White Paper

EMC CLARiiON Virtual Provisioning

Applied Technology

Abstract

This white paper discusses the benefits provided by Virtual Provisioning. It describes how EMC® CLARiiON® implements thin provisioning using thin pools and thin LUNs, and how storage space is allocated and monitored. The paper also discusses which application types are most suitable for thin LUNs.

December 2008

Table of Contents

Executive summary

Virtual Provisioning enables organizations to reduce storage costs by increasing capacity utilization, simplifying storage management, and reducing application downtime. Virtual Provisioning also helps companies reduce power and cooling requirements and delay capital expenditures.

The EMC® CLARiiON® CX4 series introduces thin LUN technology that builds on CLARiiON's virtual LUN capabilities and seamlessly integrates with existing CLARiiON management and replication software. With CLARiiON Virtual Provisioning, you can choose between traditional LUNs, metaLUNs, and thin LUNs. The ability to nondisruptively migrate data to different LUN and disk types allows you to deploy the best solution for your changing application and business requirements without incurring downtime.

CLARiiON's thin LUN technology also supports features such as hot sparing, proactive sparing, and the ability to migrate data between thin LUNs, traditional LUNs, and metaLUNs without incurring application downtime. This flexibility separates CLARiiON Virtual Provisioning from typical thin provisioning implementations.

Introduction

One of the biggest challenges facing storage administrators is balancing how much storage space will be required by the various applications in their data centers. Administrators are typically forced to allocate space based on anticipated storage growth. They do this to reduce the management expense and application downtime required to add storage later on. This generally results in the overprovisioning of storage capacity. Overprovisioning leads to higher costs; increased power, cooling, and floor space requirements; and lower capacity utilization rates. Even with careful planning, it may be necessary to provision additional storage in the future. This may require application downtime depending on the operating systems involved.

To address these concerns, the CLARiiON CX4 series introduces thin LUN technology. This technology works with CLARiiON's metaLUNs and traditional LUNs to provide powerful, cost-effective, flexible solutions. CLARiiON thin LUNs present more storage to an application than is physically available. Storage managers are freed from the time-consuming administrative work of deciding how to allocate disk drive capacity. Instead, an array-based mapping service builds and maintains all of the storage structures based on a few high-level user inputs. Disk drives are grouped into storage pools that form the basis for provisioning actions. Physical storage is automatically allocated only when writing new data blocks.

Thin provisioning improves storage capacity utilization and simplifies storage management by presenting an application with sufficient capacity for an extended period of time. When additional physical storage space is required, disk drives can be nondisruptively added to the central storage pool. This reduces the time and effort required to provision additional storage, and avoids provisioning storage that may not be needed.

CLARiiON Virtual Provisioning is a separately priced option that requires a software enabler. It is available for CX4 storage systems running FLARE® release 28.5 and later. You can manage thin LUNs using the same Navisphere® Manager GUI and Secure CLI commands that you use to manage traditional LUNs and metaLUNs. CLARiiON replication products, Navisphere Analyzer, and Navisphere Quality of Service Manager work seamlessly across thin LUNs, traditional LUNs, and metaLUNs. EMC RecoverPoint also supports the replication of CLARiiON thin LUNs.

This white paper discusses:

- How thin LUNs deliver space efficiency and make it easier to provision storage
- How CLARiiON implements provisioning for thin LUNs
- How thin LUNs and thin pools compare with traditional LUNs and RAID groups
- The types of file systems and applications that are best suited for thin provisioning

Audience

This white paper is intended for IT planners, storage architects, administrators, and others involved in evaluating, managing, operating, or designing CLARiiON storage systems.

Terminology

The following terminology appears in this white paper:

% Full — The percentage of pool capacity that is currently consumed. It is calculated using this formula:
% Full = Consumed capacity / user capacity

% Full Threshold — A parameter that is set by the user. The system generates an alert when this threshold is exceeded.

Allocated capacity — See "Consumed capacity."

Available capacity — The amount of actual physical thin pool space that is currently *not* allocated for thin LUNs.

Consumed capacity — For a thin pool this is the space currently used by all thin LUNs in the pool. For a thin LUN, this is the physical space used by the LUN.

LUN migration — A CLARiiON feature that dynamically migrates data to another LUN or metaLUN without disrupting running applications. LUN migration has many uses, for example:

- To change the type of disk drive (for example, from economical SATA to faster FC, or vice versa)
- To select a RAID type that better matches the current performance or availability requirements
- To re-create a LUN with more disk space

After migration is complete, the destination LUN assumes the identity (World Wide Name and other IDs) of the source LUN, and the source LUN is destroyed. LUN migration can be tuned by setting it to one of four speed settings: Low, Medium, High, or ASAP.

MetaLUN — A collection of traditional LUNs that are striped or concatenated together and presented to a host as a single LUN. A single metaLUN can be striped across any number of disk spindles, providing a much wider range of performance and configuration options. Additional LUNs can be added dynamically, allowing metaLUNs to be expanded on the fly.

Oversubscribed capacity — The amount of user capacity configured for thin LUNs that exceeds the actual user capacity in a thin pool.

RAID group — A type of storage pool that contains a set of disks on which traditional LUNs and metaLUNs can be created.

Storage Pool — A general term used to describe RAID groups and thin pools. In the Navisphere Manager GUI, the storage pool node contains RAID groups and thin pool nodes.

Subscribed capacity — The total amount of capacity configured for thin LUNs in the pool. This number can be greater than the available user capacity. The user capacity can be expanded by adding drives to the pool.

Thin LUN — A logical unit of storage where physical space allocated on the storage system may be less than the user capacity seen by the host server.

Thin pool — A group of disk drives used specifically by thin LUNs. There may be zero or more thin pools on a system. Disks may be a member of no more than one thin pool. Disks that are in a thin pool cannot also be in a RAID group.

