

# 2017 State of Storage in Virtualization



ActualTech Media

By Scott D. Lowe, vExpert  
ActualTech Media

Commissioned by Tegile



## Table of Contents

<b>Executive Summary.....</b>	<b>4</b>
Key Findings .....	5
<b>Setting the Stage.....</b>	<b>6</b>
Application Mix .....	6
Virtualization Penetration.....	7
Hypervisors in Use.....	8
Storage Mix: Flash and Disk .....	10
Flash Deployment Justification.....	12
<b>Key Storage Characteristics .....</b>	<b>13</b>
Storage Performance .....	13
Virtualization's Impact on Performance .....	14
Storage Capacity and Data Reduction .....	15
Flash Storage Acquisition Frequency .....	18
<b>Other Storage Characteristics .....</b>	<b>19</b>
Storage Protocol Choice.....	19
Multiple Protocol Support .....	20
Company Size and Protocol Choice .....	21
Storage Capabilities .....	22
Storage Variety .....	23
<b>VMware VVols.....</b>	<b>24</b>
<b>Cloud-based Storage .....</b>	<b>27</b>

ActualTech Media © 2017. All rights reserved.

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

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 Figures

Figure 1: A breakdown of applications running in the data center .....	6
Figure 2: Respondent virtualization penetration .....	7
Figure 3: Hypervisors in use in respondent organizations .....	8
Figure 4: Number of hypervisors in use in respondent organizations .....	9
Figure 5: General storage environment breakdown .....	10
Figure 6: General storage environment breakdown .....	11
Figure 7: Flash deployment drivers .....	12
Figure 8: Understanding the state of storage performance .....	13
Figure 9: Storage challenges analyzed by virtualization penetration .....	14
Figure 10: Reported total organizational storage capacity .....	15
Figure 11: Insight into perceived capacity challenges .....	16
Figure 12: Data reduction ratios experienced by respondents.....	17
Figure 13: Flash capacity expansion frequency.....	18
Figure 14: Storage protocols in use .....	19
Figure 15: Number of protocols in use .....	20
Figure 16: Protocol choice as a function of company size .....	21
Figure 17: Number of protocols in use by company size.....	21
Figure 18: Capabilities inherent in existing storage systems .....	22
Figure 19: Number of storage vendors in place .....	23
Figure 20: Respondent feedback regarding knowledge of VVols.....	24
Figure 21: VVols improvements to storage environments .....	25
Figure 22: VVols improvements to storage environments .....	26
Figure 23: Use of cloud storage .....	27
Figure 24: Cloud implementation drivers .....	28



## Executive Summary

As it has been for decades, data storage is one of the most vital data center resources supported by IT. Data remains the lifeblood of the business, with increasing criticality as more and more business relies on data in order to be able to operate. From automated business processes to being able to access customer records to any number of other needs, business data is more important than ever.

At the same time that the importance and amount of data continues to increase, organizations are deploying ever more virtual infrastructure to replace physical systems. Virtualization is, by far, the standard method for new workload deployment and even the most mission critical and data-intensive applications are deployed on modern virtualization platforms.

For the second year, Tegile has partnered with ActualTech Media to research how storage and virtualization trends track with one another and how respondents view technologies such as VMware's VVols. With responses from more than 700 IT pros and IT leaders, you will discover the applications and storage characteristics that are important today and how respondents' thoughts have shifted over the past year.

### About The Author



Scott D. Lowe is the CEO of ActualTech Media, a research, content, and demand gen firm focused on the technology industry. Scott has been in the IT field for over twenty years and spent ten of those years in filling the CIO role for various organizations. Along with the team at ActualTech Media Scott has written thousands of articles and blog postings over the years and regularly creates compelling new content for clients.



## Key Findings

- Virtualized instances of SQL Server continue to increase in popularity. In 2016, 66% of respondents cited running virtualized SQL Server, a combination that has increased to 71% in 2017.
- Virtualization continues to make inroads. In 2016, just 19% of companies were 91% to 100% virtualized, a number which has skyrocketed to 27% for 2017, with 5% of respondents saying that they are running data center environments that are fully virtualized.
- Hybrid storage (combination of flash and spinning disk) gains even more traction. Hybrid storage systems are now found in 55% of respondent environments compared to just 47% a year ago. In other words, more than one-half of respondents now have some kind of flash storage in their data center, whether it's included in a hybrid system or it's an all-flash scenario.
- Occasional storage performance issues persist. 51% of respondents indicate that they experience storage performance challenges from time to time.
- For companies that have deployed flash storage (either as a part of a hybrid storage system or as a part of an all-flash array), company size plays a major role in how often they expect to expand their capacity. Just 6% of small companies—as opposed to 16% of large ones—add capacity quarterly. 20% of small companies have not yet added any additional flash capacity while only 9% of large companies say the same.
- iSCSI now exceeds Fibre Channel in deployment popularity. 50% of 2017 respondents indicated that iSCSI is one of their protocols of choice, which is a significant increase over 2016's 42% figure. Moreover, iSCSI now beats Fibre Channel, which was cited as being used by 49% of respondents.
- VMware VVols continues to be panned by customers. In 2016, 55% of survey respondents felt that they had little to no knowledge of VVols. This year, that number is 66%, a 20% increase. Moreover, those that feel that they're experts in the technology dropped from 5% in 2016 to a scant 1% in 2017.

## Setting the Stage

Before we begin to analyze the full results of the survey, it's important to understand some baseline facts, all described in the following sections.

### Application Mix

Business application needs are the primary driver behind the plans and strategies for most IT departments, inclusive of their virtualization and storage activities. In fact, in reviewing the kinds of applications that are supported by IT, overall business priorities should become clear since there should be alignment between application support and business needs.

In the 2016 version of this report, Microsoft SQL Server was the most highly deployed workload being operated in respondent's virtual environments, with 66% of respondents citing this combination. 2017 has brought even more virtualized SQL Server instances to bear. As you can see in Figure 1 below, 71% of this year's respondents are operating virtualized instances of SQL Server.

