White Paper

Storage Thin Provisioning Explained

...Managing the Storage Explosion series

A White Paper by Bloor Research
Author : Peter Williams
Publish date : January 2010

Storage thin provisioning can provide a number of ongoing CAPEX and OPEX savings as well as increasing business agility. Organisations should be looking to plan for its implementation while carefully considering the functionality of the available systems in relation to their own existing infrastructure

Peter Williams

This document is published in association with

Executive summary

This paper explains what is meant by storage thin provisioning, the way it works and the storage problems it addresses. It describes its potential benefits as well as a few pitfalls for potential users to watch out for. It describes the technological approach to thin provisioning used by one leading vendor, HP, explaining its architecture in overview. Finally, a case study demonstrates benefits that can be accrued through using this approach.

Fast facts

- A series of storage problems associated with rocketing data volumes are being experienced by organisations of all sizes, including:
  - » Ever greater pressure on IT infrastructure capacity and performance, as well as on the IT staff managing it in a climate in which there is strong emphasis on keeping costs down.
  - » Increasing equipment space, cooling and power usage, yet poor disk utilisation (typically below 30%) bringing higher operating costs (OPEX).
  - » More backup and replication capacity needed as well as longer backup and archive times, plus higher licence and maintenance charges to cover more equipment and software.
  - » Potential limits on scalability and expansion because of, for instance, finite data centre space or power capacity or systems lacking sufficient scalability to cope with data growth.
  - » IT staff experiencing ever more time-consuming capacity planning and management, exacerbated by 'silos' of data for different applications needing to be separately assessed and managed.
  - » Increasing contention and uneven and/or degraded performance, often as a by-product of server virtualisation, affecting SLAs and/or increasing the need for re-balancing of the data storage load.
  - » Increased disruption to live operations when new storage infrastructure is installed.
  - » Virtualised server flexibility constrained by using an inflexible storage infrastructure undermining IT staff's ability to rapidly respond to changed business requirements.
  - » Overall, these increase an organisation's CAPEX and OPEX for systems and IT staff while putting pressure on performance, SLA achievement and business agility but storage thin provisioning helps address these issues.
- Thin provisioning, often implemented on top of storage virtualisation, directly addresses the capacity planning and management problems as data volumes increase, automating most of the process.
- Thin provisioning uses centralised management and a single virtualised pool of data aggregated from the data silos using a multi-node SAN; this simplifies the management process.
- Storage capacity is only provisioned for an application just before it needs it, and this can push up disk utilisation above 80%, reducing OPEX (less power and space used) and deferring CAPEX on new storage.
- Higher disk utilisation means physical capacity for primary storage is reduced and, if thin provisioning is extended to the backup and replication data copies, their physical capacity requirements are also reduced.
- Benefits from adding thin provisioning not already mentioned include:
  - » Largely eliminated time and cost involved for staff in data storage capacity planning and management (the only major need being to add extra physical capacity in time for when it is needed).
  - » (in appropriate systems) elimination of the time needed by IT staff to monitor and carry out system load balancing because it is handled automatically as part of the provisioning process.

Executive summary

- Reduced and deferred ongoing equipment CAPEX and OPEX through better storage array utilisation—less power, cooling, floor space, potentially less hardware and software licences and ongoing maintenance charges, and overall cost per GB of data stored.

- Cost savings as per the above may also be achieved through the software reducing backup, replication, archiving and DR file sizes in line with the thin provisioned primary storage file sizes.

- Better business responsiveness and agility through greater storage capacity deployment flexibility reducing the cost and time to implement business changes affecting data storage; overall less time and effort in storage management (complementary to server and storage virtualisation).

- Bloor believes that thin provisioning using iSCSI will appeal to a wider range of organisations than FC and prove more cost effective, not least because IT staff will not need to acquire specialist FC skills.

- Potential pitfalls to consider before implementation include: poor performance or scalability; incompatibilities, including mixed operating systems; lack of appropriate IT skills; inability to extend thin provisioning to backup and archive disk copies; ignoring emerging trends.

- The HP LeftHand P4000 SAN/IQ architecture provides the following features and functionality geared towards scalable, thin provisioned storage:

  - Storage utilisation up to 80–90% (versus typical utilisation below 30%) saves physical equipment, space and power OPEX, with CAPEX on new capacity deferred and reduced—assisted by a pricing model that facilitates incremental capacity purchasing.

  - Its thin provisioning approach largely obviates capacity planning and management, saving IT staff time and costs.

  - Centralised management via the Centralized Management Console (CMC) and use of iSCSI connection simplify and reduce the cost of new storage set-up and management and do not require specialist FC skills.

  - Scale-out clustering achieves near linear scaling of performance as storage grows through adding infrastructure alongside disk capacity; this is needed to maintain service levels.

  - The VSA option enables virtualisation with thin provisioning to be achieved without creation of a physical SAN; potential users should weigh up costs versus performance, resilience and scalability.

  - HP LeftHand claims 'Hyper-availability' (rather than "merely" high availability) is achieved for business continuity through distributing the intelligence across the storage modules plus redundancy of all components, even beyond a system or site, enabled by SAN/IQ Network RAID.

  - Consolidated snapshots and remote replication save capacity and improve backup performance, including to remote sites, and provide a low-cost DR capability per volume; a synchronous replication option can be used to provide mirroring for full DR.

  - The thin provisioning solution with storage virtualisation, complementing server virtualisation, can assist in achieving operational agility to respond rapidly to changing business needs.