Threshold alert — An alert issued when the % Full Threshold has been exceeded.

Traditional LUN — A logical unit of storage that can span a number of disks on a storage system but looks like a single disk or partition to the server. The amount of physical space allocated is the same as the user capacity seen by the host server.

User capacity — For a thin pool, this is the total capacity available to all hosts using the thin pool. This is also referred to as "usable capacity." It is measured as raw disk capacity minus overhead (RAID overhead and mapping overhead). For a thin LUN, this is the size of the thin LUN as it appears to the host. (This term also applies to traditional LUNs, where user capacity equals consumed capacity.)

Business requirements

Organizations, both large and small, need to reduce the cost of managing their storage infrastructure while meeting rigorous service level requirements and accommodating explosive storage capacity growth.

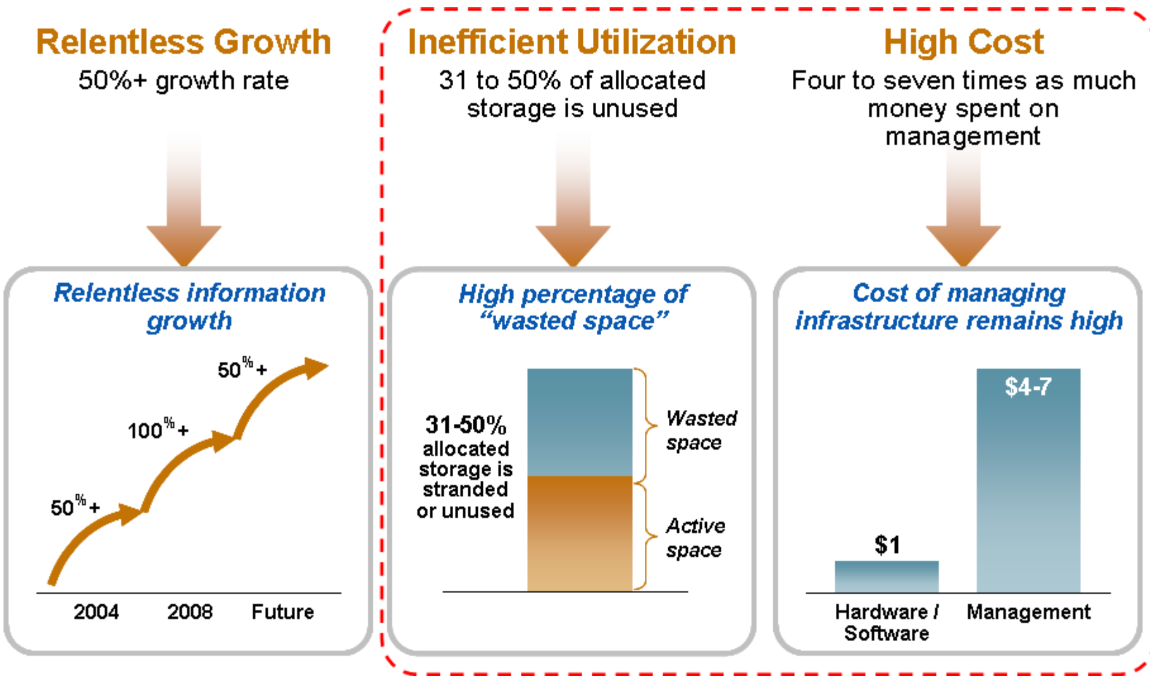

Figure 1. Information management challenges

Several business objectives have drawn increasing focus:

- *Reducing the cost of storage administration*

  "Ease of use" initiatives span multiple aspects of storage processes, including staff training, initial storage provisioning, the addition of new storage, and the management and monitoring of storage systems. Virtual Provisioning reduces the ongoing staff time required to repeatedly add storage capacity.

- *Maximizing the utilization of storage assets*

  Organizations need to accommodate growth by drawing more value from the same or fewer storage resources. Operational efficiency remains an ongoing challenge, as organizations often overallocate storage to applications to reduce the risk of outage and the need to reprovision later on.
- *Reducing capital expenditures and ongoing costs*

  Virtual Provisioning reduces capital costs by delivering storage capacity on-demand. Ongoing costs are reduced because fewer disks consume less power and cooling, and less floor space.

Comparing traditional LUNs and thin LUNs

Storage provisioning is the process of assigning storage resources to meet the capacity, availability, and performance needs of applications. With traditional storage provisioning, the host reported capacity is equal to the amount of physical storage capacity allocated. The entire amount of physical storage capacity must be present on day one, resulting in low levels of utilization. Recovering underutilized space remains a challenge. Figure 2 and Figure 3 show the differences between traditional and thin provisioning.

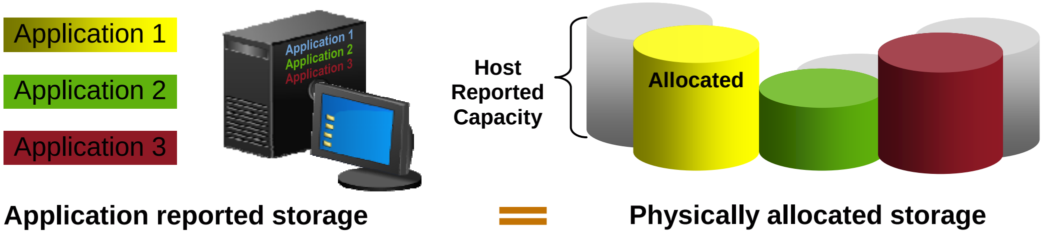

Figure 2. Traditional storage provisioning

With traditional provisioning, the storage administrator works with business people who forecast the amount of storage they will require. There is a tendency for these forecasts to be inflated. In some companies, an application administrator may monitor storage space and ask the storage administrator to provision additional storage. The storage administrator must rely on timely and accurate communications from various applications people to effectively manage storage space utilization.