What's even more interesting is the fact that the 2016 and 2017 data almost perfectly align. While there is slight variation in the magnitude, the order of commonality of virtualized applications remains unchanged over last year.

It's interesting to note that, even in the age of Dropbox and cloud-based storage services, more companies this year have virtualized their file sharing services and retained them in-house. This year, 57% of respondents say that they are running such services as opposed to 53% in 2016. ERP systems tell a similar story. In 2017, 23% of respondents are running virtualized ERP services as opposed to just 19% a year ago.

Other applications and workloads, including Exchange, SharePoint, VDI, and Oracle databases, enjoy almost identical levels of virtualized deployment as they did in 2016.

### Which of the following applications or types of applications are you running virtualized? (multiple selections allowed)




Figure 1: A breakdown of applications running in the data center

The workload mix described by survey respondents is an interesting one due to the variety of I/O patterns inherent in the aggregated list. People are running workloads that require very little peak or sustained I/O – such as file services – while also supporting workloads that can have extremely high peak and sustained I/O needs, such as VDI and database workloads.

The challenge for these organizations is in choosing and deploying storage that is affordable, has sufficient capacity, and that can meet the wide variety of I/O and throughput needs inherent in the workload mix. Just a few short years ago, supporting such workloads would have been incredibly difficult, if not impossible. Today, though, hybrid and all-flash storage systems have the capability to support a plethora of I/O needs while also providing sufficient affordable capacity to meet growing data storage needs.

### Virtualization Penetration

Although virtualization is hardly a new technology, it's still interesting to see how much of the data center is running virtually as opposed to physically. Each year, organizations continue their efforts to virtualize more workloads. As you probably know, virtualization enables a number of additional workload management, efficiency, and disaster recovery opportunities, so it's no surprise to see organizations continue to improve their capabilities in this area.

In the past year, companies have made surprising strides. In 2016, just 19% of companies were 91% to 100% virtualized. This year, that number has skyrocketed to 27%, an increase of 70%. Although not broken out in the chart, 5% of respondents are running data center environments that are fully virtualized.




Figure 2: Respondent virtualization penetration



## Hypervisors in Use

It's clear that VMware vSphere is—and will remain—the dominant hypervisor for the foreseeable future, particularly as the company continues to focus on its efforts to expand beyond the core hypervisor into storage and networking services. These expansion efforts effectively lock customers into the vSphere hypervisor, helping to protect its market share.

As a virtualization-focused report, completeness would not be attained without a look at the hypervisors in use in an organization. As storage and the hypervisor become more closely linked to one another, there will be capability differences between different systems based on hypervisor enhancements. In fact, since last year, vSphere has extended its lead in this area, jumping from 69% penetration in 2016 to 73% this year.

Likewise, Microsoft's Hyper-V hypervisor continues to make gains in the enterprise as well, moving from 34% penetration in 2016 to 37% in 2017. We fully expect to see Hyper-V's market penetration continue to increase due to the relatively recent release of Windows Server 2016. The hypervisor in this latest version of Windows Server includes a number of enhancements, making it even more compelling as an alternative to vSphere.

Oracle, on the other hand, lost ground in 2017, dropping from 16% in 2016 to 13% in 2017. A similar scenario faced Xen, which dropped from 12% penetration to 10%. KVM has stayed pretty consistent, enjoying 9% market share in 2017.




Figure 3: Hypervisors in use in respondent organizations

For this survey, the majority of respondents are running just a single hypervisor. Even in a market in which alternative hypervisors have become commonplace, the complexity introduced with a second hypervisor is largely avoided. Of course, not everyone adheres to this. 26% of respondents are running two hypervisors while 8% are running three or more in production.

**Number of hypervisors  
in use in respondent  
organizations  
(N=714)**




Figure 4: Number of hypervisors in use in respondent organizations

## Storage Mix: Flash and Disk

For all of the benefits of flash and even as prices have dropped, for the past year, hybrid storage systems—those that combine flash and spinning disk—remain the hottest sellers in the storage market. Hybrid storage systems are now found in 55% of respondent environments compared to just 47% a year ago. In other words, more than one-half of respondents now have some kind of flash storage in their data center, whether it's included in a hybrid system or it's an all-flash scenario.

The growth in hybrid storage comes at the expense of all-disk systems, which are falling out of favor, dropping from 41% of respondent environments in 2016 to 37% for 2017. In addition, this year interestingly sees a bit of a drop in those without on-premises storage environments. In 2016, 9% said that they do not have any local storage. In 2017, that figure is 6%.

All-flash penetration remains basically unchanged at 2%. However, as flash storage prices continue their downward trend and as disk-based storage systems approach the end of their usable life and are replaced, we fully expect to see this figure rise sharply.

### Which statement best describes your current on-premises storage environment?




Figure 5: General storage environment breakdown



The information provided in Figure 5 is just a snapshot of the overall aggregate storage environment. One of the goals of this survey is to better understand exactly how organizations are leveraging flash, whether that flash exists as a part of an all-flash array or it's part of a hybrid system.

To that end, respondents were asked to provide more detail about their flash deployment, the results of which are shown in Figure 6. If you compare Figure 5 and Figure 6, you may believe there to be a mismatch. In Figure 5, you learned that only 2% of respondents are running all-flash data centers, and in Figure 6, you can see that 31% of 2017 respondents have an all-flash array.

The reality is this: In Figure 5, respondents were not allowed to make multiple selections. Only if *every single storage device* in the data center was flash would that respondent be considered an all-flash data center. In Figure 4, respondents were allowed to make multiple choices. So, even those respondents operating an all-flash storage array may also have a hybrid array, meaning that, as an organization, they are hybrid. Speaking of hybrid, a full 65% of those who use flash are running hybrid systems, making it the clear leader when it comes to how flash is used in the data center. Hybrid has been shown to provide an excellent balance when it comes to managing storage costs in terms of both capacity and performance.

