



NETAPP WHITE PAPER

NETAPP THIN PROVISIONING: BETTER FOR BUSINESS

Paul Feresten and Quinn Summers, Network Appliance, Inc.
March 2007 | WP-7017-0307

HOW NETAPP THIN PROVISIONING CAN HELP YOU SIMPLIFY, STREAMLINE, AND SAVE

Thin provisioning can help businesses simplify capacity planning, streamline application and new service rollout, and reduce costs. *Storage Magazine* calls it one of this year's must-have technologies.¹ This white paper takes a closer look at the business side of thin provisioning benefits, considering how the technology can help eliminate some of the most common and taxing of storage management challenges.

As an industry leader in thin provisioning, Network Appliance brings its trademark simplicity and reliability to the technology, making it easy to set up, nondisruptive to users, and safe and predictable for use in enterprise environments. The paper provides an overview of the technology and where it is most effectively applied, descriptions of actual deployments, and a summary of the business benefits your organization can expect to achieve.

¹ "iSCSI SANs, hardware-based tape encryption, high-capacity disk drives, virtualization and thin provisioning will be must-have technologies next year." Excerpted from *Storage Magazine* on "Hot technologies for 2007," Alan Radding, Dec 2006.

TABLE OF CONTENTS

- 1 INTRODUCTION..... 3**
- 2 THE CHALLENGE OF STORAGE PROVISIONING 4**
- 3 THIN PROVISIONING: WHAT IT IS 5**
- 4 WHY THIN PROVISIONING MAKES FINANCIAL SENSE..... 6**
- 5 NETAPP'S THIN PROVISIONING ADVANTAGE..... 7**
 - THE FOUNDATION: NETAPP FLEXVOL® TECHNOLOGY 7
 - TURNING THIN PROVISIONING ON 7
 - BETTER CONTROL THROUGH POLICY-BASED SPACE MANAGEMENT..... 8
 - BUSINESS RESULTS THAT MATTER 8
- 6 PUTTING THIN PROVISIONING TO WORK:
ENVIRONMENTS THAT WILL BENEFIT MOST 9**
 - THIN PROVISIONING OF HOME DIRECTORIES 9
 - THIN PROVISIONING FOR DISK-TO-DISK BACKUP 10
 - STORAGE MANAGEMENT AND OPTIMIZATION AT CADENCE DESIGN SYSTEMS, INC. 10
 - SAP OUTSOURCING SERVICES AT SIEMENS BUSINESS SERVICES 10
 - BACKBONE FOR VIRTUALIZATION ARCHITECTURE AT A HEALTH-CARE RESEARCH COLLABORATIVE 10
- 7 SUMMARY 11**

1 INTRODUCTION

Wasted capacity, complex and protracted provisioning processes, rapid growth that exacerbates storage allocation issues—if this describes your storage environment, you should be considering thin provisioning. A well-implemented thin provisioning plan can help you save hundreds of thousands of dollars (maybe more), provision new capacity in seconds rather than days or weeks, add new users and applications without adding storage, and even boost your competitiveness by delivering more services without increasing costs.

Why doesn't everyone use it? For one thing, not every storage vendor offers the technology. Because of that, there still exists some measure of fear, uncertainty, and doubt about the maturity and efficacy of thin provisioning in enterprise environments. Also, while some businesses have been taking advantage of thin provisioning for quite some time, not all implementations are created equal. Some schemes can be difficult to understand, time consuming to set up, and expensive to manage. Successful implementation therefore relies on an approach in which thin provisioning technology is integrated seamlessly within the storage system architecture.

Industry trends towards adoption of grid architecture and virtualized server environments have been driving the need for more dynamic and flexible storage provisioning. As these technologies enable the concept of utility computing, thin provisioning supports the concept by providing storage-on-demand functionality to ensure alignment of server and storage utilization.

