White Paper

Ten Tips For Reducing Storage TCO

The Proper Design And Architecture Of A Storage Infrastructure Can Have A Significant Impact On The Total Cost Of Ownership (TCO)

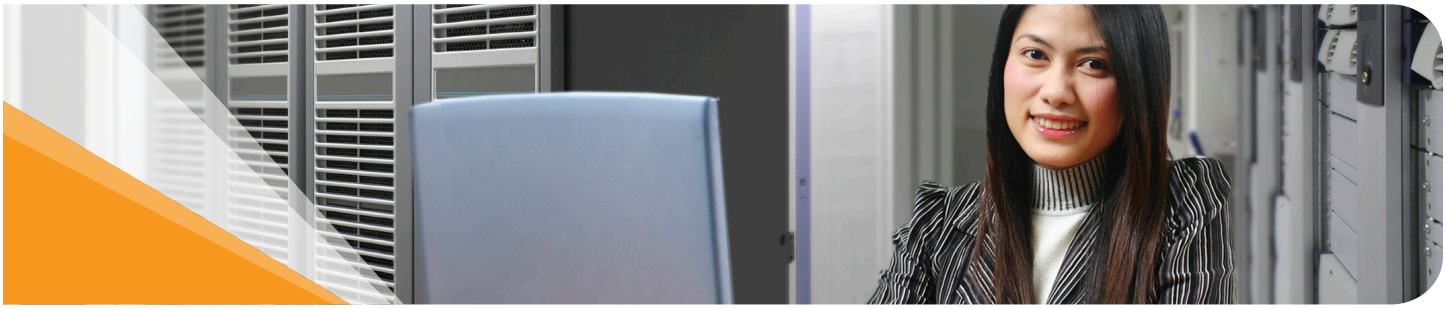

Executive Summary

As corporate storage environments continue to grow, and the data becomes more important to the organization, IT departments struggle with increasing cost and complexity. This white paper outlines the top ten strategic steps that should be taken to reduce the Total Cost of Ownership (TCO) of a storage infrastructure. By architecting the correct storage environment, significant cost savings can be recognized.

1. Use Heterogeneous Storage

The use of heterogeneous storage architectures gives users the flexibility to purchase the storage devices that provide the best price / performance ratio, rather than simply the storage device offered by their single storage vendor.

Vendor lock-in is a major contributor to the cost of storage systems. Often organizations find themselves dependent upon their storage vendor not only for the sake of providing the storage systems, but also to design the Storage Area Network (SAN). It is only natural for storage providers to recommend solutions that are profitable for them — leveraging the fact that lack of competition drives up prices.

Many vendors now propose storage virtualization solutions enabling organizations to work seamlessly with heterogeneous storage devices.

In some cases, storage vendors attempt to intimidate customers with a threat of revoking their warranty and/or ceasing to support them if they use heterogeneous storage. However, when the customer insists, an offer to use the vendor's storage virtualization solution is typically offered.

One major concern with using heterogeneous storage is the interoperability of storage services software such as snapshots and mirroring. Therefore, it is recommended to select a storage virtualization solution that also supports storage services for heterogeneous storage.

2. Use a Tiered Storage Architecture

Every organization puts multiple applications to use. Each has different storage requirements. For example, ERP software – which is the core of many businesses – may require very fast and reliable

Tiered Storage Architectures place data on the storage devices that have the most appropriate price / performance ratio. Not all data is created equally and not all data needs to reside on "tier 1" storage.

storage devices. Email systems may require mid-tier storage devices. However, development applications may use lower-tiered storage. The costly solution, which is often suggested by storage vendors and integrators, is to place all data on the high-end storage devices. The outcome is a redundant cost of storage for applications that do not require high-end storage systems. In some cases, the argument used is that "we have spare capacity on the high-end device". However, when a real need for high-end capacity arises, the extra storage space is already occupied.

The cost of storage can be easily reduced by checking the availability and performance requirements for storage – leaving you the freedom of choice, to purchase only what you need.

Since storage requirements may change over time (e.g., the Information Lifecycle Management – ILM) it is also recommended to have an on-line volume migration application for moving volumes from one storage device to another, while applications are up and running. Obviously this software must be able to migrate between storage devices from different vendors.

