

STORAGE SWITZERLAND REPORT

DOING MORE IN LESS SPACE



When it comes to storage, capacity costs money. Wasted capacity is lost money. When trying to increase efficiency, after improving staff productivity by, as we stated in a prior article, "Doing More in Less Time", the next obvious step is to do more in less space. In other words, use capacity efficiently. Space optimization is not limited to just storage. "Space" can be wasted on a virtual server cluster as well, for example, if all the hosts are not running at a high level of CPU utilization.

The challenge is that, thanks to cheap capacity (storage and CPU) or products that thinly provision, compress or deduplicate data, diligence has been replaced by under-allocating resources for the given situation. While these optimization technologies are important, their use should also be balanced with proper data center resource management. Given the current economic situation, there may no longer be budget to under-allocate. Now, data centers are being pressured to over-allocate or do more in less space.

Why has under-allocation as a method to these challenges become so common? One reason is that there is a limited amount of time that the IT staff has to address these resource allocation issues. Another is because data center management software has traditionally been too narrow in scope, too complex and too expensive. Given the steps

we referenced in our prior article, IT needs more time allocated to develop better resource utilization processes, and it needs effective tools. Data center management tools, like those from Tek-Tools are attempting to change that situation across a wide range of IT platforms.

Proper resource utilization requires that these tools provide a real-time, heads-up indication of what's going on in the environment, as well as trending capabilities going forward. Today's dynamic data center can no longer wait for a batch job or network crawl to take place. Information must be available in real time so the IT staff can react and plan accordingly.

The goal is to do more in less space. To achieve this goal you need the ability to improve compute and storage utilization rates further.

Save Space by Improving Storage Utilization Rates

Data center storage, probably more so than any other asset in the data center, is grossly under-allocated. Wasted storage capacity can appear in many areas: old data on storage that needs to be migrated, storage that is allocated but not in use and storage that was allocated but is now abandoned.

Old data that is on primary storage is potentially the simplest way to free up existing resources. Most data centers report 60% or greater the amount of data that has not been accessed in the last year. Purchasing a secondary disk tier or archive would allow this storage to be moved to less expensive disk. Most often, this secondary disk tier can compress and deduplicate data, and as a result, now offer disk capacities around \$1 per GB. A data center resource management tool, that has a storage extension to it, can analyze capacity in real time and then allow the IT staff to build a data movement policy based on age and needed capacity.

This reduction in space can come in two forms. First, additional primary storage capacity purchases can be avoided. Other than a substantial cost savings, avoiding this additional capacity purchase will save floor space as well as the time required to deploy and manage any new storage. Another space savings is in the reduction of power requirements. Many data centers simply don't have access to additional power, and they don't have space on their existing power grids for more storage.

Conversely, by identifying this old data and moving it to densely packed, higher capacity drives and optimization techniques like compression and deduplication, secondary storage tiers can save cost, floor space and improve power utilization. Further, many of these systems also employ power-efficient drives that power down or go into standby when not in use. The key is having a tool that can universally

identify this data, not just initially, but as part of an ongoing process.

The second type of space wasting is storage that is allocated but not in use. For example, this could be a 1TB LUN that was assigned to a database application a year ago because it was expected that the application would grow that large. Often, what actually happens is that utilization is about 25% of what was projected and the rest of the capacity on that LUN is 'held captive'. This is the problem that the technology 'thin provisioning' attempts to solve, by only allocating storage as it is being used.

Thin provisioning is not without its detractors. There is concern about a potential performance impact when managing thinly provisioned volumes and there is the ongoing concern of over-allocating storage to the point that volumes actually do run out of capacity, aborting applications in their tracks. Finally, if there is not budget to purchase a thin provisioning storage system, the IT staff needs tools to manage this process on their own.

Armed with the same tools that can identify old data, LUNs themselves can be monitored for actual space utilization. Further, a trend line can be built that identifies the rate at which capacity is being consumed on a LUN. Volumes that have excess capacity and a slow consumption rate can be quickly identified and, depending on the system, either reduced in size or have the active data migrated to smaller volumes.

By doing this, the storage administrator can once again save

enough storage space to delay upcoming storage purchases, as well as reap the benefits of not requiring the additional floor space and power consumption.

