



# All-Flash Datacenter – Hyperconvergence and Web-scale

COMMISSIONED BY

**NUTANIX**™



## 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

The enterprise storage market has had a relatively sheltered existence over the last few decades with few opportunities for market disruption despite the extraordinary technology developments that have taken place in this area. Competition within the storage market was limited to incumbents and startup players, which entered with the same proprietary array mindset as their predecessors and used similar pricing and partnership strategies to gain market share.

In the past three to five years, however, hyperconverged platforms and cloud services have risen to 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. Flash has become another major game changer in the storage market. It has allowed startups to upend monolithic tier one disk arrays by delivering high performance with undeniable efficiency and management benefits.

The convergence of high-performance flash storage with hyperconverged platforms is the next leap forward in the evolution of data storage. While these two technologies have already demonstrated their relative merits in the market, the combination of flash and hyperconvergence is more than just an accretive blending of benefits; it provides a transformative value boost because the strengths of flash amplify the benefits of hyperconvergence and vice versa.

The antiquated 'rip and replace' lifecycle of traditional storage is based on the three-to-five-year usable lifespan of storage arrays and is punctuated by disruptive and costly array migrations that occur toward the end of an array's productivity cycle. Hyperconverged systems augmented with flash storage provide a simpler, just-in-time productivity cycle where storage capacity can be added linearly and conveniently when needed, without the disruption of forklift storage array upgrades. The scale-out architecture of hyperconverged platforms allows IT administrators to gracefully remove older hardware from a cluster without disruption, similar in manner to the way Web-scale cloud implementations are able to leverage the latest commodity hardware improvements.

This report discusses the benefits of flash and hyperconvergence, as well as the heightened value that can be attained by blending these innovations into an all-flash datacenter.

## Flash is the choice for hot data with enterprise applications

### HIGH PERFORMANCE DRIVES DATACENTER BENEFITS

The storyline around NAND flash-based storage usually focuses on the substantial raw performance benefits of flash over hard-drive-based arrays. While a conventional hard drive can only provide 150-200 random input/output operations per second (IOPS), an equivalent flash SSD can push thousands of IOPS in the same form factor. As a result, all-flash arrays (AFA) on the market today can adequately handle even the most demanding tier one application workloads with performance requirements from hundreds of thousands to millions of IOPS.

Beyond raw performance, all-flash storage also delivers a number of key benefits that help customers justify making the leap from disk storage to flash, including:

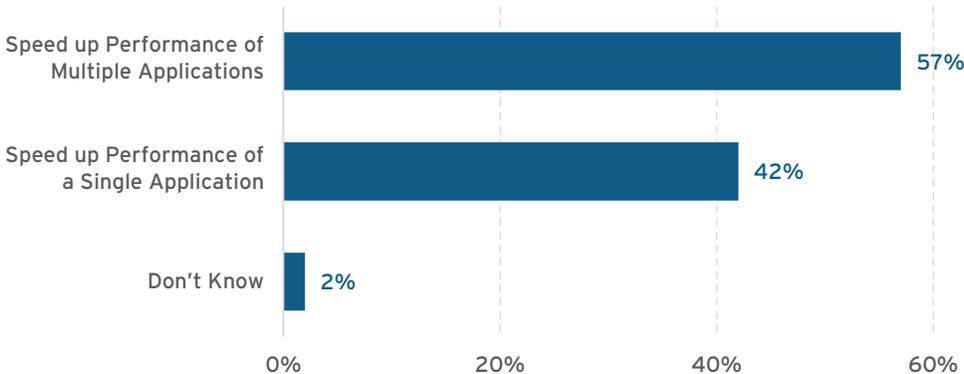
- **Power savings.** As we mentioned earlier, flash storage can deliver more performance per watt compared to conventional disk arrays. This is a key attribute for service providers and enterprises that are trying to maximize power utilization in their datacenters to avoid costly facilities upgrades to deliver more power. For organizations in dense urban areas, where additional power is not available, all-flash storage could help prolong the life of existing implementations.
- **Reduction in rack space consumption.** A single SSD can provide the performance of dozens of hard drives and, as a result, can drastically reduce the physical footprint of storage in a datacenter. It is common to see several racks of disk storage consolidated to a modest quarter rack or less of flash storage, particularly in high-performance scenarios where an IT department has deployed dozens or even hundreds of spindles to meet the IOPS requirement of a challenging workload.
- **Reduced maintenance and outages.** Unlike hard drives, which have spinning disk platters and actuators to store and access data, solid-state devices such as PCIe flash cards and SSDs have no moving parts, which allows them to be reliable, even in some environmentally challenging scenarios. A single SSD or flash card offers the equivalent performance of dozens of hard drives, so from a hardware maintenance perspective, it reduces the number of devices that can fail in an array.
- **Less application tuning and fewer code rewrites.** Because of their high performance, all-flash systems have been used to overcome application performance issues. In some cases, organizations were able to avoid sharing databases across multiple storage systems because the aggregate performance of a flash array was able to handle the entire performance workload. The high performance of flash arrays can also reduce the amount of time spent tuning applications and databases for performance, which not only saves time for internal database administrators but can also reduce the amount of money organizations invest in professional services to optimize applications.

