

Thin-provisioning options for file systems

Although there are some situations where thin provisioning of NAS devices does not make sense, most data centers can benefit substantially from the technology.

By Ray Lucchesi

January 18, 2008—Thin provisioning for file systems was first introduced in 2002 and has experienced renewed interest with the explosive growth in thin-provisioned block storage subsystems. Most major NAS vendors offered this technology by 2004 to alleviate some of the problems associated with file-system proliferation, and numerous data centers quickly discovered the benefits of thin provisioning for NAS users. In fact, one NAS vendor who was interviewed claimed a 90% usage rate over its total installed base.

Thin provisioning can be likened to an empty room with a number of inflatable castles inside. The castles each start with their initial inflation amount. As demand rises they each independently inflate to needed levels. Each castle has an upper inflation limit, and the castles in total cannot inflate beyond the room. In this analogy, the castles are file systems, and the room is the available storage.

In contrast, hard provisioning is like building a city with permanent stone castles. The ultimate size of the castle is determined upon completion of the foundation. Any subsequent change to the size would require additional, possibly major, construction.

Much like the portable inflatables, thinly provisioned file systems can automatically expand up to a pre-determined maximum limit. These pre-determined limits are defined at configuration of the file system in conjunction with an initial space allocation. Users exploring the file system under Windows would see the drive letter size as the initial allocation size.

As data is written to the thinly provisioned file system and exceeds its initial space allocation, the file system automatically expands without operator or user intervention. Subsequent write operations continue seamless expansion until the maximum limit is reached or until operator intercession.

In most thin-provisioned file systems, automatic expansion does not mean automatic contraction. In fact, only a few vendors automatically contract a file system as unneeded files are deleted. Most vendors' products use the deleted file space for subsequent file writes, leaving the "expanded" file system space allocation as is and thus, not releasing space.

Most NAS vendors with thin provisioning support both CIFS and NFS. Thin provisioning does not require special host software as it is built into the standard CIFS and NFS file protocols.

Ordinarily under thin provisioning of files, file systems are offered that auto-extend. Auto-extend features are available in advanced NAS servers and from other file-system vendors when space can be either "soft" or "hard" allocated. For example, in a 20TB NAS system, file system "A" could be initially allocated 4TB of space with an additional 4TB (8TB total) provisioned. Beyond the initial 4TB, the file system grows as needed up to the maximum 8TB.

Perhaps even more significantly, in a 20TB NAS system, further file systems could be configured up to the maximum vendor-imposed limit. In fact, the cumulative maximum limits could even exceed the actual available storage. The file systems are then allowed to expand, independently and as needed, until the actual cumulative utilization reaches maximum capacity (20TB).

