

Storage Efficiency in a Virtual Environment

Author: James Rosikiewicz
 Director of Product Management
 phdvirtual.com
 Published: April 2009

1.0 Introduction

One of the critical components of any infrastructure is the storage layer. Organizations typically invest a large amount of capital upfront for storage, often purchasing specified equipment without necessarily scoping the requirement, profiling the environment or reviewing ongoing storage management. The latter is as important as the initial scoping and design phase. Within a virtual environment, ongoing maintenance and monitoring of storage tiers is more often overlooked than in a non-virtualized environment.

To shed some light on this issue, this whitepaper discusses storage tiers as well as considerations and techniques around storage in a virtualized environment.

2.0 Wasted Space in Tiered Storage Hierarchies

Tiered storage is often described as ‘the assignment of different data categories to different storage media.’ The aim of tiered storage is to reduce the total cost of storage for an organization’s infrastructure. A tiered storage model requires data be assigned a class, which in turn dictates the tier of storage on which the data will reside.

Each storage tier is made up of a set of disks and a storage fabric. Critical data that requires constant high speed access can be kept on the fastest storage tier (Tier 1). Data that does not need such high speed access can be located on a lower cost tier (Tier 2) and so forth. By understanding data access requirements, a model can be constructed allowing organizations to purchase differing tiers of storage and connectivity fabric. This is in contrast to purchasing storage which is expensive and over-specified, or is cheaper but does not perform adequately.

Tier 1	Tier 2	Tier 3	Tier 4
Critical Data	Medium Priority	Low Priority	Archive
Fast FC (15K)	Fast FC (10K) / FATA / iSCSI	iSCSI / SATA	Tape Device
Email Database	Fileserver User Profiles	Backup Target	Regulatory Compliance Data

Figure 2-1: Tiered storage model example

With the implementation of tiered storage, the ability to move data easily between tiers allows the implementation of a data lifecycle process. This typically reflects positively on the Total Cost of Ownership (TCO) and Return on Investment (ROI) since the layered approach means data and services operate on a ‘fit-for-purpose’ infrastructure rather than the inefficient ‘one size fits all’ approach.

However, a common occurrence in this hierarchy is storage maintenance neglect. Assuming that data is kept in the correct place and moved between tiers as required, administrators may not continue to monitor the actual content within their storage. When storage tiers fill up, additional capacity is often deployed without investigating whether anything that can be done to recover and re-use any of the existing utilized space. Redundant data may sit in multiple tiers. In virtual infrastructures, there are techniques that can be employed for initial consolidation project, and periodically thereafter, to assist storage footprint efficiency.

However, a common occurrence in this hierarchy is storage maintenance neglect. Assuming that data is kept in the correct place and moved between tiers as required, administrators may not continue to

monitor the actual content within their storage. When storage tiers fill up, additional capacity is often deployed without investigating whether anything that can be done to recover and re-use any of the existing utilized space. Redundant data may sit in multiple tiers. In virtual infrastructures, there are techniques that can be employed for initial consolidation project, and periodically thereafter, to assist storage footprint efficiency.

3.0 Storage Efficient Technologies & Configuration

There are six primary technologies and configurations that can achieve efficient storage in a virtual environment. These include (1) Thin provisioning, (2) Data de-duplication, (3) Linked clones, (4) Backup target configuration, (5) VM guest disk configuration and (6) Shrinking or growing a virtual disk. Each is discussed in detail below.

3.1 Thin Provisioning

Traditional storage provisioning, sometimes known as fat/thick provisioning, generally requires disk space to be allocated beyond current needs, in anticipation of future growth. As a result large amounts of storage space are provisioned but never used. This leads to a poor ROI since actual storage usage can often be as low as 10%, due to significant amounts of unused whitespace.

Thin provisioning, however, allows higher storage utilization by eliminating the need to install physical disk capacity. When multiple applications share access to the same storage array, it provides a method of optimizing the efficiency of the available space by using only what is required when data is written.

