

2016 State of Storage in Virtualization



By Scott D. Lowe, vExpert
Partner and Co-Founder, ActualTech Media

Commissioned by



Table of Contents

Table of Contents	2
List of Figures.....	3
Executive Summary	4
Key Findings.....	5
Setting the Stage.....	6
Application Mix.....	6
Virtualization Penetration.....	8
Hypervisors in Use	9
Storage Mix: Flash and Disk	10
Flash Deployment Justification.....	12
Key Storage Characteristics	13
Storage Performance.....	13
Virtualization's Impact on Performance.....	14
Storage Capacity and Data Reduction	15
Data Reduction by Storage Type.....	17
Other Storage Characteristics.....	18
Storage Protocol Choice	18
Multiple Protocol Support.....	19
Company Size and Protocol Choice	20
Storage Capabilities.....	21
Storage Variety	23
VMware VVols.....	24
Cloud-based Storage	28
Respondent Demographics.....	31



ActualTech Media © 2015. 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.....	8
Figure 3: Hypervisors in use in respondent organizations.....	9
Figure 4: General storage environment breakdown.....	10
Figure 5: The various ways which flash storage is leveraged in the data center.....	11
Figure 6: Flash deployment drivers.....	12
Figure 7: Understanding the state of storage performance.....	13
Figure 8: Storage challenges analyzed by virtualization penetration.....	14
Figure 9: Insight into perceived capacity challenges.....	15
Figure 10: Data reduction ratios experienced by respondents.....	16
Figure 11: Data reduction ratios by type of storage.....	17
Figure 12: Storage protocols in use.....	18
Figure 13: Number of protocols in use.....	19
Figure 14: Protocol choice as a function of company size.....	20
Figure 15: Capabilities inherent in existing storage systems.....	21
Figure 16: Number of storage vendors in place.....	23
Figure 17: Respondent feedback regarding knowledge of VVols.....	24
Figure 18: VVols improvements to storage environments.....	25
Figure 19: Criticality of VVols support in a new array purchase.....	26
Figure 20: VVols criticality as a function of knowledge around VVols.....	27
Figure 21: Use of cloud storage.....	28
Figure 22: Cloud implementation drivers.....	29
Figure 23: The number of people indicating storage complexity -- by company size.....	30
Figure 24: Respondent vertical.....	31
Figure 25: Respondent title category.....	32

Executive Summary

If you've wandered the show floor at practically any industry trade show in recent years, you've seen two major data center trends playing out. The first is simply an extension of a trend that began in the early 2000s: virtualization. Modern enterprises are continuing their virtualization efforts and are moving more and larger workloads to virtual platforms. Second, you're seeing a true renaissance in the world of storage. With the rise of viable flash storage, entire new markets are being created and led by this technology. In fact, it's safe to say that the second trend, a golden era in storage, is enabling the first trend: bigger and faster virtualization.

ActualTech Media and Tegile seek to understand what's happening at this intersection. To that end, we have partnered to bring you this report. In this report, you will learn about the state of storage as it pertains to virtualized workloads. We polled over 1,000 IT pros from across the world in order to learn about what organizations of all shapes and sizes are doing with regard to storage and virtualization and to gain insight into the kinds of challenges that are being faced by these organizations and how they are leveraging such services as cloud storage, VMware Virtual Volumes (VVols), and more.



About The Author

Scott Lowe is a partner in and co-founder of ActualTech Media. 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. Scott has written thousands of articles and blog postings over the years and regularly creates compelling new content for clients.

Key Findings

- Virtual machines now power even the most mission-critical systems in the enterprise. 66% of respondents indicate that they have virtualized SQL Server. 25% have virtualized Oracle and 15% are running virtual systems to support big data/analytics needs.
- Virtualization penetration loosely correlates to storage performance challenges.
- Although all-flash storage garners major headlines, hybrid storage systems – those that combine flash and spinning disk – are, by far, the more popular choice. 61% of respondents indicate they use hybrid storage systems while only 3% indicate that they run all flash. We expect this to shift over time to favor all flash, but for the present, hybrid is the undisputed standard. Storage companies that have unified offerings in both spaces are best prepared to meet both today's needs as well as the emerging needs of tomorrow.
- Data reduction outcomes increase dramatically as organizations introduce hybrid and all-flash systems. In fact, of those reporting 10:1 or better data reduction, a scant 17% are running all-disk systems. For those who get no data reduction, 67% are running all-disk systems.
- In terms of accessing storage, Fibre Channel remains the protocol of choice while iSCSI comes in second. Usage of each protocol correlates to company size with smaller companies favoring iSCSI over Fibre Channel, while the reverse is true for large companies. 55% of respondents indicate that they are using more than one storage protocol or access method with 13% using four or more. Multi-protocol support remains a key need for storage vendors to consider in their platforms.
- VVols remains relatively unknown with 55% of respondents citing little to no knowledge of the technology. Only 5% consider themselves well-versed in VVols.
- Perception of the criticality of VVols support in a storage system is directly proportional to respondent's knowledge of VVols. As respondent knowledge of VVols increases, so does their belief that VVols is a critical storage feature.
- Cloud storage is getting good traction – 30% of respondents use cloud storage in some fashion.

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

The intersection of virtualization and storage is interesting, but in order for any analysis to be useful, there must be an understanding around business needs. Business needs are often reflected in the kinds of applications that are running in the data center. In Figure 1, you can see that databases rule the data center world, with 66% of respondents indicating a need for SQL Server and 25% running Oracle. This really isn't a big surprise given that databases are generally the lifeblood of the business.

Which of the following applications are you running virtualized?

(N=938)