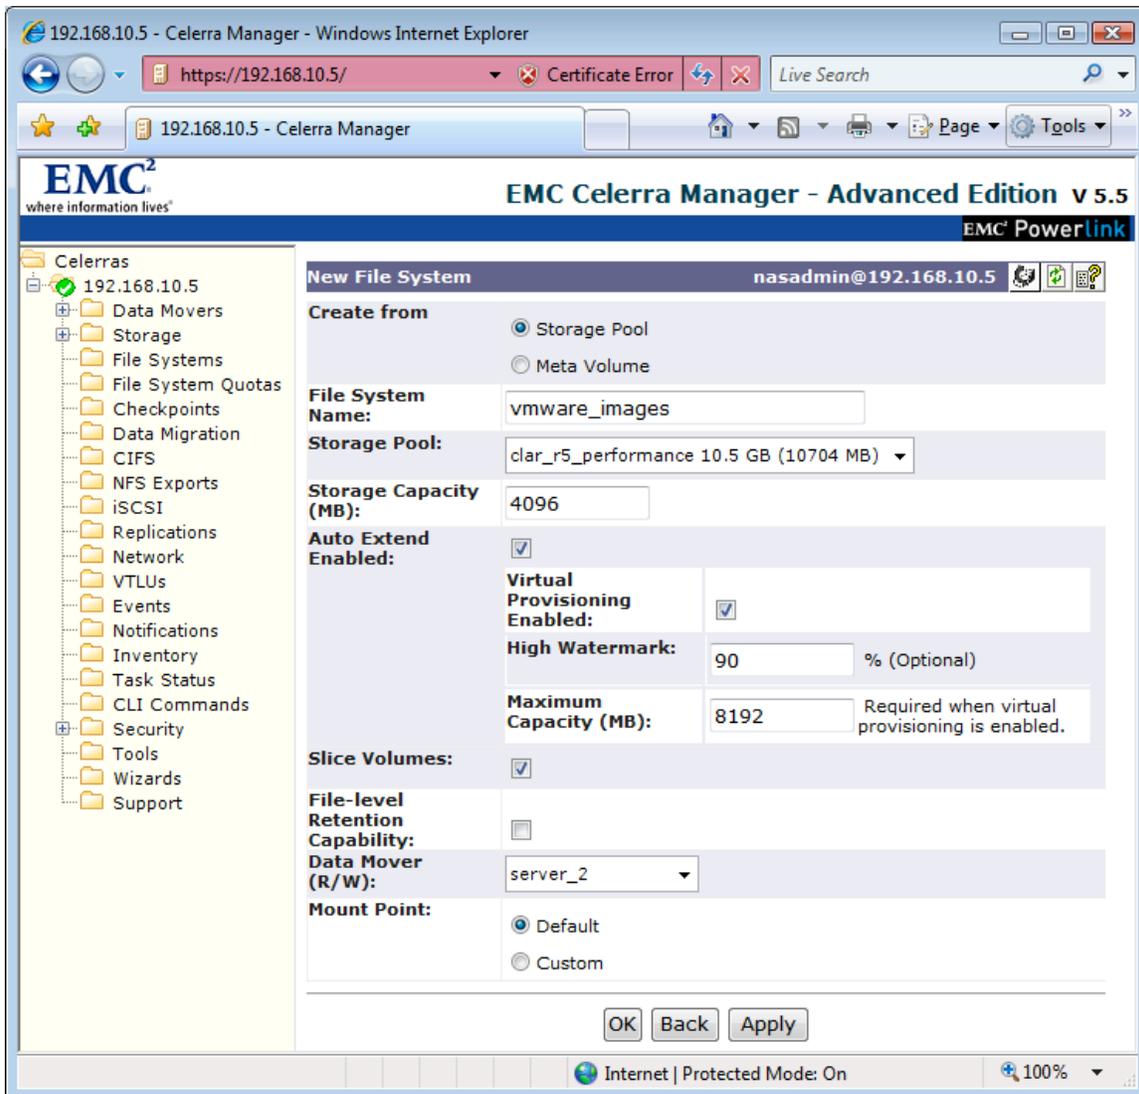
In contrast, with hard provisioning of a 20TB NAS system, a file system "B" expected to grow to 8TB over time should be hard allocated the full 8TB initially. Further adjustments of space, either up or down, would require both use/administrator intervention. In addition, only 12TB (20TB minus 8TB) of storage is available for other file systems.

Data-center benefits

The major benefit of thin provisioning of file systems is the ability to delay storage capacity capital expenditures until absolutely necessary. That is, the data center can purchase additional capacity on a just-in-time basis. In the example above, upon initial configuration the thinly provisioned NAS device had 16TB of configurable physical space versus the 12TB available in the hard-provisioned NAS device. However, even this comparison is oversimplified when one considers that a thin-provisioning NAS device can provision more maximum file system space than the actual physical storage space available.

For example, in the hypothetical 20TB thin-provisioned NAS device, four file systems could be initially allocated 4TB each. Unlike a similarly sized hard-provisioned NAS device, the file systems could each be configured to an 8TB maximum limit, potentially exceeding the actual storage capacity. In effect, thin provisioning allows each file system to grow independently with little or no operator intervention until either an individual file system reaches its previously determined maximum allocation or the total storage capacity is used.

