White Paper

The Evolution of Hyperconvergence and NetApp's Role in this Rapidly Expanding Market

Sponsored by: NetApp

Eric Sheppard
October 2018

IN THIS IDC WHITE PAPER

This IDC white paper reviews important market trends that have driven rapid growth in global hyperconverged infrastructure (HCI) deployments. This paper outlines the ways in which hyperconvergence is evolving to become an enterprise-class datacenter platform capable of supporting the stringent demands of mission-critical and hybrid cloud environments. This paper also provides an overview of NetApp's HCI portfolio of solutions and a summary of an in-depth interview IDC conducted with a NetApp HCI customer.

SITUATION OVERVIEW

Companies that want to compete in today's rapidly changing digital environment must focus on innovative use of technology that drives organizational agility and business decisions that are based on high volumes of near-real time data. Simply stated, organizational agility and intelligent use of big data have become competitive differentiators. As such, companies around the world are undertaking digital transformation projects intended to streamline costs of doing business, strengthen customer relationships, capitalize on new sources of revenue and improve workforce productivity. For companies undertaking such digital transformation, this can mean a shift from a physical product to a subscription service, targeting an entirely new market or reaching new customers that brings a different level of scale. The impact of this transformation is far-reaching and difficult to overstate. All parts of a company (e.g., executives, lines of business, facilities, IT teams, software developers, and more) are expected to be contributors to this once-in-a-generation transformation by fundamentally rethinking organizational structures, processes and tools to better deal with the ever-increasing speed of business change.

The need to digitally transform has created immense pressure for datacenter teams that are increasingly expected to be leaders in this transformation. New leaders are evolving to become cloud architects so that they can meet the considerable pressure to mimic the best of public cloud within their datacenters (i.e., provide resources that are simple to manage, allocate and consume). Indeed, IT consumers have considerable experience with public cloud services and fully expect a comparable swipe-and-go process of dialing up services nearly instantly within their own organization's IT department. Further, developers and business leaders expect their IT teams to automate the provisioning of tuned IT resources for any workload at any location whenever needed. Demand for datacenter skills is shifting from those with expertise in 3-tier SANs (server, network, storage) to skills

commonly found within cloud architect communities. Simply stated, IT teams are forced to either adapt to a changing world where public cloud simplicity is the norm or face becoming obsolete.

The rapid increase of infrastructure convergence is driven by its ability to address many of the datacenter challenges outlined above. Those that have deployed converged infrastructure are benefiting from increased IT staff productivity, improved business agility and the ability to support a modern hybrid cloud solution. Early convergence of datacenter infrastructure focused almost entirely on integration, autonomous compute, storage, and networking systems. This convergence expanded with the introduction of software-defined, scale-out solutions that drove further operational simplicity through automation and reduced complexity. These newer solutions, known as hyperconverged infrastructure (HCI) brought considerable benefits, but often came with fewer capabilities than their legacy counterparts. Only recently have we begun to see a new generation of hyperconverged solutions that combine the benefits of scale-out, software-defined architectures with mature, enterprise-class data services. IDC believes this represents a new phase, or second chapter of the HCI market. An overview of how we have gotten to this newer generation of HCI solutions, what is expected of a modern HCI solution, and the considerable benefits they offer, is provided in the sections that follow.

The Evolution of Datacenter Infrastructure Convergence

Datacenter teams are finding that the levels of scale, simplicity and agility needed to support complex digital transformation initiatives simply cannot be achieved using their long-standing practice of buying individually managed silos of datacenter resources, or just following the “check box” next to things like APIs. As a result, IT organizations around the world are shifting their resources away from standalone servers, networking, and storage systems in favor of converged infrastructure solutions that can be centrally managed with tools that offer new levels of automation. IDC has been covering the converged systems market for nearly a decade. During this time, solutions in this market have proven to be a very effective way to improve datacenter resource utilization, increase IT staff productivity, improve business agility, and reduce time-to-services.

Converged systems represent a consolidation of core datacenter technologies (servers, storage systems, networking and management software) into a single system that can be deployed, managed, and supported more efficiently than buying and building these technologies separately. Converged systems help to remove complexity and risk associated with managing enterprise-grade datacenter infrastructure so that IT teams can confidently focus their time on higher-value projects and tasks.

