TANEJA GROUP
TECHNOLOGY ANALYSTS

TECHNOLOGY PROFILE

HYPERCONVERGED INFRASTRUCTURE POWERED BY PIVOT3

Benefits of a More Efficient HCI Architecture

MARCH 2016

THE NET NET

Hyperconverged infrastructure is one of the fastest growing IT enterprise technologies and has already achieved nascent, yet significant market penetration that continues to accelerate. IDC defines hyperconverged systems as *an emerging technology that natively collapses core storage, compute, and storage networking functions into a single software solution or appliance.* Gartner predicts that HCI will represent more than 35 percent of total integrated system market revenue by 2019, building towards 67 percent of a $35 billion market in 2021. The widespread adoption and success of virtualization technologies has been driven by an insatiable desire to make IT simpler, faster, and more efficient. IT departments can no longer afford the time and effort required to create custom infrastructure from best-of-breed DIY components as they are increasingly required to support business outcomes.

With HCI, the traditional three-tier architecture has been collapsed into a single system in which the hypervisor, compute, storage, and advanced data services are integrated into x86 industry-standard building blocks. The immense success of this approach has led to increased competition in the HCI space, forcing customers to sort through the various offerings and analyze key attributes to determine which are critical to business success.

One of these competing vendors, Pivot3, was founded in 2002 and has been in the HCI market since 2008, well before the term "hyperconverged" was used. For many years, Pivot3's unique architecture has provided the most efficient scale-out Software-Defined Storage (SDS) system available on the market. This efficiency is attributed to patented design innovations. Pivot3's extremely efficient and reliable erasure coding technology called Scalar Erasure Coding can deliver significant capacity savings depending on the level of drive protection selected; in a full 16 node array the system achieves an industry-leading 94 percent usable installed capacity effectively available for storage. Conversely, many leading HCI implementations use replication-based redundancy techniques which are heavy on storage capacity utilization.

Scalar Erasure Coding also helps Pivot3 to unify hardware resources – RAM cache, SSD and HDD – across nodes and wide stripes data across all unified hardware to form a cross-cluster virtual SAN, which Pivot3 calls the HyperSAN. Optimized for performance, all IOPS and bandwidth are aggregated as new nodes are added. In addition to performance, the HyperSAN ensures maximum protection and resilience are always maintained, so during a failure event, a Virtual Machine (VM) migrates to another node and continues operations without the need to divert compute power to copy data over to that node. One of the founding principles of Pivot3's architecture is the concept that all nodes should form part of the array and contribute to overall performance with the lowest overhead possible. A low CPU overhead is critical to hyperconvergence as it is required to implement the SDS features and other VM-centric management tasks. Implementation of the HCI software uses the same CPU complex as business applications; this additional usage is referred to as the HCI overhead tax. HCI overhead tax is important since the licensing cost for many applications and infrastructure software are based on a per CPU basis. Even with today's ever-increasing cores per CPU there still can be significant cost savings by keeping the HCI overhead tax low.

The Pivot3 family of HCI products deliver high data efficiency with a very low overhead and are an ideal solution for storage-centric business workload environments where storage costs and reliability are critical success factors. One example of this is a VDI implementation where cost per seat determines success. Other examples would be capacity-centric workloads such as big data or video surveillance that benefit from Pivot3's leading storage capacity and reliability. In this paper we compare Pivot3 with other leading HCI architectures based on data extracted from other HCI vendor's reference architectures for VDI implementations. Using real world examples, we have demonstrated that with other solutions, users must purchase up to **136 percent more raw storage capacity** and up to **59 percent more total CPU cores** than are required when using equivalent Pivot3 products. **These impressive results can lead to significant costs savings**.

