



The Compelling Price / Performance Economics of the Nirvanix Storage Delivery Network Versus In-house Storage

Marc Staimer, President & CDS of Dragon Slayer Consulting
marcstaimer@comcast.net 503-579-3763

Introduction

Cloud Storage has emerged as an increasingly viable storage option for IT departments and data center managers. If you are not considering Cloud Storage for some of your data, you are missing a huge opportunity to reduce storage costs and need to take another look.

One such Cloud Storage provider has optimized its network to handle the demands of large files including media, is Nirvanix (www.nirvanix.com). Nirvanix has built a global cluster of storage nodes collectively referred to as the Storage Delivery Network (SDN), powered by the Nirvanix Internet Media File System (IMFS). The SDN intelligently stores, delivers and processes storage requests in the best network location, providing the optimum user experience. This document will examine how a multi-geographical, multi-file copy, compares in terms of price and performance with the traditional model of building and maintaining corporate data centers (or hosting equipment at facilities).

o “Cloud Storage” In Brief

Cloud storage is cost-optimized networked data storage where the file-based storage exists in the “cloud” over the public Internet instead of the users’ premises on dedicated servers. The storage is rented or leased per GB on a monthly basis following a service model much like a utility. The data can be accessed and controlled through an API or through virtualized file servers that appear to the user as a large local file server. The Cloud Storage provider operates in the background providing unmatched scalable resources transparently to the customer as required. These resources may physically span across multiple servers or in the case of Nirvanix, the resources may span across multiple data centers or even continents.

Cloud Storage is not the best choice for high performance applications such as online transaction processing (OLTP). On the other hand it is ideal for storing and delivering vast amounts of unstructured file based data such as including video, music, eBooks, podcasts, webinars, etc.; Web 2.0; or even online backup and archiving services.

This leads to the question of how Cloud Storage compare on a cost / performance basis with the customers’ in-house storage? Let’s look at cost first.

Determining the Cost of In-house Storage

Over the past twelve months, Dragon Slayer Consulting has been querying end-users on how they calculate the total cost of storage. The sample size has been 67 end users ranging in storage budget size from approximately \$1 million to nearly \$1 billion. Industries surveyed included financial services (banks, securities brokers, commodities brokers, etc.) Insurance, Energy, Entertainment, Manufacturing, Life Sciences,



Pharmaceuticals, Government (Federal, State, and Local), Aerospace, and Online backup services. The data has been quite surprising and eye opening.

The first piece of information that is not reported by the trade press or tracking firms, is that most of these organizations convert total cost of storage ownership to burdened cost per GB of storage.

The conventional wisdom is overall storage costs per GB are rapidly declining. This is based on the common knowledge that hard disk drive (HDD) cost per GB has historically decreased at approximately 8% per quarter. (Although, over the past two quarters that has decreased to approximately 4%.)

In fact, conventional wisdom is completely wrong, as it often is. Overall the burdened storage cost per GB in most data centers are actually increasing or at best staying flat. This is because the HDD acquisition cost is a very small part of overall burdened storage costs.

The Makeup of Burdened Storage Cost

Burdened storage cost is measured as the annual total cost of ownership per GB of managed storage. The costs generally include:

- Acquisition cost of the storage system(s)
 - Including the acquisition cost of the HDDs
- Annual maintenance beyond the basic warranty for the storage
 - Usually 7x24 onsite within 4 hours premium maintenance
- Upfront licensing of the storage software for the storage system(s)
- Annual licensing of the storage software for the storage system(s)
- Percentage of the storage network infrastructure required for the storage system(s)
- Percentage of the upfront licensing of the storage networking infrastructure software licensing required for the storage system(s)
- Percentage of the annual licensing of the storage networking infrastructure software licensing required for the storage system(s)
- Raised floor space (cost per square foot) for the storage system(s)
- Percentage of building maintenance, utilities, physical security, and taxes allocated per square foot of space utilized for the storage system(s)
- Vertical rack space required for the storage system(s)
- Cooling for the storage system(s)
- Power for the storage system(s)
- Amount of UPS (Uninterruptable Power Supply) required for storage system(s)
- Power generation in the event of power failure
- Annual storage administrator time in hours multiplied by the cost of the administrators per hour (including salary, benefits, perks, etc.) for
 - Implementations
 - Operations
 - Management
 - Moves, adds, and changes



All of the costs are assigned, summed, and then divided by the total number of GBs to determine the annual burdened storage cost per GB. It doesn't stop here. Since there are usually two or three different tiers of storage based on price performance, this exercise is performed on each Tier.

This is an enormous amount of work to calculate every time a storage purchase is made. Most organizations work a year in arrears and instead calculate what that burdened storage cost was for the previous year. They then use those numbers for the current year budget allocations.

Please note that every organization accounts differently for burdened cost per GB of storage. They are however, generally quite similar. Below are the results:

Annual burdened storage costs per GB & Storage System Tier Context

○ *Tier 1 Storage*

- Annual burden storage costs ranged from \$40 to \$48, with the median being **\$43**.
- The Tier 1 storage systems are the organization's primary storage. They are configured with Fibre Channel or SAS drives at 15,000 or 10,000 RPM. This is because Tier 1 storage systems are used as the highest performing, most scalable, and most reliable storage. It is used for business critical applications and data. Some of the more popular vendors models the queried IT organizations considered Tier 1 included:
 - BlueArc Titan 3100, 3200
 - EMC Symmetrix DMX-4
 - HDS Universal Storage Platform V or VM
 - HP StorageWorks XP24000/XP20000
 - IBM Shark DS8000
 - NetApp FAS 6000 / V6000 series with ONTAP GX
 - SUN StorageTek 9990V System
 - 3PAR Data S400 or S800
- To a lesser extent Tier 1 storage sometimes includes:
 - EMC CLARiiON CX3-80
 - HDS AMS 1000
 - IBM DS4800
 - NetApp FAS 6070 / V6070
 - SUN StorageTek 6540 Array

