



Considerations for the Next Phase of Hyperconverged Infrastructure

MARCH 2016

COMMISSIONED BY





ABOUT 451 RESEARCH

451 Research is a preeminent information technology research and advisory company. With a core focus on technology innovation and market disruption, we provide essential insight for leaders of the digital economy. More than 100 analysts and consultants deliver that insight via syndicated research, advisory services and live events to over 1,000 client organizations in North America, Europe and around the world. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

© 2016 451 Research, LLC and/or its Affiliates. All Rights Reserved. Reproduction and distribution of this publication, in whole or in part, in any form without prior written permission is forbidden. The terms of use regarding distribution, both internally and externally, shall be governed by the terms laid out in your Service Agreement with 451 Research and/or its Affiliates. The information contained herein has been obtained from sources believed to be reliable. 451 Research disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although 451 Research may discuss legal issues related to the information technology business, 451 Research does not provide legal advice or services and their research should not be construed or used as such. 451 Research shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.

NEW YORK

20 West 37th Street
New York, NY 10018
+1 212 505 3030

SAN FRANCISCO

140 Geary Street
San Francisco, CA 94108
+1 415 989 1555

LONDON

Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 207 426 1050

BOSTON

One Liberty Square
Boston, MA 02109
+1 617 598 7200

Executive Summary

Hyperconvergence has been receiving a tremendous amount of attention because it represents the next step in the evolution of IT resource delivery. This technology takes the idea of integrating compute, storage and networking that started with converged systems design and has improved on those architectures by adding deeper levels of abstraction and automation. Hyperconverged infrastructure (HCI) vendors promise simplified operation and the ability to quickly and easily expand capacity by deploying and launching additional modules; simplicity has been the key selling point for the HCI pioneers.

As HCI ventures even deeper into the enterprise and cloud environments, the architectures will need to become more efficient, agile and adaptable to help IT professionals shoulder the burden of rapidly growing data sets and workloads. This report discusses the benefits of HCI and the enhancements that must be made to expand HCI deeper into the mainstream enterprise datacenter.

Introduction: Market forces are driving hyperconverged infrastructure forward

Over the last few years, storage and other infrastructure teams have been asked to do more with fewer resources at their disposal. While data at many organizations continues to grow at an alarming rate, budgets are not growing at the same pace, which indicates that storage professionals can no longer rely on bulk purchases of traditional disk systems to get by.


In the past three to five years, hyperconverged platforms have become legitimate alternatives to traditional storage systems. They have shown enterprises, midsized companies and service providers that storage does not have to be confined to proprietary external arrays. The following factors are driving organizations toward the next generation – HCI:

DATA GROWTH IS OUTPACING STORAGE BUDGET INCREASES.

In our Voice of the Enterprise Storage (VotE), Q4 2015 report, 56.1% of respondents expected to see their total deployed storage capacity increase by 25% or more, with an alarming 7% of respondents expecting a 100-200% increase in 12 months (See Figure 1). Although respondents are expected to see storage budget growth in 2016, only 17.8% of respondents will be getting a significant increase beyond 25% (see Figure 2). To get to a sustainable storage environment, organizations must look beyond their existing storage and infrastructure strategies to find more efficient ways to contain data growth while accelerating the delivery of storage services to stakeholders.

Figure 1: Expected Storage Capacity Growth for 2016


Q. By what percentage do you expect your organization's total deployed raw storage capacity (as measured in terabytes) to change in the next 12 months? Please include capacity deployed both on-premises and in third party clouds. n = 574



PATHFINDER REPORT: CONSIDERATIONS FOR THE NEXT PHASE OF HYPERCONVERGED INFRASTRUCTURE

Figure 2: Storage Budgets changes 2016 vs 2015

Q. By what percentage do you expect your organization's storage budget to change in 2016 compared to 2015? n = 366



STORAGE PAIN POINTS


Organic business growth continues to be the largest driver for storage growth at organizations, with 55.4% of respondents listing it as the top storage pain point (see Figure 3). Storage infrastructure complexity is also a growing concern considering that 35.7% of VotE storage respondents are managing three or more primary storage tiers in their environment, in addition to their backup and disaster recovery infrastructures. Because of data growth and increasing complexity, the management of storage has become an acute pain point for organizations, with 34% struggling with capacity planning/forecasting, 16.6% having difficulty juggling storage silos and 14.2% complaining about a lack of skilled staff.