CONSIDERATIONS FOR IMPLEMENTING A HYPERCONVERGED INFRASTRUCTURE

Figure 1: Traditional Three-Tier versus HCI Architectural Benefits

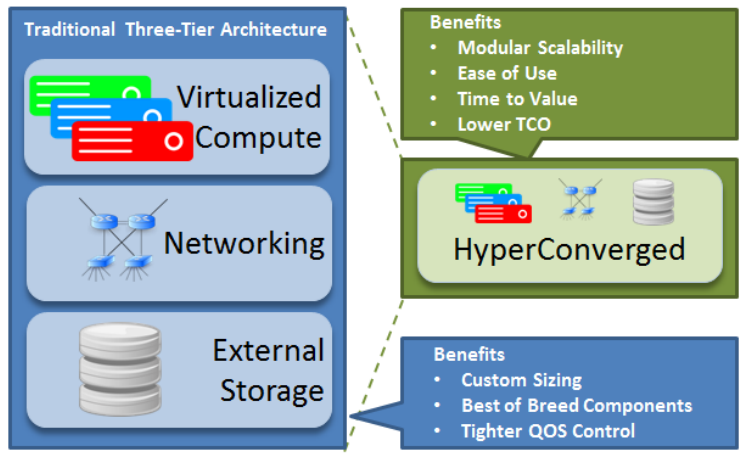

As the HCI market matures, each vendor will attempt to create new product attributes to differentiate themselves. This is natural in a maturing market and forces businesses to prioritize the most important features for ensuring a successful HCI deployment. Figure 1 outlines the core benefits of using an HCI architecture over traditional three-tier architecture. The assumption is that the reader has already selected the HCI approach over a traditional three-tier approach. Therefore, when selecting between the numerous HCI vendors/offerings, Taneja Group believes products should first meet a set of essential requirements. After verifying those requirements are met, we suggest that businesses dig deeper and evaluate a second set of requirements which ensure a more robust and cost effective solution. Both of these important considerations are discussed below.

Ensure your HCI solutions meet these essential requirements:

- **Time to value via modular scalability:** Upgrades and expansions must execute in hours - not days, weeks and months. Modular scalability saves businesses time and money and removes financial risk associated with large capital purchases.
- **Performance that scales linearly:** Performance sizing must not be a mystery. Predictable performance gains must come concurrently with adding capacity.
- **Flexibly-tuned modular scalability:** Ideally, HCI solutions should support business application workloads of all shapes and sizes. The system must be flexible enough to optimize capacity and performance in the CPU, memory and storage dimensions. For example, a VDI solution requires a different infrastructure optimization than a big data solution.
- **VM-centric ease of use:** HCI solutions that provide a VM-centric management approach should approach self-service ease of use, removing the need for IT specialists to manage the system.
- **Seamless upgrades and technology refreshes:** Ensure that the HCI solution allows for upgrades and technology refreshes that are seamless with no downtime.
- **Single point of support:** HCI solutions providers should be able to offer a single support experience and be knowledgeable regarding both the HCI hardware and the end-to-end virtualization software infrastructure that is delivered with the product.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

Additional HCI considerations that should be evaluated:

These advanced HCI considerations are typically nuances that each business must decide whether they are critical to the success of their deployment. Many differences between solutions will be exposed in the scale-out storage technology. The following list of considerations is not intended to be exhaustive but should give a sense of areas to explore before making purchase decisions.

- **Support for heterogeneous virtualized environments:** There are many virtualized environments that support HCI solutions. Some vendors focus on just one virtualized environment and others support multiple hypervisors. Prioritize how important it is to support VMware, Microsoft and KVM environments. Additionally, evaluate whether hypervisors can be mixed or changed.
- **All-inclusive hypervisor support:** Some vendors are embracing an open source version of KVM in order to integrate a full end-to-end approach and eliminate the purchase of additional hypervisor technology. Evaluate the cost savings of this approach and whether particular application workloads are appropriate for KVM.
- **Storage capacity efficiency and reliability:** Evaluate efficiency of the underlying storage service. Does the software rely on a RAID controller for hardware reliability? How efficiently does the storage software convert raw capacity to useable capacity and at what level of a failure domain? Erasure coding not only provides significant gains in capacity and efficiency, it also ensures consistent higher performance than a simpler replication-based scheme.
- **Advanced data services:** Several HCI vendors provide advanced data services of which deduplication, replication and integrated backup are examples. Evaluate the capabilities and cost of these integrated features against alternative approaches provided by other software providers.
- **HCI overhead tax:** HCI solutions inherently use compute resources that could also be used to run business applications. Evaluate whether overhead consumed by the HCI infrastructure is worth the extra features. Much infrastructure software (VMware and vSphere, for example) is licensed on a per CPU basis. Evaluate whether those costs would be significantly higher considering the overhead taxes of particular HCI solutions.
- **Shared storage performance and data locality approach:** Evaluate how performance of shared storage changes as VM's are moved across the cluster. Some HCI approaches take a significant performance hit if VM data is located on certain nodes. Other HCI approaches keep performance consistent across all nodes, but might not optimize performance when data is local to a particular VM.