However, in term of raw growth, the deployment of all-flash arrays is accelerating. In 2016, just 23% of respondents had an all-flash array. For 2017, almost one-third (31%) of respondents are operating all-flash systems. This growth appears to be coming at the expense of PCIe-based flash and DIMM-based flash modules. The reason is multi-fold. First, in many ways, PCIe-based flash was a stopgap that appears to be falling out of favor somewhat. Second, with the rise of NVMe-based flash storage, there isn't as much of a need for PCIe- or DIMM-based flash systems since NVMe leverage directly the PCIe bus.

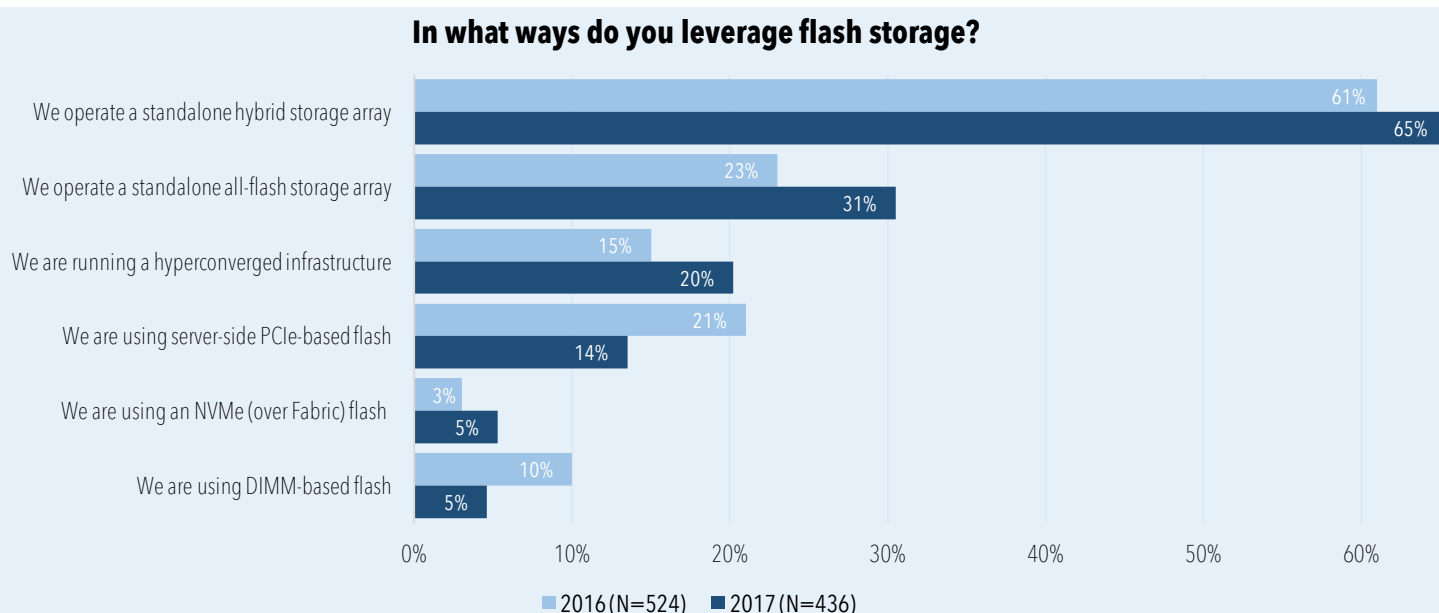


Figure 6: General storage environment breakdown

Figure 6 also shows that 20% of 2017 respondents indicate that they are running some level of hyperconverged infrastructure in their data center as well. This is up from 15% in 2016 and demonstrates increasing deployment of this technology. Most hyperconverged infrastructure systems take a hybrid approach to storage, but an increasing number of these systems are now shipping as all-flash units as well.

### Flash Deployment Justification

As the cost of flash continues to decrease, justifying the expenditure is becoming easier, but as is the case with everything in IT, there still needs to be clear business justification for investments. Of course, when it comes to storage, there is often an existing depreciation and replacement schedule that helps to drive new investment.

That said, the decision to adopt flash storage generally revolves around a few key drivers that really don't change that much year over year (Figure 7). As you can see, the need to continue to feed the application performance beast reigns supreme, with 72% of respondents citing improving application response times as a key driver for deploying flash or hybrid storage in the data center.

The most striking year-over-year change is demonstrated in the business growth and expansion driver. For 2016, 38% of respondents cited business growth as a flash driver whereas, in 2017, 46% of respondents indicate the same. This could be a simple function of replacement cycles or it could signal a new surge in IT spending if economic factors have generally improved.

### What drivers led you to deploy flash in your storage environment?




Figure 7: Flash deployment drivers

## Key Storage Characteristics

Storage capacity (MB, GB, TB, PB) and performance (IOPS, latency, throughput) are the common metrics by which storage is measured to ensure that workloads are able to operate. There are a number of different factors that impact these metrics, but in order to start with a common understanding of the challenges that respondents face in their storage environments, we asked some general questions about these two high level metrics.

### Storage Performance

Although flash storage media has effectively solved storage performance issues for many organizations, the fact is that many companies still don't have a flash strategy. Moreover, as is the case with all technology, flash will eventually be driven to obsolescence by a newer technology as newer workloads demand increasing performance.

For now, though, we sought to understand respondents' thinking around the current state of storage performance in their organizations. As shown in Figure 8, you can see that 49% of respondents never have a storage performance issue. On the other hand, this means that 51% of respondents do have storage performance challenges, with 3% of all respondents indicating that they face consistent performance challenges related to storage.