- A live case study shows how thin provisioning built on storage virtualisation—combined with other features such as very high availability, central monitoring and control and scalability—can bring hard (financial) and soft (operational flexibility) benefits to an enterprise

The bottom line

Storage thin provisioning using an approach such as the HP LeftHand P4000, provides a string of ongoing CAPEX and OPEX savings and increases business agility. Savings come in less storage equipment and lower staffing costs. Overall, Bloor believes that storage thin provisioning will soon be widely demanded for data storage, with iSCSI SAN connection also overtaking FC long-term; so organisations should be looking to plan for its implementation while carefully considering the functionality of the available systems in relation to their own existing infrastructure.

The storage explosion: problems to address

The amount of data being stored on disk and tape continues to rocket. As this happens, there is ever more pressure on the capacity and performance of the IT storage infrastructure and the IT staff doing their best to manage it. The present economic climate means organisations have to apply strong internal pressure to hold down or even reduce costs while, as a minimum, maintaining system performance. Organisations of all sizes are affected; smaller businesses' data quantities are often equivalent to those of large enterprises a couple of years ago and larger enterprises' storage requirements can run to petabytes (PBs).

More data means more and larger disk arrays to store it, in turn taking more space and power and producing more heat that needs cooling to compensate—yet more power consumption. If a data centre is used, the finite limits on space and power capacity can easily become critical and a constant management headache; alternatively the current systems may lack sufficient scalability to cope with the data increase going forward.

This problem increases with primary storage needing ever-larger back-up and remote replication copies—the latter for disaster recovery (DR), updating for branch offices in some cases, and archiving. More equipment also means higher licence and maintenance costs while more data increases backup and archiving times. These factors all push up capital and operational expenditure (CAPEX and OPEX).

The problem is typically further exacerbated by 'silos' of physically separate storage in which the data is dedicated to specific applications or is remotely located, or both. Silos can occur through takeovers or mergers between companies using different applications and data sets, or as the result of branch office autonomy. Overall disk storage capacity utilisation is also low, typically below 30%, while spare capacity needs to be allocated for growth for each silo; this multiplies the time spent by IT staff on capacity planning and management.

Server virtualisation is now commonplace. Resulting benefits can include consolidation that reduces the number of physical servers and added business agility because new virtual servers can be quickly created or collapsed as needed to meet urgent business requirements. However, there are downsides. A physical server supporting multiple virtual servers may experience a performance slowdown through increased contention, and this is especially true if each virtual server is attempting to access its own silo of data storage. This flexibility in creating virtual servers tends to mean they rapidly proliferate, which only increases this problem.

Conversely, an increase in total storage can give rise to uneven or unacceptable server access performance that may affect service level agreements (SLAs) or mean applying frequent re-balancing of the stored data load is needed. Moreover, such server flexibility is constrained by the inflexibility of the storage it uses, undermining an IT department's ability to rapidly respond to changed business needs.

These factors, in general, push up the amount of storage hardware, and the networking for attaching it, in order to get performance, increasing management difficulties. Moreover, there is typically disruption whenever new storage equipment is installed and primed for use. Multiple instances of software involved in operation and management of the data silos also adds to licensing and maintenance costs.

In summary, organisations are faced with an urgent need to tackle some or all of these storage problems—and the now commonplace server virtualisation can actually add to the pressure on the storage infrastructure and its management. Anything that can alleviate one or more of the capacity, performance and management pressures without significantly increasing CAPEX or OPEX deserves careful consideration. This is where thin provisioning comes in.

What is thin provisioning?

Thin provisioning is invariably discussed in the same breath as storage virtualisation because it is commonly implemented on top of virtualisation. However, they are different and separate (which is why they are the subject of separate papers in this series).

Thin provisioning is designed specifically to address some of the problems inherent in standard capacity planning and ongoing capacity management as the volume of data increases. It uses centralised management that can greatly simplify the management process and is also effective in reducing primary storage capacity by boosting utilisation (even above that achieved by storage virtualisation); dependent on the choice of system, this can be extended to also save disk capacity for backup and replication copies of the data.

The thin provisioning system (with complementary storage virtualisation) will bring together the data silos by adding (or extending) a storage area network (SAN) with multiple nodes—aggregating the silos' physical data and spare disk space to create a single virtual pool of data and spare disk capacity. This is then able to be managed as one from a central console. However, it is thin provisioning alone that addresses, and to a large extent obviates, the need for IT staff to carry out capacity planning and ongoing capacity management.

The thin provisioning concept is very simple with the architectural complexity hidden from the user during day-to-day operation. Whereas each server application has allocated to it a set amount of storage, thin provisioning does not provide that in practice; however, whenever an application approaches the physical limit of the storage already 'provisioned' for it, extra storage is provided from the central pool. By this means, IT staff no longer need to concern themselves with whether any individual application's storage capacity will exceed its physical limit; the system handles this automatically (autonomously).

The overall capacity will still need monitoring in order that extra equipment is added when needed, but the thin provisioning system will aid this by maintaining a displayable status on current overall capacity usage and back that up with a warning mechanism if the physical capacity limit is approaching.

As capacity is provisioned only as needed, overall disk array utilisation can be allowed to increase to 80–90% provided suitable performance can be maintained, and the approaches to this vary between the different thin provisioning solutions. The high utilisation will reduce an organisation's total equipment and OPEX for the same data capacity as well both deferring and reducing the CAPEX for new equipment.