This ability to over-provision a thin-provisioned file system provides another significant benefit: File-system space estimates can be less precise. In fact, with thin-provisioned file systems, all file systems could be configured to the maximum vendor-supported size.



Screen from EMC Celerra Manager shows thinly provisioned file system (“vmware_images”). Initial allocation is 4GB with a max limit of 8GB, with auto-extension, and virtual provisioning

Thin-provisioning concerns

Even with its potentially significant cost savings benefit, thin provisioning is not for every IT data center. Thin-provisioning file systems must be vigilantly monitored by operations staff. An inability or unwillingness to maintain this monitoring can have disastrous effects. All future write operations are terminated if the thin-provisioned file systems are allowed to consume all available physical storage space. As such, when file storage consumption reaches 50% to 80% of the available capacity more storage should be considered.

Another, less problematic, concern advanced by some vendors that do not provide thin provisioning is the overall need for multiple file systems and, conse-

quently, thin provisioning. These vendors argue that one file system and one storage pool eliminates the need for thin provisioning. While this argument may have merit, the alternative—multiple file systems—also has merit. In fact, multiple file systems can enhance some IT activities such as directory lookup, organization chargebacks, and space management.

Hybrid approaches

Some file systems support thin-provisioned block storage. Although file systems must be hard allocated to thinly provisioned LUN storage, the total physical storage actually used in these systems could be almost the same as thinly provisioned file systems. However, to retain the maximum benefit of thin provisioning (i.e., conserving actual physical storage space), the hard allocated file system must re-use deleted file space before consuming new storage space.

Is it warranted?

In a few specific instances, thinly-provisioned file systems are definitely unwarranted. These include the following:

- *IT centers without constant operations monitoring.* Because of the automatic file-system expansion of thin provisioning, data centers can be easily lulled by the smoothness of operations. If thin-provisioned file systems are allowed to consume all of the available physical storage, all write operations are stopped;
- *High-performance file systems.* In certain IT environments, some file systems are highly optimized by directing certain file-system traffic to high-performance disk space. However, thin provisioning is often not location-sensitive (i.e., all disk storage is maintained as a single pool available for file expansion). As such, the high-performance file system may quickly become less optimized as new data is routed to lower-performance storage;
- *File systems with heavy sequential write workloads.* The automatic file-expansion feature of thin provisioning does have some operational cost associated with it. Each time the file system needs expanding, thin provisioning requires some processing overhead. As such, file systems with heavy sequential-write workloads may suffer some negative throughput performance;
- *Extremely large file systems with extremely large files and/or huge data workloads.* Data centers, like large research labs or intelligence-gathering facilities, process huge amounts of homogeneous data on a daily basis. In such situations, multiple file systems would be cumbersome. For these environments, one very large file system is more optimal and thus, thin provisioning provides no advantage; and
- *Data centers charging back IT costs to other operations.* In most cases, users of thin-provisioned storage are only cognizant of actual space in use. However, many IT data centers charge back the cost of the maximum file space allocation. Because these numbers may be very different, organizational conflicts may arise. In response, some vendors have modified reporting to optionally reflect maximum rather than actual space (**see screenshot, below**).

Thin provisioning of file systems can provide an excellent alternative to today's increasingly cost-conscious data center by allowing storage purchases to be delayed and capacity to be purchased just in time. However, if the thin-provisioned file system is not vigilantly monitored, the benefits of thin provisioning can be lost. Other isolated, more unique file system requirements may also override the benefits of thin provisioning. Overall, however, data centers that can commit to constant monitoring should embrace thin provisioning for its substantial benefits. Examples of vendors that provide thin provisioning for file systems include EMC, Network Appliance (and IBM, which OEMs NetApp systems), BlueArc (and Hitachi Data Systems, which OEMs BlueArc systems), and Symantec.

© **InfoStor January, 2008**

About the author:

Ray Lucchesi is president of *Silverton Consulting*, a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community.

<mailto:info@silvertonconsulting.com>

<http://www.silvertonconsulting.com>