

From the Field: Software Defined Storage and Hyperconverged Infrastructure in 2016

January 5, 2016



ActualTech Media

By **Scott D. Lowe, vExpert**
Partner, ActualTech Media

In partnership with

ATLANTIS 

Table of Contents

Table of Contents	2
List of Charts	3
Executive Summary	4
Introduction	5
Technical Foundation	6
Technology Domain Knowledge	6
Virtualization Penetration	7
Hypervisor Usage.....	9
Remote Office and Branch Office Support	14
Storage Characteristics.....	15
Storage Capacity	15
Data Growth	20
Flash Storage Deployment.....	21
Storage Performance	22
Current Storage Systems.....	23
The Future of Storage	25
Adoption Plans/Criteria.....	28
The Form Factor Debate	30
Business Workload Support Needs	31
Software Defined Storage and Hyperconverged Infrastructure Deployment Intent.....	32
Software Defined Storage and Hyperconverged Infrastructure Deployment Experience	33
Software Defined Storage or Hyperconverged Infrastructure Deployment Timeframe	34
Software Defined Storage and Hyperconverged Infrastructure Decision Criteria	36
Appendix: Demographics.....	39
About.....	42
About Atlantis Computing	42
About ActualTech Media.....	42

ActualTech Media © 2016. All rights reserved.



Under no circumstances should this document be sold, copied, or reproduced in any way except with written permission.

All ActualTech Media research must be cited verbatim and commercial citation of research is prohibited without our express written permission. All citations must be limited in scope; full reproduction of ActualTech Media research is prohibited. Graphs can be briefly summarized when no corresponding text is available, within the appropriate context of and reflecting the spirit of the intended research.

All citations for this research must be attributed in the following manner:

Source: From the Field: Software Defined Storage and Hyperconverged Infrastructure in 2016, ActualTech Media, January 2016

The information contained with the document is given in good faith and is believed to be accurate, appropriate and reliable at the time it is given, but is provided without any warranty of accuracy, appropriateness or reliability.

The author does not accept any liability or responsibility for any loss suffered from the reader's use of the advice, recommendation, information, assistance or service, to the extent available by law.

List of Charts

Figure 1: Knowledge levels of various data center technologies	6
Figure 2: Virtualization penetration for servers and desktops	7
Figure 3: Virtualization penetration for servers and desktops – SDS/HCI adopters.....	8
Figure 4: Hypervisor in use in respondent organizations	10
Figure 5: Current and future hypervisor/container breakdown by product.....	11
Figure 6: Current and future hypervisor/container breakdown by product – SDS adopters	11
Figure 7: Current and future hypervisor/container breakdown by product – HCI adopters	12
Figure 8: Hypervisor breakdown by vertical	13
Figure 9: Number of remote sites supported	14
Figure 10: Storage capacity in primary and across remote locations	15
Figure 11: Storage capacity by company size (primary location only).....	16
Figure 12: Storage capacity by vertical (primary location only)	17
Figure 13: Storage capacity by vertical (remote locations only)	18
Figure 14: Storage capacity for SDS/HCI adopters (remote locations only)	19
Figure 15: Respondent annual storage capacity growth rate.....	20
Figure 16: Respondent annual storage capacity growth rate - by vertical.....	20
Figure 17: Flash storage deployment statistics	21
Figure 18: Performance of SDS/HCI solutions (adopters only)	22
Figure 19: Types of storage systems currently in use.....	23
Figure 20: Types of storage systems currently in use (By vertical).....	23
Figure 21: Types of storage systems currently in use (By company size)	24
Figure 22: Respondent future storage plans	25
Figure 23: Respondent future storage plans (By vertical)	26
Figure 24: Respondent future storage plans (By company size).....	27
Figure 25: Storage adoption intent by vertical	29
Figure 26: Respondent thoughts on form factor	30
Figure 27: Top use cases broken down by analysis segment.....	31
Figure 28: SDS/HCI deployment potential	32
Figure 29: SDS/HCI deployment potential (By vertical)	32
Figure 30: SDS/HCI deployment potential (By company size)	33
Figure 31: SDS and HCI deployment experiences	34
Figure 32: SDS/HCI deployment timeframe – Aggregated, by Vertical, and by company size.....	35
Figure 33: SDS and HCI adoption decision criteria – All respondents	37
Figure 34: SDS and HCI adoption decision criteria – Small companies.....	37
Figure 35: SDS and HCI adoption decision criteria – medium companies	38
Figure 36: SDS and HCI adoption decision criteria – large companies	38

Executive Summary

Software defined storage and hyperconverged infrastructure are technologies taking the data center by storm. ActualTech Media, in partnership with Atlantis Computing, surveyed more than 1,200 IT pros from 53 different countries to gain their insight into these technologies and how they might fit with current and upcoming data center plans. This executive summary provides some highlights from our research.

53 countries



1267 respondents



24+ market verticals



71% have either adopted or are targeting adoption of SDS/HCI



27% of respondents have adopted either SDS or HCI
44% of non-adopters indicated that they are either considering or definitely adopting SDS and/or HCI

Disk is dying! 19% of respondents intend to decommission all-disk systems



Introduction

IT budgets are shrinking. Demands on IT are increasing. Data center technology has become a quagmire of complexity. Traditional storage has struggled to keep pace with workload demands. With these challenges, CIOs, technical decision makers, and IT staff members are looking for ways to continue meeting critical business needs with solutions that stabilize data center costs while also being simpler to manage. Perhaps the biggest challenges facing the data center today revolve around storage. It's expensive. It's complex. And, until flash became more common, it suffered a great deal with regard to performance.

Both *software defined storage (SDS)* and *hyperconverged infrastructure (HCI)* have emerged as solutions intended to solve the storage problem. They have entered the market mainstream as forceful options for consideration. Both bring heretofore unheard of levels of simplicity while also helping to turn the data center economic picture on its head. Rather than buying three to five years' worth of storage, data center administrators can take more of a "just in time" approach to storage thanks to the easy scalability opportunities that present themselves with these architectural options.