PIVOT3 VSTAC TECHNOLOGY FOR HCI SOLUTIONS

Pivot3's HCI products are designed and optimized for specific business workloads. Each Pivot3 HCI appliance is pre-configured with the HyperSAN – the critical storage virtualization software. This software enables direct attached storage resources to join together to create the equivalent of an external shared SAN array. Pivot3 was one of the first companies to focus on storage efficiency by developing erasure coding for its core storage foundation. Erasure coding eliminates the need for replication or hardware-assisted RAID for core reliability by enabling a flexible software-driven protection scheme. The erasure coding scheme works across a variety of disk technology and server nodes with ultra-efficient storage overhead. Pivot3 has spent years optimizing their patented technology called Scalar Erasure Coding. This gives them a competitive edge for storage-centric HCI workloads such as video surveillance where storage costs greatly impact solution affordability.

Figure 2: Pivot3 HCI HyperSAN

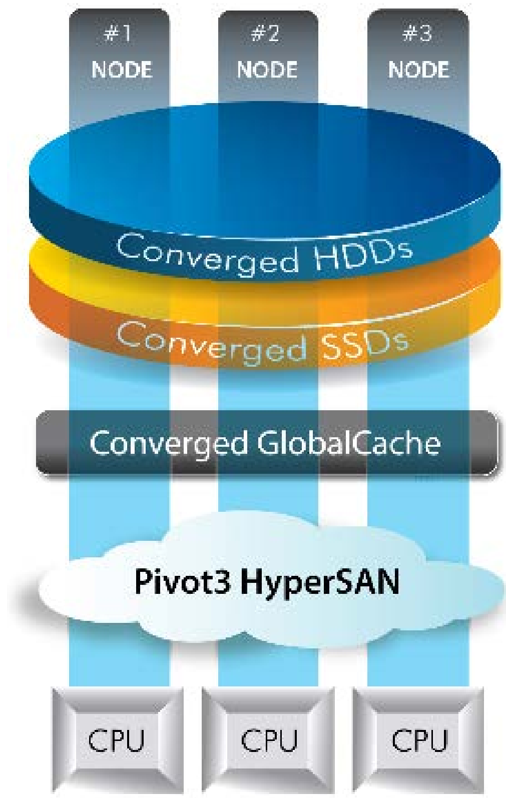

The other area of focus for Pivot3 was increasing storage performance while minimizing CPU overhead required to run the system. Pivot3 optimizes performance through a technology they call Converged GlobalCache. This innovative caching scheme can converge server DRAM with flash and is fault tolerant on writes in a cluster-coherent way. This novel approach allows for very consistent performance across the entire compute cluster. Other performance enhancements are an automatic load-balancing algorithm that combines with an ability to bypass the hypervisor-emulated IO. This creates a direct pipeline to local disks, resulting in a very high-performing shared storage solution requiring very little CPU overhead.

Storage efficiency using Scalar Erasure Coding in depth

Erasure coding is not a new technology – it has been used for over 50 years. Erasure coding was traditionally associated with an error recovery technique called Reed-Solomon which was and is still used in many removable storage devices to correct for inherent media defects. At its core, erasure coding is a method for maintaining data reliability in which information is broken into fragments, expanded and encoded with redundant data pieces, and stored across a set of different physical locations, such as disks, nodes or even geographic locations. Erasure coding enables data that later becomes corrupted to be reconstructed by using information about the data that has been stored elsewhere in the storage failure domain. Erasure coding can be used as an alternative to traditional HW RAID because of its ability to reduce the time and overhead required to reconstruct data. The key benefit over HW RAID is the fact that the failure domains can be software defined and thus failure can occur across disparate technology and failure locations. In addition, the failure domains can be optimized for storage efficiency and/or hardware repair rebuild time. Hardware-based RAID has a localized failure domain and rebuild times are proportional to the performance of individual disks which can be quite long. The historical drawback of erasure coding for primary storage is that it can be CPU-intensive, and that can translate into increased latency.