Unfortunately, the storage-capacity challenge will only get worse as the years go by. Meeting disaster-recovery requirements (29% of respondents) and meeting backup windows (17.1% of respondents) were both listed as storage pain points, which shows that data protection continues to be a growing concern. Compliance regulations are forcing organizations to retain data for longer periods and preventing storage professionals from simply deleting old data to make way for new content. This suggests that organizations will need to leverage storage-reduction technologies such as deduplication and compression to boost the efficiency of their primary and secondary storage systems (such as backup and archives) to keep data growth manageable.

PATHFINDER REPORT: CONSIDERATIONS FOR THE NEXT PHASE OF HYPERCONVERGED INFRASTRUCTURE

Figure 3: Top Storage Pain Points

Q. What are your organization's top three pain points from a storage perspective? n = 639



HCI PIONEERS FOCUSED ON SIMPLICITY AND IMPLEMENTATION SPEED

Although few would deny the impact HCI pioneers have had in terms of thought leadership for infrastructure modernization, at this point, only a relatively modest number of organizations have deployed these technologies, with only 23% of respondents currently having HCI in use in their environments (see Figure 4).


The initial deployments of HCI have been largely focused on the midrange segment of the market, where IT organizations often lack storage expertise and are not bound to a specific storage supplier. For these early customers, HCI's ability to deliver key storage functions such as snapshots/cloning, replication and flash acceleration without the need for storage SAN expertise has been a game changer, especially with IT professionals being pushed to broaden their skill sets to handle increasing workloads in datacenters.

Virtualization administrators have a growing influence on storage purchasing decisions; our Wave 19 Storage survey found that 61% of respondents reported this. This is a key point, since virtualization and cloud administrators will likely be the key IT stakeholders that push HCI into the core of the enterprise datacenter, and beyond midrange space where most of the early HCI offerings were deployed.

PATHFINDER REPORT: CONSIDERATIONS FOR THE NEXT PHASE OF HYPERCONVERGED INFRASTRUCTURE

Figure 4: HCI is Still an Emerging Space

Q. Which of the following types of storage systems and related products does your organization currently use? n = 535



Key requirements to make HCI more agile, efficient and adaptable

While there has been a large number of early success stories with small and mid-sized deployments, we are still in the nascent stages of HCI adoption, and there is more work to be done before HCI can become the dominant infrastructure standard for modern enterprises. Specifically, HCI must become more:

| | |
|------------------|---|
| EFFICIENT | To efficiently deliver resources. |
| AGILE | To keep the infrastructure manageable at scale. |
| ADAPTABLE | To match the changing needs of customers. |

EFFICIENT

Efficiency must be continually enhanced to make HCI a worthwhile investment relative to traditional infrastructure. With data growth rampant and likely to get worse, the storage efficiency of HCI must evolve to meet the challenge. Beyond scaling capacity, the performance of HCI must also improve in terms of scalability and granularity to make sure key workloads are not starved of resources. Last but not least, HCI management efficiency must take networking into account not only to ensure that HCI is simple to deploy, but also to ensure that this infrastructure will be able to optimize and enhance networks when there is a surge of consumption or if there is an unplanned outage.

STORAGE EFFICIENCY – usually in the form of data reduction – is a key requirement for organizations given that most budgets are not increasing fast enough to match data growth. Deduplication and in-line compression capabilities allow companies to store more data in the storage footprint by eliminating redundancies as data is being written to disk or flash. Deduplication works well for reducing VM images and files and has moved into the primary storage space after first becoming popular in the backup and secondary storage markets. Because deduplication does not work well on database workloads, compression has become necessary for reducing application workloads, and organizations should look for HCI offerings that not only have both of those capabilities, but also the intelligence to automatically apply the right data-reduction technology based on the workload being stored.

PERFORMANCE EFFICIENCY ensures that the expensive flash and CPU resources organizations have aggregated are not going to waste in their HCI deployments. Given the relatively high cost of flash on a dollar-per-GB basis, deduplication is an important capability for reducing costs and maximizing the utilization of flash invest-

PATHFINDER REPORT: CONSIDERATIONS FOR THE NEXT PHASE OF HYPERCONVERGED INFRASTRUCTURE

ments. The deduplication and compression capacities that we discussed in the previous section also have a major role to play in performance efficiency because they allow HCI nodes to cache more data within their expensive flash SSDs and PCIe cards.

Beyond data reduction, HCI platforms should have automated data-placement intelligence to stripe data across multiple nodes to avoid IO bottlenecks and to ensure that performance and storage capacity can scale out in a linear fashion. Data caching should also be integrated into the platform to ensure write and read operations are concentrated on high-performance flash media, while idle data is seamlessly destaged using sequential writes to low-cost hard drives. Although hard drives are much slower than flash for transactional workloads, their performance is adequate for capturing sequential write streams and for quickly recovering data in the event of a node failure.