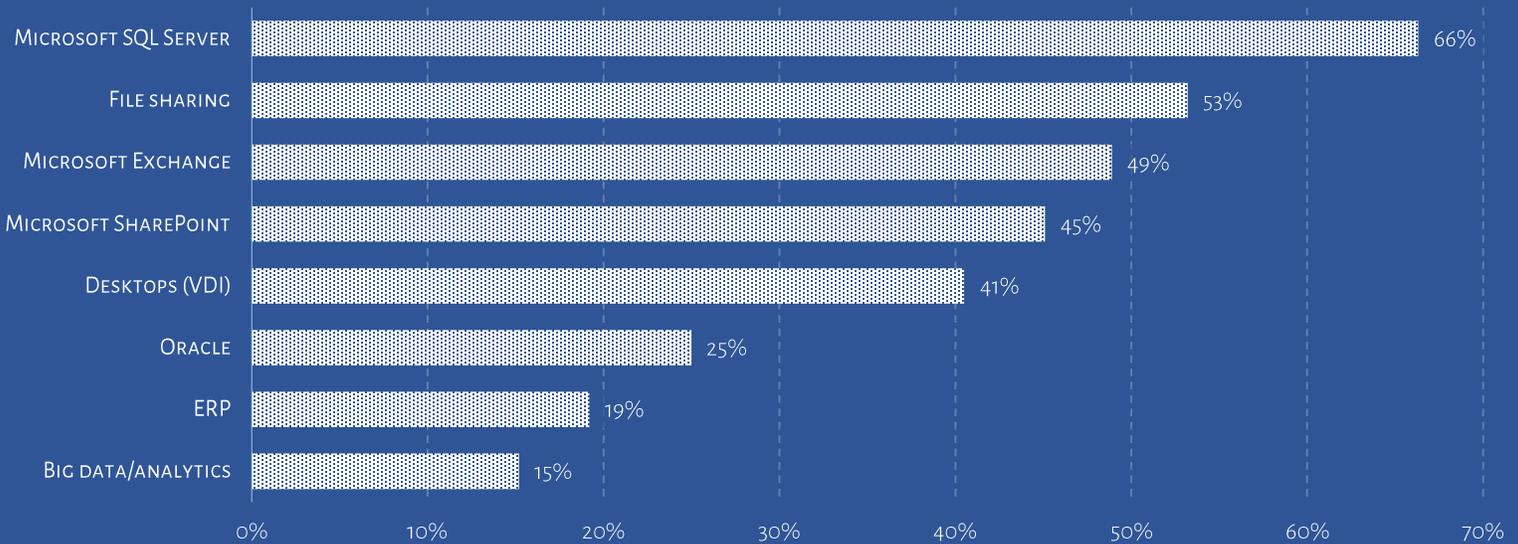


Figure 1: A breakdown of applications running in the data center

You can also see that there are other major needs in the data center, including file sharing, as well as a significant number of organizations running services including Microsoft Exchange and SharePoint. In addition, 41% of respondents indicated that they are running VDI at some level in the data center.

This workload mix is interesting primarily due to the fact that has such a variety. At first glance, you may wonder what variety there is, and in what amount, in relation to common workloads. The variety comes into play when you consider how these services interact with and leverage the virtualization and storage environments. File sharing, for example, is far from being I/O-intensive, but it can require vast amounts of capacity as users store anything and everything “on the network.” VDI, on the other hand, is one of the most I/O-intensive applications provided by IT. However, it isn’t necessarily capacity-heavy since there is ample opportunity for leveraging capabilities, such as linked clones, in order to reduce capacity needs.

Exchange, SharePoint, SQL Server, and Oracle all have very random I/O and capacity needs which vary dramatically from organization to organization.

It is this workload variety that used to stymie storage administrators. When we were in the era of the disk-based SAN, this mixing bowl of I/O was a tremendous challenge that had the potential to derail projects such as VDI. In recent years, however, flash has emerged as *the* solution for even the most vexing performance problems and has allowed data center administrators to bring virtualization to even the most I/O-intensive applications. At the same time, hypervisor improvements have enabled the creation of massive virtual machines that can be assigned sufficient resources to run big applications.

Virtualization Penetration

This report discusses storage as it pertains to virtualization, so it's important to understand where respondents stand with regard to virtualization. Figure 2 provides you with a look at this key statistic. Although this is a basic question, it's been interesting over the years to watch the bars in this chart continue to shift to the right and more and larger workloads are virtualized. Respondents for this survey are relatively heavily virtualized; more than 60% are more than one-half virtualized with over 19% getting close to being fully virtual.

How heavily virtualized is your current environment?

(N=1049)

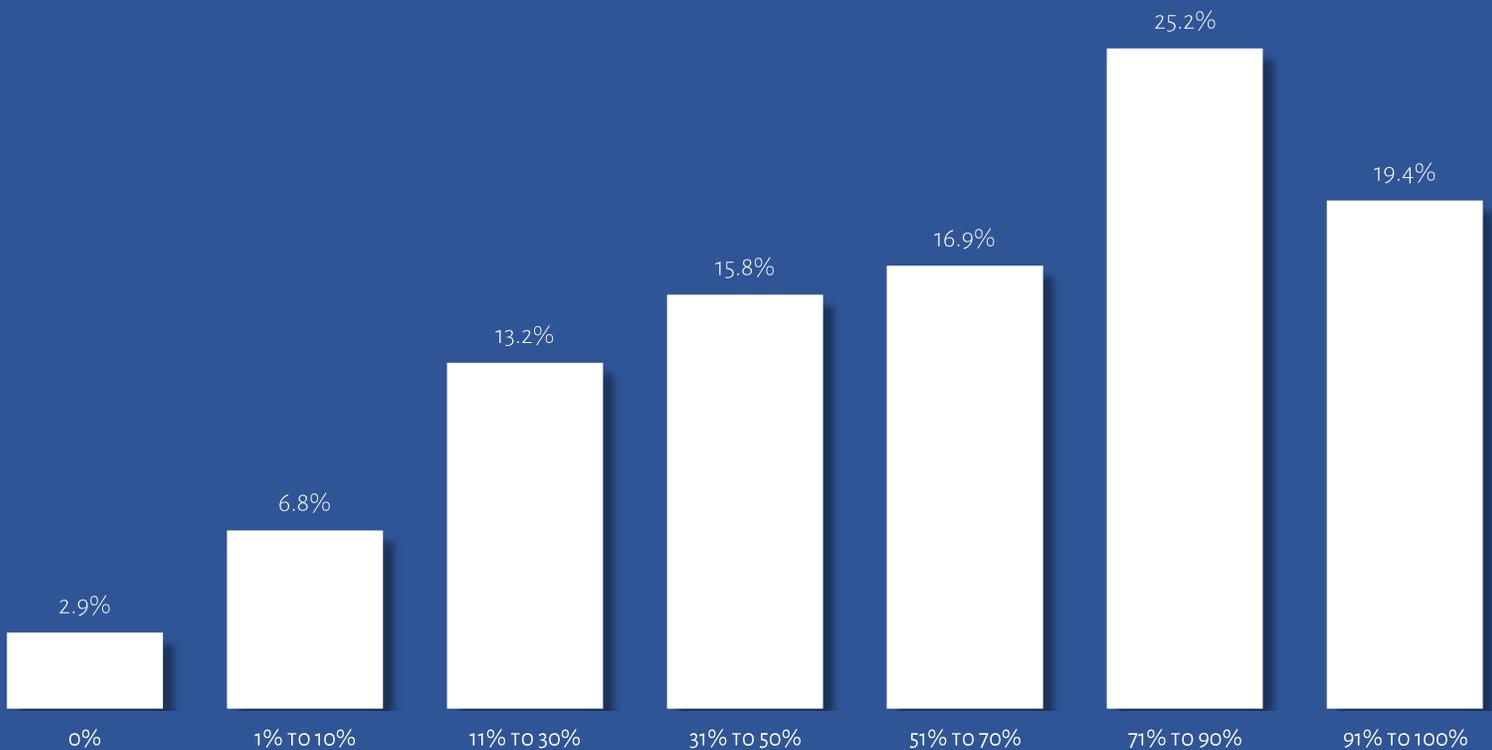


Figure 2: Respondent virtualization penetration

Hypervisors in Use

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.