Much remains misunderstood about software defined storage and hyperconverged infrastructure, though. There is often confusion about what these terms even mean. In short, software defined storage leverages a software layer to provide storage services separate and apart from what is provided by the hardware. While it is possible to build a brand new software defined storage architecture, many organizations add software defined storage tools to supplement the capabilities of their existing storage devices in order to breathe new life into them by adding modern data services (i.e. deduplication and replication). To expand capacity in a software defined storage system, administrators can either add more nodes (scale out) or add more storage to existing nodes (scale up), making such systems easily scalable.

Hyperconverged infrastructure takes the data center to new levels by eliminating the array altogether and combines storage and compute into single nodes. In both cases, growth is achieved via *scale out* mechanisms. As more capacity is needed, administrators need only to add another node to the storage or hyperconvergence cluster.

With great interest in these technologies, we sought to understand what businesses think of each. To that end, we surveyed more than 1,200 IT professionals and business decision makers to get their thoughts around these technologies and how adopters are using them. This report is the culmination of that effort. We provide herein an analyses of the total population as well as breakdowns by verticals, company size, and SDS/HCI adoption as appropriate.

Technical Foundation

We begin our analysis with a view of the respondent organization technology landscape.


Technology Domain Knowledge

We begin our analysis with a look at how the IT pros that responded to our survey view their own knowledge of various data center elements. As is very obvious, the only area in which respondents believe that they have expert level knowledge is server virtualization, with 55% responding as such. For two primary emerging technologies – software defined storage and hyperconverged infrastructure – 12% and 18%, respectively, of respondents feel that they have little to no knowledge of the subject matter. Only 18% of respondents feel that they have expert-level mastery of each these topics. Given the relative age of these technologies when compared to other data center technologies – server virtualization, datacenter networking, and enterprise storage – it's not that surprising that knowledge level is quite a bit lower. Over time, we expect to see people's comfort level with software defined storage and hyperconverged infrastructure approach that of enterprise storage, which 39% of respondents have mastered. You will notice that, overall, people are more comfortable with software defined storage over hyperconverged infrastructure. 12% say that they have no knowledge of software defined storage while 18% say the same about hyperconverged infrastructure. This is likely due to the fact that many software defined storage systems more closely resemble traditional storage arrays whereas hyperconverged infrastructure is quite different.

FIGURE 1: KNOWLEDGE LEVELS OF VARIOUS DATA CENTER TECHNOLOGIES

Knowledge level around data center technologies (N=1267)

■ Expert level knowledge ■ Some knowledge ■ Little or no knowledge



Virtualization Penetration

Particularly with hyperconverged infrastructure, virtualization penetration is a key indicator for just how much of the existing environment can be migrated. Hyperconverged infrastructure deployments require that applications run virtualized. With that in mind, gaining an understanding for a respondent's level of virtualization is important to learn just how successful that deployment might be. We learned from respondents that most are at least 71% virtualized on the server front, but that desktop virtualization is truly still in its infancy or, at the very least, not of interest to many organizations. Only 19% of respondents are more than one-half virtualized on the desktop.

For those considering software defined storage rather than hyperconverged infrastructure, virtualization levels aren't really all that important except for the fact that virtualization is just another workload type to support.

Virtualization penetration rate - servers and desktops (N=1267)




FIGURE 2: VIRTUALIZATION PENETRATION FOR SERVERS AND DESKTOPS

In Figure 3, you can see the virtualization penetration rate for those that have deployed either software defined storage or hyperconverged infrastructure. The results aren't radically different, but you can see that 75% are at least half virtualized. The most interesting item here really revolves around the desktop. In the total population, a full 23% have done no VDI. For those that have deployed either software defined storage or hyperconverged infrastructure, only 10% have not deployed VDI. This suggests that virtual desktops are of more interest to SDS/HCI adopters.




FIGURE 3: VIRTUALIZATION PENETRATION FOR SERVERS AND DESKTOPS – SDS/HCI ADOPTERS

Hypervisor Usage

Given the magnitude of virtualized applications – people are virtualizing more and bigger workloads all the time – hypervisor choice is a critical issue. Not every hyperconverged infrastructure solution is able to support every hypervisor available on the market. It's with that in mind that it comes as no surprise that VMware vSphere remains the dominant choice in the hypervisor market (Figure 4). It's also no surprise to see that, over the next 24 to 36 months, many vSphere administrators intend to migrate to the latest version of VMware's hypervisor. Hyper-V will be the likely recipient for much of vSphere's loss. XenServer 6 looks to hold pretty steady as well. However, for those on XenServer 5, it looks like they will abandon the platform for other options.

We were surprised to see that KVM did not increase in perceived future market share. In fact, based on our results, KVM's share of the market will actually decrease a little. There are a number of hyperconverged infrastructure solutions on the market that use KVM as their core. With that in mind, we believe that, rather than a decrease, we will probably see KVM adoption *increase* over the next few years. Here's why: the hypervisor layer has achieved commodity status. For many, the actual hypervisor in use really doesn't matter as long as the solution meets all needs. With the KVM-based hyperconverged infrastructure options on the market, users may not focus as much on what they're actually running. When we ask them to focus on the hypervisor, KVM doesn't stand out, but in practice, many may not really care, especially in smaller organizations.

We were not surprised to see Docker more than doubling in adoption in the next few years. Container technology is getting more attention and, much as was the case in the early days of virtualization, we expect to see container adoption start to become more interesting to people as they learn more about the technology and as it expands to support more workload types.

Hypervisor deployment statistics




FIGURE 4: HYPERVISOR IN USE IN RESPONDENT ORGANIZATIONS

What is surprising is what happens when the hypervisors are grouped into their respective vendor categories and then analyzed. As you can see in Figure 5, VMware will maintain its market share. This may seem insignificant, but when considering the market as a whole, there is a story there, especially as the virtualization market is likely to continue to grow overall. What that means is that VMware is likely to simply maintain share in a growing market while those that are abandoning other platforms – such as Citrix – are more likely to jump to Hyper-V rather than VMware.