Storage QoS is another key attribute that is emerging in HCI systems. This feature allows administrators to prioritize sensitive applications and guarantee performance for these workloads. In the cloud-provider world, storage QoS has also been used to ensure that low-priority workloads do not consume excessive resources, and do not become 'noisy neighbors' that disrupt workloads that are residing on the same hardware.

NETWORK EFFICIENCY is often overlooked as a key component for HCI, but to push these infrastructures into enterprise and cloud scale, networks must become more efficient both in terms of performance and manageability. Just as we saw with the earliest forms of HCI, ease of deployment will continue to be a key attribute for platforms. Today, HCI does a good job of seamlessly blending storage and server hardware into a shared pool of resources. The next step is that integration of networking components – both internal to the cluster and outbound – must become easier to configure and automated. Software-defined networking could be leveraged to simplify and automate processes to accelerate and standardize deployments.

In normal operation, the network of a HCI deployment has to run in the same low-maintenance manner as the nodes. This means that the network management software has to provide a unified view of the network and minimize any requirement to individually manage the devices that form the network. The ability to identify and track cluster traffic can both reduce the workload of administrators and enhance the understanding of overall cluster health. Greater network visibility adds much greater depth to cluster management.

Given the highly distributed nature of enterprises today, tighter integration and optimization of WAN network links will also be key elements for HCI. Nodes sent off to a remote office should be able to launch and configure themselves, without requiring a networking expert to be physically present. From an efficiency standpoint, WAN-optimization technologies should also be implemented to reduce the amount of data that has to be sent over the WAN and ensure that HCI clusters are synced with central offices to guard against data loss in the event of a disaster.

AGILE

Efficient scale-out architectures not only provide an easy means for adding resource capacity, but they must also help organizations easily manage additional hardware without requiring an increase in staff. To facilitate this, HCI platforms must become more intelligent and need to integrate more tightly with existing resources in the datacenter. To become more agile, next-generation HCI platforms must have:

COMMON MANAGEMENT AND ORCHESTRATION. HCI management tools today do a relatively good job of managing the resources within a cluster, but for HCI platforms to take a leap forward, integration with other existing datacenter platforms, such as traditional and converged infrastructure, is necessary.

From the context of datacenter networking, one of the challenges with hyperconverged systems is that the management software, as it simplifies the operation of the cluster, can hide the state of the network that is the core of the cluster. Having sophisticated network management software as part of the networking component of a hyperconverged cluster can greatly improve its long-term reliability and simplify operations. It can provide visibility into the performance of the network and activity levels of the system components.

Integration with well-known management tools, such as VMware's vCenter, will help simplify management by eliminating the need to add new management consoles. To simplify troubleshooting, HCI platforms should also help customers create a common control plane for servers, networking and storage to centralize logs and error reporting within existing tools.

PATHFINDER REPORT: CONSIDERATIONS FOR THE NEXT PHASE OF HYPERCONVERGED INFRASTRUCTURE

SCALE-OUT ARCHITECTURES. With the constant growth of data and applications, scale-out architectures are ideal because they allow infrastructure professionals to add capacity non-disruptively while minimizing the management impact of adding nodes. Another key benefit of scale-out architectures is that they allow organizations to start out with a small configuration and gradually grow the storage infrastructure to match the needs of their workloads.

Although scale-out is a common capability in existing HCI, most implementations on the market have a rigid architecture that forces customers to add compute, memory and storage as blocks of resources. Future generations of HCI should allow customers to add these resources independently and in a granular fashion, since the current deployments ultimately wind up creating inefficient silos with unused processing and storage resources.

ADAPTABLE

IT infrastructure technologies and services are in a constant state of evolution, and the HCI deployment must be adaptable to ensure that the organization is taking advantage of the latest technology enhancements and to meet the rapidly changing requirements for stakeholders' applications.

To become more adaptable, next-generation HCI must be able to:

INTEGRATE WITH APIs. Although the concept of using APIs for management and provisioning is new for enterprise storage environments, public cloud environments have been using APIs for these purposes for many years now, which makes this a key requirement for organizations building a private cloud environment. Through the use of API-based management, infrastructure professionals and their colleagues can create service catalogs to allow clients to request resources and have them automatically provisioned through the APIs. APIs can also be used to facilitate rapid cloning of VMs for test/dev and other purposes.