As you can see in Figure 3 (left hand side), VMware vSphere remains the undisputed leader in the hypervisor wars with 69% of respondents using it in their companies. Taking second place is Hyper-V, with 34%. In the last year, Oracle's hypervisor solution has gained steam and runs in 16% of respondent companies. Xen, KVM, and miscellaneous other tools exist in a small number of environments.

At the right hand side of Figure 3, it's obvious that the majority of respondents (65%) prefer to operate with just a single hypervisor. 27% run two hypervisors while 8% run 3 or more.

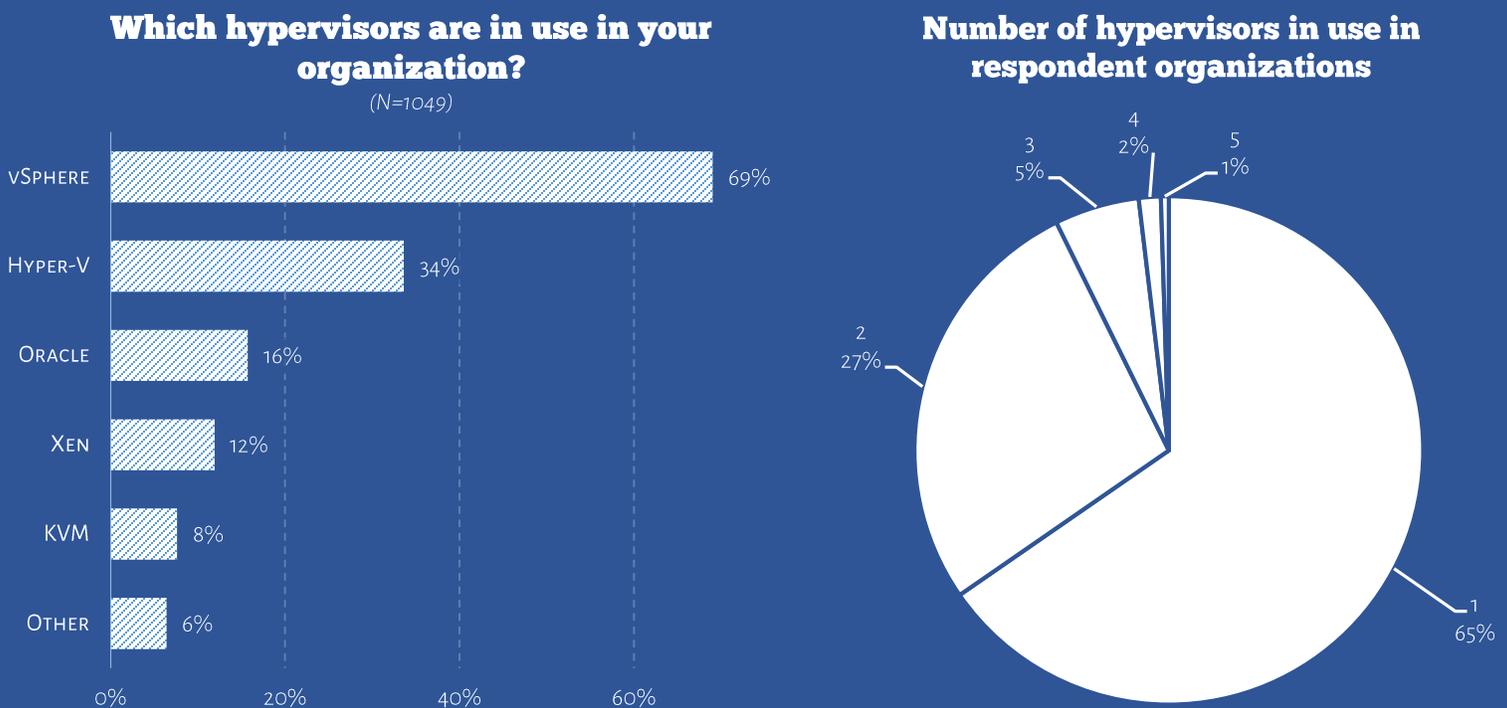


Figure 3: Hypervisors in use in respondent organizations

Storage Mix: Flash and Disk

On the storage front, many businesses have begun to add flash to their data center environments, although there remain a great many that have stayed with traditional hard disks, too. In our survey, we asked respondents to generally describe their storage infrastructure. We found that 9% of respondents do not have local storage — they may be leveraging cloud as their storage environment, while 41% of respondents continue to operate disk-only storage systems. However, a full 50% of respondents are using flash in some way. Only 3% indicated that they are all-flash, with 47% using a mix of flash and disks in a hybrid scenario.

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

(N=1049)