**Which statement best describes the state of storage performance in your environment?**  
(N=761)

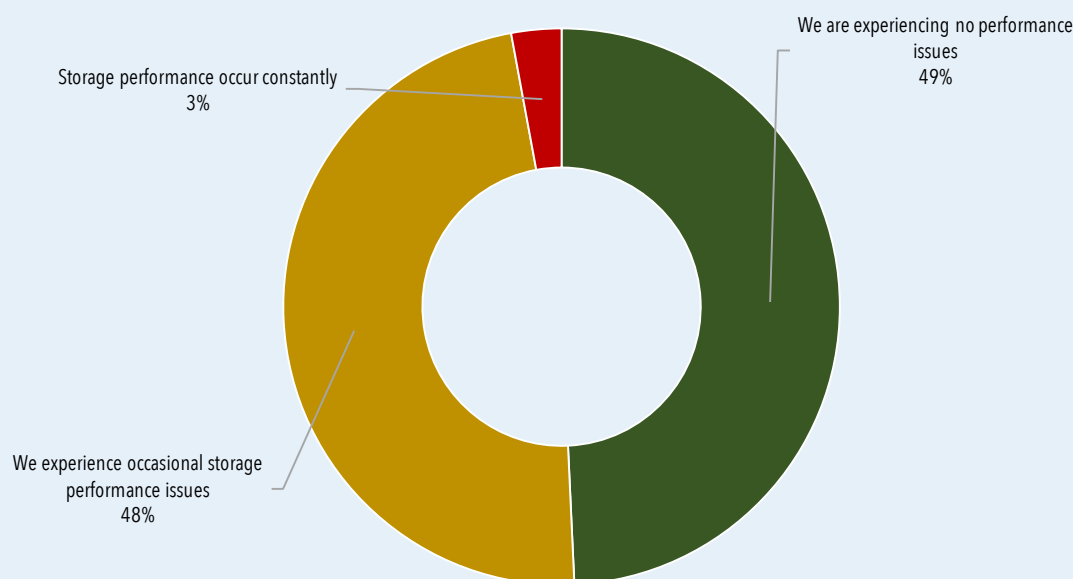


Figure 8: Understanding the state of storage performance



## Virtualization's Impact on Performance

Although crippling storage performance problems are relatively rare, the fact is that a great number of respondents do face occasional storage performance issues that can be impactful on the business. Virtualization often contributes to this problem by virtue of the fact that virtual environments mix all kinds of workloads together and expect the storage environment to deal with the mess. This so-called "I/O blender" is pervasive, but modern storage systems do a far better job at handling them than legacy systems do, but challenges remain.

As we analyzed the survey data, it became clear that there are two stories. First, those that are doing very little virtualization have far fewer storage performance problems than those that are doing more. However, those that are *fully virtualized* actually have far fewer storage performance challenges than those that are not as highly virtual.

Figure 9 demonstrates this finding. As you can see, those with no virtualization and those that are 100% virtualized have far fewer storage performance challenges than any other group. These trends seem to indicate that those who are the most mature with regard to virtualization have also implemented storage systems capable of handling the loads placed on them, which makes sense. On the low end of the spectrum, we suspect that the respondents have so little virtual infrastructure that the impact on storage simply isn't apparent.




Figure 9: Storage challenges analyzed by virtualization penetration

## Storage Capacity and Data Reduction

There was a day when raw storage capacity was all that people considered when it came to data. Today, it's just one part of the story. You've already seen respondents' experiences with storage performance, so let's take a look at capacity. Before we do, though, understand that "capacity" can mean very different things to different people. For those that have very simple storage systems, capacity is basically the usable storage space on the array. However, once you get into systems with data reduction capabilities—data deduplication and compression—capacity analysis gets a bit more challenging.

To get some sense of comparison between respondent groups, we asked respondents to tell us approximately how much storage capacity they have across all of their data center environments. Figure 10 demonstrates that respondent capacity figures are very evenly distributed from the very small to the very large.

**How much total storage capacity  
do you have across all  
of your data centers?  
(N=753)**




Figure 10: Reported total organizational storage capacity

For 2017, 49% of respondents indicate that they have all the capacity that they need for today and for the foreseeable future. While that's a good number, it does mean that 51% of respondents have storage capacity challenges. Forty-seven percent (47%) say that their capacity challenges are only occasional, while 4% say that they are regular and severe.

Now, let's look at capacity – 58% of respondents indicate some level of capacity issues with 42% saying that all is well on the capacity front (Figure 11). Another 6% of respondents indicate that their capacity issues are regular and severe.




Figure 11: Insight into perceived capacity challenges

As indicated above, data reduction technologies are those that can help organizations reduce the overall amount of storage space that is consumed by their data. These services can help companies save money on storage by reducing the amount of bits and bytes that actually have to be saved.

Data compression reduces disk capacity needs by making individual files smaller, much like a ZIP or file compression process on an operating system. Data deduplication works by eliminating identical blocks of storage. For example, if a storage system has 500 identical blocks, the storage array will store just one copy, thereby eliminating the need to store the other 499 copies.

Data reduction technologies have increased in significance in recent years with the rise of flash. Bit for bit, flash remains more expensive than spinning disk, although price parity is getting closer every year. However, when data reduction is implemented on the flash media, there can be significant capacity savings that can offset the increased cost of the raw storage.



For our 2017 data, we changed the methodology by which we requested that people provide information about the data reduction ratio they experience in their storage environments. The data reduction ratio is a way to easily convey how much capacity savings is being experienced. For example, a 2:1 reduction level effectively cuts storage capacity costs in half (you're storing 2 bytes for every 1 byte of physical space) and a 3:1 reduction reduces it by 66% (you're storing 3 bytes for every 1 byte of physical space). The higher the ratio, the higher the savings.

In Figure 12, 12% of respondents say they experience relatively low levels of data reduction, with only 2:1 reduction. There are a whole lot of reasons why people may see relatively low reduction:

- They may have a storage array that doesn't do a great job with data reduction
- They may be using an array that provides compression only and does not offer deduplication
- They may be running workloads that do not lend themselves to high levels of data reduction

Another 73% of respondents are experiencing data reduction levels of between 3:1 and 6:1, with another 17% achieving excellent outcomes at 7:1 and higher. Data reduction ratios beyond 5:1 are considered excellent.

