

Thin Provisioning

Focus on Compellent Dynamic Capacity

By Tony Asaro
Senior Analyst

October 2005

Thin Provisioning

ESG believes that one of the more important but often misunderstood advancements within data storage is thin provisioning. This report focuses on Dynamic Capacity, Compellent's implementation of thin provisioning.

Various operating systems support a feature referred to as dynamic volume expansion, which is sometimes confused with thin provisioning. Dynamic volume expansion allows customers to increase the size of their volumes as an online process. For example, if the size of the original volume is 100 GB but needs to be increased to 200 GB, the system administrator can use an online process to do this, using tools provided by the operating system. This is a valuable capability that addresses the problem of data growth by enabling system administrators to add more capacity transparently without having to re-boot servers. However, dynamic volume expansion does not support shrinking the size of a volume. System administrators typically plan ahead for growth and provision more storage than is needed for various applications. An application may take up 100 GB of actual data but, based on an analysis by or a request from a specific business unit or department, the system administrator preemptively creates a volume of 500 GB of capacity, allowing for data growth. When the volume is created, this 500 GB is allocated and dedicated to that application and no other application can use it. However, in many cases, that full 500 GB is never used and is essentially wasted. This is a major problem when managing storage capacity and is often referred to as stranded storage.

The inefficiencies of traditional storage provisioning can negatively impact capital costs and storage administration resources. The most obvious inefficiency is the amount of storage that goes unused and therefore increases the total cost of ownership. Additionally, since this allocated but unused storage capacity cannot be reclaimed for other applications, customers have to buy more storage capacity as their environments grow, increasing cost even further. At some point customers may actually be required to buy a completely new storage system in addition to the one they have in place. ESG conducted a survey in June 2005 of 20 customers who focused on the limitations of traditional storage provisioning methods. The highlights of that survey included:

- 54 percent of the customers were aware that they had stranded and unused storage capacity due to inefficient provisioning methods.
- 55 percent of these customers had between 31 and 50 percent of stranded and unused storage. For example, if they provisioned 10 TB of storage capacity, then 3.1 TB to 5 TB of that was stranded.
- 45 percent of the total users surveyed had to buy an additional storage system because they could not utilize their stranded storage. This means that, although these customers had unused storage capacity they've already paid for, they needed to buy a new storage system to meet the needs of their business.
- 30 percent of users were planning to buy an additional storage system in the next twelve months because they could not access their stranded storage. Combining the 45 percent of customers who purchased an additional storage system with the 30 percent who were planning to buy, then the number of users who needed to purchase additional storage to keep their business running increases to 75 percent.
- 80 percent of users felt that storage provisioning was a time and resource drain on their IT organizations.

The Value of Thin Provisioning

There are six ways that the Compellent Dynamic Capacity thin provisioning feature can save customers money:

1. Less storage is required initially when purchasing a new storage system.
2. Since there will be no stranded storage capacity, less storage is required over the life of the storage system.
3. Additional storage systems will not be required based on having stranded storage.
4. More servers per storage system provide greater levels of consolidation.
5. The time and resource required to perform storage provisioning tasks is reduced.
6. The performance of various data management services is improved.

Less storage is required initially when purchasing a new storage system. Using traditional storage provisioning methods, customers have no choice but to buy additional capacity up front. However, with the Compellent Dynamic Capacity, this is not required since capacity does not have to be pre-allocated. Customers can buy what they need with a minimal amount of spare capacity for short-term growth. When a pre-defined threshold is reached, more disks can be added to the storage pool. The amount of money saved can range, depending on the environment. For example, customers can start out with 25 or 50 percent of the capacity they normally would purchase with storage systems that don’t support thin provisioning. A customer who previously had to buy 10 TB of capacity because of the limitations of traditional provisioning can instead acquire 2.5 TB or 5 TB with the more efficient thin provisioning method. The point is simple – Dynamic Capacity provides a method for “just in time” provisioning that allows system administrators to be more operationally efficient and lower capital costs.