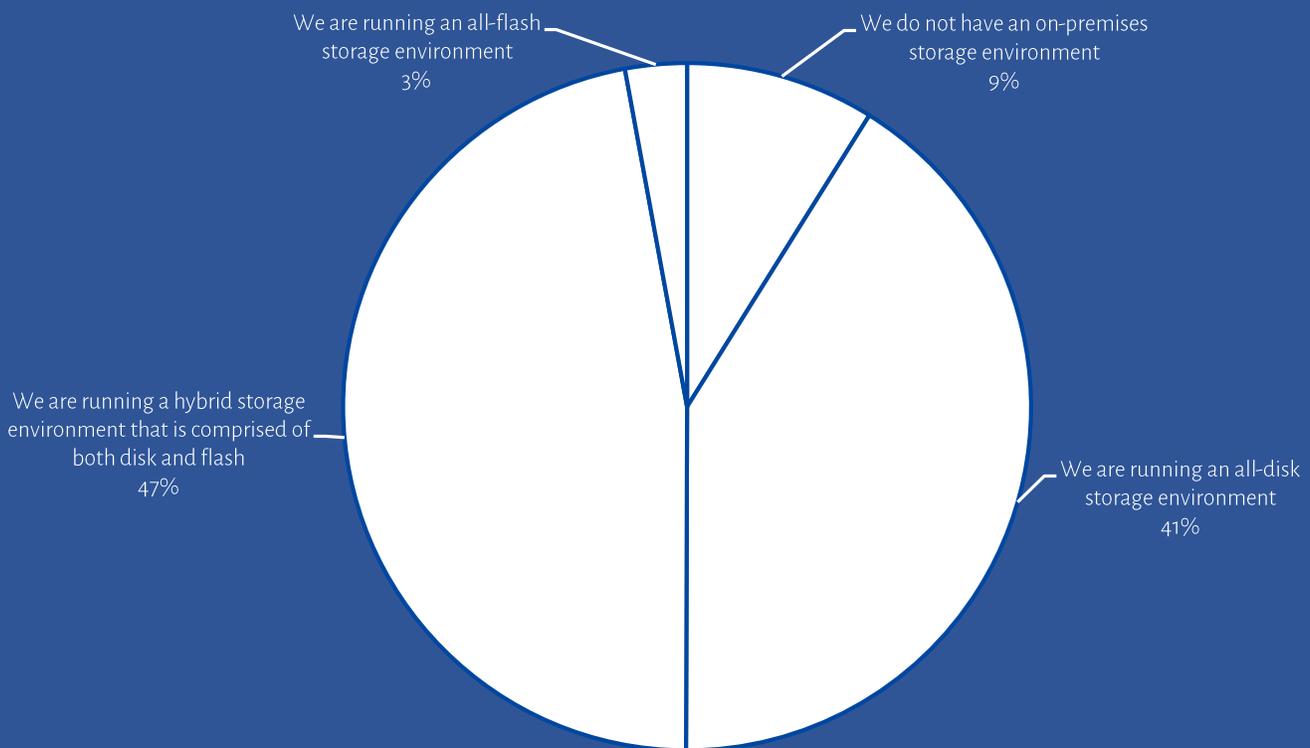


Figure 4: General storage environment breakdown

To really understand how flash storage is being used in respondent organizations, we asked for a bit more detail. For those who are using flash in some way, we asked them to tell us exactly *how* they're doing so, the results of which are shown in Figure 5. If you compare Figure 4 and Figure 5, you may believe there to be a mismatch. After all, in Figure 4, you learned that only 3% of respondents are running all-flash data centers, and in Figure 5, you can see that 23% of respondents have an all-flash array. The reality is this: In Figure 4, 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 5, 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 61% of those who use flash are running such 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.

In what ways do you leverage flash storage?

(N=524)

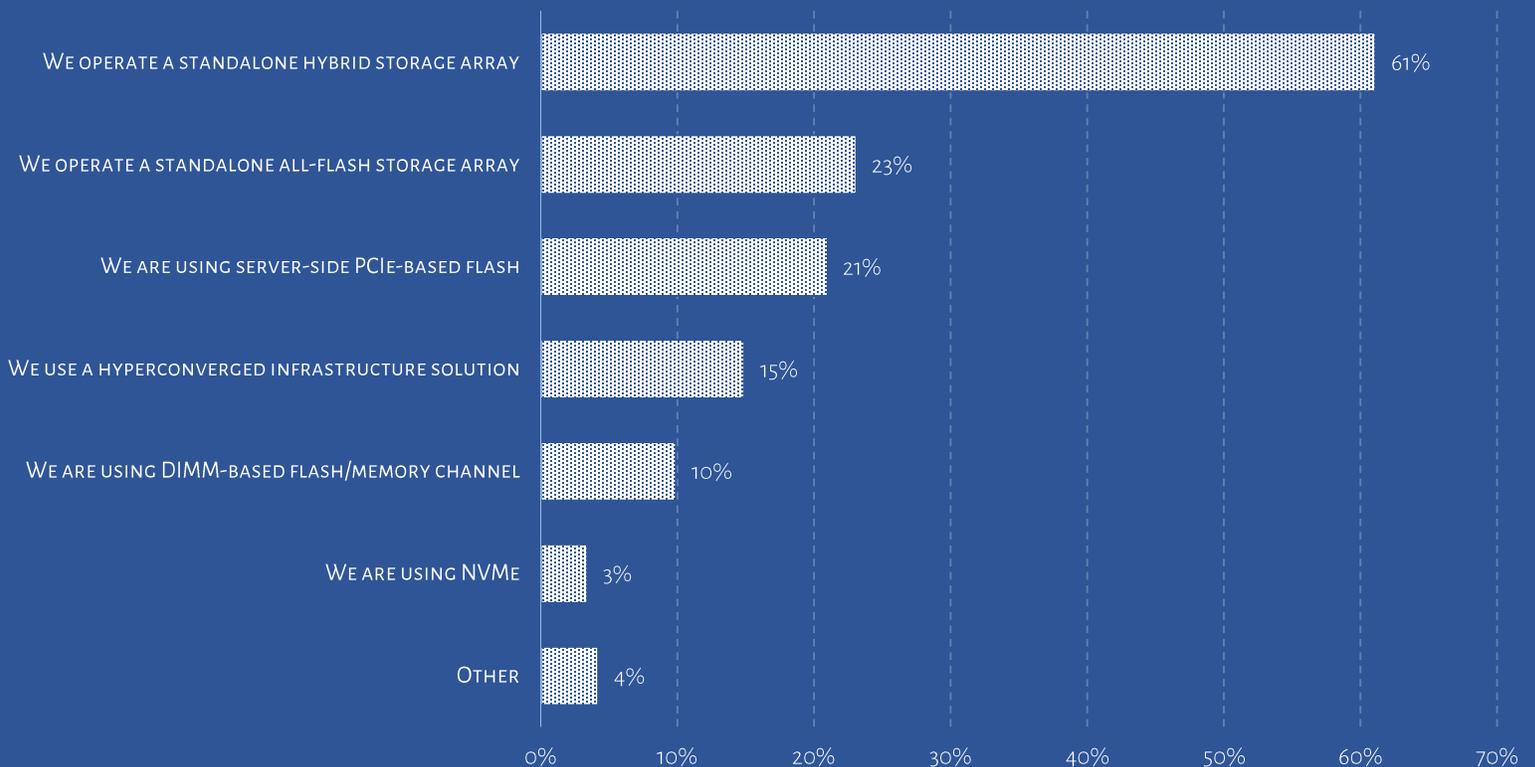


Figure 5: The various ways which flash storage is leveraged in the data center

There are a lot of ways to deploy flash beyond using array-based systems, too, and many of these options are reflected in survey responses. Just under one-quarter of flash users are using PCIe-based flash storage cards in their systems and a full 10% are using flash storage that ships in a DIMM-based form factor, and 3% are using NVMe.

Figure 5 also shows that 15% of respondents indicated that they are running some level of hyperconverged infrastructure in their data center as well. Most hyperconverged infrastructure systems take a hybrid approach to storage, but an increasing number of these systems are now shipping as all-flash units.

