Exploring The True Cost of Converged vs. Hyperconverged Infrastructure

FlashStack delivers all flash performance at a cost comparable to Nutanix Hybrid HCI

A DeepStorage Technology Report

Prepared for
Pure Storage

About DeepStorage

DeepStorage, LLC. is dedicated to revealing the deeper truth about storage, networking and related data center technologies to help information technology professionals deliver superior services to their users and still get home at a reasonable hour.

DeepStorage Reports are based on our hands-on testing and over 30 years of experience making technology work in the real world.

Our philosophy of real world testing means we configure systems as we expect most customers will use them thereby avoiding "Lab Queen" configurations designed to maximize benchmark performance.

This report was sponsored by our client. However, DeepStorage always retains final editorial control over our publications.

Third Edition September 2017

Contents

The Bottom Line

Over the past three to five years, Hyperconverged Infrastructure (HCI) solutions have gone from a new idea strictly for VDI and remote offices to become yet another option in the enterprise customer's bag of tricks. Like HCI, converged infrastructures simplify the design and procurement process, by pre-engineering server, storage and network components into a single item a harried IT manager can order on demand.

HCI vendors and proponents have promoted HCI solutions as being less expensive, more efficient and easier to manage than more traditional converged infrastructure solutions.

In this Technology Report we set out to see if the HCI camp's claim of lower costs actually hold water, especially as cluster sizes exceed four to six nodes. To test that hypothesis, we priced-out clusters of Nutanix NX-3060 nodes, with hybrid storage (using both SSDs and spinning disks), from four to thirty-two nodes. We compared them to Cisco and Pure Storage FlashStack systems that combine Pure's all-flash FlashArray//M storage array with Cisco UCS C220 1U servers.

The FlashStack solution was comparable in cost to the Nutanix hybrid solution, even when we assumed the Nutanix solution could match FlashStack in data reduction efficiency. When we adjusted our configurations to reflect the FlashArray's more effective data reduction and the CPU consumption of the Nutanix CVM, the FlashStack solution was lower in cost -- up to 40% lower.

The system configurations represent common configurations across the Nutanix users we interviewed:

- *Equal server configuration*
  - Xeon E5-2650 4v
  - 512GB DRAM
  - Dual 10Gbps Ethernet
- *Nutanix media*
  - 2 – 960GB SSDs
  - 4 – 2TB [HDDs]
- *Prices compared include 3yrs on-site support and maintenance*
- *Typical large customer discounts included*

Introduction

Hyperconverged infrastructure (HCI) represents a sea-change in data center design. By building scale-out storage into the same servers used for compute, HCI creates systems that scale from small ROBO or SMB solutions to conceptually whole data centers from just a few building blocks. As with all too many new technologies today, HCI's advocates call themselves disruptive.

In the past, we participated in several waves of what were believed to be revolutionary, disruptive technologies. We sold hard drive subsystems in the 1980s and were pleased to find that industry, and era, as one of the examples in Clayton Christensen's seminal The Innovator's Dilemma.

As the technology industry has come to worship the concept of disruption (though all too frequently without having read Christensen) technology vendors and their champions, including several hyperconverged infrastructure players, are all too eager to adopt the mantle of disruptive technologies.

> The key question is:
> "If hyperconverged infrastructure is truly disruptive, is it actually less expensive?"

One key aspect of the sort of disruption Christensen observed was that disruptive technologies had to enter the market at a lower price than existing solutions. Even if the new technology was initially inferior, it could address a market below the price point of existing technology as it matured. The market is finally disrupted when the new product becomes good enough, and then superior to the older technology while maintaining its lower cost.

We concede that HCI, by reducing the number of decisions and options, simplifies the acquisition process. We'll even concede that reducing server, hypervisor and storage management to a single console could add value by freeing up staff to work on other tasks—or provide quantifiable savings by enabling staff reductions. However, quantifying the precise operational and financial impact involves a level of prestidigitation beyond the skills of these humble authors, and is left as an exercise for the reader, who understands their costs much better than we ever could.

To finally settle the matter of whether HCI is less expensive than converged infrastructure, we set out to configure a hyperconverged solution and a solution based on discrete servers and storage that would be as directly comparable as possible. We then set out to calculate the acquisition cost of solutions that delivered comparable amounts of CPU to run and storage capacity to house those virtual machines.

For the hyperconverged solution we selected Nutanix's NX-3060-G5 in part because the four-node, 2U package is archetypical of Nutanix, and in part because it was the most popular node among the Nutanix customers we surveyed.

Those customers also told us they bought the four-node appliances full, and expanded their environments four nodes at a time.

For the servers-and-storage model we used Cisco and Pure Storage's FlashStack. This Converged Infrastructure teams Cisco UCS servers (in this case Cisco's C220 1U rack-mount server) with Pure Storage's FlashArray //M. By selecting a converged system that's orderable as a single unit we match HCI's advantage around ease of purchasing.

We found that the Nutanix hybrid systems were at least as expensive as the FlashStack all-flash solution especially when we tried to approximate real world data reduction rates.

Configuring The Alternatives

When building our configurations we tried to make the two systems as directly comparable as possible. In cases where we believed there were multiple options that could be reasonably argued, we chose the option most favorable to Nutanix, for the base model, as a balance to Pure Storage's sponsorship of this report.

Additional viable options are described, and where possible quantified, in the detailed discussion below.

The Comparison

The specific configuration of our Nutanix nodes is representative of what our Nutanix customers reported they were buying. The vast majority of the systems they were buying had 10 to 12 core processors and 256 to 512GB of memory with one or two terabyte hard drives.

We selected a configuration on the larger end of the range. We equipped each of our nodes with dual 12 core Xeon E5-2650 processors, 512GB of memory, 960GB SSDs and 2TB hard drives. More specifically, that's Nutanix's NX-3460-G5-22120: a block of four nodes in a 2U enclosure.

To get the server and storage configuration, we first calculated how much useable storage capacity the Nutanix cluster would deliver (see "Calculating Usable Capacity" section for details). We configured a Pure FlashArray to deliver similar capacity. We then calculated the acquisition cost of that FlashArray with an appropriate number of Cisco UCS servers.