In a virtual environment, a virtual machine (VM) can provide a thin-provisioned disk of 20GB. Initially, only 5GB may be used, occupying only 5GB of space on the storage tier. The guest O/S (operating system) in the VM will see 20GB and will be able to use it, growing the VM's virtual hard disk when required. This technology allows over commitment of storage, consequently slowing the requirement to purchase additional storage whilst allowing growth to occur in line with 'actual' business requirements. Whereas more traditional systems required large disk space to be allocated up front due to the complexity of growing disk volumes in operating systems, thin provisioning allows large disks to be allocated without utilizing all of that space upfront, thereby resolving complexity issues.

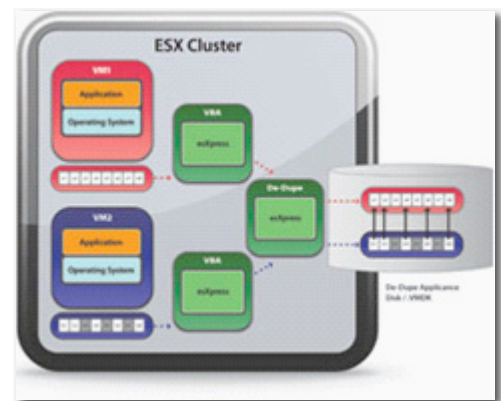


Figure 3-1: The esXpress Data De-duplication model

3.2 Data De-duplication

Data de-duplication is mostly found on higher end storage systems, which can potentially comprise Tier 1 storage. Sometimes referred to as 'Single-Instance Storage' or 'Intelligent Compression', data de-duplication strips redundant data (such as duplicated blocks of data) ensuring only a single copy is kept. Where duplicate data would have been stored a pointer is created instead, referring to the unique data copy, and thus reducing storage requirements. For example, an email system may contain 10 instances of an email, the email being 1MB in size. Normally, 10MB of space would be required to store the 10 copies of the message, however, with data de-duplication, only 1MB of space is occupied. Data de-duplication technologies also exist in some backup products. For example, esXpress by PHD Virtual, a leading VMware backup product performs data de-duplication at the source, during backup. This means

that only the required blocks of data are sent to the backup destination which saves time, bandwidth and disk space. It is important to understand that products touting data de-duplication technology should be fully investigated. Currently, there are only a few that remove 'whitespace', and this is not considered data de-duplication.

For more information on Data De-duplication and the esXpress backup product please see the PHD Virtual website at: <http://www.phdvirtual.com>

3.3 Linked Clones

Linked Clones are a more advanced version of standard cloning technologies. Clones are essentially copies of existing parent VMs which remain independent of the parent VM, whereas linked clones, are known as dependent since they utilize the parent VM for its base while still maintaining separate storage to write personalized data.

3.4 Backup Target Configuration

Storage is often configured as per a stated company standard. When storage standards are considered, resilience is normally top of the agenda. However, in certain situations, not only is resilience not required, but it can also have a detrimental effect on performance. Organizations that utilize a level of RAID for resilience will witness a trade off in performance against, say, Striping. Certain types of data are transactional and therefore do not need resilience. If data files, which are backed up to disk for near-line recovery, are to be moved on to tape, then there is no need for resilience on the backup destination disk set. This disk set should employ a RAID 0 configuration, where data is striped across all the disks in the disk set without parity, achieving the best performance possible. An organization's backup destination may also not need to be located on its Tier 1 storage - often a lower tier can be used, utilizing less expensive disks and fabric.

Information on RAID: http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks

Information on the PHD Virtual esXpress Backup Product: <http://www.phdvirtual.com/products/esxpress-virtual-backup>

3.5 VM Guest Disk Configuration

Depending on the VM guest requirements, there are various configurations that can be used to improve storage efficiency. Often, multiple virtual disks are presented to a virtual machine; the first disk stores the guest operating system and the 2nd, 3rd and subsequent disks, store application data. Taking this methodology further, a virtual disk could be added to store the guest's virtual swap file. This 3rd virtual disk, for example, could be excluded from maintenance operations and backup tasks, as the virtual memory file is often not required following a VM restore (the virtual swap file is automatically recreated by the operating system if it's missing at boot up).

Additionally, data from deleted files is frequently left on the virtual disks. Removing this data, by effectively zeroing out the whitespace on the disk is a good practice. Microsoft has a free utility called 'SDelete' that can do this for Windows guests.

