White Paper

Reviewing the Current State of Hyperconvergence and Real-World Benefits of VMware Virtual SAN Deployments

Sponsored by: VMware

Eric Sheppard
July 2016

IN THIS WHITE PAPER

This IDC white paper reviews the trends driving organizations to deploy hyperconverged systems within their datacenters. It also provides an overview of VMware Virtual SAN and the results of a recent IDC survey of VMware Virtual SAN customers. The survey results offer insights into the most common areas VMware Virtual SAN customers are seeing the benefits of the solution and the amount of capex and opex savings experienced.

SITUATION OVERVIEW

The datacenter infrastructure market is undergoing a once-in-a-generation change that is driving a fundamental need to rethink how systems are designed, deployed, managed, utilized, and refreshed. While enterprise infrastructure has seen steady improvements around resiliency, density, efficiency, and performance with each new product refresh, the impact of these improvements rarely had datacenterwide implications. That is because many datacenters were built with silos of infrastructure (e.g., storage, compute, networking) that were managed independently by silos of experts. Such a structure worked well in the past but is leaving an increasing number of IT departments struggling to keep up in a world where major digital transformations are all too common. These digital transformations are driven by a group of new technologies that, when combined, provide CIOs a brand-new platform to create new value in their marketplace that would have been impossible a few years ago. The core technologies that make up this new platform (which IDC calls the "3rd Platform") are focused on cloud, mobile ecosystems, big data, analytics, and social technologies. Companies are leveraging these 3rd Platform technologies to drive new revenue streams, to drive deeper connections with their customers, or to simply compete more effectively in their markets. Importantly, these 3rd Platform technologies are fundamentally changing the way applications are written and the way datacenter infrastructure is designed, purchased, and managed. Datacenter staff are looking for infrastructure that can:

- Reduce the amount of time datacenter specialists (e.g., storage administrators) spend conducting low-value tasks by automating where possible
- Simplify management tasks that cannot be automated in order to allow IT generalists and application owners to perform an increased amount of management and provisioning
- Efficiently scale to accommodate vast new levels of users and devices by seamlessly incorporating public or private cloud

- Eliminate LUN-level storage policies and management in favor of virtual machine (VM)–level infrastructure management and SLAs
- Leverage software-defined storage features to help reduce or eliminate the use of storage area networks and silos of storage arrays
- Leverage x86-based servers as a standard building block for compute and storage services
- Incorporate scale-out architectures that allow users to start small when needed and scale without forklift upgrades or complex migration projects

The Adoption and Evolution of Converged Systems

Many organizations have been deploying converged systems to help deal with these important datacenter changes. Such systems are providing tight integration between core datacenter infrastructure while offering centralized management and increased levels of automation. At a broad level, converged systems represent a consolidation of disparate datacenter technologies that can be acquired, deployed, managed, and supported as though they were a single system. Converged systems are designed to be deployed quickly using a modular building block. Real-world deployments typically include servers, storage systems, networking equipment, and core system infrastructure software for centralized management and automation. The adoption of integrated systems has grown quickly over the years to become a multibillion-dollar market, which can be seen as a clear indication of their effectiveness; total spending on converged systems exceeded $10.5 billion in 2015. That said, most of the converged systems deployed to date have focused on tightly integrating well-established infrastructure systems that have long been available as standalone servers, storage systems, or networking equipment. Arguably, the use of these systems as core building blocks represented the first phase of the converged system market, which has lasted for more than five years. Like today's IT departments, the converged systems market is rapidly evolving. An important component of this evolution is the recent emergence of hyperconverged systems, which IDC considers a subset of the converged systems market and the next phase of the market's life cycle.