Our solution needed Ethernet switches, so we included a pair of the Cisco Nexus 92160YC-X 25/10Gbps switches that we recently added to the DeepStorage Labs network. These 48-port 25Gbps network switches provide enough capacity to endure multiple server and storage upgrades while still consuming only a tiny fraction of the budget.

Comparing The Contenders

Following our rule about giving every arguable configuration choice to Nutanix, we first compared the systems assuming that both systems could reduce data equally. We also used the same number of servers with the FlashArray as nodes in the Nutanix cluster.

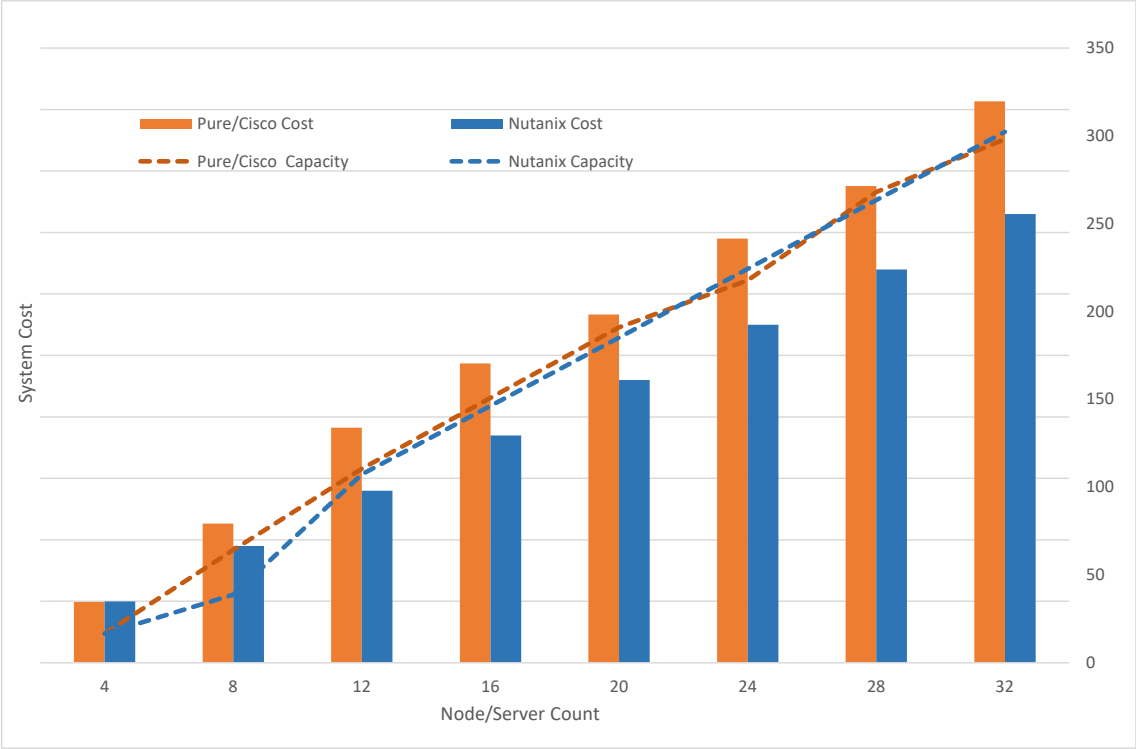

Chart 1 – Nutanix vs FlashStack Equal node count, data reduction (2:1) and useable capacity

Note that the Nutanix cluster's capacity increases dramatically when expanding from 8 to 12 nodes and the cluster size increases. Small Nutanix clusters must replicate all their data to three nodes to maintain our required degree of resiliency (see "Data Resiliency Considerations" below). As the cluster size increases Nutanix uses 6D+2P erasure coding, which is much more efficient.

This reduced efficiency in small clusters makes the Nutanix solution more expensive than the FlashStack alternative in four- where it must use 3-way replication to meet our resiliency requirement. Once the Nutanix can take advantage of the more efficient erasure coding, the cost of the flashstack system ranges from 10 to 25% higher.

The FlashStack solution delivers consistently high performance from an all-flash backend. Our Nutanix configuration has two 960GB SSDs and four 2TB hard drives per node, or about 20% of its total capacity in flash, which is slightly more than the average for hybrid systems.

While hybrid solutions like the Nutanix can deliver good performance for consistent workloads, their performance, or more specifically latency, drops significantly when use patterns change and applications access data from the spinning disks. Best-case latency for a read from the 7200 RPM disks is 10ms and that will increase dramatically if the I/O load exceeds the 400-500 IOPS/node the hard drives can deliver.

Users of hybrid systems typically see average latencies in the 2-5ms range with a significant amount of variation, especially as workloads shift and access colder data from disk as may happen when running end of period reports or seasonal applications. Users of all-flash systems like the Pure FlashArray in our comparison see not just lower, but more also more consistent latency.

Average latencies provide a decent snapshot into storage performance but a real user's performance experience is more determined by the 90th or 95th percentile latency. When workloads are shifting this 95th percentile latency is typically 10-20ms on hybrid systems like the Nutanix and under 1ms on all flash systems like Pure's FlashArray.

Basic Assumptions

We made several basic assumptions to make our hyperconverged and conventional configurations as comparable as possible:

- *The user organization will use VMware vSphere regardless of which underlying hardware model they choose.*
  - At the recent Nutanix NEXT user conference the company revealed that only 23% of Nutanix customers were running Acropolis. Because the majority of Nutanix customers use vSphere, we will too.
  - The vCenter server is the common integration point for many backup, management and other tools. Our system administrators will have some familiarity with the vCenter management console and we expect they will continue to access vCenter occasionally to manage these tools even if a simplified user interface is provided.
  - vSphere simplifies the comparison because we use the same software configurations on both systems.

Discounting

Of course only a fool, a state government, or an analyst looking for a headline would pay list price when spending over a million dollars on an IT system. Enterprise IT buyers expect significant discounts.

We applied what we believe are appropriate discounts based on the actual quotes we've seen for each vendor's products when selling to Fortune 1000 customers.