With thin provisioning, the *host reported capacity* (storage perceived by the application) is *larger* than the actual allocated space on the storage system. This simplifies the creation and allocation of storage capacity. Thin LUNs can be sized to accommodate growth without regard for currently available assets. Physical storage is assigned to the server in a capacity-on-demand fashion from a shared storage pool. The storage administrator monitors and replenishes each storage pool, not each LUN.

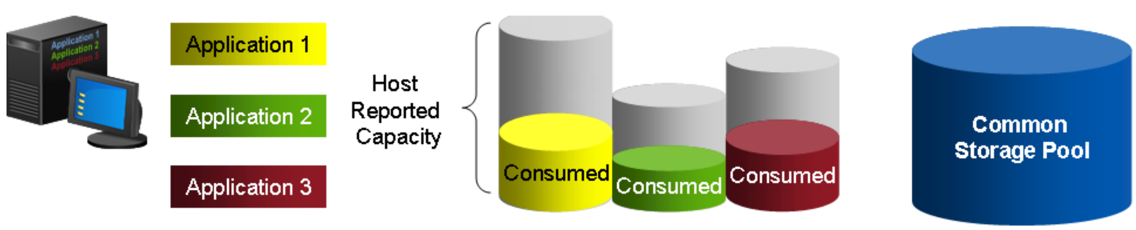

Figure 3. Thin provisioning

In either case, storage utilization must be carefully monitored. CLARiiON provides the storage administrator with specific monitoring and alerting capabilities to help manage these risks.

Thin pools

A thin pool is a collection of disks that are dedicated for use by thin LUNs. A thin pool is somewhat analogous to a CLARiiON RAID group. Thin pools can contain a few disks or hundreds of disks, whereas RAID groups are limited to 16 disks. Thin pools are simple to create because they require only three user inputs:

- Pool Name: For example, "Test & Dev Pool 2"
- Resources: Number of disks
- Protection level: RAID 5 or 6

Thin pools are flexible. They can consist of any supported Fibre Channel or SATA disk drive. A CLARiiON CX4 can contain one or many thin pools. RAID 5 is the default selection for thin pools and is recommended for most situations. The smallest pool size is three drives for RAID 5 and four drives for RAID 6. However, recommended pools sizes are multiples of five for RAID 5 and multiples of eight for RAID 6. Actual limits such as disks per thin pool and thin pools per system are listed in "Appendix A: Virtual Provisioning limits" on page 18. Also refer to the *EMC Virtual Provisioning Release Notes* for current information.

For very large pools of SATA drives, RAID 6 may be a better choice. Thin pools are easy to expand by adding drives to the pool. EMC recommends adding disks in multiples of five drives for expanding RAID 5 thin pools and multiples of eight drives for expanding RAID 6 thin pools

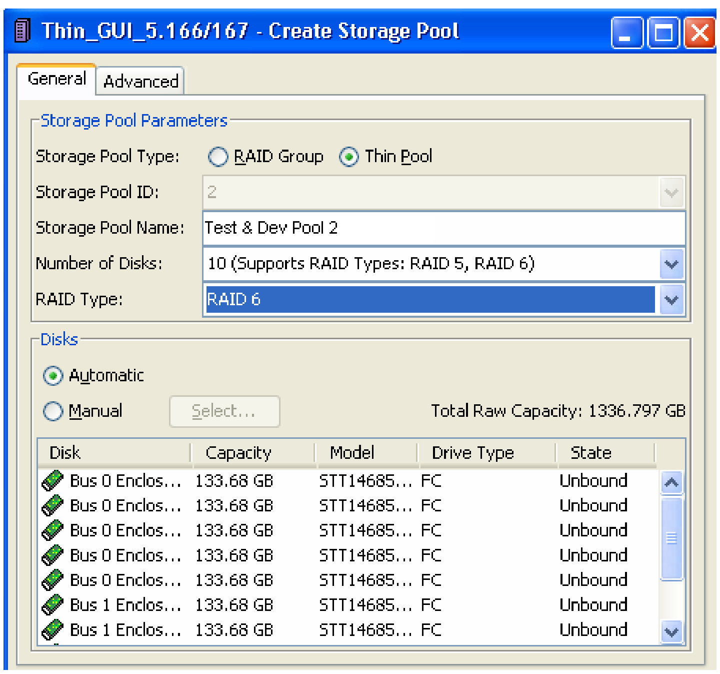

Figure 4. Create Storage Pool dialog box

There are two restrictions on what type of disks can be used in a thin pool:

- Vault drives (the first five drives in a storage system) cannot be part of a thin pool. Navisphere dialog boxes and wizards do not allow you to select these drives.
- Enterprise Flash drives (EFD) cannot be used in a thin pool.

Although it is not a restriction, it is a best practice to limit each thin pool to one type of disk drive. Fibre Channel and SATA drives should be deployed in separate pools. Drives within a thin pool should be rated at the same rpm. These recommendations are similar to recommendations for traditional RAID groups.

Disks in a thin pool should be the same size to maximize space utilization. When using different capacity drives to create a thin pool, it is best to do this in stages. For example, if you have ten 450 GB drives and five 300 GB drives, first create the pool selecting only the ten 450 GB drives, and then expand the pool by adding the other five 300 GB drives.

Monitoring, adding, and deleting thin pool capacity

User capacity is the total physical capacity available to all thin LUNs in the pool. *Consumed capacity* is the total physical capacity currently assigned to all thin LUNs. *Subscribed capacity* is the total host reported capacity supported by the pool.