SECURE SENSITIVE DATA. The data in companies' environments is rarely homogeneous, and the HCI platform must be able to deliver the appropriate level of protection to match the business and compliance requirements for each workload. While disk encryption has become more common in enterprise environments, this level of security only provides protection against hardware theft. By providing a deeper level of security, next-generation HCI platforms should offer file-level encryption as an option for securing compliance-sensitive data residing in a system. Given that data is also at risk in flight, VPN integration is also necessary for securing replication streams whether they are pushing data to remote sites or to public clouds. Beyond encryption, HCI platforms should have comprehensive auditing capabilities to track the source of breaches and data-corruption events within the four walls of the organization itself.

ASSIMILATE NEW TECHNOLOGIES RAPIDLY. As storage and HCI gradually move beyond the proprietary appliance models that have dominated for years, more infrastructures will become software-defined and leverage commodity hardware to lower costs. Hardware innovation in the commodity market for processors, memory and solid-state storage such as flash is happening at an extremely rapid pace. The transition to software-defined datacenter infrastructure will allow organizations to innovate at their pace, in contrast to proprietary array and HCI architectures where companies have to wait for vendors to approve hardware such as new flash drives.

Another emerging technology innovation that requires improved and adaptable HCI is composable hardware. This new architecture allows organizations to treat infrastructure as code and disaggregates compute, storage and networking resources to optimize them to meet application and workload demands. For example, for performance-sensitive workloads, specific processors and memory resources can be dedicated to a workload to ensure that transactional performance and latency is able to match stakeholder needs. Likewise, for data-heavy workloads that are not compute-intensive, such as a media or backup archive, hard drives and flash can be dedicated to the workload, while unused compute and memory resources can be redirected to performance-sensitive workloads.

PATHFINDER REPORT: CONSIDERATIONS FOR THE NEXT PHASE OF HYPERCONVERGED INFRASTRUCTURE

Use cases

EASY TO SET UP IN DATACENTER OR IN REMOTE LOCATIONS (DISTRIBUTED ENTERPRISE USE CASE)

Distributed organizations are commonplace (see Figure 5) today, and there is a growing trend toward edge computing. Edge computing pushes applications, data and infrastructure (services) away from a centralized, more-controlled datacenter environment. When you take remote locations into account, all of the major storage pain points we discussed earlier – capacity management, the need for disaster recovery, security and risk management and the lack of skilled staff (Figure 3) – become more difficult the further away a site is from the control and IT expertise.


For remote offices and edge computing to flourish, the infrastructure at remote sites must also be able to start out with modest configurations to keep costs down and the scalability to grow in pace with the organization's remote operations. The ability to do remote installations is key not only because it eliminates the need to send an IT staffer to a remote site, but it also ensures that resource delivery is consistent regardless of whether a workload is run at the central office or the most remote office in the organization. Specifically, remote sites will likely be staffed with an IT generalist who doesn't have the SAN networking and storage array expertise required to maintain and troubleshoot a traditional infrastructure consisting of distinct compute, storage and networking hardware.

As organizations continue to become more distributed and new applications are implemented as part of the Internet of Things, a greater amount of data will be created away from central offices. HCI is likely to be deployed in these locations, and this data must have the same level of protection and accessibility for enterprises to be successful. Key enterprise-class data-protection features such as snapshots and deduplication/compression efficiently manage data and keep applications online. Given that organizations cannot afford to lose data created at remote sites, WAN-efficient replication technologies must be deployed to keep data protected and facilitate application failover to a secondary site in the event of a site-level disaster. Remote management and monitoring will be required to provide rapid problem resolution.

HCI's ability to standardize IT resource delivery and make it service-like, combined with its integrated data-protection and data and workload-mobility enhancements, make it a powerful infrastructure option for remote locations.

Figure 5: Enterprise Organizations are Highly Distributed

Q. How many offices do you have nationally (within your headquartered country?) (Please select one) n=2027



HYBRID CLOUD - MOVE WORKLOADS BETWEEN DATACENTER AND PUBLIC CLOUD

The future of IT infrastructure is hybrid clouds blending on-premises infrastructure seamlessly with public cloud compute and storage services. In the VotE storage survey from Q4 2015, we found that 79.3% of respondents were planning to increase their spending on third-party cloud storage services in 2016. In this same survey, we found that respondents were spending 57.5% of their budgets on capital expenditures for items such as array and software purchases, while only spending 42.5% on operational expenditures such as maintenance and IT staffing.