All of this can be done without purchasing a new storage system and without some of the potential pitfalls of thin provisioning. Also, the amount of this rebalancing that has to occur, and the performance characteristics of the workloads, can be measured and tracked. This information could even lead to a faster justification of a storage system that includes thin provisioning. Now, the storage administrator would be armed with a tool that allows the data center to enjoy the benefits of thin provisioning while avoiding its pitfalls.

The final aspect of doing more in less space through better storage practices is the identification and reallocation of orphaned or near-orphaned LUNs. An orphaned LUN is a block of disk storage that was allocated to a server or group of servers that is no longer in use. The storage module of a data center management tool can identify which servers are attached to these LUNs and how much traffic, if any, is going back and forth between the servers and the LUNs.

In a situation where there is no server attached or traffic going to the LUN, a quick confirmation can lead to releasing that LUN back to the storage pool. In a situation where there is minimal traffic or minimal use of the LUN, the remaining server volume can be migrated to another LUN with the then-freed LUN being released back into the storage pool.

It may seem surprising to some that orphaned or underutilized LUNs are as common as they are, but in almost every medium and larger data center they exist. The problem has actually grown as a result of server virtualization. When a server is migrated to the virtual infrastructure, its original volume is typically never backtracked to be returned to the global storage pool.

Reducing the space used by storage is a critical and obvious first step in doing more in less space. Today, with an easy-to-use data center management tools like those available from Tek-Tools, ongoing procedures can be put in place to identify space reclamation opportunities. Space reclamation does not end with storage. Thanks to server virtualization, there is a potential gold mine of space reclamation available in that environment.

Save Space by Improving Server Virtualization Rates

Server virtualization projects typically kick-off with a goal to reduce physical server count and to increase overall compute utilization. As part of that kick-off, typically, new and more powerful servers are purchased and deployed as virtual hosts. Then, a comparatively small number of very safe workloads are virtualized. The result is that while server count is typically reduced, the compute utilization is not, since these new servers have significantly more power than the ones they replaced. And there is capacity for more servers to join the cluster, but with the next wave of workloads more mission-critical, careful planning becomes a requirement.

If more physical servers can be converted to virtual servers, the space savings in the data center can be tremendous. There is the hard reduction in the number of servers, which leads to a reduction in the power and cooling required for those servers. This, in turn, may lead to more servers being managed by the same number of admins. The challenge is making these upcoming migrations predictably safe.

Using tools like those from Tek-Tools, the current, stand-alone server environment can be measured in real time. The details of this analysis can then be compared to the virtual environment, and the systems can be prioritized to decide

which are best suited to be virtualized. In fact, simulations can be run to predict the impact of adding a workload to the virtual environment.

Doing more in less space in server virtualization can lead to the powering down and disposal of servers. This can lead to a significant reduction in the number of physical servers deployed and the space and power they consume. The denser these virtual machine populations can be made, the greater the space savings. Using tools that can accurately predict the result of this increased density and then continuously monitoring the environment for greater virtualization can be done safely.

These types of tools allow for greater virtual machine densities per virtual host while maintaining the right sense of balance for virtual machine migrations and disaster recovery.

The virtual environment can also be a potential storage waster. Often, virtual machines are created with templates. Templates are valuable aspects of server virtualization that enable rapid deployments of new servers. Typically, the storage in allocation for these templates is set to a default number. Most virtualization administrators set this to a "safe" default size which, on an individual virtual machine basis does not appear to waste much space. Virtual environments never stay at just a few virtual machines, they grow rapidly. As a result, the cumulative waste of each virtual machine can result in TBs of total wasted disk capacity.

Identifying and resizing these virtual machines and potentially changing the template deployment policy can lead to a tremendous space savings and garner all the other storage allocation benefits that were described in the previous section.

Greater space utilization, either for the compute environment or the storage environment, leads to cost reduction from delayed purchases but also reductions in floor space, power and cooling. With the reported cost of a data center floor tile running as much as \$10,000 per month, the return on the investment in a proper data center management tool can be almost instantaneous.

About Storage Switzerland

Storage Switzerland, is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>.