Those companies providing what are considered “next generation” options – such as Docker – will also rise significantly in popularity in the next few years. For today – and likely for your next replacement cycle – VMware remains the clear leader in workload support, but over time, as more hypervisor and container options grow in usage, expect to see more hyperconverged solutions that provide comprehensive support for these products. While most people don’t care whether or not a solution will support multiple hypervisors, they *do* care whether or not a solution supports the hypervisor or technology in use in *their* organization.

Hypervisor deployment statistics - by vendor




FIGURE 5: CURRENT AND FUTURE HYPERVISOR/CONTAINER BREAKDOWN BY PRODUCT

Now, let's look at the same information, but this time just for those that have already adopted software defined storage (Figure 6). The information here suggests those that deploying software defined storage will do so at VMware's expense, with that company dropping from 81% of market share to 76% among SDS adopters. Likewise, Citrix will drop from 30% share to 24% and Oracle will lose 4% of its share. The gainers will be Microsoft, KVM, and Docker. Microsoft is poised to gain 6% share among SDS adopters, while KVM will see 3%, and Docker a large 7% increase, almost doubling its market penetration.

Hypervisor deployment statistics - SDS adopters (N=254)




FIGURE 6: CURRENT AND FUTURE HYPERVISOR/CONTAINER BREAKDOWN BY PRODUCT – SDS ADOPTERS

Among hyperconverged infrastructure adopters (Figure 7), the trends are similar, but with a somewhat different magnitude. Here, VMware's market drops from 85% to 77%, a full 8% drop, which is substantial. Microsoft's Hyper-V starts today at 42% and is expected to jump to 47% among our HCI adopter population. Citrix only loses a single point of their adopter base, and KVM jumps a full 3%, to 18%. We do expect to see an increase in KVM adoption among hyperconverged infrastructure users as the KVM-based HCI options continue to penetrate the market. Among HCI users, Oracle usage is poised to drop 6%, which is interesting since Oracle has their own converged infrastructure solution. And, again, Docker looks to gain significant followers as that product continues to improve.

Hypervisor deployment statistics - HCI adopters (N=193)





FIGURE 7: CURRENT AND FUTURE HYPERVISOR/CONTAINER BREAKDOWN BY PRODUCT – HCI ADOPTERS


Breaking down the hypervisor by vertical reveals some interesting details. Most notably, KVM deployment is quite a bit lower than the general population and it's particularly low in the finance and government verticals. Further, Oracle support in government and healthcare is much lower than found in other verticals as well. The news for Citrix is not positive in the government vertical, although that sample size is relatively small.

Government deployments of Hyper-V look like they might take off, too, with a 5% increase in deployments. But, the real interest here is Docker, especially in education and finance, which looks to triple its deployment in those verticals. Education is often willing to experiment with newer technologies such as Docker and, in finance, where there is a growing need to deploy programmable infrastructure and gain more efficiencies from infrastructure.


Education (N=124)



Finance (N=116)



Government (N=76)



Healthcare (N=126)




FIGURE 8: HYPERVISOR BREAKDOWN BY VERTICAL

Remote Office and Branch Office Support

One of the key use cases that has emerged for both software defined storage and hyperconverged infrastructure is supporting remote office and branch office (ROBO) environments. These technologies are very well-suited to ROBO needs and are emerging as a leading way to support ROBO environments. Figure 9 indicates that 9% of respondents have just one remote site. 15% of respondents have more than 50 sites to support.

These are not insignificant numbers. Supporting these environments can be a major challenge using traditional systems and for a number of reasons:

- Configuration drift. Even with the best laid plans, many organizations' ROBO environment suffer from "configuration drift." This means that, over time, systems begin to vary from their established baseline configurations, resulting in a variety of configurations to support.
- Hardware and software variety. Many times, ROBO locations are not one-size-fits-all and have a wide variety of hardware and software.
- No personnel. Not all ROBO sites are staffed with IT pros.

**Number of remote sites supported
(N=600)**




FIGURE 9: NUMBER OF REMOTE SITES SUPPORTED

In particular, hyperconverged infrastructure is uniquely positioned to help ROBO-laden organizations solve these challenges by allowing them to deploy infrastructure in small, consistent blocks that can be easily and centrally managed.

Storage Characteristics

Both software defined storage and hyperconverged infrastructure are intended to solve the most vexing storage challenges inherent in the modern data center.

Storage Capacity

Being able to support workloads means having sufficient storage capacity in your organization across both your primary location as well as any remote or secondary locations. Both hyperconverged infrastructure and software defined storage solutions have the capability to support both very small as well as very large deployment scenarios and either one can support centralized or distributed storage needs. As you can see in the chart below, storage capacity varies widely and there are substantial storage resources housed at remote locations. From this chart, you can see that about 16% of respondents are running 20TB to 50TB of storage at their primary location. The most surprising piece of information here is just how much storage is present across remote and distributed sites. Only 18% of respondents indicate that they have no storage outside the headquarters location.

Storage resources available in the organization
(N=1267)




FIGURE 10: STORAGE CAPACITY IN PRIMARY AND ACROSS REMOTE LOCATIONS

It probably comes as no big surprise to learn that overall primary location capacity changes with company size. In Figure 11, you can see that smaller organizations tend to have less overall storage while large companies tend to have much more. While this is common knowledge, our data absolutely reinforces it.