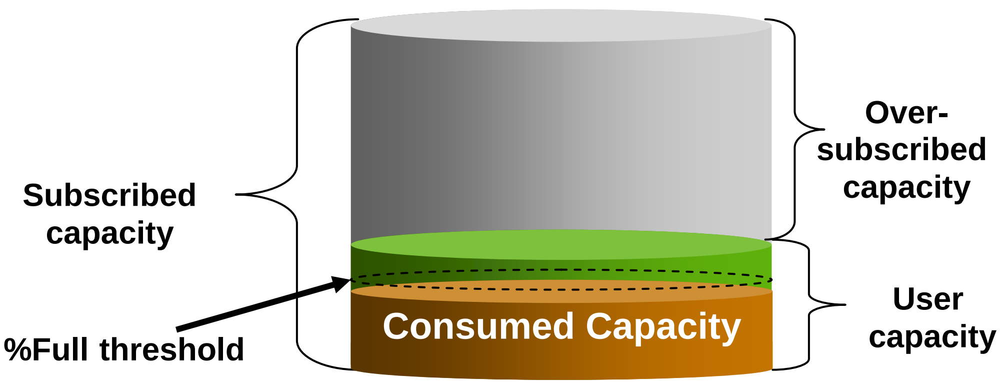

Figure 5. Pool % full threshold

Thin pools are monitored using the **% Full Threshold Alert** setting and storage capacity reports. You can specify the value for **% Full Threshold** (Consumed capacity/User capacity) when a thin pool is created. It can be changed at any time. It ranges from 1% to 84%. When **% Full Threshold** is exceeded, an alert is triggered at the *warning* severity level. When the **%Full** reaches 85%, the pool issues a built-in alert at the *critical* severity level. Beyond 85%, the built-in alert mechanism continues to track the actual **%Full** value as the pool continues to fill.

Alerts can be monitored via the **Alerts** tab in Navisphere. In Navisphere's Event Monitor wizard, you can also select the option of receiving alerts via e-mail, a paging service, or an SNMP trap. Table 1 lists the thresholds for different types of alerts.

Table 1. Threshold alerts

Threshold Type	Threshold Range	Threshold Default	Alert Severity	Side-effect
User settable	1%-84%	70%	Warning	None
Built in	n/a	85%	Critical	Clears user settable alert

Allowing consumed capacity to exceed 90% of user capacity puts you at the risk of running out of space and impacting all applications using LUNs in the thin pool.

The **Thin Pool Properties** dialog box (Figure 6) shows parameters such as subscribed capacity, over-subscribed capacity, and percent subscribed. Oversubscribing capacity is the key to reducing disk storage costs. There is no hard and fast rule for how much or how little one should oversubscribe. However, you should never subscribe storage beyond the maximum capacity of the thin pool[1].

[1] The limits for CLARiiON Virtual Provisioning can be found in the Appendix section.

Adding drives to the pool non-disruptively increases user capacity for all thin LUNs in the pool. Consumed (allocated) capacity is reclaimed by the pool when LUNs are deleted.

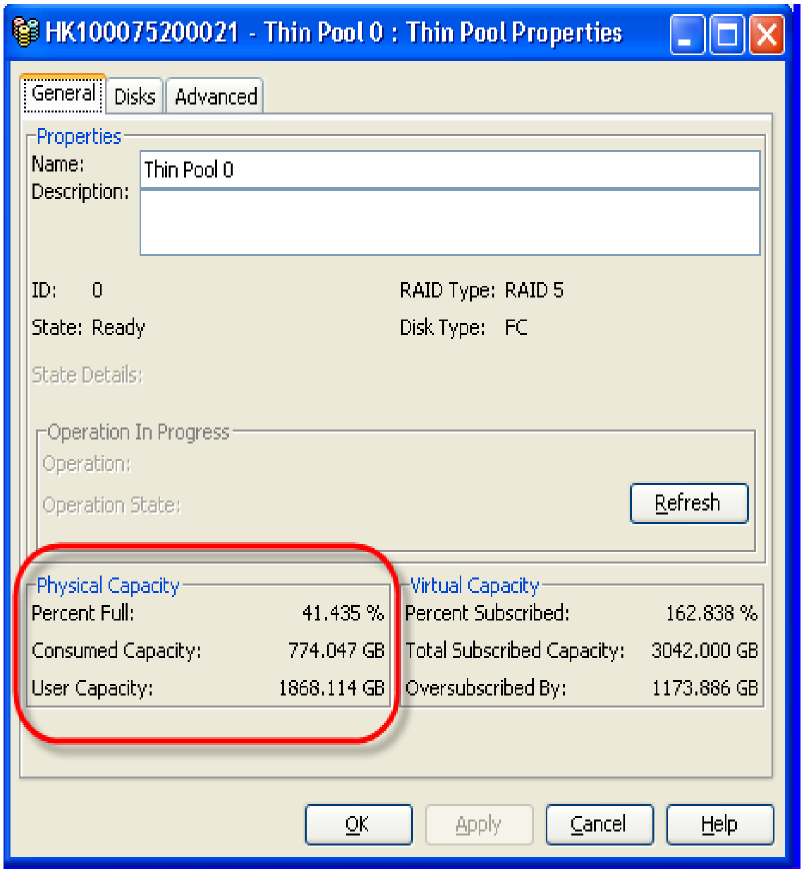

Figure 6. Thin Pool Properties dialog box

Thin LUNs

A CLARiiON thin LUN is similar to a traditional LUN in many ways. CLARiiON customers use many of the same Navisphere Manager GUI operations and Secure CLI commands on thin LUNs and traditional LUNs. Most user-oriented FLARE functions work the same, including underlying data integrity features, hot sparing, LUN migration, local replication, and LUN properties information. Thin LUNs are available as choices in Navisphere Taskbar wizards, Navisphere Analyzer, and Navisphere Quality of Service Manager. Features such as hot sparing and proactive sparing operate in the same manner. It is also possible to migrate a traditional LUN (or metaLUN) to a thin LUN and vice versa. This flexibility separates CLARiiON thin provisioning from typical thin provisioning implementations.