IDC believes hyperconverged infrastructure (HCI) is helping deliver the proven benefits of the first generation of converged systems, but doing so through a clustered, scale-out architecture that is built on x86 servers. Hyperconverged infrastructure leverages software-defined storage technologies to provide the core data services normally found on discrete shared storage systems. Users can choose to purchase hyperconverged solutions in one of two ways: as a complete appliance with all features included or as a software-based solution, which then can be deployed on precertified servers. A key characteristic of hyperconverged systems that differentiates them from traditional converged infrastructure is their ability to provide all compute, storage, and networking functionalities through the same server-based resources (or nodes). These nodes are usually deployed as clusters with three or more nodes per cluster, though some solutions do support two-node clusters. Each node within a cluster contributes all of its resources to an abstracted (i.e., virtualized) pool of capacity, memory, and compute resources. This pool of virtualized resources provides the foundation for all server-centric workloads (e.g., VM, hypervisor, and applications) as well as storage-centric workloads (e.g., data persistence, data management, replication, snapshots, and deduplication).

Benefits of Hyperconverged Infrastructure

Broadly speaking, hyperconverged infrastructure deployments can help drive lower capital costs, increased operational efficiency, reduced risk, better alignment of IT skill set, increased datacenter agility, and lower datacenter facilities costs associated with power, cooling, and floor space. The following is an overview of how this is achieved:

- **Lower capital costs achieved through the elimination of SAN-based storage solutions in favor of industry-standard servers that offer fully virtualized compute and storage services.** The scale-out architecture of hyperconverged solutions further lowers capital costs by helping 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 overprovisioning and elimination of storage silos also helps lower costs of power, cooling, and floor space within the datacenter.
- **Reduced risk of downtime associated with common life-cycle management tasks (e.g., firmware upgrades, system refreshes) because of the highly virtualized and clustered architecture of hyperconverged systems.** This architecture also enables faster system resource provisioning and scaling.
- **Integrated management software that automates many of the complex tasks needed during initial deployment, which also reduces 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 leverage IT generalists for low-value tasks, thus freeing up time for infrastructure specialists to work on more innovative projects.

OVERVIEW OF VMWARE VIRTUAL SAN

VMware has become an important supplier of hyperconverged technology and solutions over the past few years. The company works with many other technology suppliers and channel partners to offer multiple hyperconverged solutions, all of which are built on VMware Virtual SAN. VMware's hyperconverged software (HCS) solution that includes vSphere, vCenter, and Virtual SAN is designed and optimized for vSphere environments.

Virtual SAN was developed as a full software-defined storage solution with separate control and data planes. The control plane offers highly automated storage management that focuses on the unique needs of the application rather than the traditional hardware-centric approach of external storage systems. The data plane provides a scale-out architecture built on x86 servers and a virtualized SAN with flash acceleration. Data persistence is based on a distributed object store that supports VMDK-level management. Virtual SAN pools the storage resources within a cluster into a single shared data store that is abstracted and made available to any of the VMs running in the cluster. Flash-based storage and tiering are critical parts of the Virtual SAN architecture. Virtual SAN supports hybrid configurations (using HDDs for data persistence and SSDs for caching) and all-flash configurations (using SSDs for both caching and data persistence). Both configurations have a two-tier approach where writes are always written in flash first and then destaged to the persistent/capacity tier. Reads in a hybrid configuration are also cached in SSD, but reads in an all-flash configuration are taken directly from the flash resources within the persistent/capacity tier.

Review of Key VMware Virtual SAN Features

VMware Virtual SAN became generally available in February 2014 with vSphere 5.5. Those were the very early days of the hyperconverged market, and not surprisingly, the offering was initially targeted for use in noncritical or test/dev workloads. Much has changed since its initial release, thanks to the considerable resources VMware has put into the continued development of Virtual SAN. VMware Virtual SAN has matured significantly and now boasts a robust feature set designed for a wide range of enterprise workloads and use cases. VMware Virtual SAN is now able to provide customers all the benefits highlighted in the previous section (reduced capital costs, increased operational efficiency, reduced risk, better alignment of IT skill set, increased datacenter agility, and lower datacenter facilities costs). IDC has highlighted some of the more recent features that have helped increase VMware Virtual SAN's ability to drive new levels of efficiency and resiliency. These features include:

