodating explosive growth in their capacity requirements. Several specific objectives have drawn increasing focus:

- *Reducing the cost of storage administration*
Improving the ease of using storage reduces the associated time and labor expense. “Ease of use” initiatives span multiple aspects of storage processes, including staff training, initial storage provisioning, the addition of new storage, and management and monitoring of storage systems. Symmetrix Management Console (SMC) is being deployed by organizations worldwide to provide intuitive graphical controls and monitoring that help reduce training requirements and accelerate storage management processes. A particular area of concern is the need to dedicate ongoing staff time to manually provision new storage capacity as requirements for existing applications grow over time.
- *Maximizing the utilization of storage assets*
Organizations need to accommodate growth by drawing more value from the same or fewer storage resources. Techniques employed today include TimeFinder[®]/Snap, which creates space-saving snapshot images that typically require less than 30 percent additional capacity than full-volume replicas. Inefficiency remains an ongoing challenge as organizations often overallocate storage to applications in order to reduce the risk of outage and reduce the need to reprovision later on.
- *Maximizing the performance of storage assets*
Improving performance helps meet increasingly rigorous business requirements. Performance can be improved on multiple dimensions, including creating logical storage devices, mapping devices to ports, and replicating devices to local and remote locations. In each of these areas, EMC continues to

streamline and accelerate processes within the Enginuity™ microcode. Organizations remain focused on improving performance and reducing application response times.

Benefits

When used appropriately with the right applications and workloads, Symmetrix Virtual Provisioning can complement existing measures of addressing each of the challenges outlined in the previous section.

- *Ease and speed of provisioning*

Virtual Provisioning allows storage to be provisioned independently of the physical storage infrastructure. By creating a thin device that initially is larger than required by the application, organizations can reduce the need to re-provision new storage later on, as underlying physical storage consumption increases automatically as needed. These benefits are achievable regardless of application and workload type, but will need to align with organizations' existing planning and management practices. Without the need to re-provision additional storage frequently to accommodate growth, there is also no need for any associated application outages, thereby increasing the application's availability. Virtual Provisioning also has greatly accelerated the provisioning process. Creating and configuring 1 TB of storage is now more than six times as fast as one year ago (with Enginuity 5771), and more than twice as fast as with HDS USP V.
- *Improved capacity utilization*

When used with the appropriate applications, Virtual Provisioning can reduce the amount of unused physical storage by having multiple thin devices share a single thin pool, drawing physical storage from it only as needed. Storage administrators can avoid pre-allocating physical storage to applications, which helps reduce storage costs and energy consumption.
- *Improved performance*

Wide striping across the data device storage pool distributes I/O across more physical spindles than is the case with many implementations of standard provisioning. This can reduce disk contention and enhance application performance for certain workloads (as outlined in the "Performance" section).

Application capacity growth: Virtual Provisioning vs. standard provisioning

The following example compares setup and growth phases for two similar applications. The first application uses standard provisioning and the second uses Virtual Provisioning.

The first phase is setup. With standard provisioning, the user must create the device, map, mask, and perform host discovery and LVM. Virtual Provisioning requires a similar number of steps – the user creates data devices, adds them to the thin pool, then creates the thin device and binds it to the pool. As with standard provisioning, the user then maps and masks the thin device, and performs host discovery and LVM. With standard provisioning, the host-perceived capacity equals the allocated capacity. With Virtual Provisioning, a large device can be perceived, but capacity is allocated only as needed.

The second phase is adding capacity to accommodate growth. And Virtual Provisioning requires far fewer steps, making it simpler and easier to grow. While standard provisioning requires the same steps to be repeated – that is, creating a new device and presenting it to the host, and so on – Virtual Provisioning only requires that data devices be created and added to the pool, as shown in Figure 2. The thin device and its relationship to the application remains the same (provided it was made large enough at the outset):