Since there is no stranded storage capacity, less storage is required over the life of the storage system. It is the experience of ESG that customers have a large amount of capacity that is stranded and unused due to the inefficiencies of traditional provisioning methods. For example, the ESG survey indicated that between 31 and 50 percent of customers’ storage was stranded and unused due to poor provisioning. Over the life of a storage system, thin provisioning can save tens of thousands or even hundreds of thousands of dollars, depending on the amount of capacity required. To quantify this, at \$20 per GB, 5 TB of stranded storage equates to \$100,000. of **wasted cost** to the customer. For 10 TB, the waste becomes \$200,000 and so on. This wasted cost of stranded storage is eliminated with thin provisioning. Additionally, companies can delay the acquisition of disk drives and take advantage of the continuous and inevitable reduction of cost and increase in capacity. The Compellent storage system can support any disk type, allowing customers to use inexpensive and dense storage as needed.

Additional storage systems will not be required based on having stranded storage. There is a finite amount of storage capacity that any storage system can support. When the optimal capacity has been reached, customers may need to implement an additional storage system. Referring back to the ESG rapid survey, 45 percent of total users were forced to buy an additional storage system because they could not utilize their stranded storage while another 30 percent were planning to buy a new storage system. Acquiring a completely new storage system because of stranded storage is extremely wasteful. This is especially true since it is more than likely that a substantial amount of unused and inaccessible capacity already resides within the existing storage system. The additional costs associated with buying a new storage system include the up front price of the system, all of its associated software licenses and maintenance fees, and administration costs.

More servers per storage system provide greater levels of consolidation. As described above, since capacity doesn’t have to be dedicated and fixed on a per volume basis with thin provisioning, customers can create more volumes. As a result, more servers can be attached to a single storage system. Storage consolidation is a major initiative and thin provisioning can enable even greater levels of consolidation.

The time and resource required to perform storage provisioning tasks is reduced. ESG regularly hears from system administrators that one of the most mundane tasks they are required to perform is provisioning storage. Thin provisioning can essentially eliminate the need for storage provisioning. Customers can create large logical volumes without any cost to them, allowing the application to keep growing and growing as needed. ESG has spoken to customers who have implemented thin provisioning and they report that they spend literally no time provisioning storage except when creating new volumes.

The performance of various data management services is improved. When system administrators create a remote mirror of specific volumes, the data from the volume is not the only thing being copied and transferred. The stranded unused capacity is actually copied and transferred, as well. Processors, cache, and bandwidth are used to transfer an entire volume, consisting of both data and possibly large numbers of empty blocks, over expensive WAN links. This limitation of traditional storage systems also holds true for other services such as RAID rebuilds and local full volume copies. Thin provisioned volumes only contain real data and the resources needed for remote and local copies as well as RAID rebuilds are always used for real data and not empty blocks.

Thin provisioning is “just in time capacity” that essentially eliminates allocated but unused or stranded storage. It also greatly simplifies storage provisioning tasks, reducing administration costs. Customers should look at storage systems that support thin provisioning differently than those that do not when evaluating the initial capacity configuration, budgeting, capacity growth planning over time, and the total cost of ownership. **The economic impact of thin provisioning can be substantial, saving literally thousands, tens of thousands, and potentially hundreds of thousands of dollars, depending on the scope of the environment.**

Thin versus Traditional

There are a number of benefits provided by thin provisioning compared to traditional methods. With thin provisioning, customers will be able to:

- Create more and larger volumes.
- Buy less storage capacity to start with and only add disks as needed.
- Spend less time provisioning storage.
- Potentially support a greater number of servers.
- Experience much better capacity utilization rates.

This section compares traditional versus thin provisioning methods. The first example illustrates how thin provisioning is more efficient by providing a larger pool of shared capacity at the outset for multiple applications, as compared to traditional provisioning methods.

