Measuring and Improving Storage Utilization

A best practice approach to data storage management

By Stephen Foskett and Jonathan Lunt

Table of Contents

Introduction

The increasing importance of the economic use of assets has lead to the increasing use of asset utilization statistics. However, this generic term is of little use without careful definition and explanation. Additionally, misunderstandings of utilization metrics in storage discussions can result in flawed cost cases and business justifications. StorageNetworks proposes a new metric, Total Storage Utilization, to reflect all aspects of asset utilization efficiency.

This document defines key storage utilization metrics, describes best practices necessary to increase capacity utilization, and presents a case study illustrating the increase in Total Storage Utilization enabled by StorageNetworks at a large customer in the financial industry.

Storage Metrics

Before utilization can be measured, basic metrics must be defined for storage use. These definitions should be simple enough to apply to the wide variety of data storage configurations available. The following nomenclature will be used in relation to storage utilization metrics:

- **Raw** - Disk space before overhead
- **Usable** - Available logical storage space, after overhead

- **Used** - Disk space in use for a specific purpose

The three metrics are appropriate for describing storage on both hosts and arrays, as illustrated in Table 1.

	Array	Host
Raw	**Array Raw** Physical disk space installed in the array	**Host Raw** Space "claimed" by the host Operating System
Usable	**Array Usable** Space available for use after overhead	**Host Usable** Usable space in a host volume or file system[1]
Used	**Array Used** Usable storage presented to hosts	**Host Used** Logical space containing data

Table 1 - Array and Host Utilization Definition

Figure 1 on page 3 illustrates these storage metrics on a block diagram of a storage array and host. Each of the six metrics must be the same as, or less than, the previous metric.

Storage Metric Definitions

Array Raw

- The total physical disk space installed in a storage system; The sum of the sizes of all installed disks[2]
- Collected using array management tools (e.g. EMC ControlCenter)

Array Usable

- Storage available for use by hosts ("hypervolumes", "splits", "LUNs")
- Collected using array management tools

Array Used

- Usable storage assigned or allocated to hosts
- Determined with array or host tools

Host Raw

- Total storage visible to the host Operating System
- Determined with host-based tools

Host Usable

- Logical storage in a host volume or file system made available for application use
- Determined with host-based tools

Host Used

- Logical space containing actual host data
- Determined with host-based tools

[1] Raw database partitions can also be counted as host usable. These represent a challenge in calculating utilization since visibility from the operating system regarding utilization is limited. Automatic utilization agents cannot determine the usage of these areas.

[2] Note that physical disk sizes are often specified as decimal units rather than binary units. Most other metrics are expressed as binary powers. For example, a 30GB disk is 30*10^9 bytes (30,000,000,000), not 30*2^30 bytes (32,212,254,720)