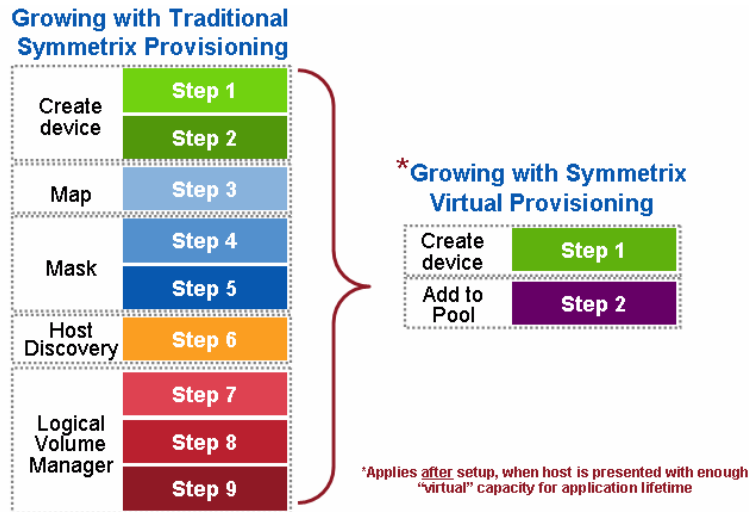


Figure 2. Required steps to add capacity with standard provisioning and Virtual Provisioning after initial setup

Users can further simplify the growth process by creating command-line scripts that automatically add data devices to a thin pool in response to alerts that are emitted when pool capacity exceeds a percentage threshold. In either case users should closely monitor the space consumption of the thin device and thin pool.

Use cases

Virtual Provisioning is most appropriate for applications that have the following characteristics.

- Predictable and controllable growth*

Applications with well-understood capacity growth requirements will help organizations avoid unexpected consumption of all space in the pool, and enable accurate initial provisioning of thin devices.
- Virtual Provisioning “friendly” environments*

If the organization seeks to improve capacity utilization with file systems, they should apply Virtual Provisioning only when files are not added or deleted frequently. Many file systems do not reuse the space associated with deleted files, resulting in used space and reducing or eliminating the capacity utilization benefits of Virtual Provisioning.

Organizations also should carefully weigh the space consumption characteristics of databases before using Virtual Provisioning in order to improve capacity utilization. Some databases preallocate space and write zeros to it. This preallocated, but unused, space cannot be shared in a thin pool, reducing or eliminating the capacity utilization benefits.

Independent of these factors, Virtual Provisioning still can improve ease of use and improve performance via wide striping in some cases. But capacity utilization will depend on the allocation methods used by specific file systems and databases as outlined previously.
- Flash drives*

Flash drives can offer both unparalleled response times as well as significant IOs/GB. Unlike traditional drives, throughput and response time do not steadily decline as utilization of the drive increases. Virtual Provisioning is therefore a natural fit to ensure that premium Flash drive resources are highly utilized.
- “General purpose” performance requirements*

Virtual Provisioning is appropriate for applications for which some performance variability can be tolerated. Some workloads will see performance improvements from wide striping with Virtual

Provisioning, as outlined in the “Performance” section. However, when multiple thin devices contend for shared spindle resources in a given pool, and when utilization reaches higher levels, the performance for a given application and thin device can become more variable.

- *Environments that need flexible provisioning (that is, test and development)*
Virtual Provisioning can be an effective means of improving ease of use and capacity utilization for lower storage tiers such as test and development.
- *Document repositories*
Document repositories can be a strong fit for Virtual Provisioning, provided the environments meet the performance, growth, and other criteria outlined previously, as rapidly rising capacity requirements for these repositories create the opportunity for improved capacity utilization.
- *Software development/source code*
Many organizations will see an opportunity to lower TCO by improving ease of use and capacity utilization for storage associated with software development activities, which often can tolerate some level of performance variability.

Virtual Provisioning on Symmetrix: Implementation overview

Symmetrix Virtual Provisioning uses logical devices called *thin devices* that can be used in many of the same ways that Symmetrix devices have traditionally been used. Unlike traditional Symmetrix devices, thin devices do not need to have physical storage completely allocated at the time the device is created and presented to a host. The physical storage that is used to supply disk space to thin devices comes from a shared *thin storage pool*, which is comprised of devices called *data devices* that provide the actual physical storage in the pool.

When a write is performed to a portion of the thin device, the Symmetrix allocates a minimum allotment of physical storage from the pool.

When a read is performed on a thin device, the data being read is retrieved from the appropriate data device in the storage pool to which the thin device is bound. When more storage is required to service existing or future thin devices, data devices can be added to existing thin storage pools. New thin devices can also be created and associated with existing thin pools.




Figure 3. Thin device and thin storage pool containing data devices

In Figure 3, as host writes to a thin device are serviced by the Symmetrix array, storage is allocated to the thin device from the data devices in the associated storage pool. The storage is allocated from the pool sequentially, striping the data across all data devices in the pool. Pool storage is allocated sequentially, regardless of whether data is written sequentially or randomly to the thin device itself.