This whitepaper provides an overview of thin provisioning, with an emphasis on how the NetApp integrated implementation ensures both management simplicity and reliability in demanding enterprise environments. Real-world deployments in production environments offer glimpses into the varied application settings where the technology delivers impressive time and dollar savings.

To best leverage thin provisioning in your environment, you must qualify your vendor, scrutinize the technology implementation, and apply it in environments where you'll get the most payback—that is, where you have large server and application deployments combined with rapid change and growth.

“Thin provisioning is extremely valuable and every storage system should support it. It is a form of virtualization that allows you to utilize and manage your physical storage capacity far more efficiently than traditional provisioning.”

**– Tony Asaro
Enterprise Strategy Group**

2 THE CHALLENGE OF STORAGE PROVISIONING

By now most everyone understands all too well the challenges of managing storage resources in a direct-attached storage (DAS) environment—one server can be bogged down with inadequate capacity while storage on a server two feet away sits idle. In that environment, administrators work endlessly to move and match storage resources to applications and user needs. The problem is compounded with every addition of a new user group, application, or server.

Most people would also agree that when compared to a DAS implementation, a network-based storage model improves business productivity and reduces total cost of ownership (TCO). But while network storage implementations improve disk asset utilization, industry estimates continue to show that disk storage utilization rates average between 25% and 40%. That means 60% to 75% of all storage capacity that is being powered goes unused. Not only is such a low utilization rate a waste of disk storage, it is a waste of power and physical space.

In most systems, storage administrators allocate and dedicate storage space to a particular volume or LUN at the time of its creation. This creates two significant administrative challenges: (1) once a volume is created on physical storage, its size is extremely difficult to change, and (2) once storage is allocated to a particular application, it is not available for another use. Because it is difficult to predict actual storage requirements, application administrators typically request much more space than they think they will need to protect themselves in case they need more storage down the line. This common practice guarantees overallocation. This is particularly important in environments supporting large numbers of business users with diverse application and capacity needs. In those settings, it is extremely difficult to predict exact usage requirements up front.

An optimal solution would provide the flexibility to provision storage without advance knowledge of exact application usage requirements. Additionally, there is the requirement to reduce or eliminate “stranded” storage by allocation of any unused capacity to applications with the most need, and to do this on an “on-demand” basis.

The increasing need to adapt to rapid growth and organizational change has driven the adoption of various virtualization technologies that allow more flexible redistribution of physical assets. Even with the most advanced understanding of these virtualization technologies, however, storage administrators can still face provisioning challenges that have the potential to negatively affect business. Poor asset utilization overspending on storage, power, and space; slow deployment of new applications; and limited growth of new accounts are just a few of the costly consequences. As a result, storage administrators and business users still express the following concerns:

- ***How can I simplify storage-capacity planning?***
- ***How do I reduce the cost of acquiring storage?***
- ***How can I respond to user requests for new capacity when my budget is limited?***
- ***How can I reduce time to market and improve competitiveness?***
- ***How can I improve service levels without increasing cost?***

The requirement for faster, simpler provisioning becomes increasingly critical with the size and complexity of the IT infrastructure. Traditional methods of allocating capacity are too unwieldy for enterprises moving to grid computing architectures or delivering Web hosting and Internet services that require allocation of storage across a vast base of users and clients with diverse and hard-to-predict capacity demands. What is clearly needed is a more dynamic storage provisioning model that enables maximum utilization of storage resources without placing an undue and unrealistic burden on storage administrators.

3 THIN PROVISIONING: WHAT IT IS

The first step in the implementation of any storage system is the allocation of capacity to servers and applications. Thin provisioning separates the logical representation of storage from the underlying physical disk arrays, making it possible to allocate more storage capacity to applications than is physically installed.

With traditional storage provisioning, disk capacity is allocated to applications regardless of how much data is actually being written. The problem is that early in a deployment, space requirements are often not known. To make sure that their applications will have adequate capacity, users often present overly high estimates of capacity requirements, and as a result, administrators are forced to overprovision the larger storage infrastructure.