The adoption of converged systems has grown quickly during the past decade, with annual spending on converged systems now more than $13 billion, with the market expanding faster than the larger datacenter infrastructure market. This type of market expansion is a clear indication of the degree to which converged systems can drive real benefits within the datacenter. That said, most of the traditional converged systems in use today are built with well-tested, but discrete server, storage and storage networking as their core building blocks. Such systems have undeniably helped push datacenter infrastructure convergence into the most demanding, mission-critical applications in use today. The use of discrete systems as core building blocks has also kept average selling prices (ASPs) of traditional converged systems relatively high. IDC data shows ASPs of converged systems built with discrete server, storage and storage networking components above $350,000 for the complete solution. While such ASPs have not impeded the continued expansion of converged systems, they have held back adoption within many organizations looking for smaller starting points and more granular building blocks.

Datacenter Convergence Evolves and Expands to Include Hyperconvergence

Like today's IT departments, the converged systems market is rapidly evolving. An important element of this evolution is the relatively recent emergence of hyperconverged infrastructure (HCI), which IDC considers a subset of the $13 billion converged systems market and the next phase of the market's lifecycle. HCI solutions deliver the proven benefits of traditional converged systems, but do so through a software-defined, scale-out architecture. HCI solutions are built as clusters of commodity servers (x86) that provide an abstracted pool of capacity, memory, and CPU cores that are used as the foundation for server-centric workloads (the hypervisor, VMs and applications) as well as storage-centric workloads (e.g., data persistence, data access and data management).

HCI deployments are driving benefits in the following key areas:

- **Lower capex.** Lower capex can be achieved through the elimination of SAN-based storage solutions in favor of industry standard servers that offer fully virtualized compute and data services. The scale-out architecture of hyperconverged solutions further lowers capital costs by helping to reduce the need to overprovision resources. Instead, customers can buy only the nodes required at the time of initial deployment and scale later as needed.
- **Reduced opex.** Reduced overprovisioning and elimination of storage silos have positive impacts beyond capex. In fact, these benefits can directly lead to lower costs of power, cooling and floor space within the datacenter. HCI solutions often integrate management software that automates many of the complex tasks needed during initial deployment while also reducing the number of steps required to provision new workloads. The result is improved IT staff productivity and increased agility within the datacenter. These same solutions also help IT departments to leverage IT generalists for low-value tasks, thus freeing up time for infrastructure specialists to work on more innovative projects.
- **Reduced risk.** The highly automated nature of HCI solutions also helps to reduce risk of downtime associated with common lifecycle management tasks (e.g., firmware upgrades, system refresh). The scale-out, software-defined nature of HCI solutions help to eliminate the need for complex and risky forklift upgrades, which have become all too common within the datacenter. Many companies leverage hyperconverged solutions as a way to improve their DR/HA processes and costs in ways not possible just a few short years ago. Lastly, HCI solutions allow users to reduce the number of technology suppliers involved within a full solution, which helps to better coordinate patches and upgrades while also reducing the number of support calls needed for the solution.

Evolution of Hyperconvergence

First made popular by public cloud operators like Google, Microsoft and Amazon, hyperconvergence emerged as a modern converged infrastructure architecture for next-generation applications. Smaller companies that were "born in the cloud" ignited further demand for scale-out, software-defined HCI systems. Meanwhile, established enterprises around the world began demanding the benefits that an HCI approach can provide – that of scale, automation and agility. On prem, HCI solutions found early success within midsize companies, remote office/branch office (ROBO) environments and virtual desktop infrastructure (VDI) projects due largely to their ability to eliminate the complexity, expense and latency often associated with SAN-based storage solutions. Awareness of the benefits these solutions bring to the table has risen steadily among enterprise IT teams, resulting in increased market adoption. Once deployed, IT teams frequently expand the set of workloads running on HCI. This considerable growth of new hyperconverged deployments and expansion of workloads running on these systems has helped to drive global sales of hyperconverged solutions (including hardware and

software) beyond $4.9 billion during the twelve months ending June 30, 2018, up 72.3% year-over-year. Importantly, the expansion of workloads running on hyperconverged solutions has brought with it new requirements. Today's HCI solutions are expected to perform predictably while supporting a larger number of workloads with differing I/O profiles. Today's HCI solutions are now expected to support mission-critical applications (e.g., SAP or Oracle), which translates into application resiliency at the system and multi-site level. A more complete overview of the requirements expected of a modern HCI solution is provided in the next section.