FIGURE 11: STORAGE CAPACITY BY COMPANY SIZE (PRIMARY LOCATION ONLY)

When breaking the data down by our four primary verticals, it's really easy to see that the 20 TB to 50 TB storage range is the sweet spot for our overall respondent group (Figure 12). It's also easy to see that different verticals have somewhat different average storage needs. For example, only 4% of those in the education vertical are running 200 TB to 500 TB of storage whereas 21% from finance have that level of capacity. Given the data-driven nature of financial companies, this comes as no big surprise, but is nonetheless interesting. By comparing the individual bar sizes in Figure 11, you can begin to see where each vertical ranks with regard to storage capacity. Here are the major ranges for each vertical (again, this is storage capacity at the primary location only):

- Education: 20 TB to 50 TB
- Finance: 200 TB to 500 TB
- Government: 20 TB to 50 TB
- Healthcare: 50 TB to 100 TB

Storage capacity by vertical (primary location only) (N=412)




FIGURE 12: STORAGE CAPACITY BY VERTICAL (PRIMARY LOCATION ONLY)

Now, let's look at the storage capacity breakdown across the aggregate of all respondent remote sites. Figure 13 excludes storage at the primary location. The data here is slightly more mixed than we see with capacity figures at the primary location, with a large number of respondents having no remote storage capacity. However, for those that do have storage resources in remote sites, the 20 TB to 50 TB range is once again the leader of the pack, but we also see a jump in the number of overall organizations that have more than 1 PB spread across remote storage systems. As mentioned earlier, this situation reinforces the need for hyperconverged infrastructure and software defined storage solutions that focus on ROBO use cases. Here are the major ranges for each vertical (this time, this is storage capacity at remote sites):

- Education: 20 TB to 50 TB
- Finance: 20 TB to 50 TB and More than 1 PB
- Government: 5 TB to 10 TB
- Healthcare: 50 TB to 100 TB

Storage capacity by vertical (remote locations only) (N=412)




FIGURE 13: STORAGE CAPACITY BY VERTICAL (REMOTE LOCATIONS ONLY)

With ROBO being a key use case for hyperconverged infrastructure, we wanted to look at overall capacity at remote locations for organizations that deployed one of these technologies. There were a total of 342 respondents that have undertaken such deployments. In Figure 14, you can see the remote storage capacity breakdown for each technology. Earlier, we learned that storage capacity and company size are linked to one another; bigger companies have more storage. From Figure 14, it's clear that some very large companies have deployed both software defined storage and hyperconverged infrastructure since the choice "More than 1 PB" garnered the greatest number of respondents.

Capacity at remote locations for those that have deployed hyperconverged infrastructure or software defined storage (N=342)

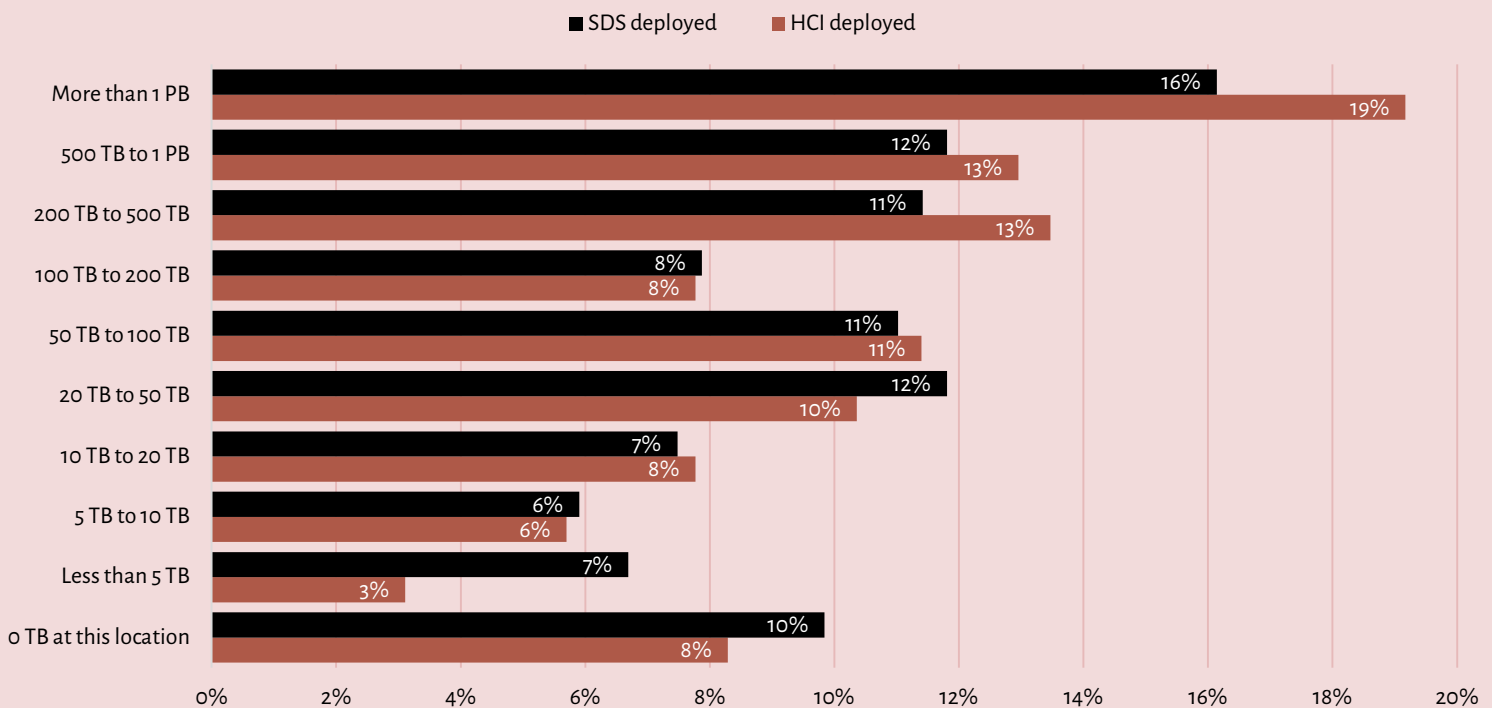


FIGURE 14: STORAGE CAPACITY FOR SDS/HCI ADOPTERS (REMOTE LOCATIONS ONLY)

Data Growth

Perhaps one of the most serious technology challenges facing organizations is keeping up with the sheer growth of data. Figure 15 shows you that most organizations are seeing a 10% to 30% annual data growth rate. However, a number of companies see much higher rates, even 50% or 100%. For these respondents, finding a storage solution that can scale easily and inexpensively is absolutely critical to maintaining reasonable level of expense and application availability.