### Workloads moving to flash for consolidation

With the growing popularity of flash-based storage systems, enterprises, midsize companies and service providers have begun moving applications to flash. In talks with early all-flash adopters, we found that databases and performance-sensitive applications were the primary workloads being moved to flash storage. VDI continues to be a key application. Analytics is another workload for flash storage consideration, and we believe these use cases will only grow in importance. Beyond the workloads themselves, what is even more important as customers look to make a large investment in flash is the fact that workload consolidation is now the norm, in place of single application acceleration implementations. This transition is an important one because it signifies that AFA deployments are moving past niche use cases to general mainstream usage.

**Figure 1: Flash Provides Consolidation – All-flash Array: Purpose AFA moving from niche to general usage**

Q: If you answered “in use” for all-flash array, which of the following best describes your use case? n=53



Source: 451 Research's Customer Insight, TheInfoPro Storage 2015

### Flash performance has a business impact

Although flash-based storage clearly has important facilities and maintenance benefits for datacenter environments, early adopters have found that the business performance benefits that are possible with the high performance of flash can actually dwarf the cost of upgrading and transitioning to an all-flash datacenter.

Early adopters have found the following benefits from their flash storage deployments:

- **Faster analytics and decision-making.** The accumulation of data alone will not generate a positive business value if it takes too long for organizations to reap actionable insights from the information residing in their repositories. We anticipate that the use of flash will increase for analytics and real-time applications (such as trading applications and the emerging space for streaming analytics) because faster insights will provide a competitive benefit for facilitating timely decision-making.
- **Meeting performance SLAs.** In financial services and shipping and logistics markets, flash has been deployed in use cases where the existing disk-based storage infrastructure dragged down performance below the SLAs. In these cases, the SLA misses would have had led to serious ramifications, including lost clients, a degraded reputation for the organization and costly penalties. In one extreme logistics example, the organization would have gone out of business if it were unable to boost performance to acceptable levels.
- **Increased worker productivity.** The performance of flash can also raise the productivity of knowledge workers. In some developer organizations, the deployment of flash arrays has reduced the amount of time required for code recompilations, which ultimately gave developers a few additional hours of productivity in a day. Besides increasing productivity, in some scenarios, flash deployment led to faster software releases and increased software sales. Worker productivity increases were also seen in VDI deployments, where faster logins and data access provided a superior experience for clients, helped to reduce helpdesk calls and increased satisfaction levels of customers.