Although not widely realized, the hyperconverged market is entering a new phase of maturity. In contrast to the early years of deployments, investment in HCI solutions are increasingly targeting mission-critical workloads, are larger in scale, and are in use by a wide range of companies. Based on past IDC surveys, business applications (e.g., ERM, CRM, supply chain management, financial management, payroll/accounting) were the most common workloads running on HCI solutions. Other common applications running on HCI include collaborative and content applications, structured data analytics and structured data management applications. Moving forward, IDC expects the use of HCI solutions for mission-critical business applications to expand considerably.

Requirements of a Modern Hyperconverged Solution

While the types of workloads running on hyperconverged solutions are a good indication of how far the HCI market has come since its early days, there is more that must be offered to drive further expansion of this market. Modern HCI solutions must close important feature/capability gaps that exist with traditional datacenter infrastructure, enable an IT departments shift towards private cloud infrastructure and support an increasing need for organizational transformation within the datacenter. The following bullets offer an overview of the most important attributes and capabilities a modern HCI solution must provide to ensure relevance as users push for expanded HCI use cases within their datacenter:

- **Eliminate "noisy neighbors."** As outlined earlier, organizations are expanding the use of HCI solutions. With the expanded use of hyperconverged infrastructure comes the need for HCI solutions that support an increased density of primary workloads. Thus, one of the most critical feature gaps that HCI solutions must address is support for predictable and guaranteed performance levels within such environments. With increased workload density comes the increased risk of resource contention. For hyperconverged solutions to thrive during this next phase of hyperconverged adoption, they must consistently deliver sub-millisecond response times and support hundreds of thousands to millions of IOPS. This is creating the need for a workload protection engine to address "noisy neighbor" issues and ensure that performance can be consistently delivered as required, no matter what else is going on in the system.
- **Integration with multiple public clouds.** Modern HCI solutions must provide a private cloud platform that can be easily integrated with hyper-scale public cloud providers to create a seamless hybrid, multi-cloud experience. Features required to achieve this include:
    - Use of service catalogs that support on-prem resource and workload deployment
    - Support for highly portable workloads and automated bursting
    - Ability to seamlessly move data between on-prem, private clouds and public clouds to support the ever-changing needs of a diverse workload portfolio
    - Seamlessly collapse management complexity to better support the lifting and shifting of applications from private clouds to trusted hyper-scalers and back again as business requires
    - Support for containers to accelerate development of new services

    - Support for chargeback of resource utilization

- **Common data fabric for private and public clouds.** IDC believes hyperconverged solutions have become an ideal platform for on-prem, private cloud deployments thanks to their scale-out, software-defined, highly automated architecture. But it is important to recognize that the role of customers' cloud-based IT solutions is to ingest, deliver, and exploit data no matter where that data is created or lives. Thus, a truly optimized hyperconverged-based private cloud should be considered a component of a broader, hybrid public/private cloud ecosystem that provides a "lingua franca" (or data fabric) that supports common data services for efficient placement, movement and use of data across a true hybrid multi-cloud environment. Such a data fabric allows IT departments to incorporate newly deployed HCI solutions in their datacenters without creating additional management silos.
- **Enabler of organizational transformation.** IDC notes that modern hyperconverged solutions are also increasingly looked upon to support organizational transformation within the datacenter. Therefore, modern HCI solutions must support consolidation of datacenter jobs/roles so that one administrator can take on the responsibility for virtualization, compute and storage. This consolidation of roles will help to free up time within the IT team for more innovative projects and ultimately help shift human capital to other critical parts of the datacenter that drive innovative, new revenue streams and customer touchpoints rather than just maintaining status quo.
- **Independent resource scaling.** Many HCI solutions scale by requiring an additional node that would increase compute and storage resources together. This is commonly referred to as linear scaling. The problem with this is no datacenter grows in a linear way across all the pools. There are times when there is need for additional compute or additional capacity, but not both. To address these increasingly diverse needs, HCI solutions must evolve into a more elastic solution. If an environment becomes "storage intensive" or "compute intensive," the infrastructure must be able to support this without stranding resources or requiring users to add a full node with more compute or capacity than is required.

OVERVIEW OF NETAPP'S HCI PORTFOLIO

NetApp announced its entrance into the rapidly growing HCI market in June of 2017, with its solutions becoming generally available in October of the same year. NetApp HCI represents a portfolio of fully integrated HCI solutions that support a wide range of configurable CPUs, GPUs, memory and storage capacity through compute and storage nodes. Based on SolidFire and its Element software first released in 2012, NetApp HCI was built in the cloud, for the cloud.