#### What data reduction ratio do you typically see in your environment? (N=281)

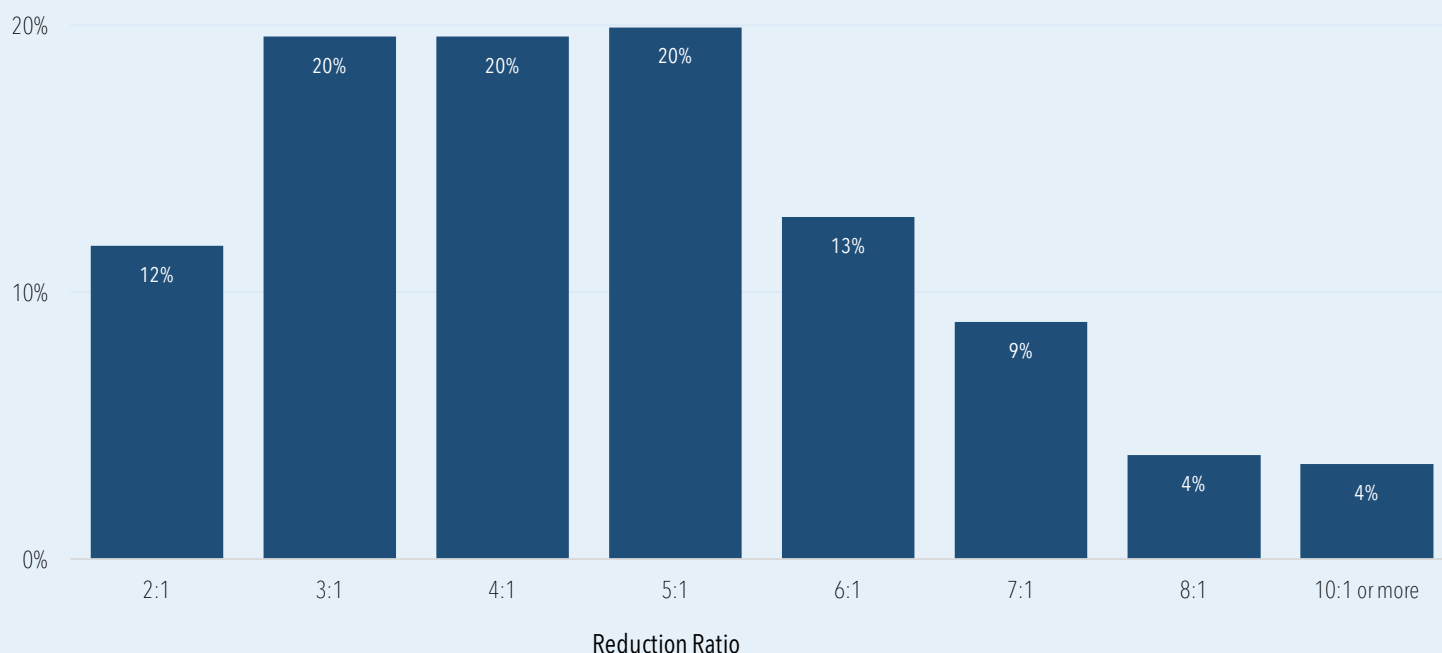


Figure 12: Data reduction ratios experienced by respondents


## Flash Storage Acquisition Frequency

Regardless of current raw storage and reduction ratios, eventually, most organizations need to add additional storage capacity to the environment in order to keep up with increasing business demands. At present more than one-half of respondents (53%) acquire more flash storage at least annually, with 10% saying that they do so on a quarterly basis (Figure 13). Fourteen percent (14%) say that they have not yet added any more flash capacity.


As expected, company size seems to play a role here, with smaller companies far less likely to perform regular capacity expansions. Just 6% of small companies—as opposed to 16% of large ones—add capacity quarterly. 20% of small companies have not yet added any additional flash capacity while only 9% of large companies say the same.

The main consistency comes at the “once a year” mark where around 26% to 30% of companies, regardless of company size, add capacity.


### At what frequency does your organization acquire new flash storage (N=435)



#### Small (0 to 499 employees)



#### Medium (500 to 4999 employees)



#### Large (5000+ employees)




Figure 13: Flash capacity expansion frequency

## Other Storage Characteristics

Besides capacity and performance, there are other interesting details of storage environments that are interesting to understand. These are presented in the following sections.

### Storage Protocol Choice

Storage protocol debates continue, but there are some clear changes in the popularity of some of the protocol choices, with significant changes over our 2016 results (Figure 14). First, iSCSI has captured the top spot, with 50% of 2017 respondents indicating use, which is a significant increase over 2016's 42% figure. Moreover, iSCSI now beats Fibre Channel, which was cited as being used by 49% of respondents. Although it's a close race, iSCSI's growth in the last year far surpassed Fibre Channel. Fibre Channel is often considered to be far more complex than iSCSI. As organizations continue their simplification efforts, we expect to see more move to less complex protocols, such as iSCSI.

Direct-attached storage has also enjoyed a rise, moving from 25% to 28% penetration since last year. We suspect that hyperconverged infrastructure deployment has had some impact on this figure. Note that, for the purposes of this survey, we consider direct-attached to be a protocol choice.

As was the case in 2016, block-based protocols in 2017 far outpace file-based ones. This year, 50% of respondents are using iSCSI and 49% have Fibre Channel deployed. NFS takes third place with 35% penetration. The various Microsoft-focused protocols (SMB/CIFS and SMB 3.0) get 28% in total.




Figure 14: Storage protocols in use

### Multiple Protocol Support

Quite often, organizations deploy multiple storage systems and storage protocols to support a wide array of workloads. As you just saw, while iSCSI and Fibre Channel are the clear frontrunners when it comes to overall protocol deployment, the fact is that 60% of respondent organizations actually run multiple storage protocols.