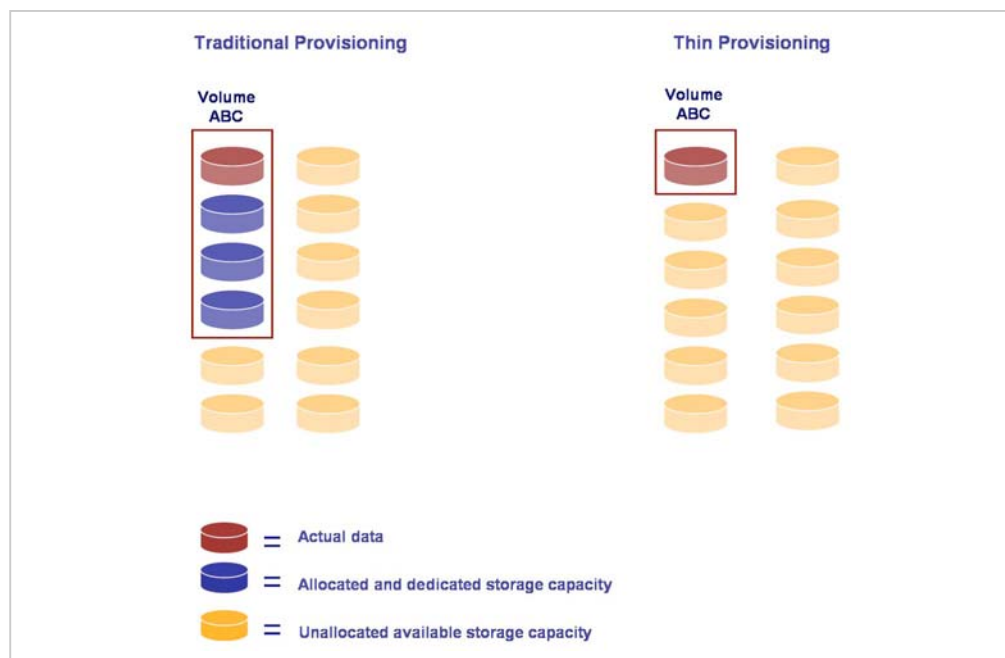
Traditional:

1. The customer's storage system has 3 TB of capacity.
2. The system administrator provisions Volume ABC with 1 TB of capacity.
3. Initially, 200 GB is stored on Volume ABC.
4. The system administrator sees that 2 TB of capacity is left for other applications, although only 200 GB of actual data has been stored.

Thin Provisioning:

1. The customer's storage system has 3 TB of capacity.
2. The system administrator creates Volume ABC with 1 TB of capacity.
3. Initially, 200 GB of data is stored on Volume ABC, as above.
4. The system administrator sees that 2.8 TB of capacity is available.

Figure One: Initial Capacity Requirements




In this example, thin provisioning offers an additional 800 GB of capacity to use for other applications over the traditional method. Depending on the number of volumes on the storage system, the delta in available capacity between traditional and thin provisioning can increase dramatically.

In Figure Two, three volumes are created using traditional provisioning methods. The total usable storage capacity of the storage system is 2 TB.

- The total volume size of Volume 1 is 500 GB, of which 100 GB is actual data and 400 GB is allocated unused capacity.
- The total volume size of Volume 2 is 550 GB, of which 50 GB is actual data and 500 GB is allocated unused capacity.
- The total volume size of Volume 3 is 800 GB, of which 200 GB is actual data and 600 GB is allocated unused capacity.

In total, the storage system stores 350 GB of actual data, 1.5 TB of allocated unused capacity, and only 150 GB of available capacity for other applications.

Figure Two: Traditional Versus Thin Provisioning Allocation



If we look at the same 2 TB storage system with thin provisioning implemented, the differences are quite dramatic. Although the system administrator creates the same size volumes for Volumes 1, 2, and 3, there is no allocated unused capacity. In total, the storage system with thin provisioning stores the same 350 GB of actual data as the other storage system, with 1.65 TB of capacity available for other applications, versus only 150 GB in the traditional storage system. **Thin provisioning effectively increases the amount of available capacity by 11 times.**

Consider the financial impact of having 11 times more available capacity over traditional storage systems, at storage totals of 10 TB, 100 TB, 500 TB, and even 1 PB of storage. As an example, consider a storage system that cost \$20 per GB. With a traditional storage system with 1.5 TB of allocated unused capacity, the cost for total storage would be \$30,000. Using it's the Compellent Dynamic Capacity thin