Erasure coding has become a popular technique when implementing data protection for scale-out distributed object storage systems. However, few vendors have been able to implement erasure coding effectively for primary storage as the engineering complexity combined with the mathematics-based protection formulas decrease CPU efficiency. Most scale-out primary storage vendors have fallen back into the less complicated but wasteful protection approach of replicating data. Pivot3 has been implementing erasure coding for primary storage since 2003. Not only have they enhanced it through the years but they have also been able to implement the algorithms in a CPU efficient manner.

Pivot3 has defined 3 levels of EC protection summarized in the table below:

Pivot3 Scalar Erasure Coding Protection Levels	
Protection Level	**Event Recovery Description**
Scalar Erasure Coding Lvl.5	- 5 simultaneous disk events or - 2 drives and an entire appliance
Scalar Erasure Coding Lvl.3	- 3 simultaneous disk events or - 1 drive and an entire appliance
Scalar Erasure Coding Lvl.1	- 1 simultaneous disk events or - 1 drive and an entire appliance

Scale-out systems ideally should recover from both an individual drive failure and an instance of the whole appliance node going down. In order to do this efficiently, Pivot3 has virtualized the failure domains across both individual disks and nodes in such a way that the entire process is flexible and dynamic. A business can choose one protection level for a data volume and instantly change it to another based on changing business needs. Pivot3 has built in a virtual spare drive so that rebuilds can start immediately upon failure. Pivot3 also provides a predictive failure mechanism so that a rebuild can start before an actual drive failure. The table below illustrates the usable terabytes of various protection schemes and drive capacities by node quantities. Also included is a virtual spare drive.

Pivot3 HCI Appliance Usable Capacity Chart (12 Drives per Appliance)																
Scalar Erasure Coding Protection:																
	Lvl.1					Lvl.3					Lvl.5					
# of Appliances / Drive Size	1TB	2TB	4TB	6TB	8TB	1TB	2TB	4TB	6TB	8TB	1TB	2TB	4TB	6TB	8TB	
1	9.8	19.8	38.5	59.2	77.0	8.9	18.0	35.9	53.8	71.8	--	--	--	--	--	
3	22.5	45.3	90.6	135.8	181.2	20.2	40.6	81.3	121.8	162.6	18.3	36.8	73.7	110.4	147.4	
4	34.2	68.8	137.8	206.6	275.6	30.6	61.7	123.5	185.0	247.0	27.8	55.9	111.8	167.6	223.6	
5	46.0	92.5	185.3	277.8	370.6	41.2	82.8	165.9	248.6	331.8	37.3	75.0	150.1	225.0	300.2	
6	57.8	116.3	232.9	349.2	465.8	51.7	104.1	208.4	312.4	416.8	46.8	94.1	188.5	282.6	377.0	
7	69.6	140.2	280.6	420.8	561.2	62.3	125.3	250.9	376.2	501.8	56.3	113.3	226.9	340.2	453.8	
8	81.5	164.0	328.4	492.4	656.8	72.8	146.6	293.5	440.2	587.0	65.9	132.6	265.4	398.0	530.8	
9	93.4	187.9	376.2	564.0	752.4	83.4	167.9	336.2	504.0	672.4	75.4	151.8	303.9	455.6	607.8	
10	105.2	211.8	424.0	635.8	848.0	94.0	189.2	378.8	568.0	757.6	85.0	171.0	342.4	513.4	684.8	
11	117.1	235.7	471.9	707.6	943.8	104.6	210.5	421.5	632.0	843.0	94.5	190.3	381.0	571.2	762.0	
12	129.0	259.6	519.7	779.2	1039.4	115.2	231.9	464.2	696.2	928.4	104.1	209.5	419.5	629.0	839.0	

Figure 3 graphically outlines efficiency growth as nodes are added to a Lvl.1 protection scheme. This demonstrates how virtualizing the erasure coding across nodes and disks can improve efficiency as more appliances are added.