FIGURE 15: RESPONDENT ANNUAL STORAGE CAPACITY GROWTH RATE

In the four charts below, we can get a look at the data growth patterns for the four primary verticals under scrutiny for our survey. As you can tell, in general, the data growth patterns are all pretty similar; most organizations, regardless of vertical, fall primarily into the 10% to 30% data growth range and have some kind of peak around the 50% data growth rate. Here, though, finance is something of an outlier, with its “secondary peak” coming in at around 45% with a smaller third peak coming at 65%.





FIGURE 16: RESPONDENT ANNUAL STORAGE CAPACITY GROWTH RATE - BY VERTICAL

It's a similar story when considering this information just for those that have deployed hyperconverged infrastructure or software defined storage. However, while the peaks are in similar places – in the 10% to 30% data growth range, fewer software defined storage users report these levels of growth.



Flash Storage Deployment

In recent years, flash storage has taken the market by storm and is poised to eventually mostly supplant disk as prices for flash continue to decrease. As of today, though, just 1% of respondent data centers are all flash. Over 60% of respondent data centers are less than one-tenth flash based, with 21% of respondents saying that they do not yet have *any* flash deployed. Just 6% of respondent data centers are over one-half flash. For vendors that are able to provide affordable flash solutions, this is actually a good news situation as there is significant upside in the flash market.

Current flash storage deployment penetration (N=1267)




FIGURE 17: FLASH STORAGE DEPLOYMENT STATISTICS

Storage Performance

While storage capacity is absolutely critical to consider, storage performance is also a key success factor for workloads. Over the years, storage performance challenges have become severe, leading to the rise of flash-based storage solutions and a renaissance in the overall storage market. Software defined storage and hyperconverged infrastructure are two rising markets that have emerged as a part of this renaissance. But, just how well are these newer entries in the storage market meeting performance goals?

As it turns out, pretty well. Only 16% of respondents have solutions that are slower than their disk-based storage systems. A full 50% say that their solutions are faster than their disk-based systems, with 14% saying that it's as fast as an all flash system (Figure 18). Overall, from a performance perspective, these newer storage options are holding their own.

Relative performance of your HCI or SDS solution (N=342)




FIGURE 18: PERFORMANCE OF SDS/HCI SOLUTIONS (ADOPTERS ONLY)

Current Storage Systems

With 75% of respondents still running such systems, disk-based storage still rules the data center, although it is joined by hybrid storage (55%), all flash storage (21%), software defined storage (21%), and hyperconverged infrastructure (16%) solutions (Figure 19).

Types of storage systems currently in use (multiple choices allowed)






FIGURE 19: TYPES OF STORAGE SYSTEMS CURRENTLY IN USE

Studying responses by vertical reveals that education lags behind in all flash deployments and that healthcare is ahead of the pack in hyperconverged infrastructure deployments. It's interesting to see that finance verticals are so much lower than the baseline when it comes to disk-based storage systems, but are somewhat ahead when it comes to hybrid storage. The message here is that there is opportunity for SDS and HCI in all market verticals.


Education (N=124)



Finance (N=116)



Government (N=76)



Healthcare (N=126)




FIGURE 20: TYPES OF STORAGE SYSTEMS CURRENTLY IN USE (BY VERTICAL)

Now, let's take a look at the kinds of storage systems in use broken down by company size (Figure 21). It becomes very clear that larger companies are leading the way to what might be considered new or emerging technologies. Almost 1/3 of respondents in large companies are running all flash storage systems. Large companies are also far ahead of the pack when it comes to deployments of hybrid storage systems and hyperconverged infrastructure.

Given the potential operational improvements that can accompany the deployment of hyperconverged infrastructure, these results are somewhat surprising. We would have expected to see more small and medium sized companies deploying the technology. However, these organizations tend to be more averse to risk as well, and may be taking a wait-and-see approach.