Storage Metrics

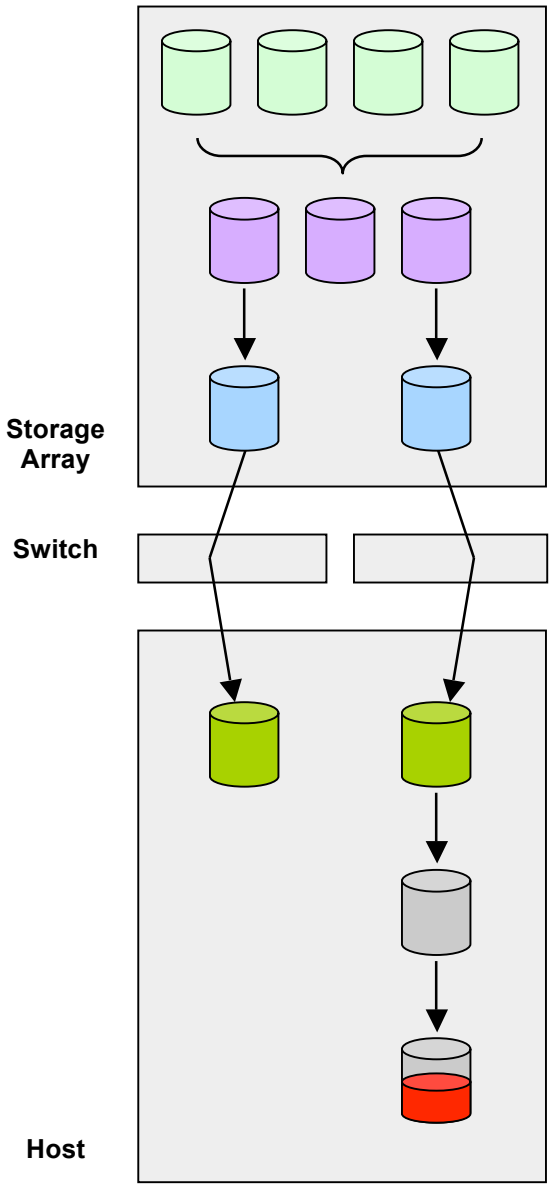

Array Raw - Physical disks in the array. In this example, each is 10GB, resulting in 40GB or **array raw** storage.

Array Usable - The amount of storage available after protection overhead. Here, RAID 5 is being used, resulting in 30GB of **array usable** storage, or 25% overhead.

Array Used - The amount of storage allocated to connectivity ports. Here, only two of the 10GB volumes have been allocated, resulting in 20GB of **array used** storage.

Host Raw - The amount of storage that the host operating system can see. In this example, both of the 10GB volumes are visible to the host, resulting in 20GB of **host raw** storage.

Host Usable - The amount of storage that the host has made available to applications. Here, the Systems Administrator has only chosen to create a file system on one of the volumes, resulting in 10GB of **host usable** storage.

Host Used - The amount of storage actually containing data. Here, only about 50% of the file system contains files, resulting in 5GB of **host used** storage.

End-to-end, 40GB of disk is used for 5GB of storage.