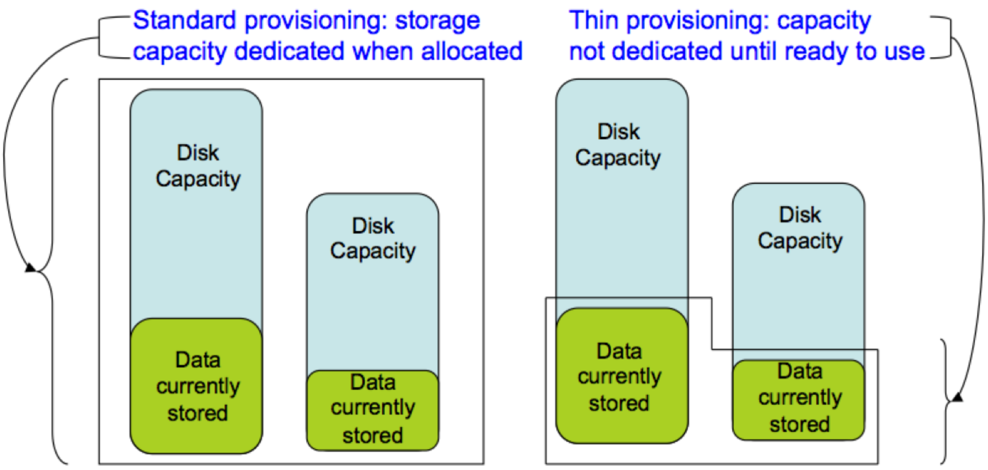

Figure 1: Thin provisioning does not dedicate storage capacity to an application when allocated but rather when it is actually needed

Why should thin provisioning be deployed? The benefits

In a traditional storage environment, IT staff must spend time calculating the disk capacity that is likely to be needed going forward (usually based on previous usage) for every separate application and/or data silo; this may be an under- or over-estimate with capacity planning often something of a hit and miss exercise. The amount of data stored for some applications will grow faster than the estimate; for others less. Nor is this a one-off requirement; once new capacity is identified it has to be requested, budgeted for, management approved, purchased and then installed—to be repeated for every data silo. Thereafter, ongoing monitoring of the capacity used is needed to ensure the next uplift purchase is made in good time.

The substantial time delay to equipment purchase approval and installation means IT staff tend to request far more extra disk capacity than they estimate will be needed soon. This reduces the frequency of purchase requests and provides capacity for handling an unexpected spike in demand by one application, but the effect is for capacity to be increased in large uplift steps even though the amount of data stored rises more steadily (albeit inexorably); so at any time there will tend to be a large amount of spinning disk capacity unused—not least because, as standard (and whether or not storage is virtualised), all the allocated capacity for each application has to be *physically* available to it all the time.

These problems are largely eliminated by thin provisioning. Thin provisioning (with virtualisation) adds storage only to top up the overall pool of capacity. Thin provisioning detects when an application actually needs extra capacity then makes the capacity available to it so that the application can continue (and regardless of whether the application is now adding more or less data than was previously estimated. Importantly, no application has any allocated storage reserved for it ahead of time; this is achieved by 'tricking' the system into thinking it has its allocated amount of storage for every application when it does not—over-provisioning—and draws on the pool of spare capacity on a first-come, first-served basis.

The aggregate of free storage space will more than meet any individual application's total needs (just not the total of all applications' future requirements), so extra physical capacity needs only to be made available when the *total* pool of storage is getting low. Thin provisioning can simply check a pre-set threshold level—and if and when the total of spare capacity drops below this threshold an automatic alert requesting more capacity to be made available is issued. By this means, the overall disk utilisation per application, and in total, can safely go much higher (up to 80% or more) because the possibility that an application will otherwise run out of capacity is removed.

In turn, storage array purchases may then be reduced in size and/or deferred (a CAPEX saving) while deferring deployment of physical capacity until it is actually needed achieves an OPEX reduction through saved running costs versus spinning disks carrying unused capacity.

Equally important is the way the thin provisioning process typically deals with new equipment being attached to the storage pool. When new physical capacity becomes available in a standard storage system, performance for an application will only be optimised if load balancing is immediately applied; even then the performance benefit will tend to diminish as more data is added unless load balancing is frequently repeated. This is time-consuming and needs controlling by IT staff.

Using thin provisioning, physical attachment of new storage capacity can be completed without disrupting live operation then, when the new storage becomes available to the system, it is autonomously and unobtrusively added to total pool of spare storage. In addition, load balancing of the whole storage data in use pool is typically an ongoing automatic background task to continually maintain optimal, evenly spread data access performance—to smooth out uneven performance and optimise it overall. This needs to be an ongoing process because the growing amount of storage is always altering the loading.

[Note: When storage virtualisation consolidates data silos this can exacerbate performance issues because there is greater contention between applications now sharing the same physical hardware for the first time; also, mission-critical on-line transaction processing (OLTP) applications do not mix well with background 'batch' tasks; so load (re-)balancing along with prioritisation of tasks is crucial. IT staff need not concern themselves with the autonomous load balancing process.]

In summary, these factors together remove virtually all storage capacity planning and management headaches and performance

Why should thin provisioning be deployed? The benefits

optimisation issues despite an ever-changing data storage picture. As physical storage capacity is only dedicated to an application when it is about to be used (rather than as soon as allocated to it), disk storage utilisation can rocket to over 80%, and disk arrays not being used do not need to be using energy and giving out heat and, if not delivered until needed, it will also save space. This all reduces OPEX. If physical capacity is incrementally increased and/or the new hardware purchases are deferred, this further reduces and defers CAPEX. The overall effect is also more environmentally friendly. Finally, this automated provisioning helps an organisation to be more responsive to rapid changes in storage capacity needs for particular applications, assisting business agility.