Flash Deployment Justification

As much as IT tries to be an agent of change, it really takes a business justification to make investments, and that goes for storage just as much as it does for anything else. In the business world, applications are king. Therefore, it comes as no surprise that almost 72% of respondents deployed flash in order to improve application response times.

Unhappy end users make for unhappy businesses. To this end, as shown in Figure 6, 48.3% of those that have deployed flash listed improving end-user satisfaction as a key driver for doing so. Happy users also make for more productive users, which can improve the bottom line. Technology investments are often made to either boost the potential for revenue or to reduce expenses. In our survey, about 38% of flash deployments were undertaken as a way to help support business growth or expansion while 25% of deployments were intended to help reduce operating expenses.

If you're wondering how flash can help on the economic front, bear in mind that flash storage, when compared to hard drives, often results in less overall power usage and can reduce cooling needs since it may not generate as much heat. Further, there was a day when people bought entire shelves of disks just to get more storage performance. When it comes to raw speed, a single flash disk can replace dozens of hard disks.

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

(N=524)

71.6%



IMPROVE EXISTING
APPLICATION RESPONSE TIMES

48.3%



IMPROVE END-USER
SATISFACTION

37.8%



BUSINESS
GROWTH/EXPANSION

25.4%



REDUCE OPERATING
EXPENSES

23.3%



DATA CENTER
CONSOLIDATION

2.7%



OTHER

Figure 6: Flash deployment drivers

Key Storage Characteristics

Performance and capacity are two critically important metrics that determine the overall suitability for a particular storage system to meet the needs of an organization. Although there are a number of factors that play into how well these metrics are achieved, we first sought to learn about how well people feel that their current storage environment is operating.

Storage Performance

Because performance has become such a critical issue in recent years (which is one of the main reasons behind the rise of flash) it's really interesting to see that a full 57% of respondents indicate that they experience either ongoing or occasional performance challenges. This is outlined in Figure 7. Meanwhile, 43% of respondents say that they have no storage performance challenges in their organizations.

Which statement best describes the state of storage performance in your environment?

(N=1049)

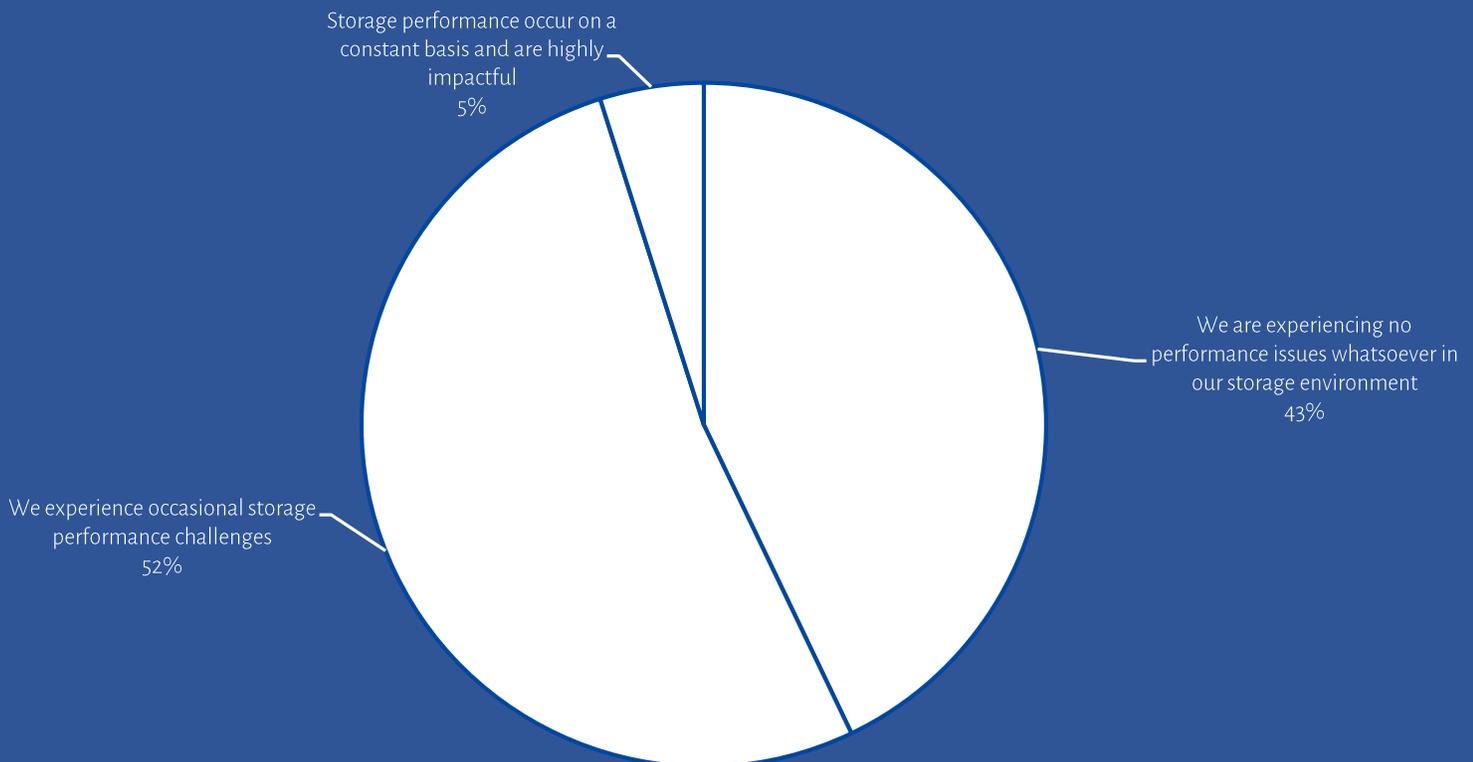


Figure 7: Understanding the state of storage performance

Virtualization's Impact on Performance

It's been said that virtualization, while it's been incredible in terms of workload management and mobility, has introduced storage performance challenges. Figure 8 seems to validate this fact, at least to a point. You can clearly see that those who are doing no virtualization have far fewer storage performance challenges with just 23% reporting issues. Bear in mind that this group represents only 3% of the respondent population, though, so it may not be as indicative of a trend as it would if there were more people included.

As soon as virtualization is introduced, 53% of respondents indicate that they experience storage performance issues with that figure increasing to 64% for those who are 31% to 50% virtualized. Beyond 50% virtualized, storage performance issues begin to slowly subside. However, those who are 91% to 100% virtualized indicate the *fewest* performance problems among all groups, except for those doing no virtualization.

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.

How heavily virtualized is your current environment? Which statement best describes the state of storage performance in your environment?