Figure 1: Storage Metric Diagram

Utilization Ratios

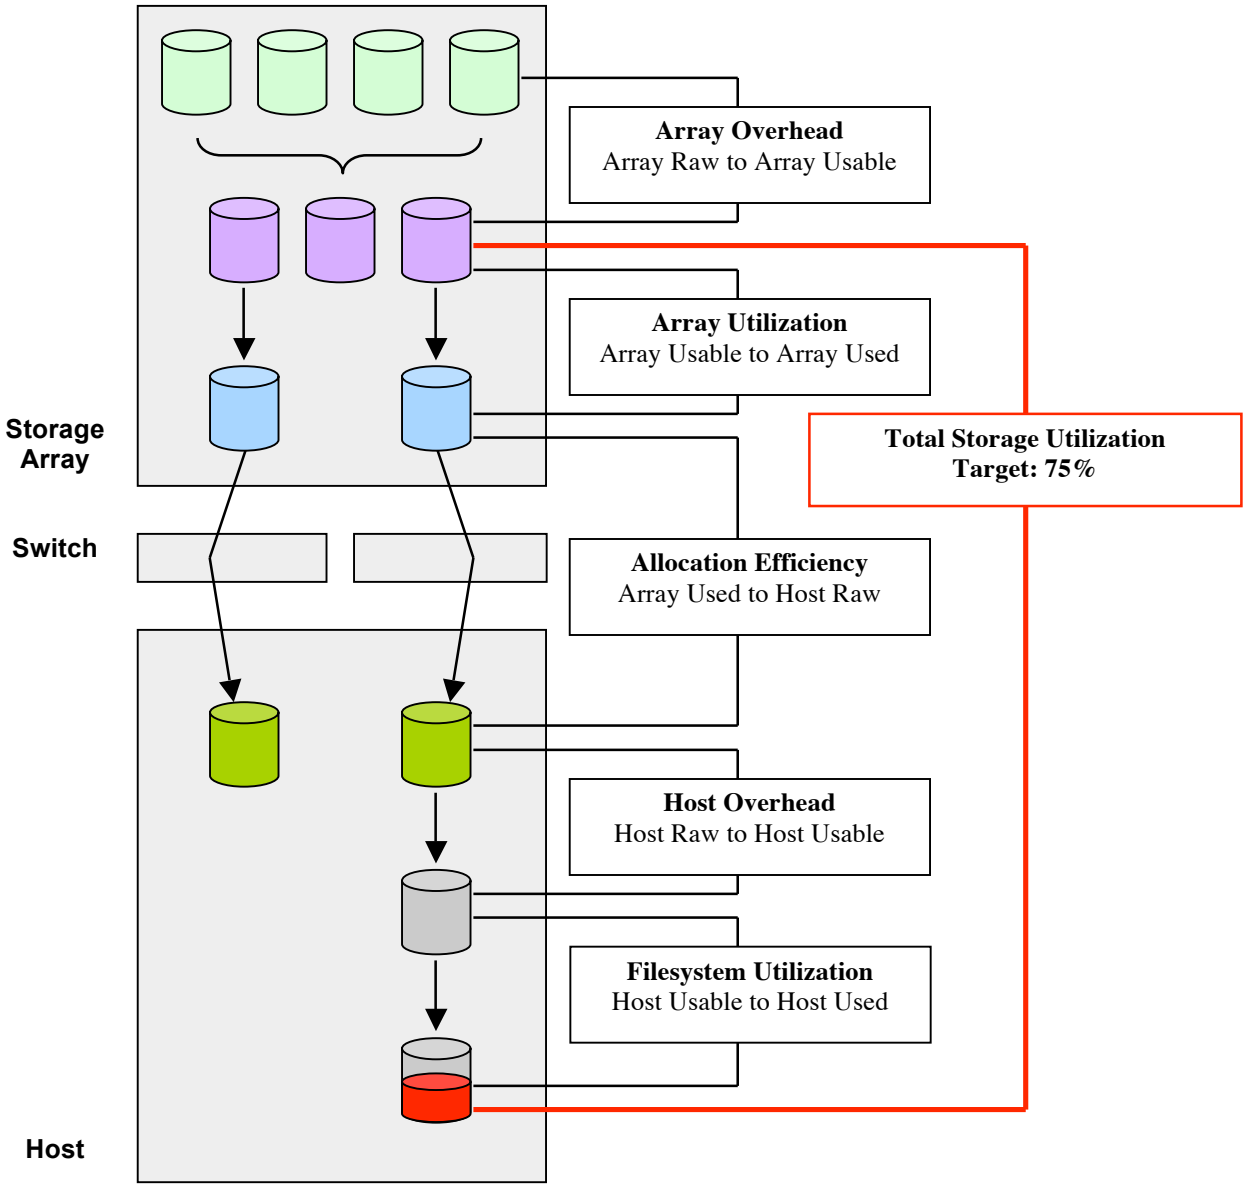

Figure 2: Utilization Metric Diagram

Storage Utilization

Once standard metrics are defined, ratios can be used to describe the storage environment. StorageNetworks has identified a number of key storage utilization ratios, as illustrated in Figure 2 on page 4. These ratios were chosen because each can indicate specific storage management problems in a concise manner.

A utilization ratio is obtained by comparing one metric with another to obtain a percentage. Inefficiencies (deliberate or otherwise) and overheads occur between each of the levels. The cumulative effect of these inefficiencies can result in very poor overall utilization as illustrated on page 11 in the section titled *Utilization Cost Savings*.

The following section describes each utilization ratio, provides reasons for high and low values and indicates achievable real world metrics.

Array Overhead

(Array Raw to Array Usable)

Array Overhead is the percentage of installed storage capacity that is not usable. Dividing Array Usable by Array Raw and subtracting that number from 100% yields the percent of overhead. Overhead here is usually due to the desired level of data protection (e.g. RAID, mirroring) rather than to poor management. These metrics must be collected directly from the storage array.

Table 2 shows the typical overhead associated with various RAID levels.

Raid Level	0	1	0+1	4	5	EMC Raid-S
Typical Overhead	0%	50%	50%	20%	20%	33%

Table 2 - RAID Level Overheads

Low Utilization

High overhead (i.e. low utilization) is normally associated with the following scenarios:

- The use of a high level of data protection (e.g. RAID 1, replication, space drives) resulting in a high protection overhead
- A purchasing policy where fully-configured arrays are purchased and deployed without a requirement for all of the storage, leaving unformatted drives in the array
- Measurement during a deployment process

While not a reflection of poor utilization, many enterprises over-engineer their storage environment, perhaps using RAID 1 where RAID 5 would perform adequately. Substantial cost savings can be achieved by rationalizing protection levels to well-defined "Classes of Service" definitions. These service levels would specify data protection and pass storage costs on to end users.