Organizations are looking to move beyond the traditional capex-heavy model of infrastructure hardware acquisition that, in the case of storage, forces customers to inefficiently purchase large sums of capacity up front. Subscription-based pricing

PATHFINDER REPORT: CONSIDERATIONS FOR THE NEXT PHASE OF HYPERCONVERGED INFRASTRUCTURE

for HCI hardware is still an evolving business, but the advent of software-defined storage together with the increasing use of commodity hardware will provide organizations with pricing flexibility beyond what is available with proprietary storage silos today.

More organizations are also transitioning to a DevOps methodology of continuous delivery. DevOps teams reduce capex and opex by consolidating technology silos – collapsing development, test, quality assurance and production systems onto common infrastructure with unified management and automation. This requires support for virtualization, containers and bare-metal environments. To benefit from this disruption, on-premises infrastructure must become more cloud-like, and it must do a better job of facilitating the movement of data and workloads both to and from public cloud environments.

As we discussed previously, HCI's ability to standardize and automate the delivery of IT services are key capabilities that can make infrastructure more cloud-like. With this automation, clients will be able to request and utilize storage and compute resources with the appropriate levels of capacity, performance and resiliency needed for specific workloads. This requires the ability to specify infrastructure operational policies for network, storage and compute elements of the physical and virtual infrastructure by directly mapping the application intent with the infrastructure policy required – for example, policies around persistent storage, volume allocation and snapshotting. This allows companies to achieve more efficient shared infrastructure for various containerized applications.

Workload migration is the next major frontier for hybrid cloud adoption, and this is an area where disaster recovery as a service (DRaaS)-enabling technologies are helping organizations convert and move workloads seamlessly between public and private cloud environments. Beyond the core data-movement technologies we have discussed – such as replication, deduplication, and WAN optimization – organizations will also need to ensure that their HCI platforms have cloud orchestration capabilities to allow them to manage and launch cloud services when needed to handle a workload, while also ensuring end-to-end security for the workload. While DRaaS technologies today are more focused on moving mission-critical workloads to ensure availability, future workload migration will become a common capability that will allow companies to tap into a public cloud when on-premises resources are not available.

The Next Generation

CUSTOMER EXPECTATIONS ARE CHANGING; YOUR INFRASTRUCTURE MUST EVOLVE

With the digital transformation that is taking place at businesses across the globe, customer expectations for information access, service availability and rapid delivery are driving organizations to revamp their infrastructures. Now more than ever, when it comes to provisioning and delivering services, time is money, and any inefficiency that emerges in business processes will become more scrutinized because of the impact to the bottom line. This market dynamic should be a driving factor in pushing infrastructure changes forward, and it will push traditional IT environments to deliver cloud-like resource-delivery models for their clients.

The increased efficiency, agility and adaptability of next-generation HCI architectures can help organizations deliver a number of key business benefits such as:

- **Meeting performance and uptime SLAs.** In the all-flash array (AFA) market, we have seen multiple customer deployments where the costly upgrade to AFA was justified by painful SLA penalties that would have hit the organizations if they did not upgrade their performance capabilities. As HCI gets deployed in more business-critical and performance-sensitive use cases, meeting SLAs will become an important part of the justification for this next-generation infrastructure. Failing to meet SLAs can also have an adverse effect on the perception of the organization and can quickly lead to lost customers. Performance is also a major factor in key areas such as VDI, where a nonresponsive virtual desktop can adversely affect worker productivity and lead to poor customer support in business-critical VDI deployments in healthcare, government and financial services organizations. Ensuring uptime is critical in many use cases where HCI is commonly deployed, such as hospitals and clinics, as well as retail and manufacturing locations.
- **Provisioning acceleration.** With the advent of cloud storage and compute services, business stakeholders have expectations for nearly instantaneous access to resources and have no patience for the hours, days and sometimes even weeks it takes for traditional IT to provision resources. The standardization and automation of resource provisioning, which next-generation HCI provides, can help organizations not only create applications faster, but scale them up to match rising demand.

PATHFINDER REPORT: CONSIDERATIONS FOR THE NEXT PHASE OF HYPERCONVERGED INFRASTRUCTURE

- **Faster insights.** Although many organizations are actively hoarding data with the hope that it will create business value someday, an inefficient infrastructure that takes too long to process data and provision resources will marginalize the value of that data. The agility of a next-generation HCI will allow decision-makers to quickly process the data at their disposal to gain insights in a timely fashion to either gain a tactical advantage over competitors or find a new opportunity to drive business growth.
- **Customer experience.** Digital transformation of businesses and cloud services has raised customer expectations. In the e-commerce space and in social networks, high latency and downtime have a direct impact on customer retention and completed transactions. Future infrastructure must be able to not only handle the daily flow of customer requests but also must be able to efficiently scale up when there is a surge in demand.