- **Data deduplication and compression** are helping drive new levels of storage efficiency within a VMware Virtual SAN all-flash cluster. Combined, these two features allow users to eliminate costs associated with wasted capacity and help drive down network traffic. Deduplication and compression have become critical features for all-flash deployments and ultimately help users offset the relatively higher-dollar-per-gigabyte costs of flash media.
- There is **support for erasure coding** within all-flash configurations. VMware Virtual SAN now allows customers to manage resiliency levels and capacity usage at very granular levels by choosing between mirroring or erasure coding at the VMDK level. This provides the considerable flexibility and level of granularity needed to further drive out wasted capacity by letting customers fine-tune the amount of capacity used to ensure data availability of a virtual machine. Rather than forcing users to pick a single storage policy for an entire cluster or pool of capacity, VMware Virtual SAN lets users assign a unique storage policy to each VMDK deployed.
- **Stretched clusters** help drive continuous application availability in the event of a host failure or a complete site failure. A hyperconverged solution protecting against VM-level failures is important but not groundbreaking. VMware Virtual SAN's stretched cluster capabilities go much further by protecting against the failure of an entire host or even an entire site when used in conjunction with VMware's vSphere HA and vMotion. Users can also leverage Virtual SAN's stretched cluster feature with vMotion to perform site maintenance without taking a virtual machine offline.
- The **quality-of-service (QoS)** feature provided within VMware Virtual SAN can limit the amount of IOPS used by each virtual machine. This allows customers to maximize their resources by running workloads with differing I/O profiles and can also ensure no one workload will consume enough IOPS to negatively impact the performance of the system.

QUANTIFYING REAL-WORLD BENEFITS OF VMWARE VIRTUAL SAN DEPLOYMENTS

This section provides insights into a sample of VMware customers currently running VMware Virtual SAN within their organization. A survey of 82 customers was completed during the first half of 2016. Results from other IDC surveys are also included in this section to help provide historical context.

A selection of respondent demographics for the 2016 *IDC Survey of VMware Virtual SAN Customers* includes:

- Size of companies:
  - Fewer than 500 employees = 39% of respondents
  - 500 to 4,999 employees = 29% of respondents
  - 5,000+ employees = 32% of respondents
- On average, respondents had 33.4 physical servers deployed in their organizations.
- On average, respondents had 426 virtual machines running in their organizations.
- 71 of the 82 respondents were using VMware Virtual SAN in production environments.

Survey Results

As noted previously, hyperconverged infrastructure is considered a subset of the broader converged systems market, which has been around for more than five years. Looking back at past IDC survey data, one can see a material shift in the major challenges users have looked to address with converged systems. Figure 1 shows us results from two separate surveys conducted in 2014 and 2015. Figure 1 shows that IT staff productivity, agility, or credibility has grown considerably as a reason for deploying converged systems. Capex-related challenges remain an important driver of adoption too. Combined, these two challenges were called out as the most important drivers of converged systems adoption approximately 65% of the time in each survey. Recent IDC surveys and anecdotal evidence of hyperconverged users have returned comparable results. That is to say, the vast majority of users are looking to hyperconverged solutions as a way to lower the total cost of ownership over their datacenter infrastructure.

With Figure 1 showing us what drives people to adopt converged and hyperconverged solutions, Figure 2 explores why our VMware survey respondents chose VMware Virtual SAN as their hyperconverged solution. The top 3 most common reasons users had chosen VMware Virtual SAN were attributable to its tight integration with their hypervisor (57% of responses), the product's ability to be deployed quickly and easily (44% of responses), and low purchase price/total cost of ownership (38% of responses). IDC believes the high ranking of these responses is tied to VMware's ownership of the entire virtualization stack, which has allowed VMware to integrate Virtual SAN directly into the kernel of its hypervisor and to embed the features/capabilities of Virtual SAN into its vCenter management tools. This has helped ensure simplified deployments and ongoing management and eliminate learning curves for new users already familiar with vSphere and vCenter.