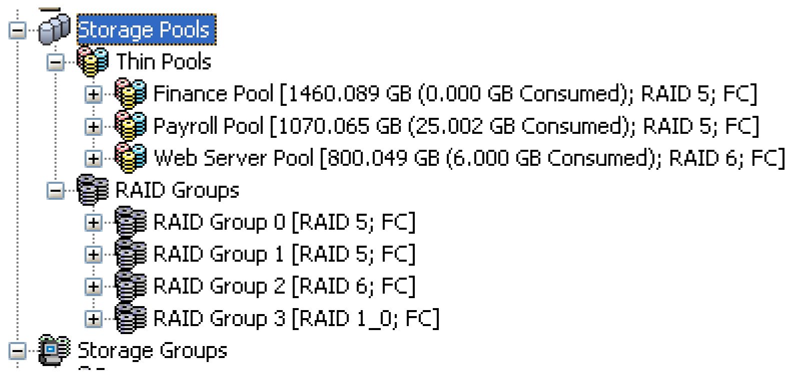

Figure 7. Storage pools in Navisphere Manager

Attributes

The primary difference between traditional and thin LUNs is that thin LUNs can consume less physical space on the storage system. Thin LUNs are simple to create, with three inputs:

- Thin pool name
- Host to provision
- Amount of host visible user capacity

Thin LUNs are easy to use because the system automatically manages the drives within the pool according to CLARiiON best practices. The Mapping Service distributes data evenly throughout the pool. Thin LUNs are easy to manage because Navisphere Manager and Navisphere CLI commands work the same for thin LUNs as they do for traditional LUNs. New property screens and reports are available to obtain information about thin LUNs.

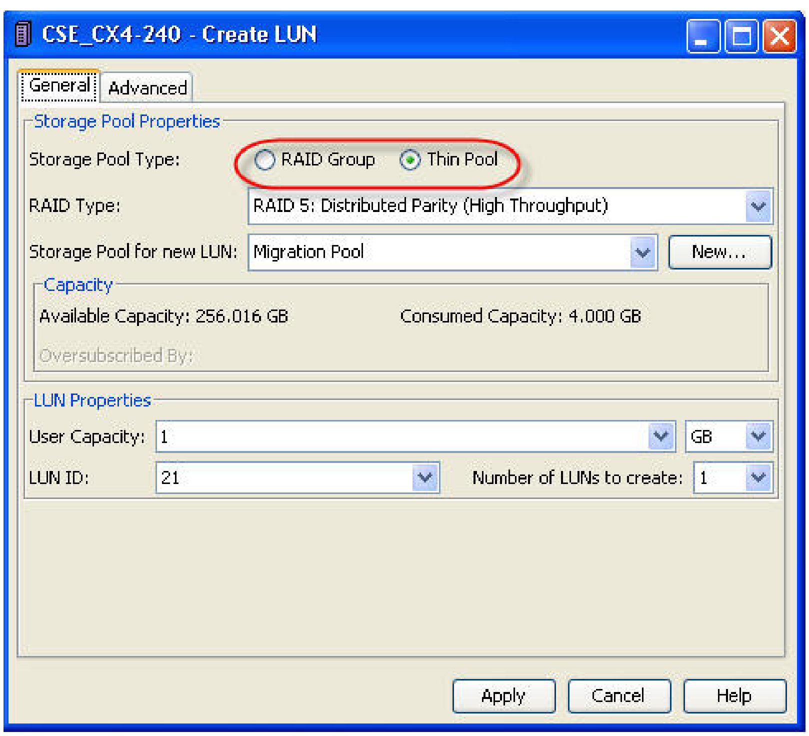

Figure 8. Create LUN dialog box

It is best to provision thin LUNs only to applications that can tolerate some variation in performance. Although thin LUNs are internally striped, performance may vary. Applications whose LUNs occupy the same spindles within a thin pool will impact one another. Use traditional LUNs for applications with stringent performance requirements. Use metaLUNs for high-throughput applications that can benefit from spreading I/O over a large number of disk drives striped over multiple RAID groups.

There are only a few restrictions with thin LUNs, and they are enforced through Navisphere:
- A thin LUN cannot be used in the reserved LUN pool (which reserves space for snapshots).
- A thin LUN cannot be used for a clone private LUN (which reserves space for clones).
- A thin LUN cannot be used as a component of a metaLUN.

Architecture and features

Specialized software known as the Mapping Service manages the placement and use of the physical storage used by thin LUNs. Data is written to 8K chunks (*extents*) and is optimally placed. This makes configuring thin LUNs easy, because the underlying software makes all the decisions about how to lay out actual storage blocks across the disks in a thin pool. Less experienced storage administrators will benefit from not having to be directly involved in the details of configuring storage. The Mapping Service performs these functions adhering to performance best practices.

When creating a thin LUN, the initial consumed space from the thin pool will be shown as 2 GB. Slightly more than 1 GB is used for metadata associated with the thin LUN, and the remaining space is available for incoming writes. As new writes come in, more physical space is allocated in 1 GB stripes. This is done in a storage-on-demand fashion. The mapping service creates embedded metadata within these additional 1 GB slices.

Using thin LUNs with applications

Thin LUN provisioning is most appropriate for applications that have the following characteristics:

- Thin Provisioning “friendly” environments

  To improve capacity utilization with file systems, use thin provisioning only when files are *not* frequently deleted. Many file systems do not efficiently reuse the space associated with deleted files, which reduces the capacity utilization benefits of Thin Provisioning.

  You should carefully weigh the space consumption characteristics of databases before using thin provisioning,. 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.

- *“General purpose” performance requirements*

  Thin provisioning is appropriate for applications in which some performance variability can be tolerated. Some workloads will see performance improvements from wide striping with thin provisioning. 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 can become more variable. Navisphere Quality of Service Manager can be used to manage resource contention within the pool as well as between LUNs in different thin pools and RAID groups.