For more information on Microsoft SDelete: <http://technet.microsoft.com/en-us/sysinternals/bb897443.aspx>




Figure 3-2: Microsoft SDelete Utility

A final maintenance task that should be periodically carried out on VM guests (assuming they are Windows guests) is to run a defrag. VMs should be treated and maintained just as if they are physical servers. The

only caveat to this is to ensure that any snapshots are committed. Because the defrag tool touches every block of the virtual disk, the snapshot will quickly grow. This may cause problems with filling up disks/ storage tiers (assuming multiple defrag tasks on multiple servers are running concurrently) or causing issues with commitment of the snapshot at a later date.

3.6 Shrink & Growing Virtual Disks

VMware provides the ability to shrink or grow a VM's virtual disk. This is an important and useful feature. In traditional FAT computing environments, if administrators want to move a server to a larger disk, they need to take an image or backup of the existing O/S and Application data, replace the hard disk and perform a restore to the new disk. Conversely, while it is not common, the ability to move to a smaller disk can significantly improve disk space efficiency.

In a virtualized environment, it can be easy to inaccurately assign the required virtual disk space for a guest VM. To achieve a better ROI, administrators need to closely monitor and maintain actual VM disk size in proportion to the actual required space of the disk. Administrators are able to shrink a disk from within the VMware tools client installed on the VM:

Start --> Control Panel --> VMware Tools --> Shrink Tab

The same can be achieved from the command line on an ESX host

To grow a virtual disk, administrators can use the VI client. With the VM powered off, edit the VM's settings and amend the size of the virtual disk (ESX 3.5 onwards). Afterwards, administrators need to extend the partition within the guest.

A third option for shrinking or growing a VM disk is to use the free VMware converter utility which is available for download from VMware's website.

For more information on VMware Converter: <http://www.vmware.com/products/converter/>

4.0 Summary

As storage requirements escalate across organizations, efficient storage utilization is needed to counter wasted whitespace and lower physical storage costs. Several tools are available to impact an organization's storage footprint efficiency: PHD Virtual helps with two of these. First, esXpress enables data de-duplication to prevent wasteful use of storage. Second, esXpress Backup can be configured so that the appropriate backup destination is selected depending on the requirement for data availability and performance.

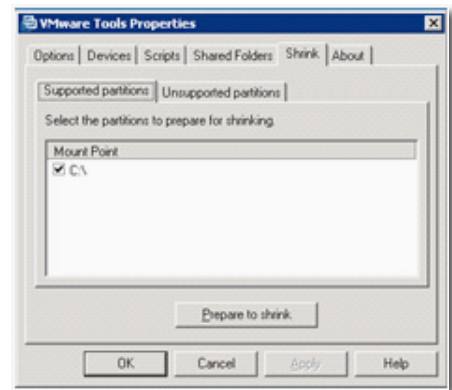


Figure 3-3: VMware tools Shrink

Scenario: Raw Disk Array 50TB (100 x 500G hard drives):

1. Subtract 1 Drive per tray for hot spare drive = 46.5TB
2. Raid 5 subtract 20% (9.3TB) = 37.2TB
3. Formatting overhead uses 3% (1.1TB) = 36.1TB
4. VMFS Working Capacity Threshold (80%: 7.2TB) = 28.9TB
5. VMDK Working Capacity Threshold (80%: 5.78TB) = 23.12TB

The total actual available storage only returns 23.1TB out of the original 50TB - a loss of 54% !

Figure 3-4: Example of Inefficient Storage Utilization

About PHD Virtual

As the pioneer of virtual backup appliances (VBAs), PHD Virtual Technologies has been transforming data protection in virtual IT environments since 2006. Its award-winning data protection solution for virtual infrastructures, esXpress, is used today by more than 1,600 enterprises worldwide to achieve unlimited dynamic growth, high availability, no single point of failure and scalable performance. PHD Virtual is committed to helping our customers and provides free, easy-to-use virtualization utilities to assist with the administration and management of virtualized environments.

North America Headquarters

111 Howard Blvd. Suite 104B
Mount Arlington, NJ 07856
Phone +1 (973) 288-7000
www.phdvirtual.com