High Utilization

Low overhead (i.e. high utilization) is typically associated with a mature

environment where all installed capacity has been configured to a required level of protection. Low overhead can also be an indication that data protection is not being performed at the array level. This could be a sign of insufficient protection, or that this protection is being handled by a volume manager at the host level.

Typical Metrics

This metric is not often used since its primary influence, the level of RAID protection, is a business decision rather than a result of poor management processes. Additionally, most utilization calculations do not include arrays that are in "mid deployment" (a frequent cause of low utilization in this category) because these immature environments would skew the results. Utilization metrics are therefore not normally quoted for this classification.

Array Utilization

(Array Usable to Array Used)

This metric is the percentage of usable array capacity that is allocated to hosts. It indicates the efficiency of storage deployment operations. This type of utilization information must be collected directly from the storage array.

Low Utilization

Low utilization is usually caused by lack of attached hosts, and is prevalent in direct-attached environments. This typically results from insufficient host ports, lack of fan-out capacity, and metric measurement early in the deployment cycle. The complexity associated with re-working previously configured storage varies from one storage vendor to another and may contribute to low utilization in this category. In particular, the complexities and restrictions of reconfiguring previously used storage in arrays may result in a disproportionately low "Array Usable to Array Used" metric.

Improving Array Utilization

The most common remedy to this problem is implementation of a Storage Area Network (SAN) to resolve a lack of host ports, a host consolidation exercise or a strategic change of storage hardware. Often, Array Utilization can be improved by offering a choice of storage hardware options, each with different storage allocation attributes.

High Utilization

High utilization typically results from a flexible storage environment in which requirements on the host side are efficiently fulfilled on the storage array side. Flexible, management tools are necessary to get close to "just in time" configuration changes. Tight integration of storage processes into project change control is essential.

Typical Metrics

	Typical	Mature	World Class
SAN	40%-70%	60%-80%	90%+
NAS	60%-80%	70%-90%	90%+
DAS	25%-40%	40%-60%	70%-85%

Table 3 - "Array Usable to Array Used" Achieved Metrics

Allocation Efficiency

(Array Used to Host Raw)

Allocation Efficiency reflects the ratio of storage presented or allocated to hosts to the amount actually seen by them. In many mature environments this ratio is near 100% (i.e. all the storage allocated is being seen), but this ratio can be extremely difficult to determine. It relies on accurate measurements of both Array Used storage and Host Raw storage, each of which is gathered by separate tools.

Low Utilization

A common reason for low utilization in this category is an inability to dynamically reconfigure storage arrays. If storage requests cannot be fulfilled during normal operation, excess storage is often assigned to hosts on the array side in anticipation of future need. Often these hosts never need the storage assigned to them, and sometimes they are never even attached to the storage array. During a new deployment, when servers are being brought on line over a period of time, some storage administrators will allocate all of the storage at once. Then, when servers are activated, their storage is available immediately.

Another reason for a low utilization in this category is consolidation or decommissioning. Here, previously used storage may still be allocated to a server that has been decommissioned, resulting in a low "Array Used to Host Raw" utilization ratio.

Mid-project measurement of these metrics and the resultant ratio will be un-characteristically low. It is recommended that metrics are not taken from arrays at this stage in the project life cycle.

Improving Allocation Efficiency

Effective storage management processes and software can bring storage deployment in line with host storage requests. If storage deployment just matches host storage needs, over-allocation is unnecessary to ensure that storage will be available when needed.

High Utilization

The majority of mature environments achieve very high utilization levels in this category. In order to sustain this, high levels of confidence are required from the storage administrator so they will not reserve large quantities of unused storage. A strong integration of storage processes, change control and project management is necessary for this confidence to be built up.

Typical Metrics

Typical	Mature	World Class
90%	95%	97%+

Table 4 - "Array Used to Host Raw" Achieved Metrics

Host Overhead

(Host Raw to Host Usable)

This metric reflects the amount of storage configured for use versus the

amount the host can see. Since the Host Raw metric is a function of the storage administration team and the Host Usable a function of the systems administration team, this metric is a useful measurement of how well the two functions are cooperating. Data for this classification is collected from the host.

High Overhead

High overhead (i.e. low utilization) is typically a result of over-allocation of storage. This is normally due either to poor storage policies or architectural constraints. In fact, the lack of mature policies frequently results in massive over-allocation with the aim of reducing administrator workload.

