



Waste-Not, Want-Not — 3PAR Builds Thin into an Information Lifestyle

Analyst: Anne MacFarland

Management Summary

In many situations, right-sizing, or good tailoring is the key to expense avoidance and environmental virtue. In IT storage, the most prominent wrong sizing usually refers to the capacity of the volumes by which storage is provisioned. This capacity addresses certain application expectations as to what will be needed – and often errs on the side of generosity. This leads to low utilization of storage arrays. In times of high-energy costs and lean budgets, this wrong sizing represents a continuing, unbearable cost.

Low utilization of storage also can be a result of high I/O demand. With server consolidation using virtual machines, the I/O demand on storage per server increases. Because the requests are coming from different applications tenants in different stages of varying processes, the I/Os are random, not sequential. IT has many ways to optimize for sequential reads – but when the nature of the aggregate demand makes all reads random, the tools for optimization are more limited.

This problem will only get worse. Databases are growing very large. The data in them is used more widely across value chains to optimize operations. New sources of data are being brought to bear on business decisions. Data retrieval becomes key to every sales engagement, to each customer care scenario, and, fundamentally, to any business strategy.

3PAR, based in Fremont, CA, developed and patented its *Thin Provisioning* for primary storage several years ago. With 3PAR, thin provisioning is a utility computing strategy, not just a tool or a check-off feature. The Thin strategy permeates the *InForm* operating system of its *InServ* arrays. It mandates the virtualization that carves every drive into 256MB virtual disks, the small 16KB units by which capacity is allocated, and the wide striping that distributes any volume across all drives of an array. As business data balloons, this tailoring adds up to real savings – and not just in capacity, but also in I/O response time (where more controllers can contribute), resilience via RAID, and in volume rebuild, where, again, a distributed approach can be taken.

Now, 3PAR now leverages its thin provisioning, not just for full, remote, and virtual copies, but also for three new Fall 2009 features – *Thin Conversion*, *Thin Persistence*, and *Thin Copy Reclamation*. These turn a highly attractive feature into the basis of a highly attractive information lifestyle. For more details about these new offerings, please read on.

IN THIS ISSUE

➤ Utility Computing.....	2
➤ Get Thin with <i>Thin Conversion</i>	2
➤ Stay Thin with <i>Thin Persistence</i>	2
➤ Building Out an Ecosystem	3
➤ Conclusion	3

Utility Computing

The tenets of *utility computing*, often targeted primarily at service providers and large IT deployments that run (and charge back) as service bureaus, has something to offer all IT installations when it comes to doing more with less waste. Utility computing is based on three concepts:

- A shared infrastructure,
- That can handle diverse and unpredictable workloads, and
- Support diverse service levels
- With high utilization, and
- Low operational overhead.

Utility computing is not just for Web Server farms, file serving, and other segregated tasks, but also for the diversity of needs and neediness that beset the everyday enterprise. By bolstering its *InForm* array operating system with the thin functionalities needed to deliver high utilization and low operational overhead, 3PAR successfully has created its own utility computing strategy for the last decade.

With 3PAR, thin virtualized storage is not pool-based. Instead, the entire array, available in granular chunks of a drive, is available to all applications. The array autonomically provisions the storage that is needed when it is needed, according to service levels defined for the situation.

There are two forces at work in providing this new functionality that generally are not described, but which help to explain how the new 3PAR features work. One is the *InForm* array operating system, described on the previous page. It does the volume management, resource management (including cache), provisioning, and scheduling of operations typical of operating systems. The other is the *3PAR third generation ASIC*, which offloads functions that are better done outside the array controller's processors— asynchronously or intensively.

New 3PAR Functionality

Now, 3PAR offers three refinements that make 3PAR arrays even more useful to an enterprise.

- **Thin Provisioning**, introduced in 2002, is enabled by the *InForm* Array Operating System, and is available on all *InServ* arrays.

- **New: Thin Conversion** is capable of converting fully allocated (fat) volumes to thin volumes
- **New: Thin Persistence** is a space reclamation capability to keep thin volumes thin over time. Both Thin Conversion and Thin Persistence are supported by 3PAR's third-generation ASIC that does the block-by-block free-space discovery. Thin Conversion and Thin Persistence are available only on newer arrays (the *T-Class* and *F-Class* models).
- **New: Thin Copy Reclamation** is a space reclamation capability to keep thin copies thin over time and is achieved not with the ASIC but rather with granular mapping with *InForm*. Thin Copy Reclamation is available on all *InServ* Models.