The chart in Figure 15 shows, in fact, that 31% of respondents are running three or more different storage protocols. What does this mean for storage buyers? To most effectively meet the needs for modern data center environments, you must purchase storage that can maximize your flexibility when it comes to protocol choice. With most organizations running storage environments that have multiple protocols, buying storage systems that can support it all is increasingly critical.

**Number of protocols  
in use in respondent  
organizations  
(N=711)**




Figure 15: Number of protocols in use

## Company Size and Protocol Choice

Organization size definitely plays a role in protocol choice. Smaller organizations tend to adopt storage protocols or services that are considered to be a bit less complex than others. As you can see in Figure 16, smaller companies have *far less* Fibre Channel uptake than larger companies and do a lot more direct-attached and iSCSI. Only 22% of small companies use Fibre Channel whereas up to 39% of larger companies do. Further, as company size increases, use of iSCSI decreases as well. Moreover, Fibre Channel over Ethernet (FCoE) does well as organization size increases as well. For those that have deployed FCoE, 49% hail from large companies of 5000 or more employees.




Figure 16: Protocol choice as a function of company size

More protocols equates to more complexity, which is probably why close to half (46%) of small companies operate with only one protocol. As company size increases, the number of protocols in use increases as well (Figure 17).




Figure 17: Number of protocols in use by company size



## Storage Capabilities

Storage isn't just storage anymore. In fact, today's modern storage systems include a multitude of features and capabilities that have become commonplace. Such features include the ability to replicate data for disaster recovery purposes, full encryption, cloud-based monitoring, and on-by-default data reduction. Figure 18 makes it clear that, while important, many of these features are still lacking—or are not in use—in deployed production storage systems.

With that in mind, Figure 18 gives you a look at the various storage features that are built into respondents' platforms. For example, 62.5% indicate that they have remote replication capability, which is a critical element of a disaster recovery plan. Plus, 59% say they have compression capability, and just under 58% report that their storage system provides data deduplication.

Only 30% of respondents indicate that their arrays provide encryption at rest. This is an increasingly important feature as enterprises strive to improve their overall security posture in response to the increasing potential for security breaches.

### Which of the following features are found in your storage environment? (N=761)

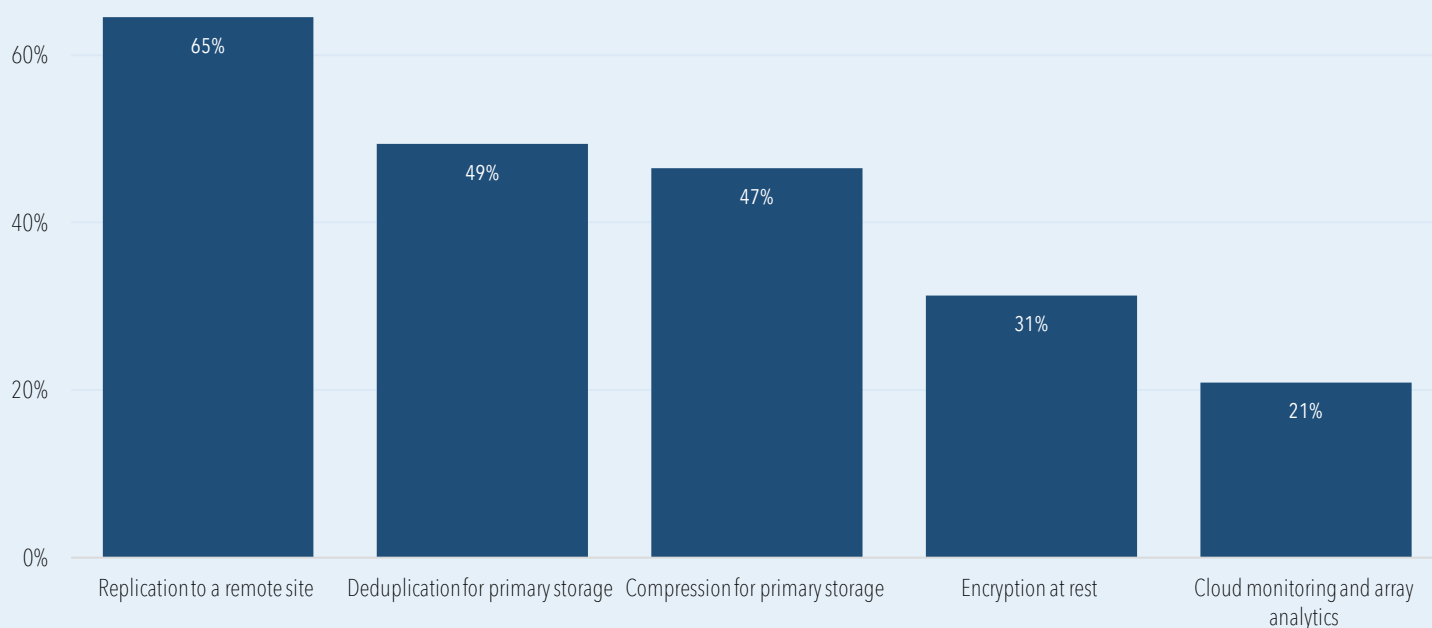


Figure 18: Capabilities inherent in existing storage systems

### Storage Variety

The more variety in the data center, the higher the administrative overhead. This is doubly true for storage, which has traditionally been among the most expensive and complex data center elements. We asked respondents to tell us how many different storage vendors they have in place (Figure 19). Close to 2/3 of respondents are supporting solutions from multiple storage vendors, which is likely to increase overall complexity.

It is likely that these organizations have not been able to deploy solutions that can meet all of their needs, hence the deployment of multi-vendor solutions. This is an opportunity for vendors, that can meet a myriad of needs to help customers simplify their storage.