Thin provisioning may also go further, as will be described. For instance, scaling performance when adding more storage can be a major headache, so the way the system is architected in this respect can be critical to longer term success. Complementary applications may also extend thin provisioning to back-up and replication copies of the thin provisioned storage to further reduce the data footprint and save on transmission and storage costs.

Other issues and potential pitfalls, and the features to overcome them

The following are some items that need to be considered before implementing thin provisioning. (It is not an exhaustive list):

- Higher utilisation of available storage capacity (a consequence of both thin provisioning and storage virtualisation) may introduce access contention problems leading to poor or uneven performance; so it is important to understand how the thin provisioning application overcomes this, including how it deals load balancing—so as to maintain optimum performance and minimum contention.

- Whatever the load balancing (and other fine tuning) capabilities, these must not be a major drain on performance, which would tend to defeat their very purpose.

- Scalability will continue to be a major issue for storage; the user needs to ensure that the solution not only allows for massive capacity growth but that this will not result in major access performance loss.

- Adding new physical disk storage capacity must not disrupt live operation in any way; ideally the software will automatically detect the added capacity then add it to the spare storage pool.

- If installing from scratch, typically alongside storage virtualisation, there will be a need to consider what storage to include in the SAN at the outset from (potentially) a mix of DAS, NAS and SAN situated centrally and remotely and what older equipment should be replaced to optimise performance.

- If the company has a mix of operating systems, typically at least Microsoft Windows and Linux and/or UNIX, the centralised thin provisioning solution needs to be able to service data access from each.

- Many organisations moving to both storage virtualisation and thin provisioning will have general IT skills rather than SAN-specific expertise (historically meaning fibre channel (FC) skills as well); the thin provisioning solution (software and hardware) needs to be relatively simple to install and use.

- Thin provisioning solutions typically require FC or, increasingly, iSCSI (but do not support both). iSCSI is likely to be the easier option for organisations without an existing FC structure and in-house FC expertise (not least because all use Ethernet).

- Other factors, such as poor network communications, can create bottlenecks that affect performance; so these need to be considered as part of the mix at the same time as planning for thin provisioning.

- VMware and other virtual server users will need to be reassured that the thin provisioning solution is fully compatible, not imposing any restrictions on storage access from the virtual server applications.

- Confirm there are no technical limitations preventing the thin provisioning solution extending beyond the primary storage to mirror imaging, snapshots and remote replication so as to further reduce overall raw capacity and (perhaps) assist performance.

- Conversely, potential users should satisfy themselves that the preferred thin provisioning solution will not, in any way, compromise the ability to recover after system failure.

Why should thin provisioning be deployed? The benefits

They should confirm how well storage capacity is self-monitored and the displays and alerts for when, for instance, spare capacity is becoming critically low.

There should also be reassurance about how the system handles quality of service (QoS), tiering of storage and mobility of data between tiers, prioritisation to individual applications (e.g. those that are mission-critical) and the ability to make changes to support business agility.

Consider the likely impact of emerging trends on infrastructure and costs longer term, for example the expansion of iSCSI and increasing availability and reduced pricing for high-performance solid state disk (SSD) storage.

The HP LeftHand P4000 approach to thin provisioning

In overview, HP's purpose-built thin provisioning solution is designed around a clustered IP SAN using HP LeftHand P4000 series hardware which also provides the storage arrays. Currently there are two main models, the P4300 with eight disk drives per storage node and the larger P4500 with 12 disk drives per node. Both can scale to over 20 nodes per cluster non-disruptively. A node consists of an Intel x86 based storage server running HP's SAN/iQ software that virtualises the RAID array-based disk storage and controls 'reservation-less' thin provisioning of storage capacity—meaning it does not reserve disk storage space for an application in advance of it being needed in practice.

The HP LeftHand P4000 SAN solution is also available as a Virtual SAN Appliance (VSA)—software only running as a virtual appliance inside VMware—which virtualises the internal and direct-attached storage of the underlying devices without needing an external SAN. HP has recently extended the maximum VSA capacity to 10GB per ESX server. Prospective users need to weigh up costs versus performance, resilience and scalability.

The system uses iSCSI connection, now by far the fastest growing connectivity method although Fibre Channel (FC) still has more SAN installations. Two obvious advantages for iSCSI over FC are its lower cost and relative simplicity to implement and these make it particularly popular in the low- to mid-range market; however, there is no obvious performance or reliability reason to prevent iSCSI being deployed widely in larger enterprises. For these reasons Bloor expects this connectivity protocol to become increasingly dominant over the next few years.

As will be demonstrated below, the P4000 architecture design has concentrated on very high availability, with no single point of failure across the SAN, and an ability to scale performance directly in line with the increase in storage capacity. These two complementary design criteria are critical since storage capacity is increasing inexorably while adding more data puts more pressure on both performance and reliability. The hardware and software is supplied as a self-contained appliance and the whole package includes utilities that extend thin provisioning to snapshots and remote replication of the data.

Product architecture

The HP LeftHand P4000 series uses rack-mounted HP SB40c storage blades in a multi-blade chassis supporting SAS and SATA disks. Its purpose-built operating software, SAN/iQ, manages the clustering within its iSCSI SAN and also provides a platform for running extra applications. This is a true cluster, which has a single IP address, and every storage server in the cluster shares equally in supporting the cluster's workload and storage capacity.