3. Allocate Storage According to Actual Requirements

The storage allocation capabilities enabled by storage virtualization can increase capacity utilization considerably. Capacity can be allocated as needed. Organizations can also purchase new storage as needed taking advantage of ongoing price reduction in storage devices.

Storage allocation is a complex task, resulting in storage administrators allocating storage based upon applications' growth estimations. This often leads organizations to purchase much more storage than is actually needed at the time of implementation. Over time, some applications may require more than was initially allocated for them and others may have extra capacity which cannot be reclaimed.

Storage virtualization solutions enable the creation of logical "storage pools" that can be expanded dynamically by adding more physical storage, and creating logical volumes from the storage pools. Those logical volumes are like any other volume to the application but can be easily expanded by the virtualization software. The result is that an organization can buy just the storage it needs, expand volumes on-demand and add physical storage only when needed. The cost of storage is lower, since organizations buy only what is required at each point in time and a very low amount of storage remains unused.

4. Implement Centralized Storage Management

The ability to manage and protect all storage devices (regardless of vendor) with a single application and from a single "pane of glass" can significantly reduce costs.

There are many software applications available to manage storage. In many cases, a single vendor may offer a few different software applications – each for different purposes (e.g., one for snapshots and one for mirroring). Then there are the cases where a user purchases multiple applications to support different storage devices. Implementing such applications requires some level of training and knowledge resulting in higher cost of maintenance. By selecting storage management software that can provide centralized management for all storage devices and services, an organization can significantly reduce maintenance costs.

5. Use Low-Capacity Snapshots to Protect Against "Logical" Failures

Using Snapshots for rapid recovery rather than tape allows users to improve their RTO from hours to minutes and similarly improve their RPO from days to hours or event minutes – as needed.

Market research shows that 93% of data loss events are from logical failures. Only approximately 7% of data loss events are due to a "physical" failure such as a disk array going offline. Among logical failures we count such events as virus attacks, corruption of file systems, or accidental file deletion by users.

Recovering from logical failures using tapes is a long and exhausting process. Tapes must be located and loaded into the tape drives. A physical location on disk must be available to restore the data. Loading the data from the tapes to the disk is time consuming. Finally, if the data cannot be found or is damaged, the process must be repeated using a different set of tapes – assuming there is a second copy of the backup.

By using low-capacity snapshots, recovery from logical failures can be done in a few minutes. A snapshot can simply be mounted to a server so that the file(s) can be copied back in place or an entire volume (or set of volumes) can be restored in their entirety. The rapid recovery capabilities of snapshots can dramatically improve the Recovery Time Objectives (RTO) of organizations. Rather than requiring hours to recover from a tape backup, organizations can use snapshots to recovery in minutes.

Low capacity snapshots consume very little storage space (an average of 5% of volume size per day). Therefore, multiple snapshots can be kept online as "recovery points". Often, users will keep five or six snapshots online per day for a five-day work week. At the beginning of the next week they replace the oldest snapshot with a new one. They keep a rolling 25 or 30 recovery points that they can use to restore data over the previous five days. If they need to restore data older than one week, they use tape. In this example, the user brings their Recovery Point Objective (RPO) to within an hour or two. By using more snapshots, they can improve their RPO to within minutes.

6. Instant Volume Replication and Low-Capacity Snapshots for Application Testing

Many organizations develop applications. In most cases, development teams need current data for development and testing purposes. Replicating a volume may require a long time to copy the entire data set. Sometimes, several teams need the data for independent tests simultaneously, resulting in multiple replications using a substantial amount of storage space.

Using snapshots for application testing gives users instant access to "current" data and also gives them the ability to roll-back to previous stages in the testing process – to review test results or to restart the test process at any stage. It also saves considerable disk space and time.

There are a few software vendors that provide an "instant copy" option, in which copies are ready for use within seconds, and the data is then copied in the background. In addition, there are low-capacity snapshot solutions that enable users to simultaneously assign the same snapshot to multiple servers for independent read and write purposes.

Using instant physical copies in conjunction with the above low-capacity snapshot solution enables the fast creation of a replication and the assignment of snapshots in parallel to multiple servers, each may be used by different testers or developers for independent purposes. Using this solution, a developer's time is saved as well as a significant amount of storage space.