To simplify our calculations, we didn't adjust our discounts with the cluster size. This makes our comparison more valid for large enterprises that may receive deep discounts even when ordering a four node cluster for a remote office than for an SMB where that same four node cluster may be essentially all of their computing infrastructure.

Readers should note that most enterprise IT vendors discount hardware and software much deeper than they discount support and maintenance while two solutions may have the same MSRP and discount rates, if one has more support costs that option will be more expensive after discounts are applied.

Data Resiliency Considerations

Since the storage admin's first commandment is, has always been and will always be "Thou Shalt Not Lose Data" providing data resiliency is any storage system's primary function.

While each administrator must determine what level of resiliency he or she requires for their workloads, as relatively paranoid types our own standard for mission-critical workloads is that the storage system be able to survive:

- A storage controller failure and a drive failure

  OR

- An HCI node failure

  AND

- An additional drive failure or an unrecoverable read error during the data rebuild
- All without interrupting data availability or data loss.

Pure's RAID-3D combines distributed double parity erasure coding with a log-structured data layout that mitigates parity RAID's write amplification to provide this level of resiliency. Because the FlashArray only supports one data protection method there's nothing to configure or tune.

To survive both a controller and drive failure, a shared nothing system, like the Nutanix cluster and almost every other HCI solution on the market, must either replicate data across three nodes, which requires three terabytes of storage for every terabyte of data, or distribute data across a large number of nodes using parity-based and more capacity efficient erasure coding.

A Nutanix system using replication factor of 3 (RF3) and erasure coding will replicate, or mirror net-new active data to SSDs on three nodes, demoting data as it cools to a capacity pool using more efficient 4D+2P erasure code. This hybrid protection model allows Nutanix to take advantage of erasure coding's greater efficiency for most data while mitigating the impact of distributed parity's higher latency on applications.

Hyperconverged and other shared nothing storage systems can offer a lower level of resiliency by mirroring data across two nodes, and/or using a single parity erasure code, which Nutanix calls replication factor 2 or RF2. While two-way replication may provide sufficient resiliency for some workloads, we are concerned some administrators may underestimate the risk of taking a compute node off line, which eliminates storage redundancy for VMs on that host.

We used this hybrid RF3 with erasure coding mode when calculating the useable capacity of the Nutanix cluster. It's the most space-efficient mode Nutanix offers that provides a level of resiliency similar to the Pure FlashArray.

Calculating Useable Capacity

While data reduction makes any discussion of storage capacity complicated we, like customers are primarily interested in how much data we can write to a storage system, or in the case of Nutanix, a subsystem. This total useable capacity, after all the data protection and data reduction technologies have done their magic is how we compared the two systems.

Because the capacity of the Nutanix cluster was our baseline, the first step was to calculate that capacity. Luckily we didn't have to build a spreadsheet that kept track of how much flash and disk each piece of Nutanix's multi-layered distributed file system consumed. Our friends at Nutanix pointed us to designbrewz.com/ an online calculator that has all the various formulas built in.

The designbrewz.com calculator doesn't include spare space into its available space calculation. Because the Nutanix file system distributes erasure coded data across the array, a Nutanix cluster can rebuild onto any available free space after a drive or node failure without dedicating drives or nodes as hot-spares.

While the Nutanix system doesn't require dedicated nodes, there does need to be enough free space for the system to rebuild onto. To ensure there was always enough free space to rebuild after a node failure we entered N-1 nodes into the designbrewz calculator, effectively calculating the capacity for an N+1 cluster where N+1 is a multiple of four.

As we saw above, the FlashArray always maintains spare space, which is accounted for in the FlashArray's useable capacity. This puts both systems on an even footing where resiliency is concerned.

Thin Provisioning and Unmap

As modern storage systems both the Nutanix distributed file system and the Pure FlashArray use metadata structures to provide highly efficient snapshots, clones and thinly provisioned datastores. The Nutanix CVMs provide an NFS interface that vSphere uses to leverage the Nutanix file system for per-VM snapshots and VM cloning. The Pure FlashArray provides a block interface, in our comparison iSCSI over 10Gbps Ethernet. Pure has recently announced support for VMware's VVOLs (virtual volumes) to provide data services at the VM level.

While a full discussion of the relative advantages and disadvantages of NFS vs iSCSI or Fibre Channel is beyond the scope of this report, the block protocols do have one advantage related to data efficiency. Thin provisioning allows data volumes to only consume storage space as data is being written to them. A 100GB thinly provisioned volume containing 20GB of data will therefore only consume 20, not 100, gigabytes of storage space.

Any thinly provisioned storage system is efficient when data is being written, the difference is how they react when data is deleted. On iSCSI and Fibre Channel datastores vSphere uses the SCSI UNMAP command to tell the storage system that some blocks of the volume are no longer in use and the storage system marks those blocks as free to be overwritten. In vSphere 6.0 and later vSphere also passes UNMAP commands through from the guest operating system releasing deleted blocks faster and at finer granularity.

NFS has no mechanism for the guest OS, or vSphere, to identify deleted blocks within a .VMDK file; therefore thin provisioned VMs on an NFS datastore like Nutanix's will grow as data is added but won't shrink as data is deleted.

Since the amount of storage capacity that will remain consumed by deleted files is dependent on the applications used, and the age of the system we did not calculate the effect of UNMAP into our capacity calculations.

Scaling

The first generation of HCI solutions were built from a single building block with a fixed amount of compute and storage. While this approach made planning easy, it limited those HCI solutions to scaling CPU and storage at the same rate even if they only needed more capacity or memory space for more VMs.

Today, Nutanix allows customers to specify the components, from CPU to SSD, that make up each node. This gives users a range of systems from eight CPU cores and 1.5TB of raw capacity to 44 cores and over 40TB. Users can even add storage-only nodes to scale capacity without the cost of supporting VMs, including the cost of vSphere licenses.

Since Nutanix relies on data locality to maintain high performance, Nutanix storage-only nodes are only used for the second or third copy of a VM's data. Storage-only nodes provide less expensive capacity but performance will be determined by the primary node's storage devices.