**Types of storage systems currently in use - by company size
(multiple choices allowed)**

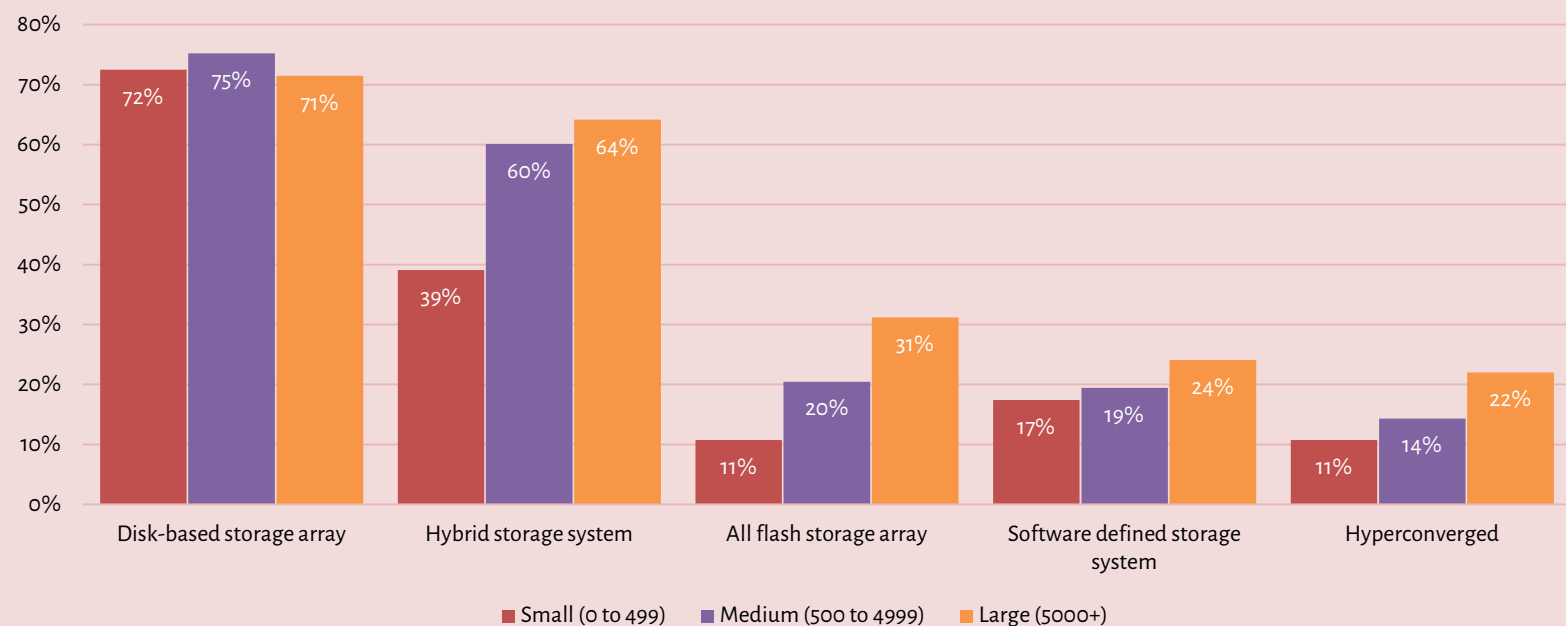


FIGURE 21: TYPES OF STORAGE SYSTEMS CURRENTLY IN USE (BY COMPANY SIZE)

The Future of Storage

When we asked respondents to tell us what their future (future = 2-to-3 years out) plans are regarding storage, the responses paint a bleak future for disk-based storage. A full 19% of respondents – almost 1 in 5 – say that they will fully decommission their disk-based storage systems over the next two to three years. The primary gainers in the same timeframe will be all flash arrays and hybrid storage arrays, but 35%+ also say that they will expand their use of software defined storage and hyperconverged infrastructure.

The future of storage in the data center (N=1267)





FIGURE 22: RESPONDENT FUTURE STORAGE PLANS

None of this comes as a major surprise. Flash storage has been plummeting in price for quite some time and is expected to hit price parity with disk within the next few years. Once raw price parity is achieved, expect to see spinning disk quickly fall off in terms of usage. Flash simply carries with it far too much performance potential when compared to disk.


Breaking the data down by vertical (Figure 23) yields some interesting results:

- Government workers hang on to everything. Across the board, their decommissioning rates are far lower than the average.
- Healthcare is ahead of the curve in decommissioning disk. At 24%, they are 5% likelier to remove spinning disk from the environment.
- Finance and healthcare are both ahead of the curve when it comes to deploying SDS and HCI. 39% of respondents in each vertical say that they will deploy more SDS. 41% of finance respondents and 38% of healthcare respondents say that they will deploy more HCI.


Education (N=121)



Finance (N=112)



Government (N=74)



Healthcare (N=123)





FIGURE 23: RESPONDENT FUTURE STORAGE PLANS (BY VERTICAL)

Large companies (5000+ employees) are absolutely the sweet spot in both the SDS and HCI markets. 46% of large company respondents say that they will add more SDS and HCI over the next two to three years. At the other end of the scale, small companies (fewer than 500 employees) are poised to hang on to what they have, with fewer than the average deploying anything new in the way of storage systems. The midmarket (500 to 4,999 employees) is average in just about every way when it comes to new systems deployment.


Small companies (N=484)

Will stay the same We will add more Will be decommissioned



Medium companies (N=380)

Will stay the same We will add more Will be decommissioned



Large companies (N=371)

Will stay the same We will add more Will be decommissioned




FIGURE 24: RESPONDENT FUTURE STORAGE PLANS (BY COMPANY SIZE)

Adoption Plans/Criteria

While software defined storage might be considered as architecturally similar to traditional storage in that storage remains isolated from compute, it is hyperconverged infrastructure – in which compute and storage are combined – that is of more interest to those considering these technologies. 27% of respondents are most likely to adopt the former while 33% plan to adopt the latter. However, those considering more traditional approaches still outweigh those looking at emerging storage approaches. 39% are considering all flash arrays while 45% are considering traditional systems, which include hybrid storage arrays (Figure 25).

50% of respondents, though, say that they intend to deploy cloud-based storage services. For the foreseeable future, we expect that most deployments will be of the hybrid variety in which organizations combine cloud-based storage with local storage. Over time, as more companies seek to further simplify their data center environments, many are turning to the public cloud, which eases deployment. However, because of security concerns, locality concerns, and even cost challenges, many companies are discovering that keeping things private makes more sense. We'll see how this plays out in the coming years, but for now, cloud is still a big plan for many.

This information would seemingly contradict what you just learned – that 19% of people currently using disk-based storage arrays intend to decommission them. However, bear in mind that, for those that intend to “remain the same” on disk-based storage, that means that they will ultimately need to replace them, which we believe is the reason that we see strong results for Traditional SAN/NAS devices in Figure 17. Also note that the response categories are slightly different, since we add cloud storage as an adoption option to this question.

In Figure 25, you can also see that we have provided a breakdown of adoption intent by vertical. It's clear that those in finance have major plans when it comes to storage in the coming years, with 67% intending to deploy all flash arrays. Finance also intends to add a lot of software defined storage (49%) and hyperconverged infrastructure (54%). We were somewhat surprised to see to relatively low software defined storage/hyperconverged infrastructure uptake intent in the education and government sectors, however. Especially in education, technology is often seen as a cost center, with the benefits of these emerging technologies helping to drive down costs.

Which technologies are you most likely to adopt?