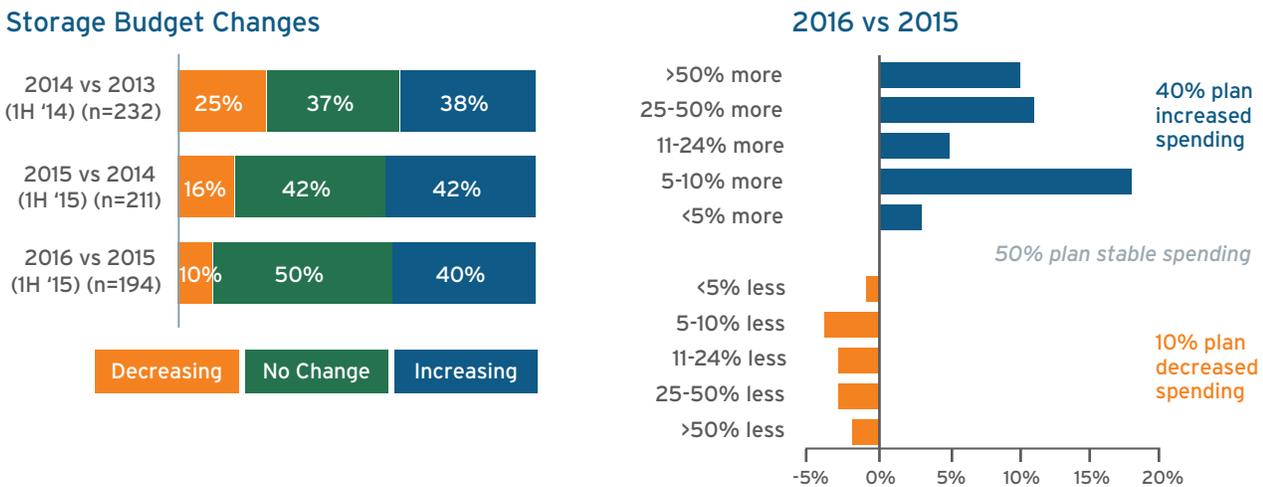
### Hyperconvergence simplifies the datacenter

Hyperconvergence platforms combine compute, storage and networking resources into modular building blocks that can be scaled out to match increasing resource demands, almost always using a software-based approach. The consolidation of these resources – especially storage – helps simplify the provisioning, management and troubleshooting of infrastructure to help accelerate the time to value for infrastructure technology investments. A number of factors are increasing the need for architecture changes, including flat storage budgets, relentless data growth, and the desire to reduce management and troubleshooting complexity and cost.

**STORAGE TEAMS ARE STRUGGLING, AND MOST BUDGETS ARE NOT INCREASING**

Over the last few years, storage teams have been asked to do more with fewer resources at their disposal. While data at many organizations continues to grow at a 30-40% annual rate, 60% of organizations reported in our Wave 19 Storage survey that their storage budgets were either staying flat or decreasing (See Figure 2). In the survey, only 15% of respondents expected to receive a sizable budget increase of 25% or more, which indicates that storage professionals can no longer rely on traditional disk systems to get by. The growing capacity and performance burden driven by business stakeholders and the performance increases required by applications and virtualization consolidation are forcing storage professionals to make a change. To get to a sustainable storage environment, organizations must look beyond their legacy arrays to find more efficient ways to deliver storage resources in a timely manner.

**Figure 2: Storage Budget Trends - 60% are Flat or Declining!**



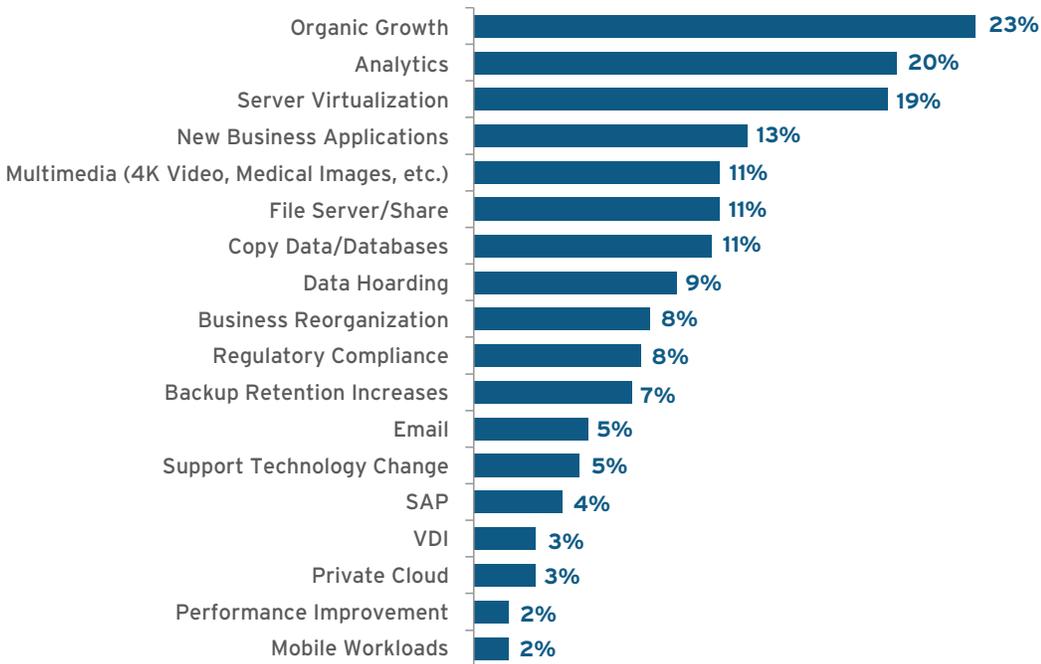
Source: 451 Research's Customer Insight, TheInfoPro Storage 2015