When we spoke to Nutanix customers, those with clusters over eight or ten nodes generally expanded their clusters with four nodes of basically the same model they were using at a time. This is, of course, somewhat related to Nutanix's packaging of four nodes in a 2U chassis for some of their most popular models.

Our comparison calculates the useable capacity of clusters of four to thirty-two nodes in steps of four. The smaller clusters of four and eight nodes have lower useable storage per node as they use three-way replication to protect their data. Once we reach 8 nodes, each additional node adds 9.75TB of useable capacity, assuming 2X data reduction, 480GB of RAM and 20 cores of compute to the cluster.

HCI supporters make a good case for their linear, scale-out model compared to servers and traditional dual controller arrays from traditional vendors.

With these traditional systems, you had to buy a new array when you reached the maximum capacity, or performance, of your array's controllers, and re-buy all the software options if you could connect your existing media. If you were unlucky you'd get caught in a hardware change and have to replace the media too, adding the joy of data migration to the big bill for the new array.

While the Pure FlashArray has a dual controller architecture similar to the traditional arrays, the folks at Pure decided to not chase the big check for a new array from their customers every three years, instead designed the FlashArray//M chassis to have a longer life than those old timers – aka traditional storage arrays. The FlashArray chassis uses shared NVRAM modules for their cache, which allows a customer (OK a Pure-trained technician because we want someone to blame), to upgrade the controllers, one at a time, without the complications of maintaining a coherent cache between dissimilar controllers.

In our comparison we specified FlashArray//M20 systems, which Pure rates at 200,000 IOPS, from 4-20 nodes and the beefier //M50s, which Pure rates at 270,000 IOPS, for the larger clusters. Under Pure's Evergreen Storage policy customers with active maintenance agreements can upgrade from the //M20 to the //M50 by paying just the difference in price between the two.

This upgradeability extends all the way to Pure's newest FlashArray//X which replaces SAS SSDs with PCIe based NVMe flash modules. They can non-disruptively upgrade //M to //X controllers and migrate data from a customer's current SAS SSDs to the NVME SSDs in the background. Pure's Evergreen Storage program also includes a trade program providing a 25% GB for GB purchase credit. The FlashArray//X can provide a big performance bump without requiring one to rebuy a whole array--or more importantly, migrating data between arrays or over the storage fabric.

The FlashArray's wide range of expandability and Pure's customer-friendly Evergreen Storage policies make the two systems roughly comparable in scalability across the range of 4-32 nodes and capacities up to 300-400TB. The Pure solution is a bit less linear, as they sell SSDs in packs of ten, making capacity steps bigger than the 9.75TB/node of the Nutanix. As an example, to match the 20-node Nutanix cluster's capacity we used a solution with 30% more capacity and the same system for the 24-node comparison in Chart 2.

Adjustments And Other Considerations

As we were building our ultimate configurations we made a series of design decisions that could have a significant effect on the cost of one or the other of the solutions.

Data Reduction

Data reduction, more specifically data deduplication, has been at the core of the FlashArray architecture from its inception. Data reduction is an always on, almost always inline process. Pure keeps a ticker on their website showing the average data reduction across their install base, which at the time of this writing was reporting just over 5:1 data reduction.

Compression and deduplication are more, though admitedly not very, recent additions to the Nutanix distributed file system that Nutanix recommends be enabled or disabled for individual workloads. Data reduction, a compute intensive process, can be challenging for hyperconverged systems where CPU must be shared with user workloads. Nutanix has addressed this with a collection of inline and post process reduction at the three-way replicated performance tier and the erasure coded capacity tier that combine to balance CPU and memory usage against storage efficiency.

Our understanding of data reduction technologies leads us to believe that the FlashArray's inline reduction will be more effective than the Nutanix methods at reducing data for several reasons.

First, a Pure FlashArray deduplicates all the data written to it as a single deduplication realm, uses deduplication block sizes as small as 512 bytes and uses a multi-stage compression mechanism, all of which will lead to very efficient reduction.

Nutanix, by comparison, uses 16KB deduplication blocks to deduplicate data across a container or datastore. Nutanix clusters with multiple containers will constitute multiple deduplication realms with data deduplicated within each realm but not across realms. As for compression, Nutanix uses a slightly less aggressive method, trading storage efficiency for CPU efficiency.

The biggest negative impact on Nutanix storage efficiency is that customers following Nutanix recommendations will only enable data deduplication for selected workloads where deduplication is especially effective, such as physical to virtual (P2V) conversions.

Of course the degree to which a system can reduce data is very dependent on the data being reduced. Thus, quantifying how much better the FlashArray will compress any given customer's data would be a challenge.

The best information we have comes from the vendors themselves. Note that Nutanix employees downplay their data reduction; one Nutanix blogger recommending that customers size systems without factoring in data reduction and take any savings as gravy. The semi-official designbrewz.com site allows you to choose 1.5:1 or 2:1 data reduction.