Zero Detection and Virtualization

These ASIC-based products are enabled by what 3PAR calls *zero detection*. Deleted file space is easily prepared for 'thinning' by simple two-step process that overwrites deleted data with zeros just as secure delete functions would do. As data is read by the 3PAR's third-generation ASIC, it detects zeros in-line at wire speeds. Instead of writing zeros to disk, the software effectively virtualizes them, eliminating the need for the underlying (empty) capacity. This capability does not require agents or specialized software and works with all servers. It requires no professional services and becomes part of business as usual (which, coincidentally, is the only way that any diet really works).

Get Thin with Thin Conversion

3PAR's Thin Conversion as described above can take data stored on other arrays, and convert it on the fly to thin volumes with minimal disruption to production workloads on the *InServ*. Particularly in contrast to other data migrations techniques that require professional services, this feature makes migration to 3PAR attractive. A terabyte-for-terabyte technology refresh is no longer a requirement where Thin Conversion can be applied.

Stay Thin with Thin Persistence

Any application that writes large portions of data, and then quickly deletes it is referred to as a transient data application. This data, once deleted, persists as capacity-squatting

ghosts in a thin volume. *Thin Persistence* addresses the need to groom out these pockets of unused capacity.

As with Thin Conversion, 3PAR's ASIC-based Thin Persistence can find and reclaim capacity – but in this case, this is not part of a data migration. It is a grooming of a production array of primary storage. While this might raise worries about performance degradation, the ASIC-based utility is able to operate at wire speeds with minimal impact to production workloads on the InServ.

Because of the fine granularity of 3PAR's software architecture, phone-home data on InServ arrays indicates that 3PAR customers average a 60% reduction in required capacity due to Thin Provisioning. *Thin Persistence*, can nominally increase this efficiency by another 8-to-12%. This is a tidiness discipline that enhances utilization. It is like minimizing your carbon footprint, but instead of saving the planet you are saving money – both in capacity that does not have to be bought and in the energy costs that can be avoided.

Thin Copy Reclamation

This extends the stay thin concept beyond primary storage to copies. Given the number of copies that are created and deleted, this is the completer initiative to curbing capacity sprawl.

A discussion of *thin* would not be complete without a mention of *compression* and *deduplication*. 3PAR's thin functionalities are a lifestyle. Compression and deduplication schemes work like a corset. Like wearing a corset, the software that supports these schemes becomes part of the status quo, for it is needed in order to expand the data so it can be used. Compression is great for sequential data streams like audio and video files. It is fine for datasets being downloaded for intensive use. Deduplication is very useful for secondary copies of data kept for recovery or for anticipated re-use. It is less useful for the opportunistic transactions of organizational and business use – data that is kept on primary storage. 3PAR's focus on 'thin' leadership mirrors its focus on primary storage applications, for which thin technologies are the only viable capacity reduction alternative in the market today.

Building Out an Ecosystem

The benefits of the 3PAR's thin philosophy can be further leveraged with the cooperation of information managers like file systems. File systems can also benefit if they have fuller coordination with array operating systems. A recent 3PAR-Symantec partnership has produced an API for file-system-to-array communication, which shares substantial commonality with the soon-to-be-ratified T10 SCSI thin provisioning standard. Both companies are promulgating the API.

3PAR has introduced a fourth product, 3PAR Thin Reclamation for *VERITAS Storage Foundation*, which provides intelligent re-thinking of Symantec file systems via the thin API. It notifies the 3PAR array of freed blocks via a standard SCSI command. The process is automatic and involves no actions by storage or server administrators. It works with VERITAS Storage Foundation version 5 MP3 and later. This institutionalizes the philosophy of thin. Such a philosophy is desperately needed in data storage, where sprawl has been the default norm.

The goal of the partnership going forward is to create an ecosystem where *True Thin* is built-in and completely automatic. Symantec (Veritas) has partnered with Hitachi Data Systems and others to promulgate use of this API. 3PAR has pressed its operating system partners to consider both the Thin API as well as the T10 standard. This expands the domain of tight collaboration between file systems and storage arrays, and adds more efficiencies to an always-troublesome link between domains. This is a good thing

Conclusion

3PAR's extensions of its Thin Provisioning turn a function into a lifestyle. If you want to get the full economic benefits of thin provisioning with the least overhead, these new 3PAR functions give you the ability to build Thin into your information access lifestyle.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Anne MacFarland is Director of Information Solutions for The Clipper Group. Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 128. (Please dial “128” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.