New storage modules may be added to an existing SAN/iQ cluster or be used to create a new one as needed. The entire operation may be controlled by one administrator from the central management console (CMC), an intuitive 'single pane of glass' GUI; this is irrespective of whether the storage is centralised or spread across multiple remote sites. This is one aspect of maintaining simple operation even when provisioning new data storage.

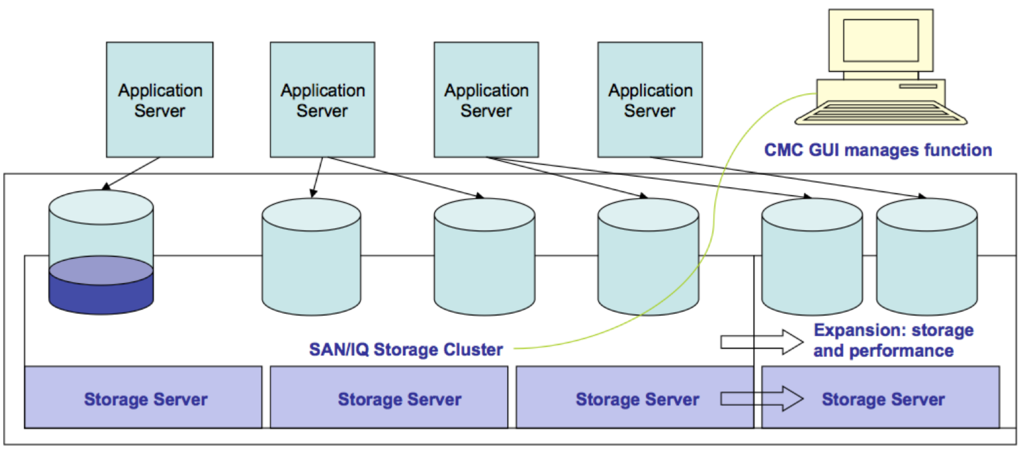

Figure 2: The SAN/iQ storage platform uses true clustering and maintains a constant ratio of processing resources to disk capacity to achieve near linear scaling as the cluster grows

The HP LeftHand P4000 approach to thin provisioning

A major reason for clustering storage is that it facilitates creation of shared storage nodes out of 'silos' of storage dedicated to individual server applications. By consolidating these silos (storage virtualisation) into a single aggregated pool of storage, all of it can be made available to all applications. To assist in achieving this, SAN/iQ configuration wizards are used during initial set-up to guide the administrator through the process of discovering existing P4000 systems and then creating new data volumes and connecting them with the appropriate live servers.

Thereafter, adding new storage to the cluster is like adding any other Ethernet device because SAN/iQ intelligence helps automate the process. Assigning new physical capacity to a virtual storage volume can either be entirely automatic or achieved by selecting from a pull-down menu; either way this takes very little time and requires no system downtime, so it allows IT staff to add and manage storage during normal business hours. (iSCSI is also inherently simpler to install and manage for SAN storage connectivity than widely-used FC.)

In a standard virtualisation process, consolidating such storage resources can lead to contention, resulting in a performance hit because the aggregated storage is no longer dedicated to a single server and application. HP's SAN solution circumvents this partly by deploying SAN/iQ Network RAID; this stripes and mirrors multiple copies of the data right across the whole cluster of storage nodes. This works in two ways:

- Since each RAID array has a separate set of read-write heads—and there may be hundreds of arrays within the cluster—each can be accessed simultaneously with all others without any contention between them, inherently reducing contention between applications.
- SAN/iQ includes virtual IP load balancing which distributes the data path load across all storage modules in the cluster to remove any performance bottlenecks—since no one device controls either management or data traffic.

SAN/iQ's Distributed Storage Matrix (DSM), together with the hardware architecture, also achieves enterprise-level high availability; thanks partly to all stored data being mirrored, Network RAID eliminates any single point of failure in the HP P4000 SAN. The supported hardware RAID levels 5, 6 and 10 are further backed by Network RAID carrying RAID levels (per volume) of 0, 2, 3 and 4; this double-level protection copes with double disk failure and system and site failures.

High availability (HA) features extend to power supplies and disk drives being hot pluggable for quick replacement without system downtime and an integrated storage controller with battery-backed DDR2 cache for every storage module. Together these provide 'hyper-redundancy' to ensure the HP P4000 SAN maintains continuous stored data availability and live application operation, even in the event of a failure of power, network, disk, controller, entire storage node or site.

Business continuity is further assisted by the HP P4000 VSA including integrated synchronous and asynchronous replication (see Supporting Products); importantly, SAN/iQ thin provisioning extends into the back-up and remote replication products to save on total backup disk capacity and assist speed of data transfer. A simple management-controlled failover and failback procedure facilitates low-cost, centralised backup and disaster recovery (DR) per volume.

This SAN/iQ Network RAID approach for the P4000 achieves a "scale-out" architecture that, unusually but importantly for thin provisioning, scales performance alongside capacity. Every time a storage module is added to the cluster, the capacity, performance, and redundancy (HA) of the entire storage solution increases. Alongside more disk capacity, resources added are: an extra storage server (CPU and main memory) share in managing the storage, I/O and SAN/iQ cluster; a hardware RAID controller and additional battery back-up for the extra

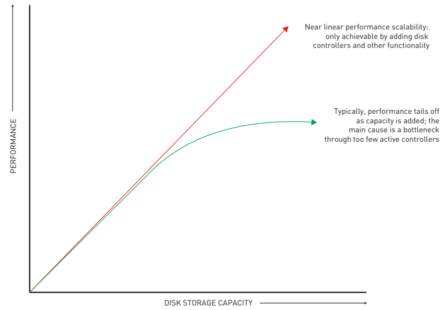

Figure 3: Scale-out architecture adds infrastructure alongside disk capacity to avoid bottlenecks that undermine performance

The HP LeftHand P4000 approach to thin provisioning

storage per storage server; further network interfaces available to operate in parallel. By this means a 30-storage server system serving 30 application servers achieves close to 50,000 I-Os per second (IOPS). This contrasts particularly with most systems which do not add extra controllers as standard, so the few controllers become an increasing bottleneck as storage capacity increases.