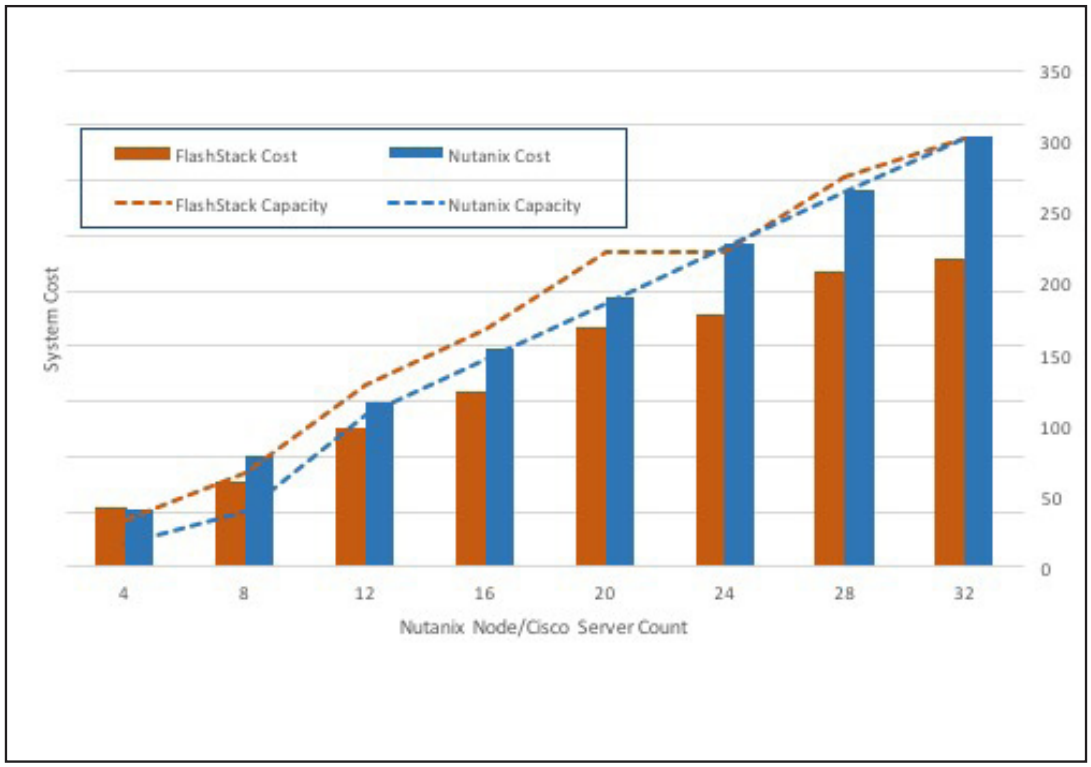

Chart 2 – Cost and Capacity with Pure Reduction at 4:1

Assuming a higher data reduction rate allowed us to configure a smaller FlashArray for each capacity point in our comparison.

As expected, increasing reduction efficiency reduced the cost of the FlashStack making the Nutanix solution consistently the more expensive solution with the exception of the smallest 4 node configuration where Nutanix has a small cost advantage. However, the

FlashStack doesn't scale quite as linearly as the Nutanix, especially in the transition from 16 to 20 nodes where the most logical step up on the FlashArray was a relatively large 22TB datapack.

Most significantly, while Nutanix does scale more linearly, the slope of the Nutanix line in Chart 2 is steeper than the FlashStack's somewhat more jagged path.

Equalizing Compute Resources

Hyperconverged solutions use the same processor to manage storage and run user workload virtual machines. This means that any CPU cycles or host memory used for storage management or data services are unavailable for user VMs. This reduces the total number of user VMs any given host can support.

By default, the Nutanix controller VM (CVM) is assigned 8 vCPUs and 16GB of RAM. When deduplication is enabled, Nutanix's official recommendation is to increase available RAM to 24GB. When we talked to Nutanix users in the field, they recommended 32GB of RAM for the CVM.

Our database of virtual server statistics tells us the average VM uses 2 vCPUs and 8GB of RAM. Our hosts with 48 execution threads (2CPUs x 12Cores x 2Hyperthreads/core) and 512GB of RAM will support 64 average VMs before overcommitting memory but only 24 average VMs before overcommitting CPU threads. Since these systems are going to run out of CPU before they run out of memory, we're going to ignore the Nutanix VM's memory impact.

Since each server/node has 48 threads and the Nutanix VM uses 8 threads the Nutanix VM consumes 1/6th of the total compute capacity of the server. To create Chart 3 below we equalized the number of CPUs available to run virtual machines by using 6 Cisco servers for every 7 Nutanix nodes.

Since our cluster sizes scale in increments of four nodes the ratio in Chart 3 below is closer to 7 Cisco servers to 8 Nutanix nodes. As you can see in Chart 3 below this adjustment is small at about 5% of the overall price.