■ All Data (Baseline) (N=1267) ■ Education vertical (N=126) ■ Finance vertical (N=78) ■ Government vertical (N=78) ■ Healthcare vertical (N=130)




FIGURE 25: STORAGE ADOPTION INTENT BY VERTICAL

The Form Factor Debate

Hyperconverged infrastructure and software defined storage solutions are available as either just software deployments or as hardware appliances that include the software. There are different solutions available depending on customer needs. Software-only solutions provide more hardware flexibility since the customer can specifically size the individual nodes. Preconfigured hardware appliances offer a bit less individual resource flexibility, but do offer a simplified deployment experience. As you can see in Figure 26 below, for those that have an opinion, most prefer appliance-based solutions, but not by a wide margin. 57% of respondents are keeping their options open and considering both kinds of solutions.

There is a lot of variety among form factor desire by vertical, but not much by company size (which, as a result, is not shown). As you can see in Figure 26, people in finance are far more likely than others to take a side in the form factor debate. Respondents in other verticals – particularly in education – preferred to consider both software and hardware-based solutions.

Thoughts on form factor (N=951)




FIGURE 26: RESPONDENT THOUGHTS ON FORM FACTOR

Business Workload Support Needs

Storage and data center infrastructure is deployed to support business workloads. We asked respondents to tell us what they want to accomplish with software defined storage and hyperconverged infrastructure. Figure 27 provides you with a look at the top three use cases identified by each segment that we analyzed for this chapter. As becomes very apparent, Test and Development is a clear top use case for those that have deployed or have an interest in software defined storage while server virtualization is, in general, a top choice for those that have deployed or have an interest in hyperconverged infrastructure. Given the highly versatile nature of software defined storage, it's not a surprise that it has use for more than virtualization tasks. Hyperconvergence, on the other hand, assumes that virtualization is the standard, and virtualized server workloads are a must on these platforms, hence respondent interest in server virtualization for hyperconvergence. Other top use cases include database workloads, VDI, private cloud, file and print and data center consolidation.

Data Segment	Type	Top use case	Secondary use case	Tertiary use case
All Respondents	SDS	Test and Development (43%)	File and print (41%)	Server virtualization (40%)
	HCI	Server virtualization (45%)	Virtual desktop infrastructure (VDI) (39%)	Database (39%)
Education vertical	SDS	Test and Development (44%)	Private Cloud (37%)	File and print (36%)
	HCI	Server virtualization (45%)	Database (41%)	Virtual desktop infrastructure (VDI) (40%)
Finance vertical	SDS	Test and Development (62%)	Database (59%)	Big Data (55%)
	HCI	Server virtualization (68%)	Test and Development (65%)	Datacenter consolidation (64%)
Government vertical	SDS	Database (46%)	Server virtualization (45%)	Test and Development (45%)
	HCI	Server virtualization (45%)	Test and Development (41%)	Virtual desktop infrastructure (VDI) (40%)
Healthcare vertical	SDS	Test and Development (42%)	Private Cloud (40%)	File and print (39%)
	HCI	Server virtualization (47%)	Database (45%)	Datacenter consolidation (42%)
Those that have deployed SDS	SDS	File and print (55%)	Server virtualization (55%)	Database (55%)
	HCI	Test and Development (47%)	Server virtualization (46%)	Datacenter consolidation (44%)
Those that have deployed HCI	SDS	Test and Development (45%)	Analytics (42%)	File and print (39%)
	HCI	Server virtualization (64%)	Private Cloud (56%)	Database (55%)

FIGURE 27: TOP USE CASES BROKEN DOWN BY ANALYSIS SEGMENT

Software Defined Storage and Hyperconverged Infrastructure Deployment Intent

As mentioned, software defined storage and hyperconverged infrastructure solutions are somewhat new entrants into the storage market and are, for many, still being proven. As they continue to prove their capabilities, more organizations will consider them for implementation. According to our survey respondents, 15% are either very likely or definitely planning to deploy such services over the next two to three years. 53% say that it's a possibility while 32% say that it's either not likely or there is no chance of deployment. In general, this is a good news story for vendors selling these solutions and is also a good indicator of interest in this technology for those considering the technology.

Deployment possibility for software defined storage or hyperconverged infrastructure (N=1011)




FIGURE 28: SDS/HCI DEPLOYMENT POTENTIAL

On a vertical basis, education remains the most resistant to moving toward one of these technologies with 39% saying that they definitely won't or are unlikely to undertake an adoption. On the other hand, at 25% saying very likely or definite, the finance vertical appears to be a sweet spot for hyperconverged infrastructure vendors.




FIGURE 29: SDS/HCI DEPLOYMENT POTENTIAL (BY VERTICAL)

As has been the case in other sections of this report, company size appears to be a good predictor for how likely an organization is to eventually deploy software defined storage or hyperconverged infrastructure. Just 17% of large companies say that they either won't or are unlikely to deploy SDS or HCI. On the other end of the spectrum, a full 40% of those from small companies say the same.

As we've said before, these results are somewhat surprising given the friendliness of HCI in particular when it comes to easing overall data center operational burden.




FIGURE 30: SDS/HCI DEPLOYMENT POTENTIAL (BY COMPANY SIZE)

Software Defined Storage and Hyperconverged Infrastructure Deployment Experience

Companies don't deploy technology for technology's sake. They deploy it in pursuit of a goal of some kind. Most often, new technologies are deployed because they either cost less or are more efficient in some way. This fact certainly holds true for software defined storage and hyperconverged infrastructure solutions. Given people's concerns around traditional storage costs and complexity, it would make sense that those that have adopted newer methodologies would do so to offset cost and complexity.

We asked respondents to tell us about their experiences with software defined storage and hyperconverged infrastructure as it relates to a number of different areas. Figure 31 provides a look at the results. In almost every area, people have had a better experience – or at least a comparable one – with software defined storage and hyperconverged infrastructure than they did with whatever they had before. The only exception is around personnel cost, which have increased for those that have deployed software defined storage.

In terms of systems performance, data center space, and power and cooling costs, there have been tremendous gains for implementers of software defined storage and hyperconvergence. On the performance front, it's more than likely that the gains have come from the fact that the previous storage was more disk-focused while the new solution is either hybrid or all flash.