If the system is being run for Microsoft Windows operating system, SAN/iQ invokes a device specific module (DSM) to work with Windows Multi-Path I-O (MPIO) iSCSI driver, with intelligence to calculate the location of any block in any virtual volume; this means it can bypass the re-direction normally invoked by the load balancing functionality so is faster (and especially so for sequential data access).

A new module can be configured dynamically from the CMC and is automatically brought into the total storage resources pool without disruption or system downtime. Applications can also remain online during all maintenance.

The P4000 thin provisioning process

Each server application has a pre-set amount of data storage allocated to it and, in standard environments, this total capacity has to be physically available. All thin provisioning approaches work using the principle that there is spare capacity within the storage allocation that does not need to be dedicated to the application initially (although the server system needs to think it is). The result is that far more storage may be allocated than is physically available (sometimes referred to as 'over-provisioning').

The basic functionality of SAN/iQ's thin provisioning approach is included within the features described above but the way it allocates capacity is extremely simple. It works as follows:

- When utilisation of a storage volume for an application reaches around 90% (i.e. there is only 10% spare), this automatically triggers SAN/iQ to reserve more space for it from the central free storage pool.
- This extra space is only added in small (128MB) increments so very little spare capacity is dedicated to each volume at any one time.

It can be seen that overall utilisation by any one application will hover between 80 and 90%* (compared with a typical utilisation below 30%). Total disk storage utilisation will be a little lower because there is a reliance on extra capacity being present in the storage pool; however, this can be kept to a minimum by only adding physical storage to the pool in small increments.

HP encourages this within its pricing model which facilitates users purchasing only the capacity they currently require; users benefit by deferring CAPEX for new capacity and avoiding power and space OPEX for under-used spinning disk. This works well because, as explained above, the extra storage capacity comes with extra systems to maintain performance.

*Note: Some thin provisioning systems allow the user to set a parameter to determine at what percentage utilisation to trigger more capacity allocation; there is a trade-off between this extra procedure perhaps allowing the utilisation to go even higher than 90% (although clearly not much more) or always leaving the system to take care of it, as in this case.

So the IT staff need not concern themselves with either storage capacity planning or management for each application. There is, however, one figure for them to monitor—the overall utilisation of the total physical capacity (spinning disks) in the cluster. At the CMC, the administrator can click on a selected cluster (where there is more than one) then instantly view a graph of how space is being used by each volume, as well as for the snapshots, and what capacity is still free.

An extra figure is provided that shows the maximum 'provisionable' space—a calculation showing what extra space could still be used by the applications in the unlikely event that all physical storage were deployed. However, there are also alerts for when physical storage is getting close to full. Among the standard alerts configured for the nodes (which mostly cover equipment malfunction) is one that is triggered when cluster utilisation reaches a percentage threshold; this threshold may optionally be varied by node and can be set by the user. These alerts are sent to the management console either by email or via SNMP. By this means, the IT staff can simply monitor usage to see when to order more equipment (i.e. without needing to spend time on complex capacity planning calculations) knowing there is extra protection from the system pro-actively warning if capacity has reached the overall cluster utilisation threshold.

The HP LeftHand P4000 approach to thin provisioning

In summary, this method means that overall disk utilisation will automatically increase up to a maximum of 80–90% and, primarily through the system's architecture (including the use of Network RAID and other high availability features), neither performance nor reliability will be seriously impacted by the growth in data storage. The HP P4000 thin provisioning solution probably automates provisioning of new storage capacity as far as can be achieved.

Supporting products

The main supporting products are provided as part of the overall solution package. These include integrated snapshot and remote copy (which uses the snapshot approach). The clustered architecture allows consolidation of snapshot copies; so snapshots can write out point-in-time copies of data from multiple storage nodes into pools of storage copies on a per-volume basis, effectively thin provisioned to save output space. Snapshots can then be used for recovery of individual files or folders or a whole volume can be rolled back to a prior state; snapshot copying is available to every volume in the cluster.

The integrated asynchronous remote replication feature extends this capacity saving to remote sites, with reduced data being transmitted across the network. This also incorporates bandwidth throttling as it replicates snapshots between P4000 SANs at primary and remote locations; copies are thin provisioned without any space needing to be reserved for the output in advance. A simple to use CMC-controlled failover and failback procedure can be deployed for low-cost, centralised backup and disaster recovery (DR) per volume.

The multi-site and DR solution feature, now incorporated into the main software, takes this further to provide synchronous replication; placing alternate P4000 storage nodes in different locations then means synchronous replication (mirroring) for DR takes place automatically.

The SAN/iQ software is modular and straightforward to port to new operating platforms. A wide range of server operating systems are supported by the HP P4000 series including latest versions of Microsoft Windows Server (including x64) and Hyper-V, Novell NetWare, SUSE Linux including Enterprise Server x86/x86-64, Red Hat Enterprise Linux and Fedora, HP-UX, IBM AIX, VMware ESX/ESXi and Site Recovery Manager, Citrix XenServer and Mac OS X Leopard.