Figure 4. Thin device and thin storage pool with added data devices

In Figure 4, as the free space in the thin storage pool becomes exhausted, new data devices can be added to the thin storage pool. The storage allocation striping continues from existing data devices onto the new data devices added to the storage pool.

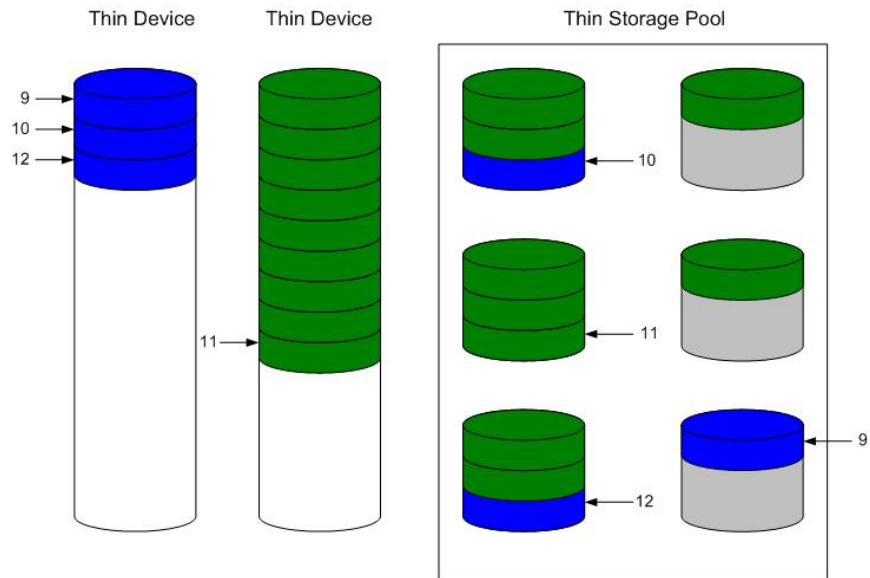


Figure 5. Additional thin device added

Figure 5 shows that new thin devices can also be added to the thin pool. Each thin device still allocates data sequentially across the storage pool, although as multiple thin devices allocate simultaneously from the pool the net effect is that data placement becomes more random.

Product features

Symmetrix Virtual Provisioning will operate on any new or existing Symmetrix DMX-3 or DMX-4 system and will integrate easily into organizations' existing knowledge and management processes for the Symmetrix platform. Configuring, replicating, managing, and monitoring thin devices and thin pools involves the same tools and the same or very similar functions as those used for years with standard provisioning. Symmetrix Virtual Provisioning will operate on typical Fibre Channel or new Flash drives, and thin pools can coexist with standard storage in the same system.

Configuration

Users have the ability to create thin devices and data devices, and to manage multiple thin storage pools within a DMX-3 or DMX-4 system. Thin devices are configured in a similar manner as standard devices using EMC Solutions Enabler or the Symmetrix Management Console, for example, using the same drop-down menus and dialog boxes in SMC. Users first create the thin storage pool, then create data devices (with specified protection schemes such as RAID 1 mirrored, or RAID 6) and add them to the pool. Once the pool and data devices are created, thin devices can be created and bound to the pool.

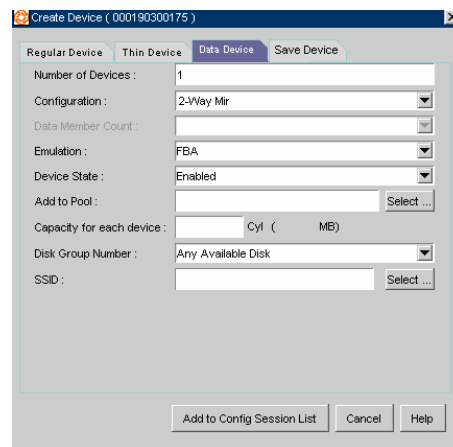


Figure 6. SMC interface for creating data devices and adding them to a thin storage pool

The thin device appears to hosts and applications as a standard device, and is mapped and masked like a standard device. Thin devices also can be used to create metadevices, provided that all members of the metadevice are thin devices. There is no mixing of thin and normal devices in the same metadevice. When a thin device is unbound from a pool or deleted, the tracks previously allocated to it on data devices in the storage pool are marked as unallocated. This means that when a subsequent thin device is bound to the data device, any reads will return zeros, preventing data from a previous thin device from being returned to the host.