FIGURE 1

Primary Challenges Users Are Looking to Address with Converged Systems

Q. What is the primary challenge that your IT organization is looking to address with converged systems?

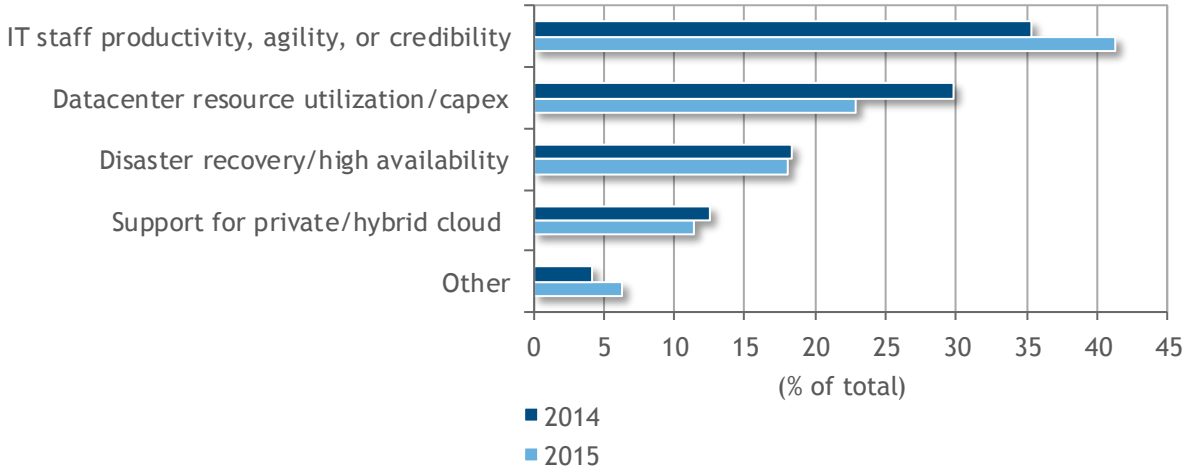

n = 289 (2014), 297 (2015)