Data center space is much improved in hyperconverged infrastructure scenarios since compute and storage are integrated together into server nodes. Further, less equipment in the data center translates to lower power and cooling costs.

You will notice that direct costs – acquisition and support costs – stay relatively constant. About the same number of respondents experienced higher costs as lower costs. While still providing a lot of value, software defined storage and hyperconverged infrastructure solutions have not yet helped companies reduce initial expenditures on infrastructure, but *have* helped when considering the total cost of ownership. This leaves a major opportunity for companies in the emerging storage space that can reduce both acquisition cost and TCO.

What has been your experience in each of the following areas following your adoption of SDS and HCI? (N=342)




FIGURE 31: SDS AND HCI DEPLOYMENT EXPERIENCES

Software Defined Storage or Hyperconverged Infrastructure Deployment Timeframe

Deployment of software defined storage and hyperconverged infrastructure is happening in waves and is more than likely taking place based on existing hardware replacement cycles. Over the next year or so, 17% of respondents say that they will undertake deployments. Over the next two years, that number jumps to a total of 62%. 27% of respondents say that they are uncertain as to their deployment plans, which could mean that they are still not sure whether they will definitely deploy or they truly don't know when they might plan a deployment.

By vertical, you can see quickly that finance and healthcare plan to move very quickly into SDS and/or HCI. In finance, 42% of those responding say that they will make a jump within the next 12 months. In healthcare, that number is 30%.

In terms of company size, 25% of small companies and large companies say that they will also deploy within the next 12 months as well.

Hyperconverged infrastructure or software defined storage deployment timeframe (N=630)











FIGURE 32: SDS/HCI DEPLOYMENT TIMEFRAME – AGGREGATED, BY VERTICAL, AND BY COMPANY SIZE

Software Defined Storage and Hyperconverged Infrastructure Decision Criteria

Most people would likely assume that cost would be the primary decision point around any new technology, but, interestingly, that's not the case. Cost is actually tied for second in terms of decision criteria. Overall performance of a solution is the key issue for many people (Figure 33, 72%), while cost is tied with availability as the second most important need (68%).

Particularly noteworthy here is that respondents rated performance and things like cost as top criteria, but did not choose as top criteria the method by which those benefits are achieved (i.e. all flash configurations and server brands). The same holds true for high availability and stretched clustering abilities. Further, features such as data reduction, which can significantly lower costs, were not rated as highly as direct cost savings. Of course, often when people think of "cost" in data center solutions, they often equate that to "price." With that thinking, it's not a surprise to see cost and data reduction considered separately. For many, features like data reduction don't change the price, but they do decrease the total cost of ownership (TCO), which is not something that is always considered when purchasing a new solution.

We mentioned that a lot of people – close to one-third – indicated that server brand is not important. In recent years, commoditization at the server level has led people to understand that the brand of the underlying hardware in many cases isn't all that significant. While there may still be compelling reasons for some to adopt server hardware solutions that may bring ancillary benefits, for many, they don't care about the brand as long as the hardware can adequately do its job.

The figures on the following pages provide you with a look at how our overall respondent pool feels about specific decision criteria. We've also broken out the responses by company size and discovered:

- In large companies, availability beats performance when it comes to decision criteria around SDS and HCI. In other words, they'd rather have something that just works.
- Cost is not the primary driver for any population (performance and availability beat it), but it is in the top 3 criteria.
- In larger companies, items such as REST APIs and stretched cluster support become more critical. While still rated at the bottom of the decision criteria list, there are far fewer people in large companies saying that these items are not important.

Key decision criteria around adoption of SDS or HCI (All respondents) (N=1256)




FIGURE 33: SDS AND HCI ADOPTION DECISION CRITERIA – ALL RESPONDENTS

Key decision criteria around adoption of SDS or HCI (small companies) (N=486)




FIGURE 34: SDS AND HCI ADOPTION DECISION CRITERIA – SMALL COMPANIES

Key decision criteria around adoption of SDS or HCI (medium companies) (N=373)




FIGURE 35: SDS AND HCI ADOPTION DECISION CRITERIA – MEDIUM COMPANIES

Key decision criteria around adoption of SDS or HCI (large companies) (N=373)





FIGURE 36: SDS AND HCI ADOPTION DECISION CRITERIA – LARGE COMPANIES


Appendix: Demographics

No analysis is done on the raw demographic data but it is included in an appendix to provide a frame of reference.


How large is your IT staff?




What is your role in the IT decision-making process?



Which phrase best describes your title?



What is your industry vertical?



About

About Atlantis Computing

Atlantis, winner of the Best of VMworld and Best of Citrix Synergy awards, offers the industry's most flexible and powerful Software-Defined Storage (SDS) platform. Atlantis delivers the performance of an all-flash array at half the cost of traditional storage. Atlantis HyperScale leverages the Atlantis patented SDS platform to deliver all-flash hyper-converged appliances that are 50 to 90 per cent lower cost than traditional storage or other hyper-converged appliances. To date, Atlantis has deployed over 52 Petabytes of storage for more than 1,000 mission critical deployments, including some of the largest virtualization deployments in the world. Atlantis is privately held and funded by Adams Street Partners, Cisco Systems, El Dorado Ventures and Partech Ventures, with headquarters in Mountain View, California and offices in Europe.

Atlantis Computing is a trademark of Atlantis Computing. All other trademarks are the property of their respective owners.

About ActualTech Media

ActualTech Media provides enterprise IT decision makers with the information they need to make informed, strategic decisions as they modernize and optimize their IT operations. Leading 3rd party IT industry influencers Scott D. Lowe, David M. Davis, and James Green and special technical partners cover hot topics from the software-defined data center to hyperconvergence and virtualization. Cutting through the hype, noise and claims around new data center technologies isn't easy, but ActualTech Media helps find the signal in the noise. Analysis, authorship and events produced by ActualTech Media provide an essential piece of the technology evaluation puzzle.

More information available at

<http://www.actualtechmedia.com/>