**How many storage vendors  
do you currently use  
for your organization's  
primary data?  
(N=715)**




Figure 19: Number of storage vendors in place

## VMware VVols

VMware's VVols has an interesting history. VVols started life with an announcement and a lot of fanfare and predictions about how it would revolutionize storage. And, then... it spent years in development while the market continued to innovate. Upon release, storage vendors could choose to implement it in whole or in part... or not at all.

Most new workloads run virtualized and with a good many reasons. Virtualization has effectively transformed physical servers to virtual ones and enabled a great deal of mobility and new data protection opportunities. However, when it comes to policy application at the storage level, there hasn't been a great deal of broad market support. Although there are some solutions to integrating storage more deeply with virtualization, there hasn't really been anything standard.

The VVol construct allows the storage and the virtualization management layer the ability to work more closely together. VVol support in a storage array makes VMware's storage policy-based management (SPBM) possible. SPBM is a key piece of the software-defined datacenter (SDDC). It's fair to say that, without SPBM, the SDDC vision simply cannot be fulfilled.

Today, VVols doesn't have the same mystique it enjoyed when it was just on the drawing board and it appears as if it's become far less deployed than once envisioned.

VMware's VVols originally aimed to simplify storage in virtual environments. However, as is the case for any technology, in order to be successful, administrators need to understand the technology. As our year-over-year results indicate (Figure 20), VVols is not doing well. In 2016, 55% of survey respondents felt that they had little to no knowledge of VVols. This year, that number is 66%, a 20% increase. Moreover, those that feel that they're experts in the technology dropped from 5% in 2016 to a scant 1% in 2017.

### What is your knowledge or understanding of VMware's VVols?




Figure 20: Respondent feedback regarding knowledge of VVols

Storage vendors make individual decisions regarding their support for VVols. VVols is a set of APIs that must be integrated by storage vendors and, once a vendor has added this API support, VMware vSphere can more directly integrate with the storage services. Deep integration can help to improve overall policy assignment with regard to virtual machine management. Policy-based administration improvements were cited in 2017 by 42% of respondents as a positive outlook on how VVols can improve the storage environment (Figure 21).

VVols can also help to eliminate silos within the IT department by enabling virtualization administrators to provision their own storage (41%) without having to involve a storage administrator. While this may sound blasphemous to certain parts of the department, this flattening of IT is a major trend that allows companies to save money and to simplify what can be complex processes.

For 2017, the biggest change in people's views around VVols is around monitoring. VVols enables deeper visibility into virtual machines and related storage and can make it easier for organizations to identify performance problems. Please note, however, that our survey methodology for this question changed in 2017. In the prior year survey, all respondents, regardless of level of knowledge around VVols, had the opportunity to answer this question around VVols benefits. For 2017, only those respondents that indicated some level of familiarity with VVols were asked to tell us how they believe VVols can improve the storage environment.

**In what ways do you believe that VVols can help to improve your storage environment?  
(multiple selections allowed)**




Figure 21: VVols improvements to storage environments

There is good news for vendors that have chosen to eschew VVols support in their storage products. As shown in Figure 22, only 6% of respondents cited VVols as a critical feature. However, 79% did indicate that it would be nice to have, although not required while 15% said that it's not important at all. These numbers are an evolution over 2016 results and do indicate that those that have some understanding of VVols see the potential for benefit, but they don't feel that the benefit would force a decision to an array with VVols support.