Architectural constraints in both the host hardware and file system layers can also result in low utilization in this category. Some operating systems cannot see newly allocated disk volumes without a reboot. Since constant 24x7 availability is often an operational priority, some server administration teams over-allocate "Host Raw" storage to avoid frequent reboots. Similarly, some file system tools cannot provide on-line file system resizing even if disk space can be dynamically allocated without a reboot. The result in both cases is a trade-off between utilization and 24x7 availability.

Another impact on this metric is the use of volume manager-based mirroring. This is unusual in enterprise-class storage, since this level of protection is more efficiently performed in the array. In some cases, systems administrators may use volume manager mirroring to provide Business Continuance Volume (BCV)[3] functionality for on-line backups. In such cases, the BCV-like volumes are best left out of the metric calculation since they have a genuine business function.

Lowering Host Overhead

Improvements can be seen in this classification by:

- Implementation of mature storage management processes to enable efficient delivery of new storage when it is required
- Use of modern file systems such as Veritas volume manager that allow on-line file system resizing
- Use of centralized, hardware mirrored arrays to remove the requirement for inefficient host-level mirroring

Low Overhead

Low overhead (i.e. high utilization) is associated with environments where the host hardware and file system layers can both support on-line changes in a safe manner. Additionally, close integration and trust between the storage and systems administration teams will exist to allow for timely delivery of the necessary storage. Visibility of file system trends is necessary, typically provided by in-house tools or specialized utilities such as StorageNetworks STORos.

[3] Business Continuance Volumes (BCVs) are “extra mirrors” of volumes, created either in the storage array or on the host with a volume manager. These spare copies of data can be separated (“split”) from the production data when needed and accessed separately as a point-in-time copy. They can be re-synchronized later while the production volumes remain on-line.

Typical Metrics

Achievable metrics depend on good business practices, host support for on-line changes, and a low overall rate of data growth, which may force systems administrators to over-allocate to meet demands without rebooting.

Typical	Mature	World Class
30%-50%	10%-30%	5%

Table 5 - "Host Raw to Host Usable" Achieved Metrics

File System Utilization

(Host Usable to Host Used)

File system utilization is the amount of available file system space that actually contains data. File system utilization is familiar to most systems administrators. This metric is often shown in simple system commands like "df" on UNIX or "dir" on Windows. Data for this classification is collected from the host.

Low Utilization

Low utilization usually indicates poor storage allocation processes, lack of forecasting and measurement tools, architectural constraints in the volume manager, or high-change file systems.

In environments where poor storage processes are employed, systems administrators or DBAs will often request excessive storage capacity, normally on the grounds that obtaining space is too difficult. This leads to low utilization in this category. Additionally, a lack of information flow between project managers and systems administration can lead to the initial creation of incorrectly sized volumes and poor forecasting of growth.[4] Many enterprises do not perform trend analysis on file system use, typically because of a lack of tools.

Some file systems are categorized as "high-change file systems" and are intended to have low "Host Used" rates. Examples include file systems for temporary files, crash dump space, print spooler files or system log files. These are expected to be able to cope with sudden, massive growth. To drive high "Host Used" rates in such areas will actually endanger safe running of the operating system.

Improving File System Utilization

Improvements can be seen in this classification by:

- Use of modern functionally rich file systems (such as Veritas Volume Manager) to allow on-line file system resizing
- Implementation of mature processes and procedures to enable on-line expansion of file systems whenever required
- Right-sizing of existing file systems[4]

[4] Note that downsizing a file system is considerably more complex than upsizing. While most systems administrators are happy to increase file systems on line, the reverse is true of downsizing. Standard procedure will be to back up the data, destroy the file system, create it and restore the data. It is therefore of critical importance that any file systems are right-sized initially.

High Utilization

This is associated with environments where the file system layers can support on-line changes in a safe manner. Additionally, close integration and trust between the systems administration, DBA and application support teams will exist to allow for timely delivery of the necessary file system space. Visibility of data growth trends is necessary, typically provided by in house tools or specialized utilities such as StorageNetworks' STOROs.