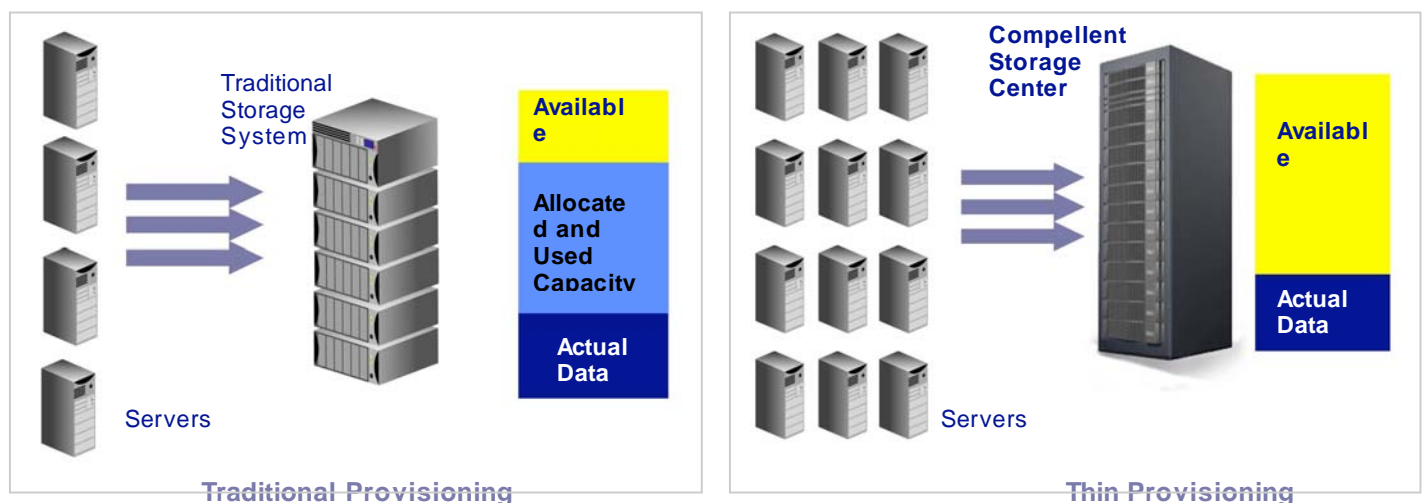
provisioning feature, there is no cost for the allocated but unused capacity. The cost savings on 10 TB of allocated unused capacity would be \$200,000.

Thin provisioning provides customers with a greater pool of storage capacity, allowing system administrators to create more volumes for different applications. As a result, thin provisioning enables customers to allow more servers to share a single storage system.

Traditional

1. The customer's storage system has 3 TB of usable capacity.
2. The system administrator uses the traditional method to provision storage. He or she creates a 500 GB volume for each of the application servers. Each volume contains 200 GB of actual data.
3. The storage system in total has 1 TB of available capacity left, 800 GB of actual data, and 1.2 TB of allocated unused data.
4. The system administrator is limited to supporting only a total of four servers in this example.

Figure Four: Server Support



Thin Provisioning

Using thin provisioning in the same example, the system administrator can support a far greater number of servers.

1. The customer's storage system has 3 TB of usable capacity, as above.
2. The system administrator uses the thin provisioning method. He or she creates a 500 GB volume for each of the application servers. Each volume contains 200 GB of actual data.
3. While the storage system has 1 TB of available capacity left, 2 TB of actual data is stored on the storage system.
4. The system administrator is able to connect 12 servers to the storage systems and has plenty of capacity to add more if required.

Another major benefit provided by thin provisioning is the ability for customers to buy less storage initially and over the life of the storage system. Traditional provisioning requires system administrators to calculate the volume size of their applications based on the actual data, anticipated growth of each volume over time, plus additional capacity as a buffer just in case the data grows faster than expected.

Thin provisioning requires system administrators to buy a pool of storage large enough to support the actual data and capacity based on short-term requirements – three months, six months, or even a year. As the storage pool is used and a predetermined threshold is crossed, the system administrator is alerted and can add more capacity as needed transparently and online.

In Figure Five, the amount of actual data and available capacity is the same with both methods but with thin provisioning no additional storage is needed for the allocated unused storage capacity. ESG has seen customers reduce their initial capacity requirements to as low as 25 percent of what they originally would have needed to acquire with traditional provisioning.

Figure Five: Capacity Requirements



Instead of anticipating the behavior of users and applications with a predetermined amount of capacity, system administrators instead add capacity as needed. This results in requiring less storage initially and for the life of the storage system.

The critics of thin provisioning are, not surprisingly, the storage vendors that do not support it. Here are some of the objections and ESG’s answers to them:

- Objection:** Thin provisioning supports bad practices of over-allocation.
- ESG’s View:** This is untrue, since system administrators can still enforce hard limits using thin provisioning. One of the major benefits of thin provisioning is to prevent over-allocation.