The benefits achieved by this thin provisioning method

- The HP P4000 thin provisioning method does not reserve storage space, only dedicating it to an application when about to be used, so storage utilisation goes up to 80–90%; so half to two thirds of physical disk capacity is typically saved over conventional storage systems—saving space and power OPEX as well as licensing and maintenance charges.

- For the same reason new disk storage capacity purchases can be deferred to much nearer the required usage date and they can be small and incremental; this defers and saves CAPEX as well as OPEX; this is supported by HP P4000 pricing model that facilitates capacity to be purchased only as needed from the outset. This also simplifies equipment purchase planning.

- It obviates the need for application- or server-level capacity management and greatly simplifies overall capacity planning; this provides a major saving in IT staff time and hence costs.

- Ease of setup and centralised management from the CMC likewise reduces IT staff time for storage set-up, management and monitoring.

- Using iSCSI connection rather than FC for the SAN lowers CAPEX, simplifies set up and management functions and removes the need for specialist FC IT skills; this further reduces equipment and staff costs when compared with FC alternatives.

- The scale-out clustering design with SAN/iQ Network RAID and adding infrastructure alongside extra disk capacity, achieves near linear scalability that will maintain performance, to maintain service levels.

- 'Hyper-availability' and business continuity is achieved through distributing the intelligence across the storage modules and redundancy of all components; backup and recovery utilities forming part of the overall solution.

- Consolidated snapshots and remote replication take copies from multiple nodes into a single storage copy pool per volume; this saves capacity and improves backup performance, including to remote sites.

The HP LeftHand P4000 approach to thin provisioning

- The CMC-controlled remote replication also facilitates a low-cost DR capability per volume while the synchronous replication option can be used to provide mirroring for full DR.

- A good thin provisioning solution sitting over storage virtualisation (as in this case) that complements server virtualisation will help in achieving real operational agility for IT to rapidly respond to changing business needs.

Case study: Mitsubishi Electric Automation

Illinois (US)-based Mitsubishi Electric Automation Inc. was formed through the merger of three of Mitsubishi Electric's automation product business units in 1997. It supplies factory automation products and solutions to the North American and Latin American markets and is one of many subsidiaries worldwide of the $40 billion Japanese Mitsubishi Electric Corporation, which provides automation products for many industrial markets.

Within its product portfolio the Mitsubishi Electric Automation supplies: programmable logic controllers (PLCs), variable-frequency drives, operator interfaces, motion control systems, computer numerical controls, industrial robots, servo amplifiers and motors, and industrial sewing machines. Such diversity is one reason why there is a huge amount of both business and technical data to store.

The company's core software, which was already installed at the outset, is SAP Enterprise Resource Planning (ERP). This was using data held on direct-attached storage (DAS) devices with a single link to the server, which was not very efficient or resilient; this was highlighted when a serious system crash left much of its data unavailable for several days. Mitsubishi then realised just what information it had stored and how critical it was to the functioning of the business. It was quickly concluded that a new storage and data protection solution was needed.

Among the criteria identified for the new system was the need: a) to simplify storage management and b) to eliminate such single points of failure as the DAS connection. Early investigation by Mitsubishi's technical team led to their ruling out FC connection on the basis of its high implementation and management costs, and instead to go for an iSCSI SAN. This limited their choices with the main iSCSI SAN contenders being HP LeftHand P4000 and Dell EqualLogic; Mitsubishi, already an HP user, chose the HP LeftHand P4000 SAN solution after having been satisfied that it would provide the high availability, performance and security required.

Mitsubishi Electric Automation now uses HP P4000 SANs running on 10 HP ProLiant storage servers for its SAP ERP storage; this is held on a central database which uses Oracle 10g release 2. This storage covers all financial, sales, services, purchasing and shipping data. This solution sits well with the company's existing server and PC standards (all HP) as it provides a consistent HP environment. The company was among the first to run its ERP system on HP LeftHand SAN architecture—and its ability to do this was a major factor in the purchase decision.

An integrated VMware server virtualisation and HP P4000 storage virtualisation strategy has led to increased efficiency and productivity, reduced costs and saved energy. The whole backup process is now centralised with all storage viewable from the CMC.

Storage consolidation through the HP LeftHand P4000 thin provisioning approach has also eliminated most potential points of failure for storage. The vulnerable direct server to DAS link has been superseded by redundant network links through redundant switches to the high redundancy HP P4000 SANs, as is vital for critical enterprise systems. Since the new software and hardware was installed, no storage failure has resulted in SAP being brought down.

Both day-to-day operational and non-routine time savings have allowed staff to be freed up to concentrate on more business-critical matters. One non-routine saving example is the procedure when changing SAP from an old to a new production system after an upgrade. Previously, with its DAS storage, four full backups were needed at different points, each lasting about three hours (so 12 hours in all). The simple-to-run P4000 snapshot capability now completes backups in seconds and, utilising the SAN/iQ thin provisioning software, the output is also thin provisioned to be extremely efficient in disk space usage. One spin-off result of this is that the company now keeps multiple versions of the same data set—yet taking up little more than the single version.

This snapshot function is also a key factor in bringing greater flexibility and business agility. The company's technical architect cited an example of when Mitsubishi's SAP consultants asked for a backup before they applied some database changes; the consultants had expected several hours wait so were shocked to find the snapshot backup was completed in seconds, saving several hours of consulting time. It also highlighted much greater flexibility for applying system patches and changes.