Source: IDC's *Converged Systems Surveys,* 2014 and 2015

FIGURE 2

Most Important Criteria for Choosing VMware Virtual SAN

Q. What were the three most important criteria that helped you choose your VMware Virtual SAN solution?

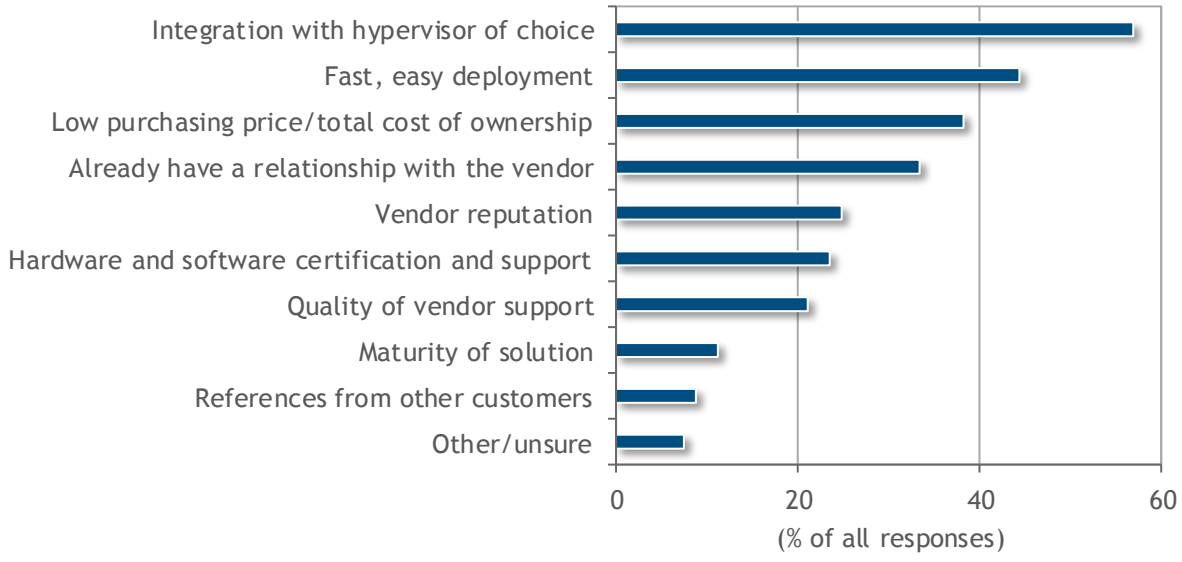

n = 82

Note: Multiple responses were allowed. A total of 220 responses were given.

Source: *IDC Survey of VMware Virtual SAN Customers,* 2016

Figure 3 reviews the general benefits our survey respondents realized through the use of VMware Virtual SAN. Figure 3 shows faster provisioning times, reduced capital spending, improved staff productivity, and improved utilization rates of storage and compute resources as being the most common benefits achieved. This aligns well with the general market drivers of converged and hyperconverged solutions and tells us that VMware Virtual SAN is able to provide the most important features/benefits to real-world users of hyperconverged solutions.

FIGURE 3

Benefits Achieved with VMware Virtual SAN Solutions

Q. Which of the following best describes how your organization has benefited (or expects to benefit) from the use of your VMware Virtual SAN solutions?

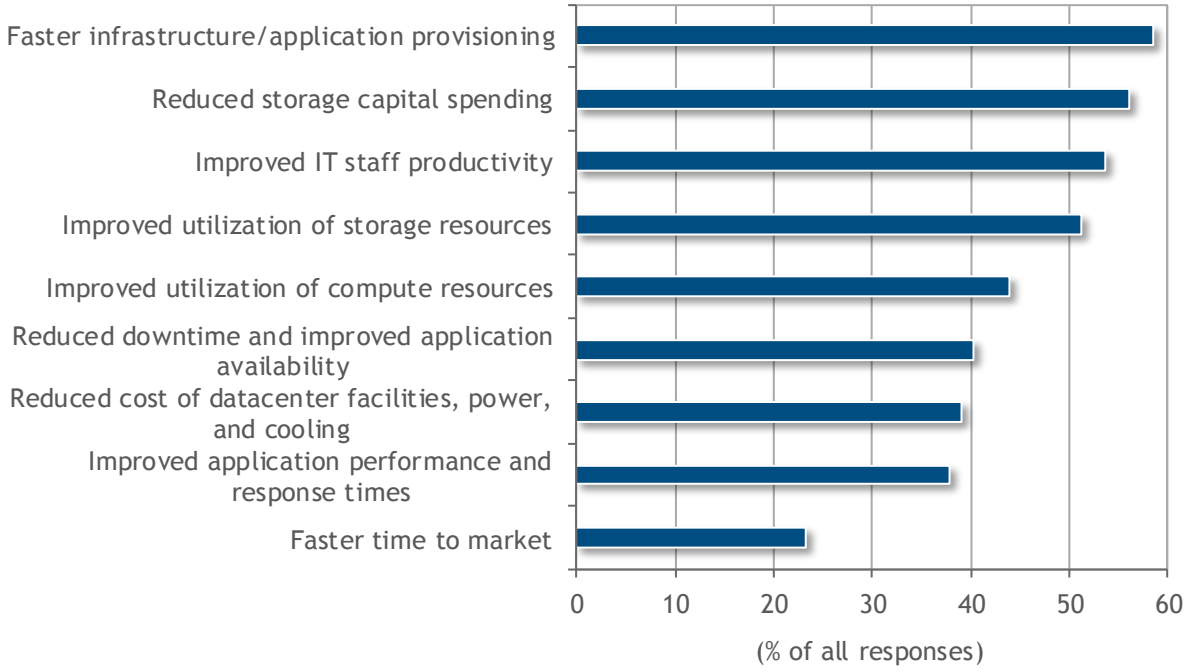

n = 82

Note: Multiple responses were allowed. A total of 335 responses were given.