With the ability to do thin provisioning, NetApp storage systems make it possible to oversubscribe free space and adapt rapidly to the changing needs of the enterprise. This is essentially an allocate-on-demand model that allows allocation of storage to application needs, whether or not all of the provisioned capacity is actually installed up front. Because you can present more storage space to the hosts or servers connecting to the storage system than is actually available, storage purchases can be deferred until real application capacity thresholds are realized. Because disk storage prices decline significantly over time, this can result in dramatic cost savings. It also allows businesses to take advantage of the latest disk technology.

To illustrate the traditional approach, Figure 1 shows a 500GB volume allocated with only 100GB of actual data, while the other 400GB allocated has no actual data stored in it. This unused capacity is assigned to a business application, even though the application may not need all 500GB until a later time. That allocated but unused 400GB of physical disk capacity is guaranteed as available but temporarily wasted, tying up capital dollars that could be well spent to meet other business needs or held in reserve to satisfy budget objectives. Furthermore, even though all of the storage capacity is eventually used, it may be years before that happens—and as a result, the business could well miss out on reduced per-gigabyte prices and the latest storage-technology advances.

With thin provisioning, the storage provisioning process starts out the same. Using the same example, the storage administrator provisions 500GB to the business application but uses only 100GB of actual data. The difference with thin provisioning is that the unused 400GB remains available to other applications. This approach allows the application to grow transparently, while ensuring that capacity and money are not wasted. Thin provisioning is really “just-in-time” (JIT) storage, because physical storage is fully allocated only when it is truly needed by the application. The rest of the storage remains in the “free pool,” to be used on an as-needed basis. Storage administrators can set thresholds that trigger alerts when more disks should be added to the pool.



Figure 1) NetApp thin provisioning (just-in-time storage).

4 WHY THIN PROVISIONING MAKES FINANCIAL SENSE

The financial benefits of the thin provisioning concept are evident in a variety of business sectors. In the telecommunications industry, for example, companies have long utilized thin provisioning models. Recognizing that it would be economically impractical to deploy worst-case bandwidth, telcos design their systems to match usage patterns. Rather than building out an infrastructure that can support the unlikely scenario of every subscriber trying to make a call at the same time, telcos configure systems to accommodate the more likely peaks in usage. Empirical evidence supports such theories and has borne out the effectiveness of the model.

The same economic principle applies to the storage industry; capacity demands vary by application type, organizational dynamics, and variable usage patterns. To understand why this principle makes sense, consider the following:

- IT managers and storage administrators routinely report using just 25% to 40% of their total disk capacity. As illustrated in Figure 2, the opportunity exists to significantly reduce the amount of disk storage that needs to be purchased to support a specific application usage level.
- Storage administrators frequently report that it is difficult to recover “stranded” storage capacity and reallocate it to applications that need additional capacity.
- Case work suggests that although a new application might seem to require a volume of at least 1TB, it is more than likely to be using only half (or less) of the original allocation a year later.
- Allocated disk capacity that is not utilized wastes money that could be invested in alternative infrastructure, services, and support staff. Also, since many software applications are priced based on capacity, additional licensing cost is associated with underutilized storage.
- Wasted disk space results in additional waste in floor space and power in the data center. In the next 18 months, increases in average storage-rack density are expected to drive average power consumption up from 2kW per rack to 30kW per rack.