A concern for some about thin provisioning is whether performance is maintained despite higher disk utilisation. In fact Mitsubishi reported that the HP P4000 SANs outperform

Case study: Mitsubishi Electric Automation

traditional SANs as well as other iSCSI SANs. A key factor in this is the P4000 scaling architecture adding a controller for each new node to handle the increased data (versus systems which experience controller bottlenecks), plus its integration with the other HP software and hardware.

Mitsubishi is not able to precisely quantify its overall financial savings but has estimated that its server virtualisation integrated with the P4000 thin provisioned virtualised storage SANs will save multiple hundreds of thousands of dollars over the next two years. To this can be added the 'soft' cost savings already described, including: reduced downtime, increased performance and improved staff productivity related to tasks such as backups and copying.

An easier part of the cost reduction calculation—which also has a "green" (energy-saving, space-saving) benefit—is that the combination of HP P4000 SANs, ProLiant blade servers and VMware server virtualisation has rendered nearly 30% of the hardware superfluous to requirements. So this equipment has been taken out of service and yielded a dramatic reduction in power consumption and heat output, while also freeing up future growth capacity.

The company's commitment to the HP P4000 thin provisioned storage SAN infrastructure is reflected in its spreading use. As well as SAP ERP, it now runs customer relationship management (CRM), business intelligence (BI) and Enterprise Portal; Microsoft Exchange Server 2003 data and much of the web infrastructure resides on the SANs while the VMware infrastructure also uses the SANs for storage.

As storage has grown, with new applications migrated to the SAN, so new nodes have been added—but only as needed as determined by the thin provisioning—while overall performance has been maintained. Such expansion would not have been achievable without this built-in scalability.

Summary and Conclusions

Storage thin provisioning using an approach such as the HP LeftHand P4000 can provide a number of ongoing CAPEX and OPEX savings as well as increasing business agility. Most visibly it will boost disk storage utilisation to reduce or defer new capacity purchases, in the meantime saving power, heat output and space. It also removes a number of IT staff headaches, the most obvious being in largely removing their need to carry out storage capacity planning and management per application, so IT staff time and costs are reduced. It is typically and logically implemented at the same time as storage virtualisation, which it complements and, with server virtualisation (now commonplace), will facilitate rapid response to changing business needs.

Bloor believes that thin provisioning will soon become a de facto requirement for data storage, with iSCSI SAN connection overtaking FC to become the most popular approach. All organisations should be planning to implement it but they need to consider the functionality of the various available systems in relation to their own existing infrastructure. Overall, Bloor believes that storage thin provisioning—in the way applied by the HP solution described above—will prove among the most cost-effective and flexible long term.

Further Information

Further information about this subject is available from http://www.BloorResearch.com/update/1063

Bloor Research overview

Bloor Research is one of Europe's leading IT research, analysis and consultancy organisations. We explain how to bring greater Agility to corporate IT systems through the effective governance, management and leverage of Information. We have built a reputation for 'telling the right story' with independent, intelligent, well-articulated communications content and publications on all aspects of the ICT industry. We believe the objective of telling the right story is to:

- Describe the technology in context to its business value and the other systems and processes it interacts with.
- Understand how new and innovative technologies fit in with existing ICT investments.
- Look at the whole market and explain all the solutions available and how they can be more effectively evaluated.
- Filter "noise" and make it easier to find the additional information or news that supports both investment and implementation.
- Ensure all our content is available through the most appropriate channel.

Founded in 1989, we have spent over two decades distributing research and analysis to IT user and vendor organisations throughout the world via online subscriptions, tailored research services, events and consultancy projects. We are committed to turning our knowledge into business value for you.

About the author

Peter Williams
Practice Leader - Information Storage Management

Peter Williams' IT industry experience spans over 44 years and he is well-respected for his breadth of both technology and business knowledge. Peter joined Bloor Research in late 2004 after four years as an IT writer and correspondent. Peter recently took on the Information Storage Management brief for Bloor.

Information storage management here refers to the management of the system-attached and network-attached storage equipment for an enterprise and for the software deployed for data backup, replication, archive, restoration and disaster recovery.

Peter's early experience was primarily with computer manufacturer NCR, working in a number of technical roles including: programmer, systems analyst, installation consultant – then graduating to project manager responsible for major software developments in the banking, retail and health sectors. In the late 1980s he became NCR's national marketing manager for the logistics and manufacturing vertical sectors. In 1993 he joined a high-tech marketing company, writing promotional and technical material for companies such as Digital (DEC), Kyocera Europe and RS Components, while he also gained a BSc in Computing part-time.

After a contract period as information systems manager for a pipework tools manufacturer he entered IT journalism with IT Week, part of VNU Business Publications, in October 2000 and was later a founding correspondent within the VNU News Centre team. In his four-years with VNU he provided some 2000 reports and analyses on both technology and business for the VNUnet web site plus IT Week, Computing, Network News and Computer Reseller News IT business publications. Peter was able to conduct in-depth interviews with some top IT industry leaders. Latterly he also wrote regular newsletters on open source and business hardware.

Away from IT, Peter plays an active role in his local church. He is an honorary vice-president of the Sussex County Table Tennis Association having long ago represented Sussex and England. Married 36 years, he has three grown up children.

Copyright & disclaimer

2nd Floor,
145–157 St John Street
LONDON,
EC1V 4PY, United Kingdom

Tel: +44 (0)207 043 9750
Fax: +44 (0)207 043 9748
Web: www.BloorResearch.com
email: info@BloorResearch.com

4AA1-6527EEW