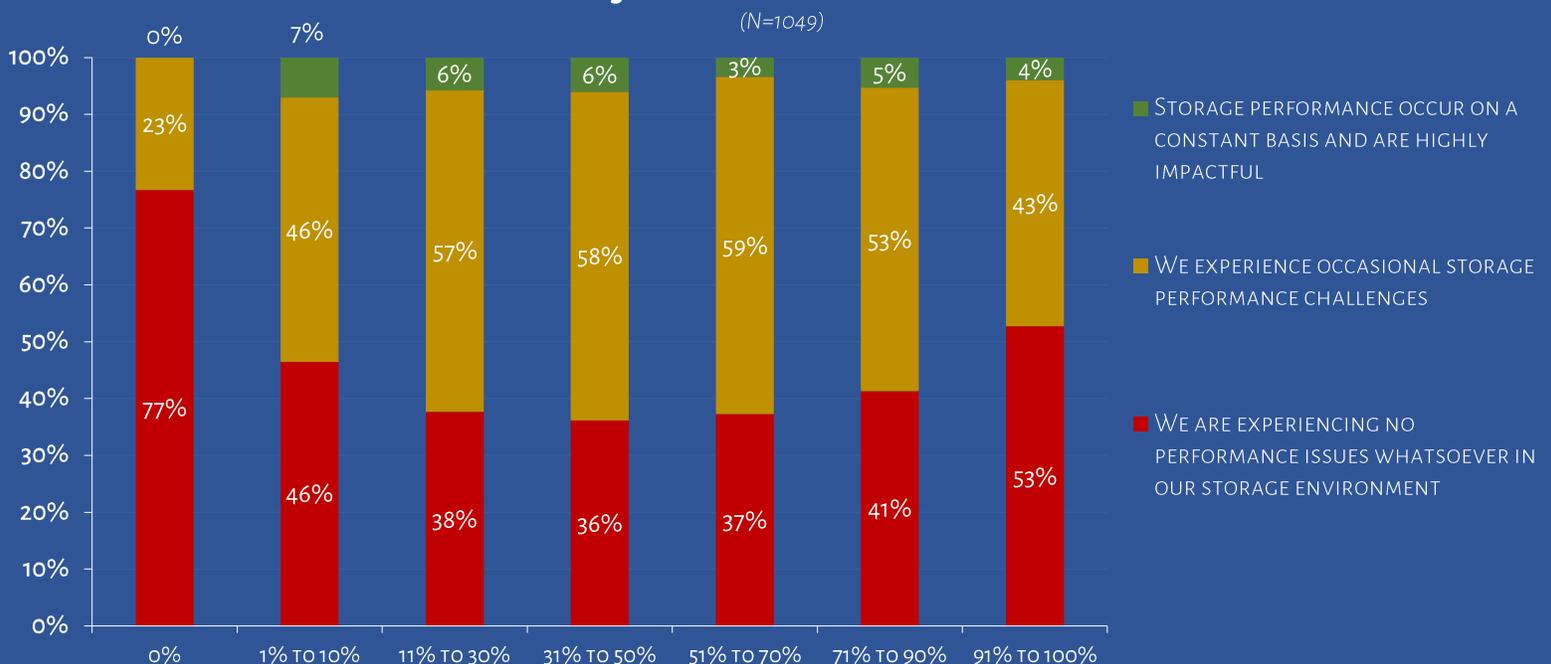


Figure 8: Storage challenges analyzed by virtualization penetration

Storage Capacity and Data Reduction

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 9). Another 6% of respondents indicate that their capacity issues are regular and severe.

Which statement best describes the state of storage capacity in your current storage environment?

(N=1049)

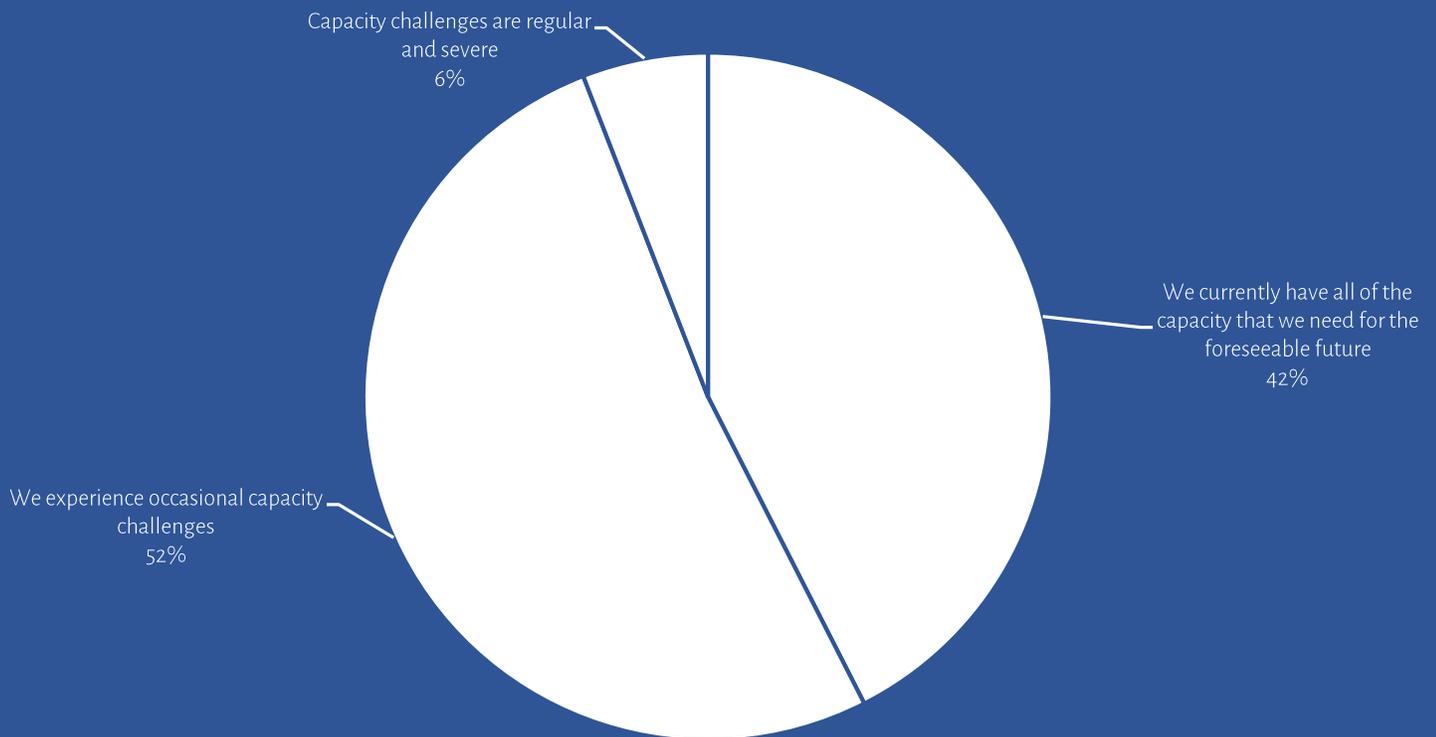


Figure 9: Insight into perceived capacity challenges

Data reduction, meaning data deduplication and compression, are two key technologies that 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. 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.