Figure 3: Efficiency versus scale

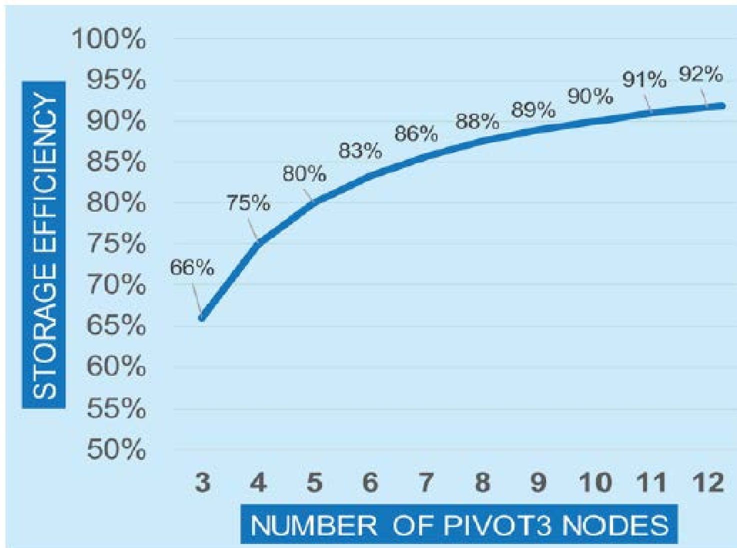

THE COST BENEFIT OF A MORE EFFICIENT HCI ARCHITECTURE

A more efficient HCI architecture can result in significant cost savings, especially as the solution grows. Pivot3 architecture efficiency saves money in two critical areas. If a solution is growing rapidly in the storage dimension then Pivot3 erasure coding reduces costs through reducing the amount of additional storage required. Alternatively, if a solution requires high VM density, as in VDI deployments, then Pivot3's lower HCI overhead tax significantly reduces the expense of adding additional CPUs. Another significant cost savings that HCI vendors tend to ignore is the cost of virtualization infrastructure software which is typically licensed on a per CPU socket level. Whenever a non-Pivot3 HCI node is added it typically requires two additional CPU sockets which require additional licenses. When a business adds HCI nodes due to an architectural limitation in storage or compute efficiency, they end up increasing expenses, not only for extra nodes but also for licensing the virtualization software on those nodes. If additional features are enabled the base-level list price per socket for vSphere enterprise rises. Another consideration is yearly software maintenance which is typically 25 percent of list price per socket.

Pivot3 Storage versus other HCI Implementations – an efficiency example

Many of the HCI architecture implementations today are using the concept of replication to ensure reliability during either a component failure or node failure. In many instances, the solution treats a component failure like a node failure in that any additional failure will bring down the solution and may result in data loss. As the entire cluster grows, component failures inherently increase in probability. Replication schemes mitigate this by increasing protection through increasing the replication factor. For example, two extra copies will cover a two node failure or two disk failures on independent nodes. However, storage efficiency will go from a maximum of 50 percent to a maximum of 33 percent to obtain this extra measure of reliability. This is

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

counterproductive when scaling as the replication protection factor is likely to be greater with larger clusters. This in turn greatly diminishes the effective capacity. The following table details the relative efficiency ratings for various protection levels for Pivot3 erasure coding versus a replication approach.

Pivot3 Erasure Coding Storage Efficiency Versus Replication Scheme					
Pivot3 Erasure Coding Storage Efficiency			Replication Scheme Storage Efficiency (best match)		
Protection Level	**Event Recovery Description**	**Efficiency**	**Protection Level**	**Event Recovery Description**	**Max Efficiency**
Scalar Erasure Coding Lvl.5	- 5 simultaneous disk events or - 2 drives and an entire node	73%*	***Replicated Copies = 3***	- 3 disk or node failures	25%
Scalar Erasure Coding Lvl.3	- 3 simultaneous disk events or - 1 disk and an entire node	81%*	***Replicated Copies = 2***	- 2 disk or node failures	33%
Scalar Erasure Coding Lvl.1	- 1 disk or node failure	92%*	***Replicated Copies = 1***	- 1 disk or node failure	50%

**Based on 12 appliances with 12 drives per appliance*

To illustrate further, assume this scenario: A business needs an HCI solution with 100TB usable storage. They believe that a 1 drive or node failure is sufficient for their reliability needs (Lvl.1). The business compares two vendors who propose a 24TB raw storage per node HCI solution. The following table outlines the number of nodes required to meet the minimum requirement of 100TB usable storage. Depending on price per node of the appliances, the resultant replication-based solution could cost an additional **50 percent** more due to the additional 3 nodes required versus the comparative 6 node Pivot3 solution.