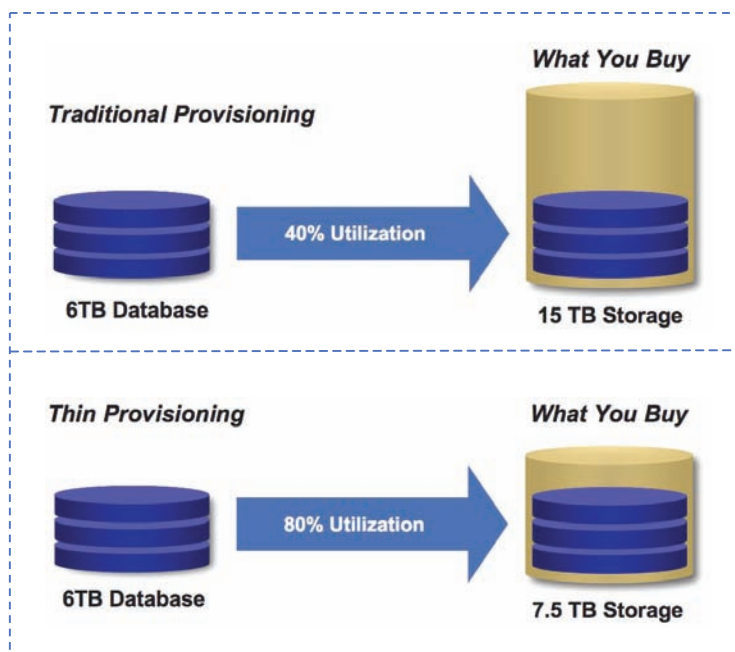


Figure 2) Impact of 2X increase in storage capacity utilization.

Thin provisioning addresses each of these issues to deliver substantial financial benefit, both in terms of capital costs and ongoing operational expenses. And when disk storage is purchased over time on an as-needed or JIT basis, for example, it is possible to leverage the ever-decreasing price per gigabyte and to take maximum advantage of increasing disk densities

5 NETAPP'S THIN PROVISIONING ADVANTAGE

Over the years, NetApp has demonstrated a commitment to simplifying enterprise data management by reducing both storage cost and administrative complexity. A key element in this commitment is the inclusion of storage virtualization capabilities at every level of the storage hierarchy. This includes the data-access level, with support for both file and block virtualization in a unified, easy-to-manage storage architecture. The result is that you can focus your attention on business productivity and growth rather than on the management of complex hardware configurations and operating procedures.

With the introduction of NetApp Data ONTAP® 7G technology, NetApp extended this advantage to areas of volume and LUN provisioning to provide unprecedented levels of simplicity, data management flexibility, and capacity utilization.

THE FOUNDATION: NETAPP FLEXVOL® TECHNOLOGY

With Data ONTAP 7G, NetApp introduced a breakthrough technology that greatly simplifies the provisioning of storage. A FlexVol volume, or flexible volume, is a logical data container that can be managed and moved independently from the underlying physical storage. NetApp FlexVol technology decouples the direct connection between volumes and their associated physical disks to vastly increase flexibility and storage efficiency. An entity termed an “aggregate” provides the connection between the logical flexible volume and the underlying physical storage. By making all disks available to all data sets through a common pool of storage, NetApp FlexVol technology opens up untapped resources to maximize both performance and capacity utilization. Because FlexVol volumes can access all available physical storage in the system, dramatic storage utilization improvements are possible. In a recent performance and usability study by VeriTest, it was determined that an OLTP database application deployed on a competitive solution required 39% more disk capacity than an equivalent NetApp FlexVol implementation².

TURNING THIN PROVISIONING ON

FlexVol technology enables superior storage utilization in its basic configuration, and this advantage can be extended through a simple setup procedure that specifies whether a particular volume is guaranteed a specific amount of physical space from the disk aggregate. If the space guarantee is set to “none,” thin provisioning is automatically set, and volumes can be sized at the full amount of virtual space required. Figure 3 illustrates this concept, showing how the total space provisioned in all flexible volumes exceeds the amount of physical disk space.




Figure 3) NetApp thin provisioning.