- Objection:** The storage administrator is essentially deceiving the application.
- ESG’s View:** The application is being “deceived.” However, we see no issue with this. It is not uncommon for various technologies to spoof applications.

- Objection:** Thin provisioning requires the system administrator to second-guess the application users’ requirements.
- ESG’s View:** This is fundamentally wrong, since thin provisioning takes away the guesswork, which is required by traditional provisioning.

- Objection:** Additional management must be provided to know when to increase actual allocation and provision more storage.
- ESG’s View:** System administrators are alerted when to add more storage capacity. However, this is a simple process. Thin provisioning actually reduces management, by simplifying the task of storage provisioning.

Compellent Storage Center

The Compellent Storage Center is one of the few storage systems that supports thin provisioning through its Dynamic Capacity feature. The Compellent Dynamic Capacity software has both a positive economic impact and simplifies the mundane task of provisioning storage. ESG believes that the Compellent Storage Center is an excellent SAN storage system due to its architecture, scalability, ease of use, features, and functions.

The Compellent Storage Center is a next generation SAN-based storage system with a modular hardware architecture. Customers can scale Storage Center by adding capacity and functionality as the demands of their environments change and still retain only a single system to support. The economics of a modular network storage system are compelling, with the potential to significantly reduce capital and operational costs. With a modular network storage system, the customer only has to add the capacity or functionality when necessary which costs far less than acquiring a whole new system, with its associated software and additional maintenance charges.

Another major benefit of the Compellent solution is that it allows customers to implement intelligent in-the-box multi-tiered storage, providing the ability to move data transparently and on-line between different tiers of storage (i.e. FC and SATA) within a storage system. The Compellent Storage Center provides this capability through its Data Progression software. System administrators can set up user-definable policies to move data based on particular criteria. For example, Storage Center can be configured to move least accessed data from higher cost, high performance FC drives to less expensive but slower SATA drives. Intelligent multi-tiered storage allows customers to respond to performance, capacity, and access characteristics based on the needs of the applications. Intelligent in-the-box tiering provides better storage utilization and cost savings. What is unique to Compellent is that its Data Progression capability works at the block level, while other SAN-based storage systems work at the volume level. In any storage system volume, there may be frequently accessed blocks and others that are never used. At the volume level, as in the SAN-based systems, the storage system identifies the entire volume as being frequently accessed. At the block level, as supported by Compellent, blocks that are never accessed will be moved to a lower cost tier of storage. The block level approach is potentially far more useful and can have a greater cost impact to customers.

The Compellent Storage Center also supports the following:

- FC and iSCSI interconnects
- Striping data over a large pool of disks for greater application performance
- Highly-efficient Boot From SAN
- Snapshot copies
- Remote mirroring – synchronous and asynchronous
- A wide range of comprehensive reporting tools

The combination of various Compellent features and functions adds value on top of value. The Compellent Dynamic Capacity (thin provisioning), Data Progression (block level intelligent multi-tiered storage), Remote Instant Replay (data replication), and scalable clustered architecture all work in concert, resulting in a powerful, compelling yet practical storage system. The value proposition of Compellent ultimately results in minimizing complexity and providing customers significant amounts of capital savings.

ESG’s View

Thin provisioning ranks at the top of the list of the most important storage system features, along with differential snapshots. The economic benefits and the reduction of storage provisioning tasks are reasons enough to evaluate storage systems that support thin provisioning, giving them a distinct advantage over products that only provide traditional storage provisioning methods. Thin provisioning is just too valuable for customers to ignore and over time as more customers become aware of its benefits, demand will drive support. Today, just a few storage vendors support thin provisioning, including Compellent. They will gain insight and experience while other vendors eventually struggle to catch up.

The Compellent Storage Center is an extremely well conceived and designed next generation SAN storage product. Thin provisioning is just one of the many valuable features that it supports as part of an excellent and advanced storage system. Compellent customers are saving a ton of money by leveraging the capabilities of Storage Center, and Dynamic Capacity is one of the key features they are embracing. There is no greater proof or validation than the fact that Compellent customers are currently saving themselves thousands, tens of thousands, and even hundreds of thousands of dollars.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.