A policy of high file system utilization can be difficult but wise. Most modern file systems have automatic defragmentation functions that require some overhead, typically 10% - 15% to operate. Driving file system utilization in a normal multi-purpose environment (i.e. not solely for db table space) beyond 85% will result in a degradation of defragmentation and thus an overall performance drop. In cases where an entire file system is dedicated to a database, this is not an issue since the defragmentation is provided by the database. In such cases, "Host Usable to Host Used" utilization can safely near 100%.

Typical Metrics

Typically, file system utilization is high (above 80%) for special-purpose file systems, such as database volumes, and much lower for many other file systems. Application and operating system areas are often utilized below 10%, and even home directories and other file storage areas can be below 50%. Overall, across all file systems on a database server, the following metrics are seen:

Typical	Mature	World Class
60%-70%	70%-80%	80%-85%

Table 6 - "Host Usable to Host Used" Achieved Metrics

Total Storage Utilization

(Array Usable to Host Used)

Finally, the Total Storage Utilization metric summarizes how well a company manages its storage assets across the entire business. This ratio is the default storage utilization metric used in publications and reflects the actual value an enterprise is deriving from its storage asset. Care is required in calculating this ratio to ensure that it accurately indicates utilization of the storage environment. Since the result of this ratio is often used in business cases and receives wide attention, it must be both logical and defendable.

Low Utilization

Low storage utilization is a result of failing to apply the procedures, technology and practices outlined in the above sections.

High Utilization

A high utilization in this metric is not easy to achieve since it requires consistent mature storage and systems management at all levels. Of critical importance is the strong integration of storage processes into both change control and project lifecycles.

Typical Metrics

This category is an end-to-end metric. Inefficiencies in each of the tiers are effectively multiplied. Therefore only obtaining average utilization at each level can result in a very low end-to-end metric.

Typical	Mature	World Class
20%-40%	40%-60%	60%+

Table 7 - "Array Usable to Host Used" Achieved Metrics

Improving Storage Utilization

Utilization Best Practices

StorageNetworks has developed a storage management lifecycle that includes the following activity areas: Forecasting, Deployment of capacity, Requisitioning (presentation of storage to hosts), and Management. Each of these plays a major role in achieving a highly utilized storage environment. StorageNetworks drives or assists in driving the following utilization best-practices targets:

- Drive Array Utilization (Array Usable to Array Used) to greater than 90% (a storage administration responsibility)
- Drive Allocation Efficiency: Bring Host Usable to be as close to Array Used as possible (a joint responsibility)
- Drive Filesystem Utilization ("Host Usable to Host Used") above 80% (a systems administration responsibility)

StorageNetworks recommends using the Total Storage Utilization ("Array Usable to Host Used") metric to create utilization targets, since it reflects all aspects of storage utilization at once. The best practice target for this ratio is 75%.

Utilization Cost Savings

Using typical industry figures, a low level of Total Storage Utilization (Array Usable to Host Used) can cost far more than a target value for world-class practices. Table 8 shows an example of how storage utilization drives storage cost.

In this example, the typical enterprise has 1 TB of production data and runs storage at 21% Total Storage Utilization. Assuming a cost of 10 cents per raw MB, they are spending 3.5 times more for storage than if they implemented world-class storage management processes.

This low utilization is due to their reliance on Direct-Attached Storage (DAS), which impedes utilization improvements. DAS leads to a lack of connectivity and over-allocation of storage. Poor storage management practices keep file system utilization below 70%, leading to the final 21% utilization rate.

	Array Utilization	Allocation Efficiency	Host Overhead	File System Utilization	Total Storage Utilization	Cost per TB
Typical DAS	55%	90%	35%	65%	21%	$500k
World-Class SAN	95%	100%	1%	80%	75%	$140k
Savings per TB						$360k

Table 8 - Cost Savings Illustration

Improving connectivity and storage allocation in a Storage Area Network with storage management best practices allows this customer to bring Total Storage Utilization much higher.

In practice, an enterprise in this situation could reorganize storage assets and defer purchases. Although this would delay recognition of the total of savings, deferred spending on storage (often a large percent of IT budgets) is a sound decision in a down economy.

Other Usage Metrics

There are a number of other usage metrics that are less-frequently discussed. These metrics are normally not used in utilization calculations, but they can have a substantial impact on the overall efficiency of a storage infrastructure.