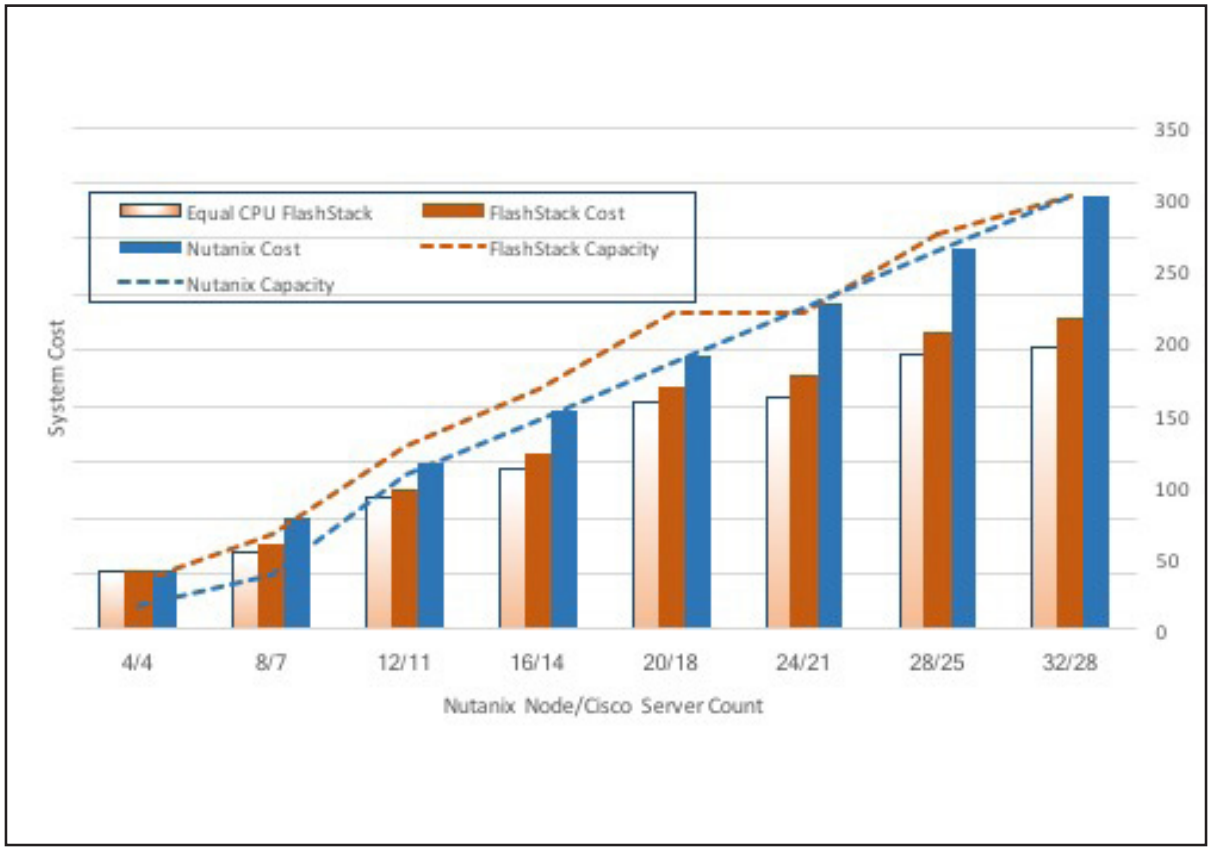

Chart 3 Cost and Capacity with Equal CPU and Equal Data Reduction

To calculate the cost of each server we added the cost of the typical software load for a server, including vSphere Enterprise Plus and Windows Data Center edition, to the cost of the Cisco UCS C220 server, as well as three years of maintenance on everything. With these well-equipped servers, the software and hardware costs were about equal after we applied typical enterprise discounts.

Making It Up In OpEx

Our analysis concentrates on cost of acquisition, so it includes three years of support and maintenance and installation services, as those are included in the initial purchase of almost all enterprise IT products.

We did not attempt to quantify how much, if any, time either system would save in regular maintenance as we've found as consumers of such research that the author's assumptions frequently have so much influence on the results as to make the process unreliable.

Basically, we believe the IT Pros that read this report know their operational costs better than we could guess at them.

Many HCI proponents argue that HCI systems are so much simpler to operate that the resulting reduction in operating expense would be so significant that customers would be foolish to continue using a server and storage model.

Sure, modern HCI solutions are significantly easier to manage than the Fibre Channel SANs and disk arrays of the early 2000s. The problem with using that fact as an argument for significantly reduced OpEx is that modern iSCSI attached storage arrays like the Pure FlashArray//M are also significantly easier to operate than the Fibre Channel Clariions of old.

The question users should be asking isn't which system is less expensive to operate but whether the difference between the two systems is significant enough to make a difference.

Some analysts argue customers should ignore OpEx savings unless the IT department is prepared to reduce headcount. That might be true in a tightly siloed organization where admins have a narrow band of responsibility and can't or won't work outside of their specialty. More progressive IT departments will benefit from OpEx savings of less than an FTE but even they will not get the additional value implied by simply multiplying the number of seconds saved performing common tasks times the task's frequency.

In this section we'll look at how both systems would be managed to see if we really can make up any cost difference in OpEx.

Integrated Host Management

Hyperconverged appliance advocates, including Nutanix, contend that HCI simplifies management by merging hypervisor, storage software and server firmware management into a common user interface. This is true. Nutanix users can manage the Nutanix hardware, from node deployment to firmware updates from within Prism. However, Nutanix has not integrated the server and storage management features of Prism into VMware's vCenter.

vSphere users, who still make up a significant majority of Nutanix users, will have to switch back and forth between the vCenter management console and Prism. Prism does provide some basic information on the vSphere environment but we don't consider it a full replacement for vCenter especially in organizations run vSphere already.

Maintaining the Cisco/Pure Storage system is a bit more complex but not the nightmare of yesteryear that involved some poor junior sysadmin standing in front of a server for an hour loading each firmware patch from a separate CD. As we discuss in A Word on Cisco UCS below, UCS Manager is the market leader in server management.

UCS Manager integrates into vCenter, giving vSphere administrators the control they need over the server hardware and boot environments in a familiar management environment. Our administrator will still need to update firmware on the FlashArray periodically. That simple task does require an additional user interface.

We assume users of either system will use vCenter and VMware Update Manager to maintain their host software environment. While some organizations may choose to use Nutanix's Prism to manage ESXi patches, we believe the effort of installing VUM to be worthwhile even if only to manage third party management agent and driver updates.

Organizations with experienced vCenter administrators will like the UCS approach as it keeps them in the console they call home. Organizations without established vCenter skills will find the Nutanix approach easier and may have to find other solutions to manage any servers they retain outside the Nutanix cluster.

As a worst case we estimate the organization deploying the Cisco/Pure solution would have to spend one person-day installing and configuring UCS Manager and perhaps as much as an hour per server per year.

Over the 3 year useful life of the proposed systems that's 8 hours for installation and thirty-three hours of additional upgrade effort a total of forty-one hours or $20,500 at the premium

Host Upgrade Considerations

Both solutions provide for non-disruptive upgrades of host hypervisors through vSphere which facilitates migrating workloads off a host before the upgrade and via DRS, rebalancing the cluster. Users on the Nutanix solution will see a greater performance impact as removing a node for maintenance reduces the total storage performance available to the cluster.

In addition since Nutanix relies on data locality for storage performance VMs migrated to hosts without a copy of that VM's data will have reduced storage performance as data populates to its new home.

In addition to reducing storage performance removing a node from the cluster temporarily to perform an update also reduces the cluster's resiliency. This is one of the reasons we recommend N+2 resiliency for all shared nothing storage systems. If a node remains offline for an extended period of time it will rebuild the missing data, consequently generating both storage and network I/O.

Integrated Storage Management

Many HCI customers cite simplification of storage management as one of the key benefits of their HCI solution. Prior bad experiences with Fibre Channel, arcane storage system CLIs or a hidebound storage administrator has put them off managing external storage. While these customers weren't looking, the state of the art for external storage management has come a long way.

One way we simplified storage management was to use iSCSI over 10Gbps Ethernet, rather than Fibre Channel to connect our servers to the Pure FlashArray in the FlashStack. Modern storage systems like the FlashArray integrate into the VMware ecosystem through APIs like VAAI and a vCenter plugin. All the admin needs to do is create

a few volumes on FlashArray and connect the ESXi hosts to them. Since the virtualization administrator has full and simple control of the storage system from inside vCenter there's no real need to access the FlashArray's user interface.