Pivot3 Erasure Coding Storage Efficiency Versus Replication Scheme Example #1		
Item	**Pivot3**	**HCI (replication based)**
Usable TB Target	100 TB	100 TB
Raw Capacity Per Node	24 TB	24 TB
Nodes Needed to Hit Target	**6**	**9**
Incremental Nodes %		**50%**
Raw Capacity Purchased	144 TB	216 TB
Storage Efficiency	81%*	50%
Storage Protection	Lvl.1	Replicated Copies = 1
Usable Storage	116.3 TB	108 TB

**Based on 6 nodes with 12 drives per node*

What if that same business needs to double the capacity from 100TB to 200TB usable capacity and wants to increase the reliability by a factor of two to cover the increased quantity of nodes and drives? The following table outlines this scenario and demonstrates a potential cost increase of up to an astonishing **136** percent due to the purchase of 15 more nodes. This scenario renders a replication based scheme almost unfeasible.

Pivot3 Erasure Coding Storage Efficiency Versus Replication Scheme Example #2		
Item	**Pivot3**	**HCI (replication based)**
Usable TB Target	200 TB	200 TB
Raw Capacity Per Node	24 TB	24 TB
Nodes Needed to Hit Target	**11**	**26**
Incremental Nodes %		**136%**
Raw Capacity Purchased	240 TB	624 TB
Storage Efficiency	80%*	33%
Storage Protection	Lvl.3	Replicated Copies = 2
Usable Storage	211.8 TB	207.8 TB

**Based on 10 nodes with 12 drives per node*

Pivot3 versus other HCI Implementations – VDI Density

Comparing the actual HCI overhead tax across HCI vendors is challenging due to the influence of many factors. Instead of that perspective, the approach taken in this analysis was to use publically available reference architectures with actual test results. These follow the Login VSI standard which most vendors use to validate their VDI implementations. A search of publically available Login VSI-based reference architectures was used to compare Pivot3 overhead with other leading HCI vendors. While data collected using this approach remains somewhat subjective due to the various test setups, VDI density is a good approximation of relative performance for each HCI architecture and the HCI overhead tax can be derived from that density. Since density is publically promoted by each vendor, this makes the comparison as neutral as possible. The following table summarizes what we found among HCI vendors.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

High Level VDI Comparisons Using Vendor's Own Reference Architecture				
VDI Optimized Solution				
Comparative Component	**Pivot3 (Med)**	**Vendor A (Med)**	**Pivot3 (Pro)**	**Vendor B (Pro)**
VDI Platform	Horizon View	Citrix	Horizon View	Horizon View
Virtualization Platform	vSphere	vSphere	vSphere	vSphere
Login VSI Load	Login VSI 4 Medium Worker (2 vCPU & 2GB)	Login VSI 3.7 Medium Worker (2 vCPU & 1GB)	Login VSI 4 Professional Worker (2 vCPU & 4GB)	Login VSI 4 Professional Worker (2 vCPU & 4GB)
Reference VDI Size	357	300	255	160
Compute Platform/Appliance	Dell 730xd (1 node per)	Proprietary (4 node per 2U)	Dell 730xd (1 node per)	Dell 630xd (1 node per)
Cores per CPU	8	8	8	12
Memory per Node	256 GB	256 GB	256 GB	512GB
CPU Speed	2.6 GHZ	2.2 GHz	2.6 GHZ	2.6 GHZ
Nodes in Solution	4*	4	4*	3*
VDI Loads per Node**	119	75	85	80
Storage Software Approach	vSTAC	Distributed scale-out	vSTAC	Distributed scale-out
Storage Protection Scheme	Erasure Coding	Replication	Erasure Coding	Replication
Storage Drive Config	SSD/HDD Hybrid	SSD/HDD Hybrid	SSD/HDD Hybrid	SSD/HDD Hybrid

** One node dedicated to VDI management overhead and not included in per node density calculation*

***Numbers based on vendor actual reference architecture testing and claims*

To add perspective, assume a scenario where a business would like to build a 2000 seat VDI implementation. The assumption is that each HCI solution will scale linearly as you add more VDI seats. The following table outlines a comparison of each of these reference architectures. For a 2000 seat VDI implementation, other replication-based HCI solutions require up to **59 percent** more nodes and up to **59 percent** more CPU cores than Pivot3.