**DATA GROWTH AND DATA-PERFORMANCE DEMANDS PLAGUE ORGANIZATIONS**

Organic business growth continues to be the largest driver for storage growth at organizations. Analytics (20% of respondents) and multimedia (11% of respondents) were also listed as major drivers for capacity growth, along with server virtualization (19% of respondents). Unfortunately, the storage capacity challenge will only get worse as the years go by, given that copy data (11%), regulatory compliance (8%) and backup retention increases (7%) were all listed as significant drivers for data growth. Regulatory compliance challenges are forcing organizations to retain data for longer periods, and they prevent storage professionals from simply deleting old data to make room for new content. The importance of copy data and backup retention 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.

**Figure 3: Drivers of Capacity Growth**

Q: What projects, technologies or initiative are most responsible for capacity growth in your networked storage? n=244



Source: 451 Research's Customer Insight, TheInfoPro Storage 2015

### IT ORGANIZATIONS ARE STRUGGLING WITH COMPLEXITY

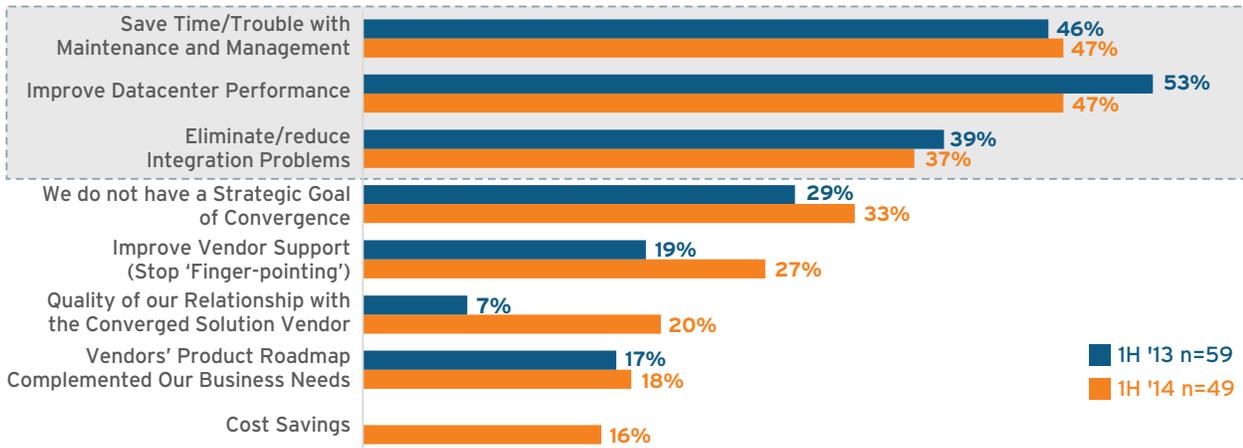
Beyond the storage silo, IT organizations are struggling with the complexity of their environments and have recently turned to converged systems to help reduce their management burden and time to value. For 47% of respondents, converged systems were seen as a means of reducing the amount of time and trouble related with maintenance and management of their systems. The elimination and reduction of integration problems was listed as a convergence driver for 37% of respondents, which is a major factor that can help organizations quickly reap the value from their costly IT infrastructure investments. Performance continues to be a major issue for organizations, and for 47% of respondents, converged offerings were viewed as a means of improving datacenter performance (See Figure 4).

In addition to leveraging flash and cloud technologies, storage professionals are also driving storage projects around storage virtualization, analytics tools and optimization technologies to keep up with their data burdens. While converged infrastructures can help accelerate time to value, the day-to-day storage management burden is still a major factor that organizations must deal with.