Source: *IDC Survey of VMware Virtual SAN Customers,* 2016

Figure 4 provides insight into how much more efficiently users of VMware Virtual SAN are able to conduct some of the most common tasks conducted by IT staff. In Figure 4, the percentages represent the amount of improvement rather than the number of responses. The savings shown in this survey are genuinely impressive for all of the tasks shown and are likely indicative of just how different VMware Virtual SAN (and hyperconvergence in general) is compared with traditional three-tier storage environments. The biggest efficiency gains can be found within tasks related to data protection or disaster recovery. The survey respondents have realized 59.5% savings in IT staff time spent on these tasks compared with pre-VMware Virtual SAN environments. This is followed by 56% savings and 53% savings on provisioning storage and virtual machines, respectively.

FIGURE 4

Quantifying the Improvements Achieved to Specific Tasks Performed with VMware Virtual SAN

Q. When thinking of the tasks that will be conducted more efficiently, please quantify the improvement against previous environments as a percentage of IT staff time.

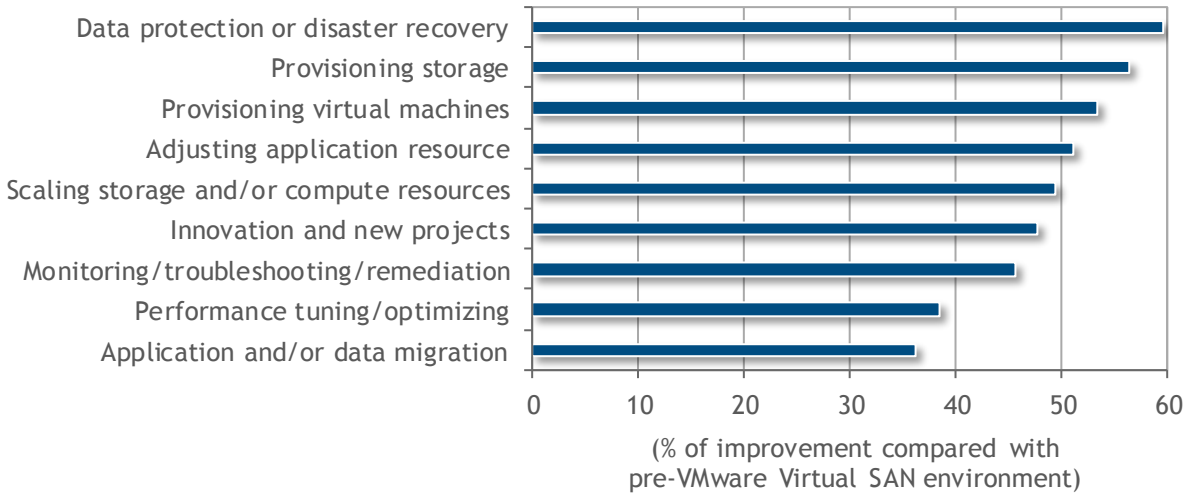

n = 82

Source: *IDC Survey of VMware Virtual SAN Customers,* 2016

Capital Expense Savings Experienced

Figure 5 provides the average capex savings experienced by VMware Virtual SAN users compared with the previous solution that it replaced. Not surprisingly, VMware Virtual SAN's ability to eliminate the need for external, shared storage by virtualizing the data services resulted in considerable reductions when reviewing capital costs of storage hardware. Here, respondents experienced a 32% capex reduction compared with the hardware of the solution it replaced.