By comparison, administrators on the Nutanix system have to log into the Prism user interface to create and manage containers. When creating those containers, administrators have to choose protection levels, and select whether features like compression, deduplication and erasure coding should be enabled on each container. Pure's FlashArray has always-on deuplication, compression and RAID-3D erasure coding. All the administrator has to do is specify the size for a LUN and create it through the vCenter plugin, which will automatically mount it on the hosts.

Both systems integrate and simplify storage management sufficiently that the total time spent performing storage management tasks for this system are on the order of 4-10 person days/year and therefore insignificant compared to the total million dollar price tag.

As this report went to press Pure's VVOL implementation was entering a public beta test cycle. The VVOL/VASA APIs allow vSphere to leverage the FlashArray's native data services like snapshots and clones on a per-VM granularity. When Pure delivers VVOLs support, storage administration will get even simpler as the entire array can be one datastore and administrators can use VMware's SPBM to manage volumes.

Rackspace Power And Cooling

Our choice of the Nutanix four-node, 2U block gave Nutanix a big advantage in density at two servers per rack unit. The Cisco C220s alone take up twice as much rack space, and the FlashArray will take up another 4-12U depending on the number of SSD shelves.The Nutanix website shows 1150W as typical consumption for the NX-3060 blocks in our comparison, while the Pure site shows a FlashArray//M50 we use in our larger configurations as using 650-1280W. Cisco's UCS Power Calculator configured like the ones in our model use 515W at 75% utilization.

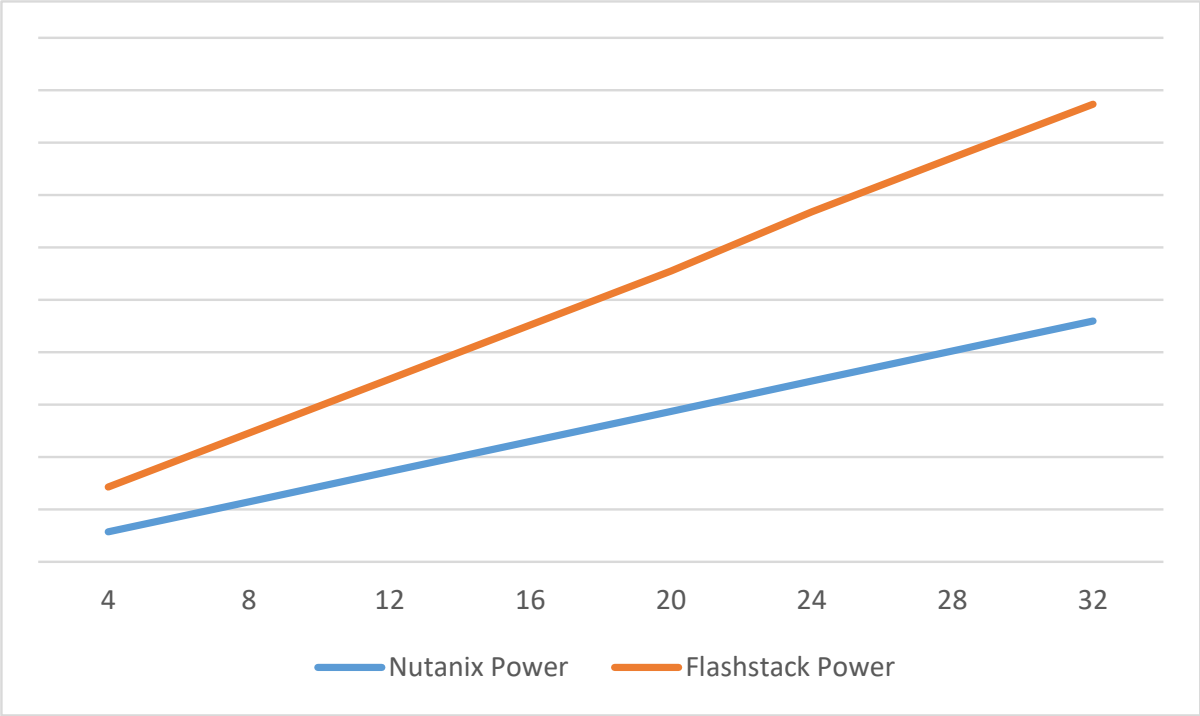

Chart 4 – Power Consumption

Our calculations show the FlashStack solution uses significantly more power than the Nutanix cluster. This gives Nutanix the advantage in both power and space.

We calculated the cost difference between the two solutions assuming the national average power cost of 11 cents per kilowatt hour and a data center PUE of 2. Over three years this possible cost saving never exceeds 4% of the initial purchase cost.

We did not include this cost of power calulation in the comparison charts 1-3 as power is not reliably a cost IT departments are concerned with. Over our consulting career we worked with clients where data center power was on the IT budget, the facilities budget or simply included in the clients rent regardless of the amount of power IT consumed.

Note: In the first edition of the report we used 1150 watts for the power consumption of a Nutanix node when it is actually the value for a four node block. We regret the error and have corrected the chart and text above.

A Word on Cisco UCS

Ever since its introduction in 2009, Cisco's Unified Computing System has led the market in server management. While other first-tier server vendors including Dell and HPE have server management suites, these tools are optional add-ons that simplify firmware updates and other day-to-day tasks for otherwise independent servers.

Service Profiles and Stateless Computing

One of the major benefits that Cisco UCS brings to the datacenter is stateless computing. Rather than having pre-configured NICs, MAC addresses, and/or WWNs, a Cisco UCS server acquires its configuration from UCS Manager through a service profile that's associated with that server's function.

This stateless configuration allows a user organization to replace failed servers or repurpose servers in a matter of seconds. Service profiles extend beyond simple BIOS settings but also ensure that network resources are properly configured to provide the connectivity a server will need in its new role.

The value of a stateless computing model becomes more apparent as the size of the UCS environment grows and servers can be dynamically assigned roles. Because the configurations we're comparing in this report are made up of a single cluster, we're not assigning any significant value to UCS management.