# PATHFINDER REPORT: ALL-FLASH DATACENTER – HYPERCONVERGENCE AND WEB-SCALE

**Figure 4: Drivers for Convergence**

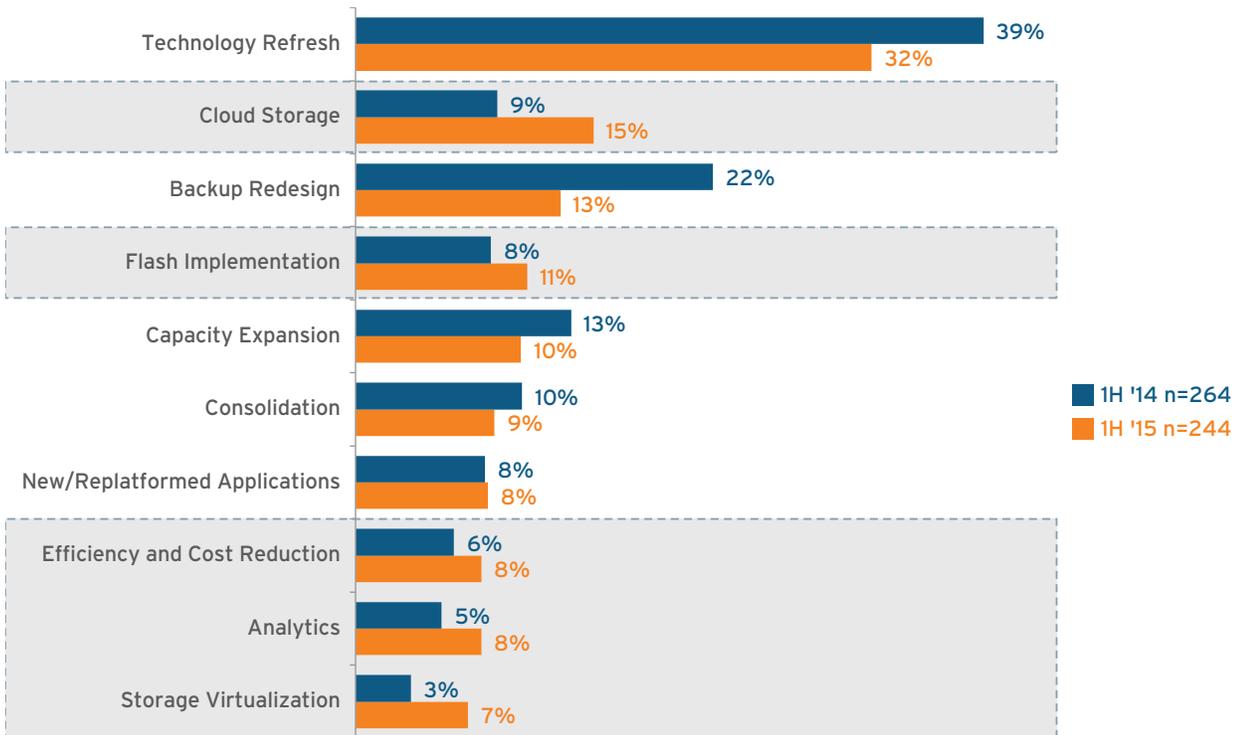
Q: What are the primary drivers behind adopting convergence?



Source: 451 Research's Customer Insight, TheInfoPro Storage 2015

**Figure 5: Storage Projects Highlight Need For Efficiency**

Q: What are your organization's top three storage-related projects in the next 12 months? List up to three.



Source: 451 Research's Customer Insight, TheInfoPro Storage 2015

**THE RISING INFLUENCE OF VIRTUALIZATION AND CLOUD TEAMS**

Virtualization and cloud teams have risen to power over the last decade and are now looking to wield their influence to boost their own management and resource utilization capabilities. For 61% of respondents in a recent 451 Voice of the Enterprise (VoTE) survey, VMware administrators have a growing role and influence in storage procurement, and we believe this trend will only increase with time. While these teams worked closely with colleagues in the storage team in the past, a number of factors have forced cloud and virtualization teams to take a more active hand in infrastructure purchasing and management decisions, including:

- **Slow resource provisioning.** The rising expectations of customers and partners has greatly increased the need for rapid resource provisioning for storage, processing, memory and networking – especially in virtualization and cloud environments. Slow provisioning continues to be a major cause of client dissatisfaction at organizations, and is usually one of the biggest drivers that business stakeholders use to move workloads to public cloud environments. To stay relevant, IT organizations must look to accelerate provisioning by using automation and integrated tool sets, which cut across the server, storage and networking silos in a traditional IT organization.
- **Lack of visibility and inefficient troubleshooting.** The narrow focus of specialized storage, networking and application management tools makes troubleshooting a difficult process. Cloud and virtualization teams want broader access to management tools to help pinpoint where problems are occurring and are no longer content to wait for other colleagues to help them.
- **Rise of virtualized business-critical applications.** With the exception of development sandbox environments, few workloads are considered non-critical in most environments today, and as such, cloud and virtualization professionals need faster and more granular recovery tools to protect their workloads. As the need for disaster recovery and business continuity continues to rise, IT organizations will need to have simplified tools that are highly integrated and optimized for the infrastructure, similar to what we see today in hyperconverged platforms or hypervisor-level data-protection software packages. Going forward, infrastructure and workloads must work together seamlessly to help facilitate workload migration, both in the event of a disaster and for the opportunistic consumption of idle resources at another facility or another cloud.

**Figure 6: VMware Administrators growing in influence**

Q: To what extent do you agree with the following statements? Please use a 1-5 scale where '1' is completely disagree and '5' is completely agree. n=80 to 83



Source: 451 Research's Customer Insight, TheInfoPro Storage 2015

**The Flash Datacenter – Hyperconvergence and Web-scale**

Currently, all-flash storage and hyperconvergence are the two top innovation trends in the datacenter, and the combination of these two disruptors will have a significant impact on IT infrastructures in the years to come. The ongoing convergence of flash and hyperconvergence will continue to evolve because these two innovation tracks complement each other. Although flash provides a significant performance boost over conventional disk-based storage while reducing management concerns, Web-scale architectures can make the all-flash performance engine even better by adding scale-out capabilities and providing VM-level granularity for provisioning and data protection. As infrastructures grow to handle the growth in workloads and data, IT organizations will need simplified management tools to deliver resources quickly and to snuff out troubleshooting issues before those issues can harm their productivity and reputation.

## IMPACT ON APPLICATION LIFECYCLES

Looking across the lifecycles of applications, especially business-critical workloads, the combination of all-flash storage and hyperconverged platforms can provide improvements over the siloed architectures found in traditional IT environments. The key value proposition of Web-scale-based hyperconvergence combined with all-flash storage is most prominent when looking at how these technologies will impact various phases of the application lifecycle.

---

## PLANNING AND DEPLOYMENT PHASE

The planning and deployment phase of infrastructure acquisition is difficult because it forces IT organizations to make hard decisions that will impact their datacenters for at least the next three to five years. Organizations often purchase far larger configurations than necessary to ensure that the performance and capacity capabilities of their compute and storage systems do not get overrun by demand before the next major refresh cycle. Databases and VDI will continue to drive up performance requirements along with new use cases such as real-time applications and analytics.

In contrast, Web-scale hyperconverged platforms with all-flash storage can provide the following benefits:

- **Scale-out architectures.** With these architectures, Web-scale hyperconverged platforms allow customers to start out with modest configurations of just a handful of nodes for an initial set of applications. They can quickly add capacity and performance to match needs without a forklift upgrade to replace existing systems. From a planning perspective, this model is superior because customers do not have to commit to making large purchases at the start, and they can defer hardware acquisitions until they actually need the additional resources, which is a major factor in helping IT organizations use their budgets more efficiently. As deployments get larger, organizations can leverage the power of their hyperconverged clusters to consolidate more applications in their infrastructure to boost the efficiency and utilization of their resources.
- **Infrastructure that is aligned with virtualization and enterprise/business applications.** Web-scale hyperconverged appliances hook into the management tools of virtualization platforms and applications, which allows these stakeholders to gain access to resources without the intervention of storage teams. From a planning perspective, this gives application and virtualization stakeholders more leeway to deploy quickly with a modest amount of resources and gradually increase to production performance and capacity levels.

---

## OPERATIONS PHASE

Web-scale hyperconverged platforms can deliver improvements for the management, performance-optimization and data-protection aspects of the operations phase.