² VeriTest, “Network Appliance FAS3070 and EMC Clariion CX3-80: Comparison of Performance and Usability: November, 2006”

To further enhance the inherent advantages of FlexVol technology, the ability to deploy a storage-on-demand environment is facilitated through the following additional capabilities:

- Simple setup. In just seconds, you can specify whether individual volumes should be thin provisioned.
- No-downtime resizing. You can grow or shrink volumes nondisruptively, without taking them offline.
- Fast capacity redistribution. Storage capacity can be resized on the fly to accommodate user, application, and business requirements.
- Easy space reclamation. Using NetApp policy-based tools (for example, volume autosizing and auto-Snapshot delete), it is possible to easily reclaim space and take advantage of thin provisioning across more applications.

BETTER CONTROL THROUGH POLICY-BASED SPACE MANAGEMENT

NetApp Data ONTAP software combines the JIT storage advantages of thin provisioning with policy-based control for optimal application and service-level flexibility. By setting thresholds or policies under which storage space can be reclaimed from deleted Snapshot copies that are no longer needed, the storage administrator can effectively leverage storage capacity without jeopardizing application performance or the availability of storage to critical programs. Because policy-enabled space reclamation and volume resizing can be set in advance, the risk of running out of space is greatly reduced. This alleviates the concern that thin provisioning could result in a shortage of physical space for mission-critical applications, because policy-based space reservations can be established on the basis of service-level requirements.

BUSINESS RESULTS THAT MATTER

As a result of implementing thin provisioning on a NetApp system, you can expect the following:

- Improved storage utilization. It's not unusual for businesses to improve storage utilization dramatically. Anecdotal evidence suggests that up to 2X improvement is not uncommon for businesses leveraging a combination of NetApp Snapshot™, FlexVol, and thin provisioning technologies.
- Reduced capital expenses. By leveraging thin provisioning, you'll be able to purchase less storage up front, as well as over the lifecycle of your systems.
- Reduced operating expenses. Because you'll use less disk capacity, you'll also use less data center space and power. Another benefit will come in the form of lower software license fees for products with capacity-based licensing. And, perhaps most significantly, with greatly simplified capacity planning and provisioning, your administrative dollars can be reduced or applied to higher-level business tasks.
- Support for more applications and users. Because thin provisioning enables a decoupling of new application deployment and the addition of new users from the asset procurement process, it is possible to react more quickly to organizational growth.
- Improved service levels and product competitiveness. The storage-on-demand model allows you to react quickly to competitive threats because products and services may be enhanced and expanded in scope without significant change in the supporting IT infrastructure. And thin provisioning can be combined with a policy-based hierarchy of service-level priorities to ensure that storage space is optimized for your business.

6 PUTTING THIN PROVISIONING TO WORK: ENVIRONMENTS THAT WILL BENEFIT MOST

So where does it make the most sense to take advantage of thin provisioning technology? In general, thin provisioning delivers maximum benefits in settings characterized by complexity, large numbers of servers and applications, service provider environments where storage must be allocated across a vast number (tens or hundreds of thousands) of users/clients, as well as virtualized server and grid computing environments where dynamic storage provisioning is essential. The technology would be particularly well suited in provisioning storage for file shares and home directories, Web hosting and Internet services environments (including photo storage and e-mail accounts), information services, and disk-to-disk (D2D) backup.

The following examples help illustrate how thin provisioning can yield benefits in specific application environments.

THIN PROVISIONING OF HOME DIRECTORIES

Many organizations have found that capacity requirements for home directories are quite diverse across the user community. Because the number of home directories can grow significantly as organizations grow, this is a natural place where thin provisioning makes sense.

To illustrate this, ESG Lab created a test scenario that simulates an IT organization providing users with 400GB of home-directory space on a network-attached shared drive. ESG Lab used both the traditional method of provisioning volumes and the NetApp thin provisioning implementation for each of the users' home directories. ESG Lab configured the home directories at 400GB, each using a pool of 2TB actual disk capacity.