In Figure 10, 13% of respondents say they experience relatively low levels of data 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 1/3 of respondents are experiencing data reduction levels of between 2.1:1 and 6:1, with another 6% achieving very good outcomes at 7:1 and higher.

A 2:1 reduction level effectively cuts storage capacity costs in half and a 3:1 reduction reduces it by 66%.

What data reduction ratio do you typically see in your environment?

(N=832)

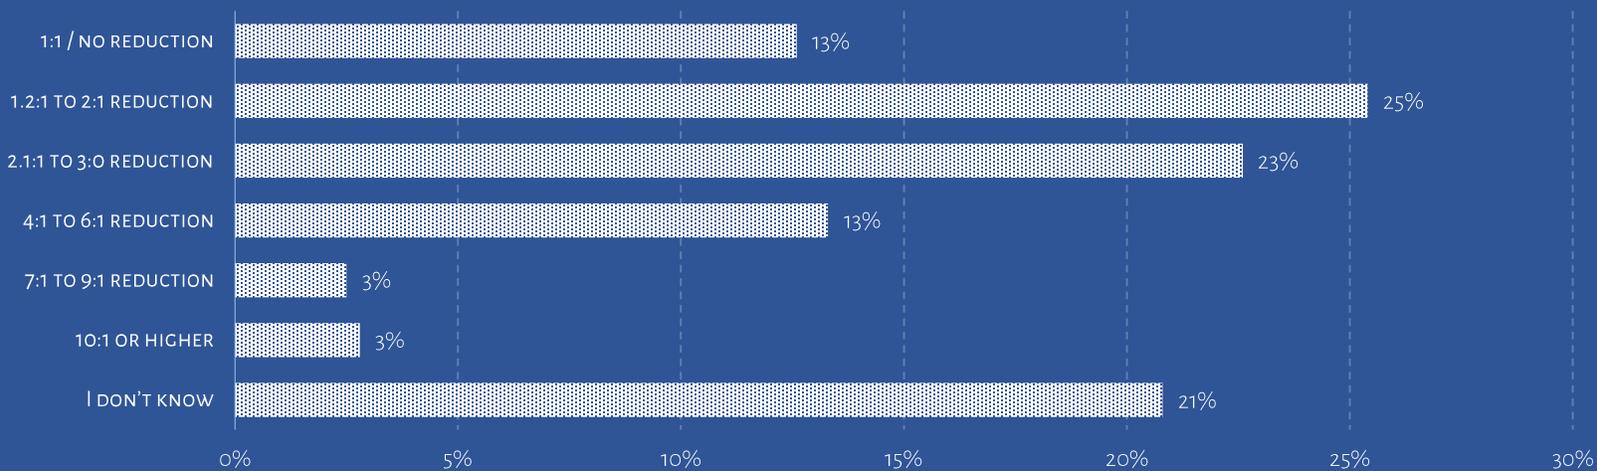


Figure 10: Data reduction ratios experienced by respondents

Data Reduction by Storage Type

If you are ever wondering whether hybrid and all flash storage can provide you with storage capacity benefits in the form of data reduction (data deduplication and compression), wonder no more! As is crystal clear in Figure 11, data reduction outcomes increase dramatically as organizations introduce hybrid and all-flash systems. In fact, of those reporting 10:1 or better data reduction, a scant 17% are running all-disk systems. For those who get no data reduction, 67% are running all-disk systems.

What data reduction ratio do you typically see in your environment? Which statement best describes your current on-premises storage environment?

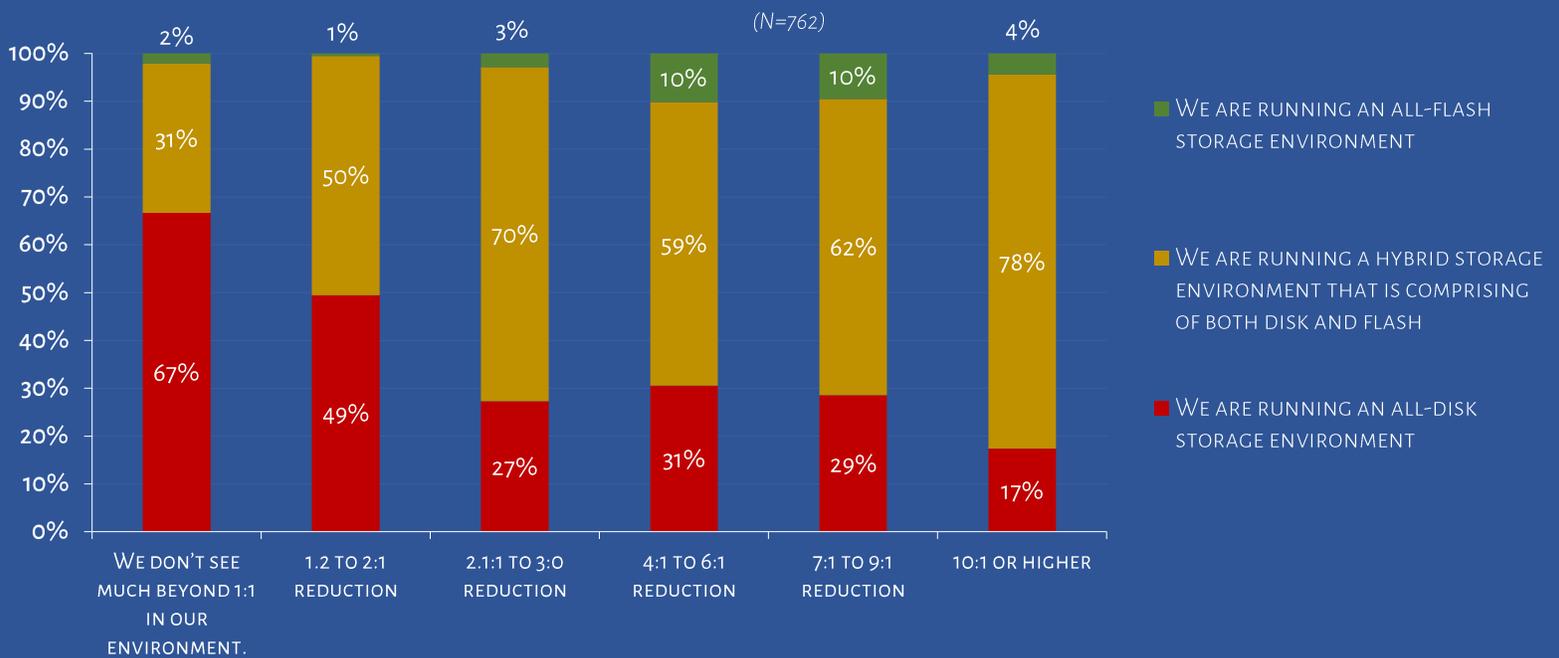


Figure 11: Data reduction ratios by type of storage

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

Back in the day — and still continuing today to some extent — people used to fight over Mac and PC. Today's storage administrators have similar disagreements. Some are in the file-based NFS corner while others won't ever give up their fibre channel environments. We asked our survey respondents to tell us about the various storage protocols that are in use in their organizations (Figure 12).