Pivot3 Lower CPU Overhead VDI Comparison Versus Other HCI Reference Architectures				
Item	**Pivot3 (Med)**	**Vendor A (Med)**	**Pivot3 (Pro)**	**Vendor B (Pro)**
VDI Seat Target Capacity	2000	2000	2000	2000

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

VDI Density per node	119	75	85	80
Nodes Needed to hit target	**17**	**27**	**24**	**25**
Incremental Nodes %		**59%**		**4%**
Number of Cores Per Node	16	16	16	24
Number of Cores to hit target	***384***	***432***	***384***	***600***
Incremental Cores %		***59%***		***56%***

Pivot3 versus other all Flash HCI Implementations – VDI Density

Many vendors including Pivot3 have turned to all flash implementations to improve VDI performance. When evaluating all-flash implementations, we evaluated Pivot3 against a leading vendor that performs inline deduplication and compression in software to determine the impact of HCI overhead tax on those data services. In this case we compared the competitive vendor's datasheet information only as we did not have access to reference architecture or detailed test information from that vendor. The following table outlines the comparison.

High Level VDI Comparisons Using Data Sheet Information		
VDI Optimized Solution		
Comparative Component	**Pivot3**	**Vendor C**
VDI Platform	Horizon View	Citrix
Virtualization Platform	vSphere	vSphere
Login VSI Load	Login VSI Knowledge Worker (2 vCPU & 2GB)	Login VSI Knowledge Worker (2 vCPU & 2GB)
Reference VDI Size	768	500
Compute Platform/Appliance	Dell 730xd (1 node per)	Lenovo x3550 (1 node per)
Cores per CPU	14	12
CPU Speed	2.3 GHZ	2.5 GHZ
Nodes in Solution	4	4
VDI Loads per Node	192	125
Storage Software Approach	vSTAC	Distributed scale-out
Storage Protection Scheme	Erasure Coding	Replication, Deduplication, Compression
Storage Drive Config.	All SSD	All SSD

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

The following table outlines a comparison of each all-flash configuration. For a 2000 seat VDI implementation, competing/or/replication-based HCI solutions require up to **45 percent** more nodes and up to **25 percent** more CPU cores than Pivot3.

Pivot3 Lower CPU Overhead VDI Comparison for All-Flash		
Item	**Pivot3**	**Vendor A**
VDI Seat Target Capacity	2000	2000
VDI Density per Node	192	125
Nodes Needed to Hit Target	**11**	**16**
Incremental Nodes %		**45%**
Number of Cores Per Node	28	24
Number of Cores to Hit Target	**308**	**384**
Incremental Cores %		**25%**

TANEJA GROUP OPINION

Hyperconverged infrastructure has become a mainstream approach to simplifying data center deployment of virtualization. These products use modular scale-out architectures, which makes them an ideal solution for many business workloads. As the hyperconverged products mature, businesses should look deeper into the attributes and value that each vendor brings. Pivot3 solutions have been keenly focused on minimizing CPU overhead in addition to making the storage tier extremely efficient, reliable, and high-performing through many years of development invested in Scalar Erasure Coding. Other hyperconverged vendors have focused on adding numerous features which erodes efficiency and depletes resources required for business applications. As demonstrated using real world examples, other HCI solutions may require up to **136 percent more raw storage capacity** and up to **59 percent more total CPU cores** than an equivalent solution using Pivot3 products.

In this maturing market, businesses need to thoroughly research and evaluate hyperconverged solutions by analyzing the true costs of each alternative. For businesses sensitive to the cost and capacity of storage or if the cost of CPU-based licensing is significant, we would highly recommend that you consider Pivot3 as one of the most efficient implementations of hyperconvergence available.

For more information please visit www.tanejagroup.com.