As a result, ESG Lab was able to set up 10 times more user home directories with thin provisioning than with the traditional method. Using traditional provisioning methods would have required 20TB of usable capacity to support 50 home directories, compared to just 2TB of capacity required when using thin provisioning.³

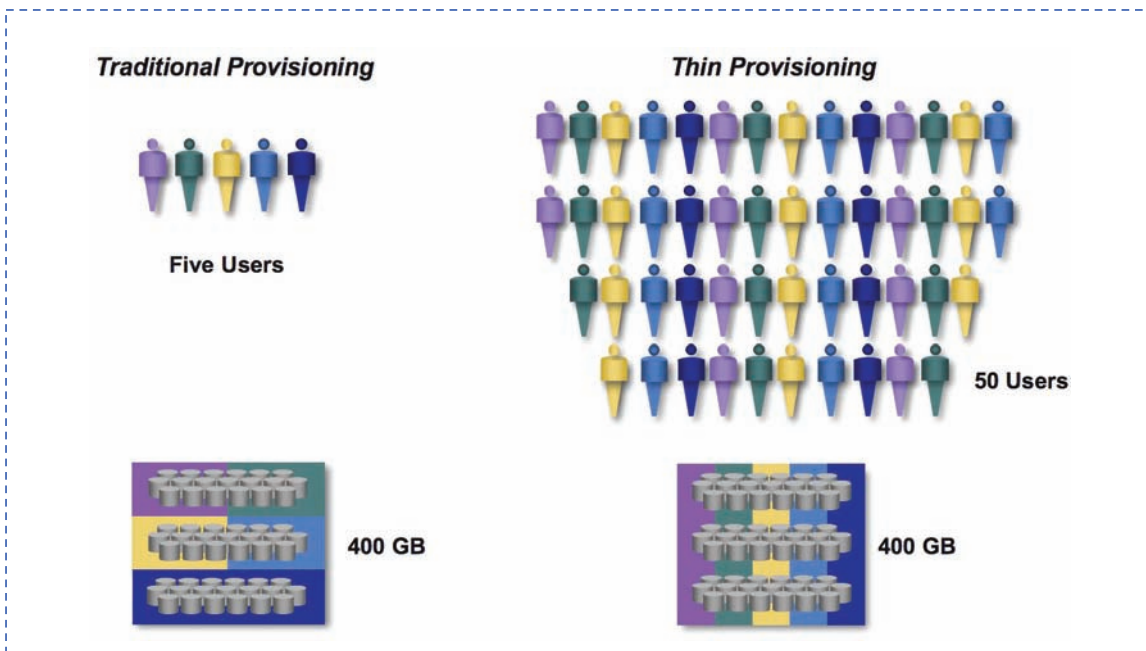


Figure 3) NetApp thin provisioning scalability.

³ Network Appliance. "Data ONTAP 7G and V-Series Virtualization, a Validation Study by ESG Lab," September, 2005.

THIN PROVISIONING FOR DISK-TO-DISK BACKUP

When it comes to disk-to-disk backup, provisioning can be even more complicated than allocating space for primary storage requirements. You not only need an estimate of the growth in primary storage usage, you also need to know the rate of change in each volume. One innovative NetApp customer, a large company that sells backup services to internal customers, overcame such uncertainties by applying thin provisioning on secondary storage for disk-to-disk backup. Over the course of a year, the customer increased primary storage capacity from 500TB to 900TB without needing any additional secondary storage.

The company's data center had been continuously pinched for floor space, power, and cooling, so this savings represents a significant benefit beyond the savings in capital outlays. The company has been able to delay the purchase of any new secondary storage for a year as a result of thin provisioning and the increased efficiency it provides. Storage utilization increased from less than 40% (due to mostly underutilized volumes) to nearly 70%.

This company also uses thin provisioning for home directories on the production side of the house. According to company practice, each of 4,500 users has up to 1GB of network file storage as a home directory, which would require 4.5TB of total storage. Using thin provisioning, the company meets this requirement with only 600GB of actual disk storage.

STORAGE MANAGEMENT AND OPTIMIZATION AT CADENCE DESIGN SYSTEMS, INC.