NetApp has designed its HCI solution to align with the evolving demands of hyperconvergence outlined above. The following is an overview of key features NetApp HCI provides:

- **Workload protection.** The underlying storage technology in NetApp HCI is Element software. Element software, core to NetApp HCI's architecture, allows for workload protection for every application consolidated on NetApp HCI through Element's QoS. NetApp HCI delivers guaranteed application performance for multiple workloads, allowing customers to consolidate many applications on it rather than using it for single workloads and creating more silos in the datacenter. Element software also brings mature data services and integration capabilities such as integrated replication, data protection, data reduction and high-availability services to HCI. Another highlight is Element software's API integrations on VMware stack, allowing simple centralized management through a vCenter plug-in that gives full visibility and control

over the entire infrastructure. NetApp HCI also supports ONTAP Select out of the box, giving ONTAP users the opportunity to continue using it on HCI. As workloads are dynamic and unpredictable, customers should use performance monitoring tools and metrics to correctly provision storage and compute to ensure performance.

- **Independent scaling.** NetApp HCI offers the flexibility of scaling compute or storage nodes independently with simplified choices of small, medium or large nodes. This allows users to decouple compute and storage capacity scaling. IDC believes this granular level of scaling helps to eliminate the overprovisioning resources common within systems that scale linearly rather than independently.
- **NetApp Data Fabric.** Data Fabric is NetApp's unified suite of data services that provide consistent data services across traditional on-premise, private clouds and public cloud platforms. Based on NetApp's well tested, enterprise class ONTAP software, this data fabric provides users with the "lingua franca" required to move data seamlessly and securely between different locations, architectures and cloud platforms — whether on-premises or within a public cloud. Expanding Data Fabric to its HCI solutions eliminates management silos common within many HCI solutions deployed alongside traditional data center infrastructure.
- **Automated deployment and management.** All NetApp HCI solutions come with a NetApp Deployment Engine (NDE), which reduces nearly 400 manual deployment steps down to less than 30 highly automated steps. Further, NetApp HCI solutions provide full integration with VMware vCenter via a plug-in, which makes ongoing management intuitive to a vast universe of VMware administrators. NetApp HCI leverages Element software APIs for highly automated integration into higher-level management, orchestration, backup, and disaster recovery tools.
- **Support for multi-hybrid cloud (Future).** Although a future feature at the time this paper was written, NetApp has announced that it will support multi-hybrid clouds through deployments of an on-premise region or a fourth hyperscaler to create a pool of cloud data services that can be accessed across hybrid multi-cloud and multiple private clouds.

Highlights from an In-Depth Interview with a NetApp HCI Customer

As previously noted, IDC interviewed a NetApp HCI customer to better understand how the solution is used in a real-world setting. The following is a short overview of that customer interview.

IDC interviewed a datacenter architect of a large U.S.-based hospital with more than 9,000 employees. There are two IT groups within the hospital that support general IT and research IT, respectively. This overview covers only the portion of the hospital supported by the general IT team. There are 175 staff members within this general IT team who manage 800 applications — of which 90% have been virtualized. These applications run on 350 physical servers and 1.4 petabytes of storage capacity. The hospital's general IT team manages two datacenters.

The hospital has deployed three NetApp clusters with a total of 42 compute nodes and 14 storage nodes. All three clusters were deployed to support the hospital's VDI. Two clusters were deployed to support 4,000 users (2,000 users per cluster); the third cluster supports 500 remote desktop session hosts (RDSH). The number of sessions supported will vary throughout the day, with concurrent sessions expected to reach peaks of 2,500 to 3,000 users. The vast majority of users supported will be a part of the clinical staff (e.g., nurses, doctors, x-ray technicians), with the IT staff also expected to make up a small of portion of sessions running simultaneously. All users are clustered into one of six types of use cases, based on their role within the organization. A total of 50 desktop applications are supported.

The hospital's initial HCI deployment was triggered by a project intended to increase the number of virtual desktops from 500 to 4,000 within the hospital. The goal of expanding the number of virtual desktops was to maximize the amount of time clinicians spent focusing on patient care. Moving 4,000 users to a virtual desktop environment greatly improved the workflow of hospital staff, who are now able to leverage features such as "batch sign-in," "follow me desktops," and "follow me printing."

This VDI expansion project provided the opportunity to explore new technology suppliers for the hospital's VDI software and hardware. The IT team considered HCI from the very beginning of this project due to the relatively small size of their datacenters. HCI proved to be an effective way to support a large number of VDI users within a very small footprint.