- *Environments that need flexible provisioning (for example, test and development)*

  Thin provisioning can be an effective way to improve ease of use and capacity utilization for lower storage tiers such as test and development.

- *Document repositories*

  Document repositories with rapidly rising capacity requirements can benefit greatly from the improved capacity utilization offered by thin provisioning, provided their environments meet the previously outlined criteria.

- *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, because these development activities can usually tolerate some level of performance variability.

Oracle[2]

Traditional database files as well as Oracle ASM include an Auto Extend feature that can take advantage of thin provisioning. Without Auto Extend, using CREATE DATABASE with traditional database files would cause Oracle to writes zeros to all blocks of a tablespace file. Oracle DBAs can elect to use traditional LUNs in small RAID groups for log files and place the Oracle database volumes in a thin pool. This places log and database volumes on separate disks (a best practice) while still providing the benefits of thin provisioning for the database volumes.

Microsoft SQL Server

Microsoft SQL Server 2005 introduced new functionality that changed how databases are created. Earlier versions of SQL Server fully initialized all data file and transaction log file components by writing to every page in all database files and log files. With SQL Server 2005, the database creation phase no longer requires initializing all data files if Instant File Initialization can be utilized. EMC recommends using Microsoft SQL Server 2005 (or later) with Instant File Initialization functionality when using thin provisioning.

Microsoft Exchange

By default, new Exchange 2007 database are between 2 MB and 4 MB in size, and grow by 2 MB as additional space is needed. This auto-extend behavior is efficient from a thin pool perspective, as only space immediately needed by the database file is allocated. Exchange environments that most effectively use Virtual Provisioning have a large space allocated for mailboxes and are slow to fill that space.

NTFS

Many New Technology File System (NTFS) operations are thin provisioning-friendly. For example, if you do a NTFS quick format, NTFS does not pre-initialize the physical space. Instead, it creates file system metadata that only consumes a few gigabytes of metadata on the thin LUN. NTFS writes new data on the LUN and updates metadata accordingly.

Not all NTFS operations are thin provisioning-friendly. If you delete a file in NTFS, the file system knows it can use those blocks to store new data. The underlying CLARiiON thin LUN, however, is not aware of this. When you create a new file on the host file system, NTFS may or may not use the space freed up by a deleted file. If NTFS writes the new file in the previously freed-up area, then the thin LUN will not grow. However, if it writes in a previously unwritten area, then the thin LUN will consume more space from the pool. You are still likely to see storage space savings with NTFS when compared to traditional LUNs, as well as other benefits from thin provisioning such as ease of management.

VMware[3]

The VMware Virtual Machine File System (VMFS) has some characteristics that are thin provisioning-friendly. First, a minimal number of thin extents are allocated from the thin pool when a VMware file

---

[2] For best practices planning for database deployments, refer to the *Leveraging EMC CLARiiON CX4 Virtual Provisioning for Database Deployment* white paper available on Powerlink

[3] For detailed implementation of Virtual Provisioning with VMware, refer to the *Implementing Virtual Provisioning on EMC CLARiiON and Celerra with VMware Virtual Infrastructure* white paper available on Powerlink.

system is created on thin LUNs. VMFS does not write all of its metadata to disks on creation. The VMware file system formats and uses the reserved area for metadata as requirements arise.

When creating a VMware virtual disk, the default type is "zeroedthick". This is the recommended virtual disk type for thin LUNs. Using zeroedthick, the storage required for the virtual disk is reserved in the datastore, but the VMware kernel does not initialize all the blocks. The blocks are initialized by the guest operating system when writes to previously uninitialized blocks are performed.

The VMware kernel provides other mechanisms for creating virtual disks that are not thin provisioning - friendly. The "eagerzeroedthick" format is not recommended for use with thin LUNs because it performs a write to every block of the virtual disk at creation and results in equivalent storage use in the thin pool. The VMware virtual disk type called "thin" is not recommended because it creates a VMware thin LUN on top of an array-based thin LUN. This results in extra management complexity because oversubscription of physical storage occurs at two independent layers that do not communicate with each other.

VMware DRS and VMotion are thin-friendly. VM Clones and Templates are problematic. VM Cloning fully allocates all blocks. There is currently no workaround for this. VMware Templates also allocate all blocks. The workaround is to shrink VMDKs before creating a template and use the "Compact" option.

Traditional LUNs versus thin LUNs

The key is to understand your application requirements and select the approach that meets your needs. If conditions change, you have the option of using CLARiiON LUN migration to migrate from thin LUN to traditional LUN or vice versa.

Use RAID groups and traditional LUNs:
- When milliseconds of performance are critical
- For the best and most predictable performance
- For precise data placement
- When you want one method for all applications and services
- When you are not as concerned about space efficiency

Use thin provisioning with thin pools for:
- The best space efficiency
- Easy setup and management
- Minimal host impact
- Energy and capital savings
- Applications where space consumption is difficult to forecast

Thin pools and thin LUNs offer many benefits, but it is important to remember that they require careful monitoring and are not suited to every application.

LUN migration

Thin LUNs can be the target or source of LUN migration operations. LUNs can be moved between thin pools or between a thin pool and RAID group.