Ongoing Support And Maintenance

Our comparison includes the cost of a three year, 7x24x365 support and maintenance contract with 4 hour response.

Most IT vendors include three years of support in their initial quote. Historically, by the end of three years advances in technology make upgrading attractive despite the pain of migrating to a new system. While newer, faster, smaller systems are attractive to the IT department, finance may want to get just a bit more life out of the organization's investment.

While legacy storage vendors frequently increase the cost of maintenance in year four or five to encourage users to buy new storage arrays, Pure Storage's Evergreen Storage program allows a customer to keep their storage system under maintenance at the same rate as long as they like and even includes free upgrades for a system's controllers with a three year support renewal.

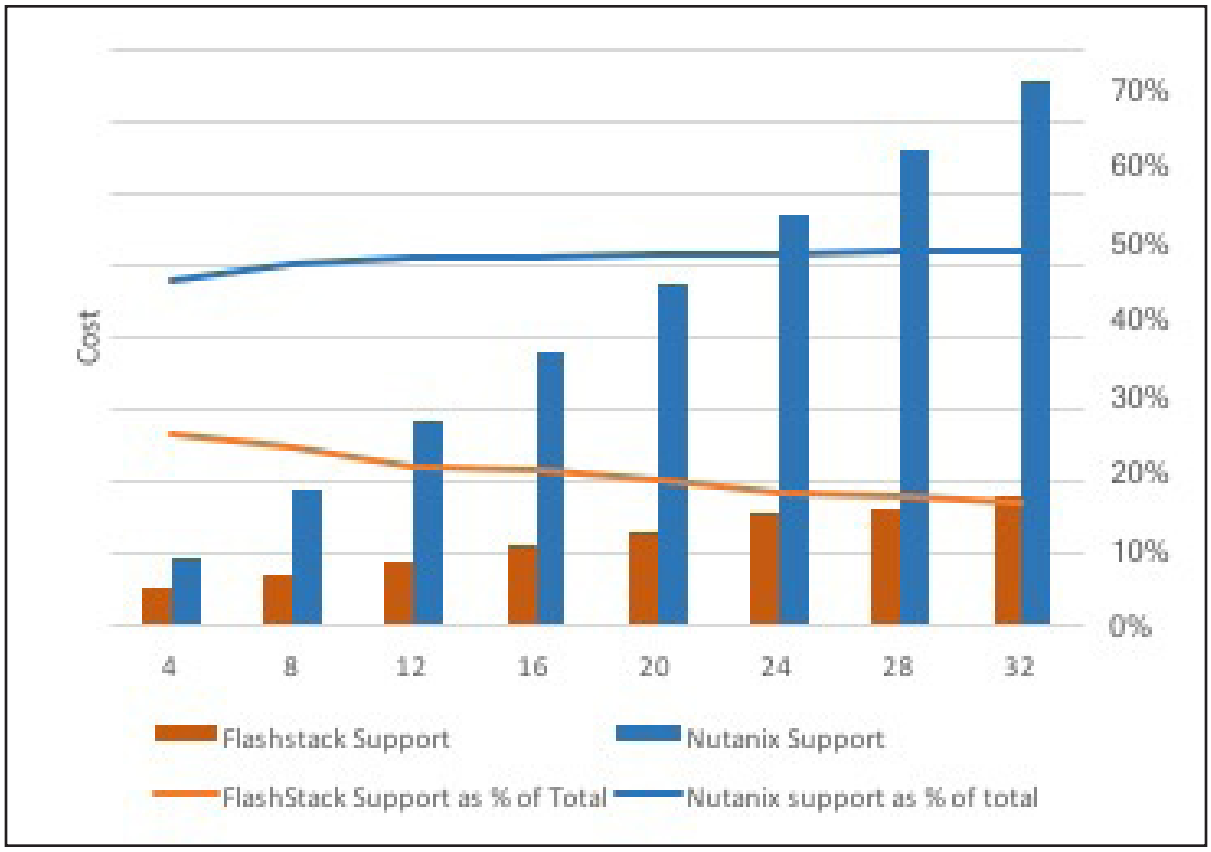

Chart: 5 Annual Support Costs

Vendors generally discount hardware and software significantly more than they do support and maintenance. As a result, maintenance makes up a higher fraction of the discounted price as the negotiated discount increases.

When comparing our two solutions, maintenance makes up almost half the total cost of the Nutanix solution and for the FlashStack solution decreases from about 25% to under 20% as the configuration grows.

Conclusions

The proponents of HCI argue that the HCI architecture leads to systems that are inevitably simpler, and less expensive, to purchase, install and manage. To support this assentation HCI vendors play up the "savings" their customers have achieved by replacing their traditional solutions with new HCI systems. It is of course easy to make any current technology look good compared to the 5 year old technology customers are retiring from their data centers.

These comparisons ignore that any current server with dual 10-core processors will replace five old servers with dual or quad-core Xeons and that the current generation of storage systems are vastly simpler to manage than the disk arrays of old.

We actually compared the cost of a typical HCI solution from Nutanix, a leading vendor in the HCI space, to FlashStack, which Cisco and Pure Storage call a next-generation converged infrastructure solution. This comparison revealed that the all-flash FlashStack solution was at worst approximatly 25% more expensive than the hybrid Nutanix alternative.

When we adjusted our model to account for the greater data reduction capabilities we believe the Pure FlashArray would provide and the CPU consumption by Nutanix' storage CVM, the FlashStack solution was as much as 40% less expensive than the Nutanix.

About this Technology Report:

This report originated with Howard Marks' suspicion that HCI solutions were actually more expensive that conventional server and software systems. Pure Storage agreed to sponsor more extensive research resulting in this report.

Both Pure Storage and Nutanix consider their pricing confidential. To prepare this report we acquired a copy of Nutanix price list and several quotes for Nutanix gear from our contacts at large enterprise customers. We used this data to calculate the list price and appropriate discounts for the Nutanix gear. The list price and discounts for the Cisco gear was also calculated from actual customer quotes.

Pure Storage provided us with hardware and maintenance list prices for the various FlashArray configurations used in our comparisons and suggested the discounts we applied as being typical for large enterprise customers.