Solutions from multiple HCI providers were considered before the IT team decided upon NetApp HCI. NetApp's solution won out over all others, in part, because of its ability to scale compute and storage independently. The hospital's VDI architecture required a large amount of compute, but a relatively limited amount of storage. With NetApp HCI, the datacenter architect stated, "when we expand, we want to be sure that we only needed to expand the resources needed." While need to expand hasn't yet surfaced, the datacenter architect expects additional compute resources to be needed far more often than additional storage capacity. Thus, NetApp HCI has provided a solution that aligns well with the expected asymmetrical growth of resources and limits the need to buy unneeded storage capacity.

Initial deployment of NetApp HCI was very quick. The high degree of automation provided by the NetApp Deployment Engine allowed the datacenter architect to deploy and configure all 42 HCI hosts during a lunch break. This included all aspects of the initial infrastructure deployment, including installing ESXi software and creating the pools of storage capacity. When referencing the ease of deploying NetApp HCI, the datacenter architect said, "I went through a wizard, clicked 'OK' multiple times and then went to lunch. When I came back, everything was deployed." He went on to say, "this type of one-touch deployment was lovely." When asked to compare this process to deploying traditional solutions, IDC was told that the NetApp HCI deployment required far less effort, time and risk. NetApp HCI automated steps that were previously highly manual and siloed by technology.

When asked about any operational benefits related to the deployed NetApp HCI solutions, IDC was told that the solutions are not yet running in full production. As such, it was still too early to provide a complete set of metrics that quantify the operational savings from on-going management of the NetApp solution. That said, the datacenter architect was able to point to considerable benefits related to the steps required to create a datastore for a new host/application. In the past, creating a new datastore required many manual steps that included the creation of a LUN, configuring storage networking, presenting the LUN to a host and allowing the host to format the raw capacity. Now, with NetApp HCI, all these steps have been eliminated. Instead, the users are presented with a wizard that asks a few basic questions and then automatically runs through the set of previously manual steps. Further, because the HCI solution is highly consolidated, there is no need to coordinate multiple teams to build out a new host and provision storage for that host. Instead, one person can utilize the NetApp Deployment Engine. When asked to quantify how much time can be saved with the automation enabled by NetApp's Deployment Engine, IDC was told: "Including the coordination of teams, we've taken a three-day process down to just three minutes."

Lastly, IDC explored the hospital's plans to expand its use of NetApp HCI. We were told that the CIO is currently considering expanding the use of VDI on NetApp HCI from 4,000 users to more than 9,000 users. We were also told that the use of NetApp HCI could be expanded to support non-VDI applications after the current deployment had operated an ample amount of time within production.

Ultimately, the datacenter architect sounded optimistic by stating: "We have the option of expanding it to quite a few other areas within the hospital, to the point that it could become the standard for us."

Ultimately, IDC was told that the IT team was very happy with the experience with NetApp HCI. The solution exceeded expectations in key areas such as the amount of compute power and storage capacity contained within the small footprint, and the ease with which an ESX could be deployed. When referencing these experiences, the datacenter architect jokingly stated, "It's *almost* made me lazy." Jokes aside, this hospital's experience with NetApp HCI is indicative of the type of benefits frequently highlighted by organizations that have started down a similar path. And while this hospital is still working through the final production rollout of NetApp HCI, its decision to deploy NetApp HCI is already driving considerable savings.

CHALLENGES AND CONCLUSION

The first generation of HCI solutions were primarily driving operational simplicity and rapid time to service. This addressed critical pain-points related to complexity and inflexibility within the datacenter. However, as with any early market deployments, users were often creating silos of HCI clusters for specific applications and use cases. Companies now want to increase the workload/VM density of their HCI clusters to get the most out of their investments. As such, HCI solutions will need to evolve to address many of the challenges that the first-generation solutions were not designed to tackle, such as workload consolidation, independent and flexible scaling, and predictable performance.

Looking forward, datacenter infrastructure will be increasingly influenced by software-defined solutions running on server-based platforms that incorporate hybrid cloud into an administrator's workflow. While this change has undeniably begun, it is still in its early days. IT teams that do not yet feel ready to deploy this type of infrastructure should keep an eye on the future and work closely with companies like NetApp to make sure today's datacenter investments prepare the datacenter for such inevitabilities.

Organizations ready to move forward with hyperconverged solutions should look for solutions that expand the traditional operational simplicity benefits to also include next-generation features outlined in this paper, such as workload protection, seamless integration with public clouds, a common data fabric between private and public clouds, and independent resource scaling.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com