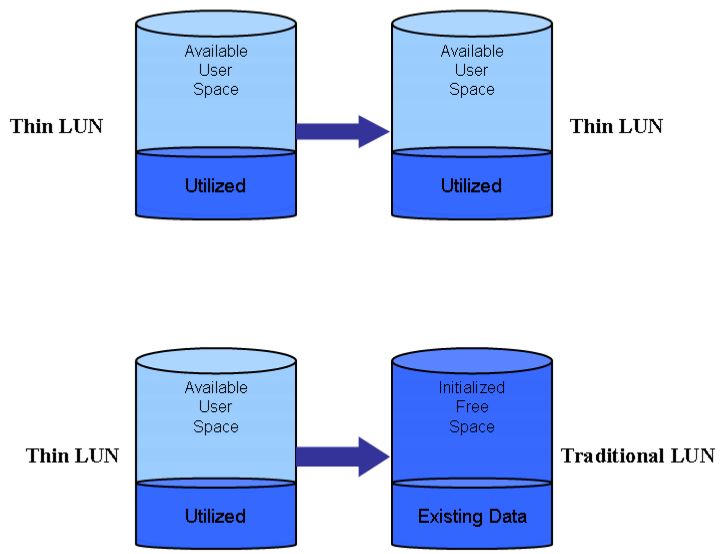

Figure 9. LUN migration examples

Migrating from a traditional LUN to a thin LUN allocates space on the thin LUN that is equal to the size of the traditional LUN, but it provides an opportunity to expand the host volume and take advantage of thin LUN characteristics. For example, consider a development database on a 100 GB traditional LUN, in which the application has consumed 80 GB of the 100 GB LUN. You can create a thin LUN with 300 GB of user capacity, and then migrate from the 100 GB traditional LUN to the 300 GB thin LUN. On Windows, the target LUN would be extended using "diskpart". This results in a thin LUN that has the same allocated capacity as the source LUN (100 GB plus a few gigabytes more for thin LUN metadata). The actual data area consumed by the application is still 80 GB, and new writes to the LUN start filling up the remaining 20 GB. Once it reaches 100 GB, it will continue to grow in a storage-on-demand basis up to the new user capacity presented to the host.

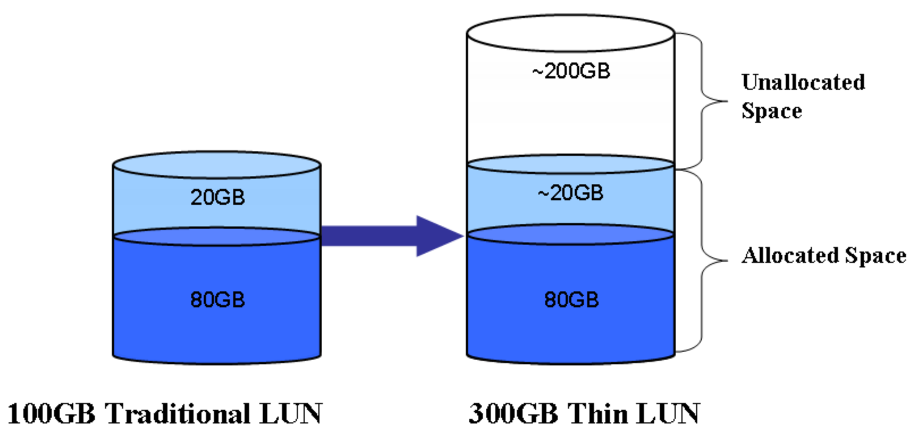

Figure 10. LUN migration example[4]

Replication

CLARiiON Virtual Provisioning supports local replication using SnapView; SnapView snapshots and clones are both supported. SnapView clones support replication from thin to thin, thin to traditional, and

[4] ~20 GB as some space will be used by thin LUN metadata

traditional to thin. When cloning from a thin LUN to a traditional LUN, the physical space of the traditional LUN must equal the subscribed capacity of the thin LUN[5]. This will result in a fully allocated thin LUN if the traditional LUN is reverse-synchronized. Cloning from traditional to thin is not recommended as the initial synchronization will force the initialization of all the subscribed capacity. Also note that thin LUNs cannot be private LUNs. This means that thin LUNs cannot be used in reserved LUNs pools for SnapView snapshots, or as clone private LUNs (CPLs) for SnapView clones.

Remote replication is supported through RecoverPoint. RecoverPoint supports replication for thin to thin LUNs and traditional to thin LUNs. Support for EMC MirrorView® and EMC SAN Copy™ will be in future releases.

Navisphere Quality of Service Manager (NQM)

Thin pools have many applications that share the same set of physical drives. NQM ensures that an application's performance is not degraded by the needs of a different application running on another LUN in the pool. NQM allows you to set and adjust storage service priorities to the different LUNs holding different sets of application data. NQM is completely supported with thin LUNs. Thin LUNs can be added to IO Classes just like any other LUN. There are no restrictions on combining thin LUNs with traditional LUNs. Both **Limit** and **Cruise Control** policies can be used to manage the storage system performance of applications mounted on thin LUNs.

Phased implementation for CLARiiON thin provisioning

CLARiiON thin provisioning will be implemented in phases. The initial phase is being released in December 2008 (FLARE release 28.5). This release includes basic thin pool and thin LUN structures as well as the ability to expand a thin pool by adding disk drives. Thin LUNs are supported for LUN migration and local replication using SnapView (snapshots and clones). Navisphere Analyzer and Navisphere Quality of Service Manager are also supported. Remote replication of thin LUNs is supported by RecoverPoint.

Later in 2009, support will be added for array-based remote replication: MirrorView/S, MirrorView/A, and SAN Copy. A later release will include the ability to shrink a thin pool by removing disk drives. The ability to create and track thin LUN space reservations will be added, as well as the ability to expand and shrink thin LUNs.