Frame utilization

This metric measures the extent to which array disk slots have been populated with disks. In the case of non-modular arrays (such as EMC's Symmetrix), customers often buy a frame and a small number of disks, intending to fill it up as needed. This purchasing decision may be legitimate but can lead to excessive cost per MB if the predictions of growth fail to materialize. A number of factors can lead to low frame utilization:

- Lack of visibility can result in the purchase of additional frames before existing ones have been filled.
- In DAS environments where array connectivity is restricted, an array will frequently run out of ports before it runs out of storage capacity. This results in a large number of free disk slots but no means of connecting a server to them.
- In recent years, most storage manufacturers have allowed customers to purchase fully configured arrays but only pay for the quantity of disk space being used. This new purchasing model, when used in a SAN configuration, has gone some way to resolving the issue of low frame utilization.
- Additionally, many manufacturers are moving to a hardware design that is more modular. Both Sun and Compaq have enterprise storage

solutions that allow both disks and array components to be added as required, offsetting the large up-front cost associated with earlier "single frame" designs.

Port utilization

This metric measures the extent to which an array's connectivity ports are used. Enterprise class arrays are provided with a large number of connectivity ports (e.g. SCSI, ESCON or Fibre Channel). In a DAS environment, storage administrators are forced to balance the total disk capacity of the array with the total number of connectivity ports. This balance almost always results in two states:

- All the storage has been allocated, but there are connectivity ports free (i.e. the ratio of storage to ports is too low).
- All the connectivity ports have been allocated, but there is storage capacity free (i.e. the ratio of ports to storage is too low).

In either case, inefficiencies are introduced. In the first case, connectivity has been paid for that cannot be used. In the second case, storage capacity has been paid for that cannot be used.

The advent of SANs and the ability to share connectivity ports between servers promises to resolve this issue. However, SAN interoperability challenges remain, leading to many Fibre Channel fabrics remaining small "SAN Islands". In DAS environments, this is still a large problem and a cause of inefficiency in the storage.

Database utilization

This metric measures the extent to which a database actually uses the table space it has been allocated by the DBA. It is a subset of the "Host Used" metric, and databases have similar utilization ratios to host storage. For example, a raw volume or file system created for database use could be thought of as "Database Raw", a database file as "Database Usable", and the data in a database as "Database Used". The level of Database Usable to Used utilization can often be seen when database exports are performed, and can be quite low. The same issues apply here as to elsewhere:

- Databases in early stages of maturity are likely to have low utilization rates
- Some table spaces need to have low utilization (e.g. sort and temp areas)

To achieve consistently high utilization, good communication and trust is needed between the Storage, System, Database and Application administrators

Storage, systems and database processes must be tied closely into change control and project lifecycle processes in order to deliver an overall high level of utilization.

Case Study

A large Wall Street financial institution had a significant direct-attached storage infrastructure in 1999. At that point, the company made two key decisions with respect to its storage strategy: to centralize storage management and to

drive costs out of the environment. This required significant organizational change, a technology shift toward Storage Area Networks, and a whole new set of storage management best practices. This company soon realized that to achieve these goals within the necessary timeframe, it needed to look for a partner with the necessary depth of expertise and technology.

Part of this high cost was a result of low asset utilization in the Direct-Attached environment and technology constraints inherent to Direct-Attached Storage. With asset utilization around 40%, this company saw significant potential for savings just by better utilizing existing storage assets.

StorageNetworks was selected as the partner to implement this strategy. Over several months, this environment evolved to support 40+ TB of primary data. The following utilization metrics, achieved at this customer's data centers, illustrate the value of improved storage management:

- Just five enterprise-class arrays deployed over the past 12 months
- Average Array Usable to Array Used ratio of 92% on arrays that have reached maturity
- The storage management group bills business units for capacity they request. This, combined with reporting tools, drives systems administrators to request only the storage they need. The result is a Host Used/Host Usable ratio of approximately 80%
- The combined effect of these two ratios is a 75% Array Usable to Host Used ratio (Total Storage Utilization).
- This improvement in utilization reduced Total Cost of Ownership for enterprise storage by 47%.

StorageNetworks, Inc.
225 Wyman St., Waltham, MA 02451
(781) 622-6700 www.storagenetworks.com