○ *Tier 2 Storage*

- Annual burden storage costs ranged from \$13 to \$18, with the median being **\$15**.
- The Tier 2 Storage systems are used primarily as secondary storage for unstructured data, multi-media libraries, backup, etc. Tier 2 Storage systems are not expected to be as fast Tier 1. As such they are optimized for capacity vs. performance. De-duplication technologies have their biggest play here. It is Tier 2 Storage that ultimately competes with Cloud Storage. They are typically configured with 7200-RPM FATA, SATA, and fat SAS drives. Sometimes Tier 2 Storage systems are configured with high capacity 10,000-RPM SAS or FATA drives. Some of the more popular vendors and models considered Tier 2 included:
 - DELL/EqualLogic PS Series



- EMC CLARiiON CX3 and AX4
 - EMC Celerra
 - HDS AMS
 - HDS WMS
 - HP EVA 4000, 6000, and 8000 series
 - HP MSA series
 - IBM DS4000 series
 - IBM DS3000 series
 - LeftHand Networks NSM
 - NetApp FAS or V 6000, 3000, and 2000 series
 - SUN StorageTek 6000 and 3000 Series Arrays
 - 3PAR Data E200
- **Tier 3 Storage**
- Annual burdened storage costs ranged from \$10 to \$14, with the median being **\$12**.
 - The Tier 3 storage systems used primarily as long-term storage used for archival, near-line, even offline storage. Performance is a non-issue so these systems are completely optimized for highly scalable capacity with very low performance. The data residing on this storage is rarely accessed. All drives are high capacity low RPMs. Deduplication and MAID technologies with spin down capabilities are often utilized. Some of the more popular vendors and models considered Tier 3 included:
 - Copan Systems Revolution
 - EMC Centera
 - EMC CDL (Virtual Tape Library)
 - HP Virtual Tape Server
 - NetApp NearStore FAS and VTL
 - NEXSAN SATABeast
 - SUN StorageTek VTL

Nirvanix's Cloud Storage Costs Comparison to Tier 2 Burdened Cost per GB


Nirvanix Cloud Storage prices on a per monthly basis ranging from **\$.25** to **\$.71** per GB per month for one to three copies of geographically distributed data. They do offer lower pricing when there is no distribution and there are fewer features; however, for this comparison, the full-featured distributed capability will be used.

Tier 2 annual burdened cost per GB of \$15 translates into **\$1.25** per month per GB and with a nominal cost of money of 8% goes up to **\$1.30** per month. And that cost is per each copy of the data. This means the cost per file stored will quickly increase when data is replicated to other locations, whether local or geographically distributed.

An apples-to-apples cost comparison shows that Tier 2 monthly burdened cost per GB ranges from approximately 500% more than Nirvanix Cloud Storage to as much as 550% more for 3 copies of the data. Even if the cost of Cloud Storage is compared to Tier 3 monthly burdened cost per GB, (**\$1** / month and **\$1.05** / month with an 8% time value of money), the Tier 3 storage is more than 400% more expensively.



That's a pretty compelling economic argument and Nirvanix Cloud Storage is clearly less expensive. But now the question moves to differences in performance.

	Nirvanix Storage Delivery Network		Tier 1 Storage		Tier 2 Storage		Tier 3 Storage	
	Per Year	Per Month	Per Year	Per Month	Per Year	Per Month	Per Year	Per Month
Single file copy burdened costs ¹	\$4.00	\$0.25	\$43.00	\$3.58	\$15.00	\$1.25	\$12.00	\$1.00
(3) geo-distributed file copy burdened cost ¹	\$8.52	\$0.71	\$129.00	\$10.75	\$45.00	\$3.75	\$36.00	\$3.00

¹Burdened costs in this comparison do not take the time value of money into consideration. Bandwidth charges are an extra consideration for all scenarios.

Table 1: Burdened Cost Comparison

Nirvanix Cloud Storage Performance Comparison to Tier 2 Storage for a Distributed Audience

The key point in the comparison is the distributed audience. To look at this on an apples-to-apples basis means that the users are accessing the data over the public Internet or private Intranet, which is a TCP/IP WAN. It does not matter how good the performance is for the storage network behind the firewall when the bottleneck is that TCP/IP wide area network.

Performance is most affected by speed-of-light latency, TCP latency, (both of which are directly correlated to circuit distance between user and files), as well as packet loss. As a result, Nirvanix's geographically distributed multiple file copy Cloud Storage has a significant performance advantage. By placing the files closer to the user, both speed-of-light and TCP latency are minimized. Packet loss is also minimized because the probability of packet loss increases as distance increases. The net effect is noticeably superior performance on Nirvanix Cloud Storage versus the homegrown storage network.

If the homegrown storage is to match Nirvanix's distributed advantage, it must also provide multiple locations with multiple copies of the files. The penalty is increased costs.

Under either scenario, it leads to a very large Nirvanix cost / performance advantage.

Conclusion

For storage Use Cases where the speed of the public networks is enough such as archiving, certain backup segments, online media storage, Web 2.0 applications and many more, Cloud Storage networks with a global footprint and a unified namespace offer a price/performance value that is hard to ignore. The Nirvanix Storage Delivery Network™ will tend towards 1/5 or less the cost of building your own Tier 2 storage infrastructure. Additionally, you will likely experience better performance results with a highly capable storage network in the end.