[5] The examples in Figure 9 also apply to SnapView clones

Conclusion

When implemented appropriately, thin provisioning can be a powerful complement to organizations' processes and technologies for improving ease of use, enhancing performance, and utilizing storage capacity more efficiently. CLARiiON Thin Provisioning integrates well with existing management and business continuity technologies, and is an important advancement in capabilities for CLARiiON customers. CLARiiON thin provisioning:

- Saves time:
  - Easy to create pools and LUNs
  - Easy to monitor and manage
- Reduces provisioning uncertainty:
  - Decisions are easy to modify
  - No impact on host servers
- Reduces upfront investment and saves energy:
  - Highly space-efficient
  - Multiple applications share resources
  - Physical storage can be added as required
- Builds on existing CLARiiON features:
  - Migration is supported between all types of LUNs
  - SnapView snapshots and SnapView clones
  - Navisphere Quality of Service Manager and Navisphere Analyzer
  - Navisphere Reports

References

- *EMC CLARiiON Performance and Availability: Release 28.5 Firmware Update — Applied Best Practices*
- *Implementing Virtual Provisioning on EMC CLARiiON and Celerra with VMware Virtual Infrastructure — Applied Technology*
- *Leveraging EMC CLARiiON CX4 Virtual Provisioning for Database Deployment—Best Practices Planning*

Appendix A: Virtual Provisioning limits

System	Max. thin pools	Max. disks per thin pool	Total disks in all thin pools	Max. thin LUNs per pool	Max. thin LUNs in all pools
CX4-120	10	20	40	256	256
CX4-240	20	40	80	512	512
CX4-480	40	80	160	1024	1024
CX4-960	60	120	240	2048	2048

Appendix B: Guidelines for thin pools and thin LUNs[6]

Thin pools

- Limit each thin pool to one drive type:
  - Fibre Channel and SATA drives should be deployed in separate pools.
  - Drives in a thin pool should be the same rpm speed.
  - To maximize space utilization, create or expand the thin pool with drives of the same capacity.

- Create or expand pools in increments of five drives for RAID 5 and eight drives for RAID 6.
- RAID 5 is recommended for most situations. For very large pools of SATA drives, RAID 6 may be a better choice.
- Avoid running pools at near maximum capacity.
- Ensure that a monitoring strategy is in place for all thin pools

Thin LUNs

- Provision thin LUNs for applications that can tolerate some variation in performance.
  - Although thin LUNs are internally striped, performance may vary.
  - A misbehaving application will impact other applications sharing disks in a thin pool.
- Use traditional LUNs for applications with stringent performance requirements.
- Use metaLUNs for high-throughput applications that can benefit from I/O spread over a large number of disk drives.
- Identify Virtual Provisioning "friendly" host operating systems and applications for thin LUN use:
  - File systems that efficiently reuse deleted space
  - Applications that reserve space but do not initialize all storage

[6] For continuous updates to best practices concerning Virtual Provisioning please use the CLARiiON best practices document *EMC CLARiiON Performance and Availability: Release 28.5 Firmware Update — Applied Best Practices*

Appendix C: Thin pool properties

Thin pool property	Description	User editable?	Value range
Name	Friendly name of the thin pool that is also its unique ID (must be unique across all thin pools on the array) Max length: 255	Yes	ASCII, all printable characters (0x20 to 0x7E) with no leading or trailing spaces
Description	User text describing the thin pool Max length: 255	Yes	n/a
ID	Numeric ID of the thin pool (unsigned 32-bit) Note: This will display as a property but is only used internally. The name is the ID.	No	0 to 4294967295
RAID type	RAID protection level of the thin pool	Only at thin pool creation time	5, 6
Disk type	Type of disks used by the thin pool	Disks selected at pool creation/expansion time	• FC • ATA • SATA • Mixed
State	Current state of the thin pool	No	• Initializing • Ready • Faulted • Offline • Destroying
State details	Optional additional information to describe the state of the pool	No	Status code from driver
Raw capacity	Total amount of physical disk capacity in the pool	Entered at pool creation/expansion time	Total capacity of all physical disks in the pool
User capacity	Total amount of storage capacity in the pool excluding RAID overhead	No	Raw Capacity minus RAID overhead
Allocated capacity	Amount of user capacity that is allocated to thin LUNs	No	0 to User Capacity
Available capacity	Amount of user capacity that can be allocated to thin LUNs	Indirectly by expanding pool	0 to User Capacity
Subscribed capacity	Total amount of storage capacity subscribed to by all thin LUNs in the pool. This number can be larger than the pool capacity	Indirectly by creating/destroying thin LUNs	0 to Total User Capacity of all thin LUNs in the pool
Operation in progress	Indicates that a long-running operation is in progress, and may show the percentage completed	No	• Initializing • Recovering • Destroying thin LUN

Appendix D: Thin LUN properties

Property	Description	User editable?	Value range
Name	Friendly name of the thin LUN (must be unique across all thin LUNs per array) Max Length: 255	Yes	ASCII, all printable characters (0x20 to 0x7E) with no leading or trailing spaces
ID	Numeric ID of the thin LUN	No	Shared with all LUN types
Unique ID	WWN	No	WWN
Thin pool	Friendly name of the thin pool that contains this thin LUN	No	Valid thin pool name
RAID type	RAID protection level of the thin LUN's thin pool	No	5, 6
Drive type	Type of disks used by the thin LUN's thin pool	No	FC, ATA, SATA, Mixed
State	Current state of the thin LUN	No	• Initializing • Ready • Faulted • Offline
State details	Optional additional information to describe the state of the thin LUN	No	Status code from driver
User capacity	Size of the thin LUN that is to be presented to the host	Only at LUN creation time	Depends on physical disks in pool
Consumed capacity	Amount of space of the thin LUN that has been used by user data	No	Depends on physical disks in pool and RAID type
Free available pool capacity	Amount of space in the thin pool that is available for thin LUNs to consume	No	Depends on pool
LUN type	Type of LUN	No	Thin
Auto assignment	Enable the array to auto-assign the thin LUN to an SP	Yes	Yes/No
Current owner	Current ownership of the thin LUN (same as traditional LUN)	No	SP A, SP B
Default owner	Default ownership of the thin LUN (same as traditional LUN)	Yes	SP A, SP B
Allocation owner	SP that owns the TLU allocation table	No	SP A, SP B
Alignment offset	Offset value to align host access	Only at LUN creation time	Same as other LUN types