FIGURE 5

Average Capital Cost Reduction Realized by Deploying VMware Virtual SAN

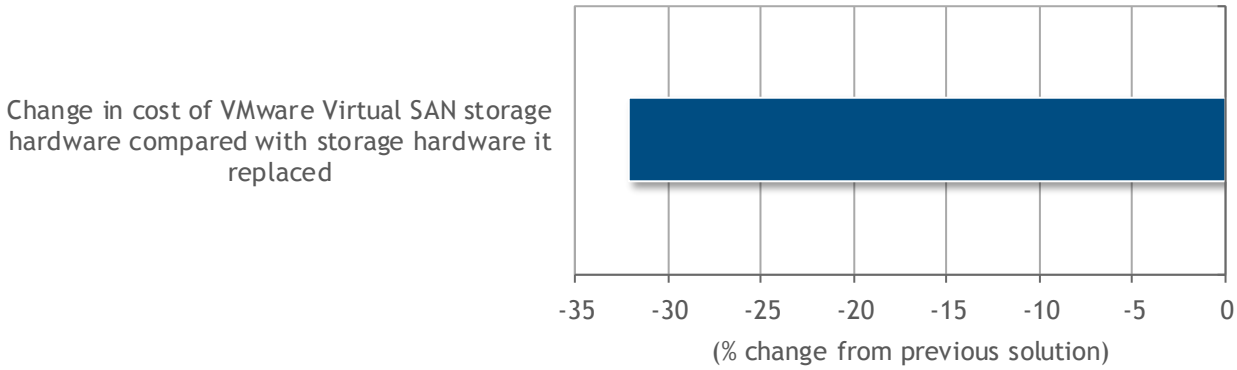

n = 82

Source: *IDC Survey of VMware Virtual SAN Customers,* 2016

Operational Expense Savings Experienced

Figure 6 provides the average opex savings experienced by VMware Virtual SAN users compared with the previous solution that it replaced. On average, customers experienced a 23% reduction in IT staff time spent on managing storage, a 12% reduction in the cost of their datacenter/facility space, a 10.4% reduction to their total cost to power and cool the solutions, and a 5% reduction in the number of full-time employees (FTEs) required to manage the storage resources. IDC would expect these savings to grow over time as users increase their use of VMware Virtual SAN and migrate additional workloads to the systems. A reduction in the number of FTEs required to manage storage is a frequently mentioned benefit for HCI users, the vast majority of whom use this reduction to redirect staff resources toward higher-value tasks. Also, it should be noted that HCI architecture often requires change management and an overall different way of thinking about infrastructure than traditional/legacy storage.