### MANAGEMENT

- **Consolidated view of processor, memory and storage resources.** From a management perspective, Web-scale hyperconverged platforms provide a consolidated view that complements existing virtualization tools such as VMware vCenter to view resource utilization and contention issues that occasionally arise in shared infrastructure environments. This particular benefit is particularly noticeable in troubleshooting sessions, where traditional environments and siloed management teams often struggle with the sharing of management data and collaboration to close out trouble tickets.
- **Rapid VM cloning to facilitate testing for patches and upgrades.** The ability to rapidly create clones and test environments is a major strength for Web-scale hyperconverged platforms. Beyond accelerating test/dev for building new applications, these capabilities can be valuable for testing new patches and security fixes before deploying them out to production systems.

### PERFORMANCE OPTIMIZATION

- **Locality benefits.** Although flash and other forms of solid-state media can deliver high performance, the response times of these technologies are limited by their proximity to server CPUs. For example, in a traditional SAN environment, a request traveling from a CPU to a SAN-attached flash array can take a millisecond to hundreds of microseconds to reach its destination. In all-flash hyperconverged systems, flash resources reside in the same physical system as the server CPU, which can help reduce latency to tens of microseconds. While this may not seem like a large amount of time, as applications are consolidated, the aggregate performance required from a cluster increases the number of transactions that need to be fulfilled and affects the latency performance of the cluster.

- **Performance scalability.** The scale-out architecture of Web-scale hyperconverged systems allows organizations to aggregate the performance of their clustered systems. As individual systems get taxed, workloads or individual VMs can not only use the local, but also the remote flash resources, or if necessary, they can be seamlessly moved to cluster nodes, which have more idle flash and RAM resources to meet the performance need. With flash in the mix, individual nodes can deliver far higher aggregate performance and can also make performance more consistent by allocating the resources based on the needs and relative value of the workload.

### DATA PROTECTION

- **VM-centric snapshots for local protection.** In contrast to volume-level snapshots, which are common with traditional storage systems, VM-centric snapshots are becoming more popular in virtualization environments because they can provide superior granularity (at a VM level). Innovations continue to occur in the space, with some vendors adding single-file-restore capabilities to their VM-centric snapshots to accelerate and simplify data recovery.
- **VM-centric replication.** Similar to snapshots, the big benefit of having VM-centric replication is the superior granularity of this deployment. In contrast to traditional storage systems, which replicate at a volume level, with VM-centric replication, organizations can choose which VMs to replicate to a secondary site for remote data protection or for workload migration purposes.
- **Deduplication.** As we mentioned earlier, space efficiency is a key requirement for organizations given that most budgets are not increasing. Deduplication and in-line compression capabilities allow organizations to store more data in their storage footprints. 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 the flash investment.
- **Disaster recovery and business continuity.** Beyond the core replication and snapshot capabilities, a Web-scale hyperconverged platform should have workload-intelligence and orchestration tools to simplify failover and failback operations. With the emergence of hybrid clouds, customers will also be able to leverage public cloud resources for data and workload recovery in the event of a datacenter outage.

### HARDWARE RETIREMENT

- **Remove hardware from clusters without disrupting performance or availability.** The data and workload-mobility capabilities of Web-scale hyperconverged systems allow customers to gracefully move workloads off older systems without suffering through planned or unplanned downtime.
- **Retire hardware without painful data migration.** A core principle of Web-scale environments is that they can continue to deliver services even when nodes fail. With Web-scale hyperconverged platforms, organizations can pull old nodes or failed nodes out of the rack and replace them with newer systems. Because these platforms use commodity hardware, the new nodes can deliver higher performance at a lower cost than their predecessors. This allows organizations to reap the benefits of X86 processor and flash innovations in a rapid and migration-free manner.
- **Eliminate disruptive and expensive three-year hardware-replacement cycles.** The consistent lifecycle of Web-scale hyperconverged platforms allows organizations to scale up their resources gracefully without wholesale disruptive hardware replacements. This allows organizations to provide continual improvements to their infrastructure without forcing stakeholders to wait for the next refresh cycle to leverage the latest technologies.

### Conclusion and Recommendations for IT Professionals

Datacenter infrastructure teams, including storage personnel, are struggling with rising performance and capacity demands, but flash has helped alleviate the problem. Flash systems are becoming mainstream and are being used for consolidation of databases and other tier one applications. The correct usage of flash is necessary to maximize the value of this precious investment. In parallel, Web-scale technologies and hyperconvergence are helping simplify datacenters by consolidating compute and storage and simplifying performance management and scalability. Consider Web-scale-based hyperconvergence with flash to improve the efficiency of tier one applications throughout their lifecycles.