At Cadence Design Systems in San Jose, California, NetApp thin provisioning, in conjunction with NetApp V-Series systems, provided an effective way to consolidate storage on IBM ESS storage arrays without the traditional need to dedicate disks to users for their file shares. According to Mike Forman, director of North American IT operations, "This gives us the ability to borrow unused capacity and give it to someone else without the customer knowing."⁴ The result has been an increased utilization of existing storage systems without the need to purchase additional storage whenever clients request it.

SAP OUTSOURCING SERVICES AT SIEMENS BUSINESS SERVICES

SAP outsourcing is one of the most important business lines at Siemens Business Services (SBS). In this arena, SBS has been taking advantage of the FlexFrame structure, a joint development of Fujitsu Siemens Computers, Network Appliance, and SAP. SBS's implementation of the FlexFrame adaptive IT infrastructure solution leverages Linux® blade servers and a 10TB NetApp storage system.

"FlexFrame enables us to offer our SAP customers an infrastructure with flexible provisioning of server and storage capacity," says Robert Pikart, manager of the SBS Outsourcing Unit. "In this application, the virtualization ability of NetApp storage solutions can be utilized especially well. Storage can be allocated in a flexible manner when it is needed, for example, to increase performance. The possibility to overprovision is particularly practical—that is, we can allocate an arbitrarily high amount of capacity without actually having to reserve fixed hard disk space. This is an extremely convenient feature that provides significant leeway in resource utilization."

BACKBONE FOR VIRTUALIZATION ARCHITECTURE AT A HEALTH-CARE RESEARCH COLLABORATIVE

Based on NetApp advantages in virtualized data management (including thin provisioning capability), replication technologies, reliability, and ease-of-management, this international research collaborative selected a high-availability Network Appliance unified storage solution. Researchers now utilize the NetApp solution to access general operations information (such as CIFS fileshares), as well as all clinical specimen data. The consolidated storage solution facilitates global access and data sharing, improved version control, and increased security and protection of information assets.

⁴ eWeek, December 13, 2004, Ziff Davis Media, Inc.

Beyond providing centralized data storage—complete with disaster recovery—the NetApp solution serves as the backbone of the organization's virtualization architecture. In less than a year, the IT team has effectively made the transition from a 30-physical-servers infrastructure to a VMware-based production environment that includes some 125 virtual servers.

After rigorously evaluating other vendors' technology, the organization concluded that NetApp thin provisioning offers unique benefits for a virtualized infrastructure. For example, although in-production capacity is 50TB, actual capacity utilized is just 3TB. When applications start to approach thresholds, adding capacity is as easy as plugging a new shelf into the SAN and running a simple command to extend the aggregate. With this technology, the organization conserves both storage and administrative resources.

7 SUMMARY

As evidenced by a wealth of real-world deployments and successes, thin provisioning can help businesses simplify capacity planning, streamline application and new-service rollout, and reduce costs. This new model of storage provisioning can help eliminate some of the most common and taxing capacity management challenges and represents one of the most significant paradigm shifts to be widely adopted by leading storage vendors.

As an industry leader in thin provisioning, Network Appliance brings its trademark simplicity and reliability to the technology, making it easy to set up, nondisruptive to deploy, and safe and predictable for use in enterprise environments. Although results are always implementation dependent, customers who take advantage of NetApp thin provisioning can expect to achieve major improvements in storage utilization, reduce storage costs, simplify capacity planning, and more quickly deploy and scale applications.

Integrated thin provisioning technology is a key component of our continuing commitment to simplification of enterprise data management. At NetApp, we take this responsibility seriously and we will continue to devote our resources and talents to the development of data management and storage technologies that will help you stay competitive and meet your business challenges.

“The possibility to overprovision is particularly practical—that is, we can allocate an arbitrarily high amount of capacity without actually having to reserve fixed hard disk space. This is an extremely convenient feature that provides significant leeway in resource utilization.”

**– Robert Pikart, Manager
SBS Outsourcing Unit**