Users can add data devices to storage pools to grow capacity nondisruptively. The ability to create and manage multiple thin pools allows workloads to be isolated from one another. This can reduce the risk of workloads affecting each other's performance or availability of storage space, and fits with strategies to tier storage in the array (see the "Tiered storage" section).

Replication

Organizations will be able to perform "thin to thin" replication with Symmetrix thin devices by using standard TimeFinder, SRDF[®], and Open Replicator operations. This includes TimeFinder/Snap, TimeFinder/Clone, SRDF/Synchronous, and SRDF/Adaptive Copy. In addition, thin devices can be used as control devices for hot and cold pull and cold push Open Replicator copy operations. When pulling data to a thin control device, all data is copied from the remote device, leaving the target thin device fully allocated physically.

Monitoring

Monitoring the space consumption of thin devices and thin pools is easy with existing tools. Typical monitoring includes tracking the thin device utilization over time relative to what was provisioned to the host. Other abilities include monitoring the utilization of each shared pool (% full) used by Virtual Provisioning. Monitoring of storage pool utilization can be performed using Solutions Enabler SYMCLI or SMC. Threshold-based alerts – for example, XX% of physical capacity for a pool – are available through either SMC or the SYMCLI.

Tiered storage

Symmetrix Virtual Provisioning adds a new dimension to tiered storage in the array, most notably (when used appropriately) by improving the utilization of storage tiers that can tolerate performance variability. Thin devices and storage pools can be used in a tiered storage configuration in the same manner as standard devices. Thin devices will adhere to priorities assigned to them by Symmetrix Priority Controls. Data devices within a thin pool, meanwhile, can be assigned to cache partitions created by Dynamic Cache Partitioning. In addition, different RAID groups can be assigned to different thin pools, in accordance with their designated tier.

Flash drives

EMC has now integrated enterprise-class Flash drives directly into the Symmetrix DMX-4 storage array, the first and only enterprise array with support for this emerging generation of drive technology. With this capability, EMC creates a new “Tier 0” ultra-performance storage tier that exceeds the limitations previously imposed by magnetic disk drives. Virtual Provisioning provides the same ease of use and capacity utilization benefits on Flash drives as on standard Fibre Channel drives.

Performance

The performance implications of the use of thin devices depend on the nature of the workload and the state of the thin device. In any thin device implementation there are response time and throughput overheads that are incurred the first time when a write is performed on an unallocated region of a thin device. In the Symmetrix implementation these overheads are quite modest. The overhead applies primarily to the first write to a thin device extent, so it will tend to disappear altogether once the “working set” of a thin device has been written to.

As already discussed, the back-end layout of the storage underlying each thin device will tend to be widely striped, typically spanning a much greater number of drives than a regular device. This will tend to make it easier to keep the workload balanced across the back end, which can improve performance by reducing the frequency of hot spots.

Most performance improvements will be seen by random read and write workloads that benefit from wide data striping across many data devices. Sequential workloads, which do not benefit from wide striping, in many cases will see roughly comparable performance when compared to standard devices, although performance will vary by workload and configuration.

Implementation considerations

When implementing Virtual Provisioning it is important that realistic utilization objectives are set. Generally, organizations should target no higher than 60 percent to 80 percent capacity utilization per pool. A buffer should be provided for unexpected growth or a “runaway” application that consumes more physical capacity than was originally planned for. There should be sufficient free space in the storage pool equal to the capacity of the largest unallocated thin device. Organizations also should balance growth against storage acquisition and installation timeframes. It is recommended that the storage pool be expanded before the last 20 percent of the storage pool is utilized to allow for adequate striping across the existing data devices and the newly added data devices in the storage pool.

Conclusion

It is critical to define the processes, applications, and workloads for which Virtual Provisioning can be used most effectively, as well as the specific benefits that can be achieved. When implemented appropriately, Virtual Provisioning can be a powerful complement to organizations' processes and technologies for improving ease of use, enhancing performance, and utilizing storage capacity more efficiently. Symmetrix Virtual Provisioning integrates well with existing management and business continuity technologies, and is an important advancement in capabilities for DMX-3 and DMX-4 customers.