If a test fails, the tester would typically like to run another test on the same data for debugging purposes. With this solution, multiple snapshots could be taken of the core data during the tests. So, if a test fails in the middle of a multi-staged process, a software tester could go back and review what the data looked like at any of the previous stages and/or continue the test from that stage, rather than restarting the test from the beginning.

7. Implement a Cost-Effective Disaster Recovery Site

Disaster recovery solutions are usually quite expensive and complex. A typical solution requires a mirrored site with the same equipment as the original site, connected via expensive communication lines.

Using a heterogeneous remote mirroring solution in conjunction with snapshot assisted mirroring technologies can make remote mirroring affordable not only for tier-1 data, but all data in a corporate IT environment.

The first step in reducing the cost of a remotely mirrored implementation is to us a solution that supports heterogeneous storage devices. Therefore, you have an expensive tier 1 storage device in your primary site, but mirror to a mid-range device. If there is a disaster, most organizations can run in a slightly degraded mode until they recover their primary site. However, there is no need to put an expensive tier 1 device, sitting idle in the remote site.

Statistically, over the life of a remote mirror implementation, the most expensive component is the network communication lines between the two sites. Therefore, if it is possible to reduce the transactions between the two sites, the TCO is reduced. There are technologies available, such as snapshot-assisted mirroring that can reduce the required bandwidth (and cost) since it transfers only the changed data, regardless of how many times it was changed during the snapshot period. Implementing these solutions can reduce the bandwidth requirements by an order of magnitude.

Using snapshots to enable "hot" backup during production hours reduces the risk of failures and can eliminate the need for off-hours staffing requirements.

8. Use Snapshot-Based Backup Processes

As the amount of data for backup increases, it becomes more costly. In order to minimize disruption to production servers, backups are executed during the night independently and without supervision. This may result in failures, or alternatively, organizations pay extra for a night shift to be present during the process. Due to increasing production requirements and increasing data capacities, backup windows are in many cases not large enough to complete the required task. This leads to huge infrastructure upgrade expenses (e.g., servers, LAN, faster tapes). By using low-capacity snapshots for backups, there is no disturbance to production servers or the LAN. A dedicated backup server can be mounted to a snapshot and perform the backup at any given time. It eliminates the need to upgrade the servers or the LAN, and cancels the need to dedicate a night shift to supervise the backups.

Snapshots can be used at the remote site to implement backups and enable development and testing processes. This increases productivity and reduces cost by turning "idle" equipment into productive environments.

9. Use of a Disaster Recovery Site for Backups

When a disaster recovery site exists, it can be used for backup purposes. In this case, there is no need for a backup server, since the DR site already has servers on call for failure and can be used for other tasks. Backup tapes are frequently shipped to a remote vault, in which case, the DR site may function as the remote vault and the shipment costs are eliminated. By implementing a tape backup at the remote location, the requirement for a "vaulted" set of tapes is achieved automatically.

10. Use of a Disaster Recovery Site for Application Development and Testing

Disaster recovery sites typically have a very current mirror of the data, a large amount of equipment (e.g., servers, terminals, networks, printers, desks) and people that are on-call "waiting" for a disaster to occur. Typically, there is not a lot of work being done at the disaster recovery site, because most mirroring solutions do not allow you to access the mirrored data without "breaking" the mirror and running exposed. However, there are solutions available that use remote snapshots to create read/write-able views of the data that can be used for development and testing purposes. Software teams may use this equipment for development and testing in order to save the cost of purchasing the same equipment for the main site.

Summary

There are a number of key steps that can be implemented to design and implement a storage architecture that will have a material impact on the cost structure. This document has outlined ten of them that can significantly reduce the cost of your storage infrastructure.

By planning ahead you can eliminate vendor lock-in, improve the productivity of your staff and capital resources and ultimately improve the success of your

sgi INNOVATION FOR RESULTS

This document is provided by SGI, an authorized reseller of LSI Corporation

For more information and sales office locations, please visit the LSI web sites at: lsi.com lsi.com/contacts

North American Headquarters
Milpitas, CA
T: +1.866.574.5741 (within U.S.)
T: +1.408.954.3108 (outside U.S.)

LSI Europe Ltd.
European Headquarters
United Kingdom
T: [+44] 1344.413200

LSI KK Headquarters
Tokyo, Japan
Tel: [+81] 3.5463.7165

LSI™