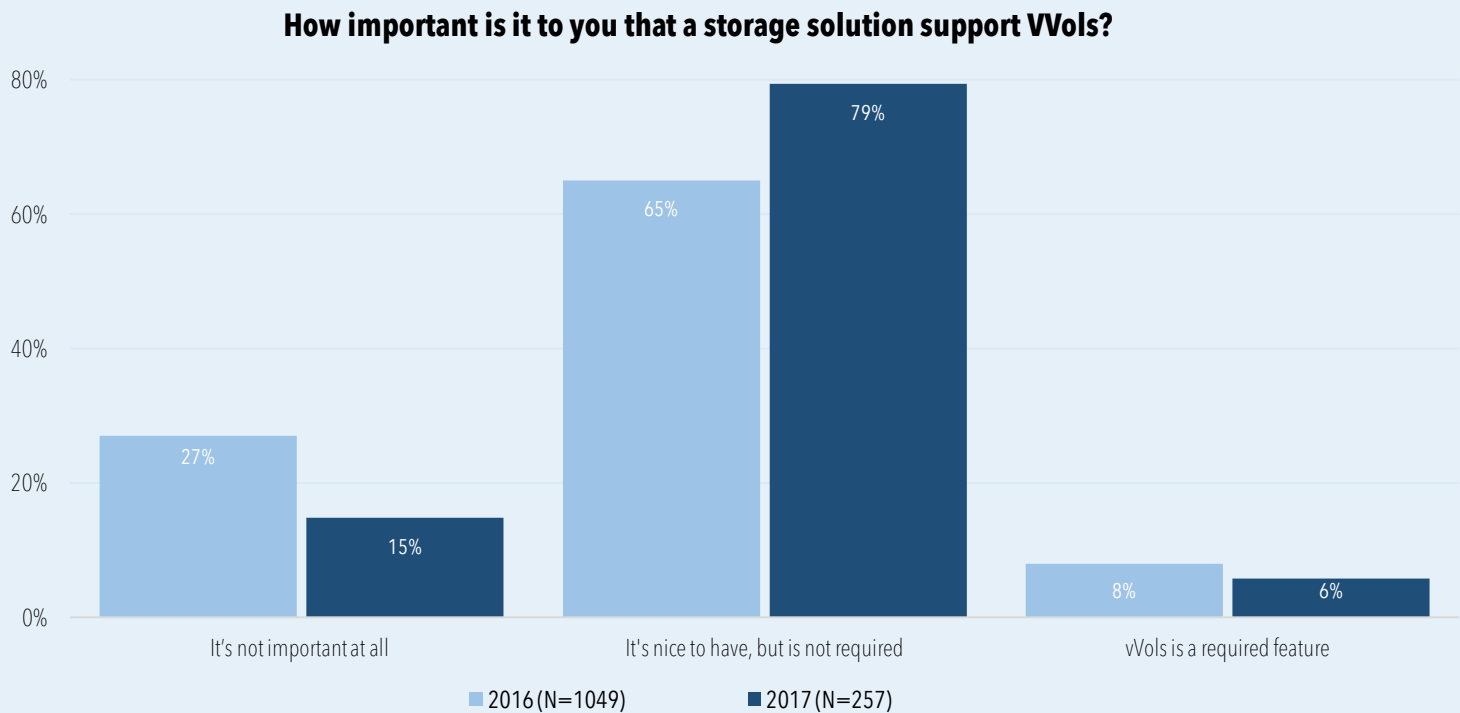


Figure 22: VVols improvements to storage environments



## Cloud-based Storage

Everything is going to the cloud! Right? Well, maybe not. Although there is definitely growth in cloud-based services, storage in the cloud may not be enjoying the same levels of growth as other services.

When asked whether or not they were leveraging cloud-based storage, just 26% of respondents said that they were, in fact, doing so with another 34% indicating that they are considering adoption of cloud-based storage in some form. A full 40% said that they were not leveraging cloud-based storage and had no plans to do so.

### Are you using cloud storage at all in your environment?




Figure 23: Use of cloud storage

Adopting cloud services is one path forward for many organizations, but what are the underlying reasons that they choose to go this route? It generally comes down to a discussion around money. Figure 24 shows that 54% of respondents cite the economic structure of cloud storage as a key driver.

Over the past year, disaster recovery as a cloud storage driver has increased pretty dramatically. In 2016, 40% of respondents indicated that they preferred to use cloud for disaster recovery rather than a secondary data center. In 2017, that number jumped to 48%.

Likewise, those looking for different ways to respond to flexing storage demand seem to be expecting those cloud to do that for them. In 2016, 34% of respondents indicated that cloud allowed them to react to changes in storage demand while, in 2017, 44% indicated the same.

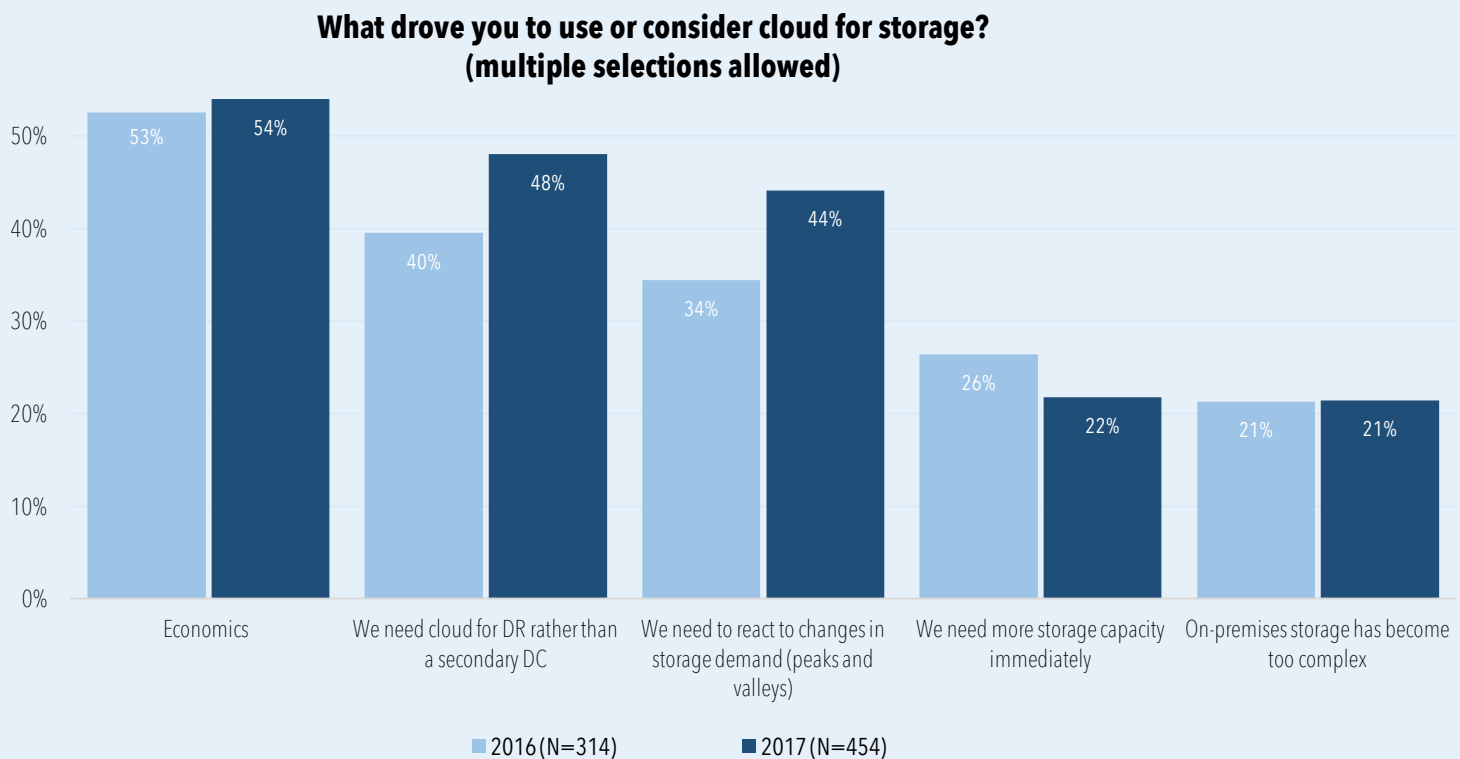


Figure 24: Cloud implementation drivers