FIGURE 6

Average Operational Cost Reductions Realized by Deploying VMware Virtual SAN

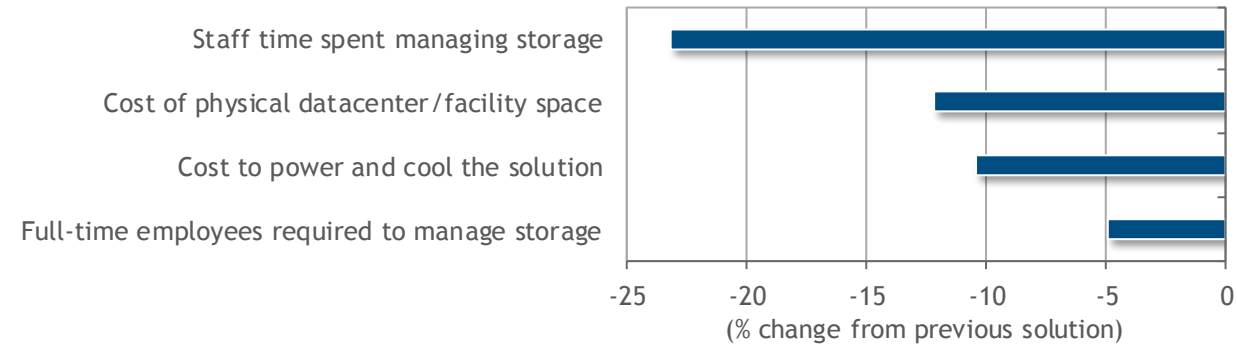

n = 82

Source: *IDC Survey of VMware Virtual SAN Customers,* 2016

High-Level Overview of Capex and Opex Savings

Quite a few factors go into datacenter infrastructure capital expenditures, which will result in differing levels of savings when moving to a hyperconverged solution. A few important factors include current use of storage infrastructure software (e.g., storage tiering software, replication software, system management software, data protection software), class of existing external storage system, amount of raw capacity, type of storage media, or type (and complexity) of the storage network being replaced. This list is not exhaustive, but it includes some of the biggest contributing factors to capital cost reductions when moving to a hyperconverged solution. It should also be noted that running software-based storage services on the same systems also hosting the applications will likely require more robustly configured servers (e.g., more memory and faster CPUs). This will move some capital costs previously associated with SAN-based storage to costs associated with servers and software. That said, the overall initial costs are often lowered by the elimination of the dedicated storage, increased levels of workload consolidation, and higher utilization of server resources. The scale-out, clustered nature of hyperconverged systems can increase initial/up-front capital savings by enabling a more granular level of resource and cluster scaling.

The degree of operational expense savings experienced when moving to a hyperconverged solution will also vary by environment. That said, organizations should expect hyperconverged solutions to drive opex savings through improvements associated with initial installation, faster resource provisioning and system scaling, reduced risk and planned downtime for life-cycle management, increased IT staff productivity, and increased business unit productivity because of reduced application I/O latency. Of course, operational expenses also include facility costs. Here, one may expect to benefit from a smaller datacenter footprint by migrating workloads from external, shared storage to hyperconverged solutions that offer software-defined storage services through the same systems hosting the workload.

Combined, the capex and opex savings achieved by migrating workloads to hyperconverged solutions can help lower total cost of ownership considerably. This is especially true when savings are measured over the typical life of a system (normally three to five years). Although the hyperconverged market is still relatively new, adoption of the technology has reached a point where companies are able to measure the savings achieved through the examples listed previously. A VMware customer that is in the IT professional services and cloud service provider market claims to have lowered its storage-related capex by as much as 65% after migrating workloads to the VMware Virtual SAN solution. A U.S. government agency claims to have reduced its storage-related capex by 33% and ultimately achieved a 50% reduction in total cost of ownership. Again, the amount of savings will vary for each deployment. That said, there is ample evidence building up to suggest hyperconverged solutions can drive real benefits and savings.

CHALLENGES, OPPORTUNITIES, AND CONCLUSION

Decades of innovation have brought us datacenter solutions that are undeniably more capable than anything offered just a few short years ago. That said, too many datacenter teams continue to buy and manage their infrastructure the same way they did 10 or 20 years ago. This process/structure has become untenable. IT departments must look to new datacenter infrastructure technologies like hyperconverged infrastructure if they want to keep up with unprecedented changes occurring all around them. Indeed, businesses of all sizes are looking to transform their company to find new revenue streams, create deeper connections with their customers, or simply compete more effectively against new competitors that are unburdened by legacy systems and practices.

The benefits and savings listed throughout this white paper are very real and indicative of those seen in other IDC surveys on hyperconverged infrastructure. And while the sample size of 82 customers is not uncommon for a market as nascent as hyperconverged infrastructure, users should couple these findings with complementary resources to get the most robust level of insight into real-world experiences.

Finally, IDC recommends a measured, common-sense approach when deploying hyperconverged solutions. Start by using VMware Virtual SAN with workloads that have already been virtualized, places where storage skills may be limited (e.g., a remote/branch office), or within environments where a three-tier shared storage system may be cost prohibitive. IDC has talked with a large number of organizations that started working with hyperconverged solutions in a comparable way. Over time, the vast majority of these companies expanded their use of hyperconvergence by migrating an increasing number of workloads onto the solution.

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