Block-based protocols absolutely win the day with 48% of respondents using fibre channel and 42% using TCP/IP-based iSCSI. NFS takes third place with 36% penetration. The various Microsoft-focused protocols (SMB/CIFS and SMB 3.0) get 27% in total.

Of course direct-attached storage isn't a protocol, but it's still a very strong contender in the data center, at 25%.

What storage protocols are you using for your virtual environment?

(N=1049)

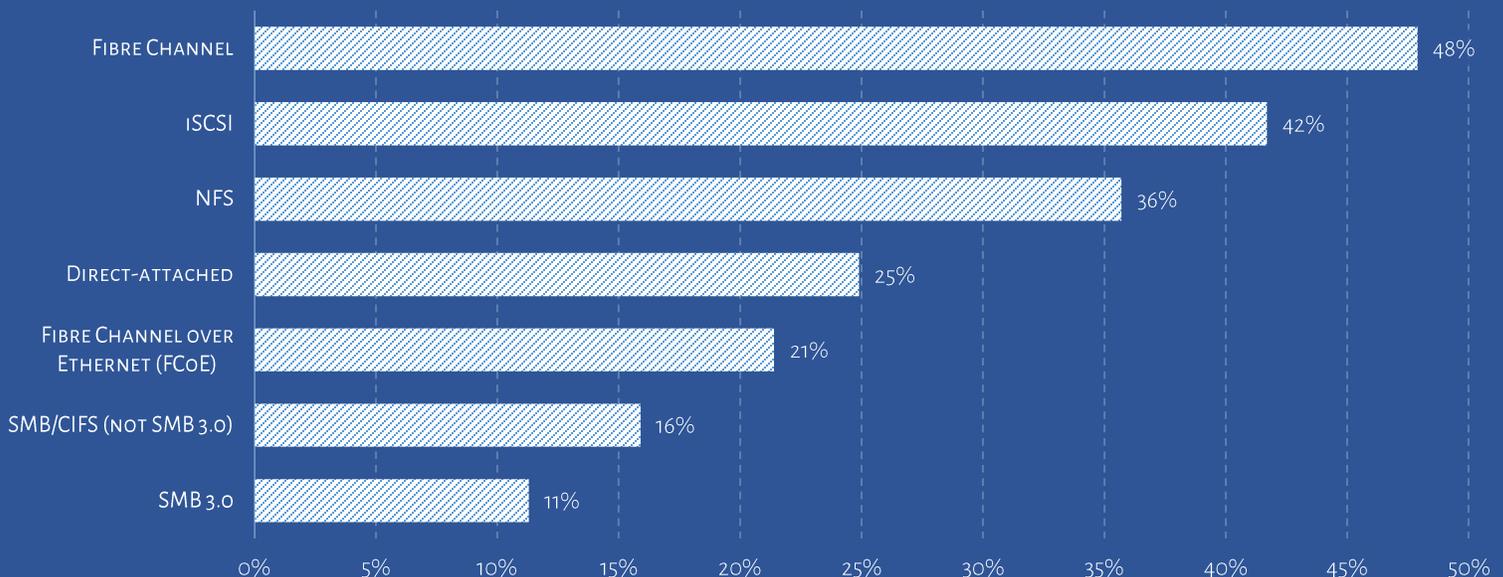


Figure 12: Storage protocols in use

Multiple Protocol Support

There are some other interesting ways to consider storage protocol data, too. We know that Fibre Channel and iSCSI are the clear winners in the protocol wars, but, as you can see in Figure 13, only 44% of respondents work in organizations that use just a single protocol. 56% of respondents are managing multi-protocol storage environments with 13% of the total using four or more protocols.

There are a number of different ways to interpret this information and to guide organizations as to what it means. It could mean that there are a lot of organizations that have too much complexity. It could mean that storage vendors must be ready to handle environments that require multi-protocol support.

The reality is that the answer is more than likely to be *both*. There is a lot of complexity out there and storage vendors must be ready to support it.

Number of protocols in use in respondent organizations

(N=1005)

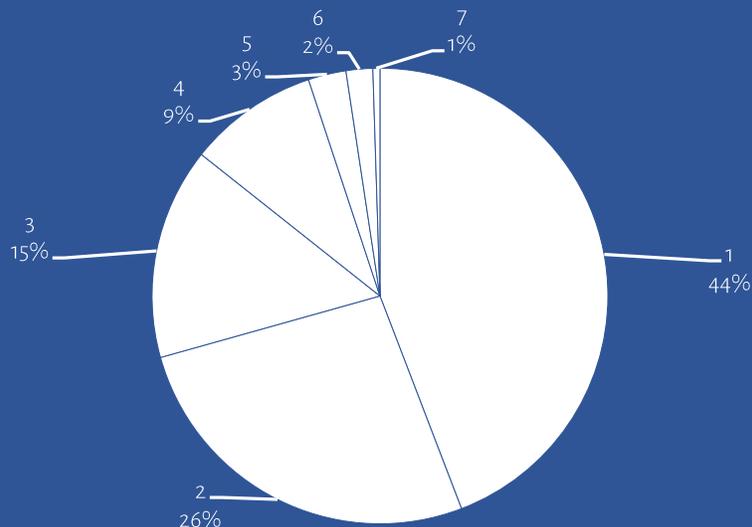


Figure 13: Number of protocols in use

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 14, smaller companies have *far less* fibre channel uptake than larger companies and do a lot more direct-attached. Only 12% of small companies use Fibre Channel whereas up to 32% of larger companies do. Further, as company size increases, use of iSCSI decreases: 25-29% for small companies and down to 13% for bigger ones. Fibre Channel over Ethernet (FCoE) does well as organization size increases as well.

How many people work in your company? What storage protocols are you using for your virtual environment?

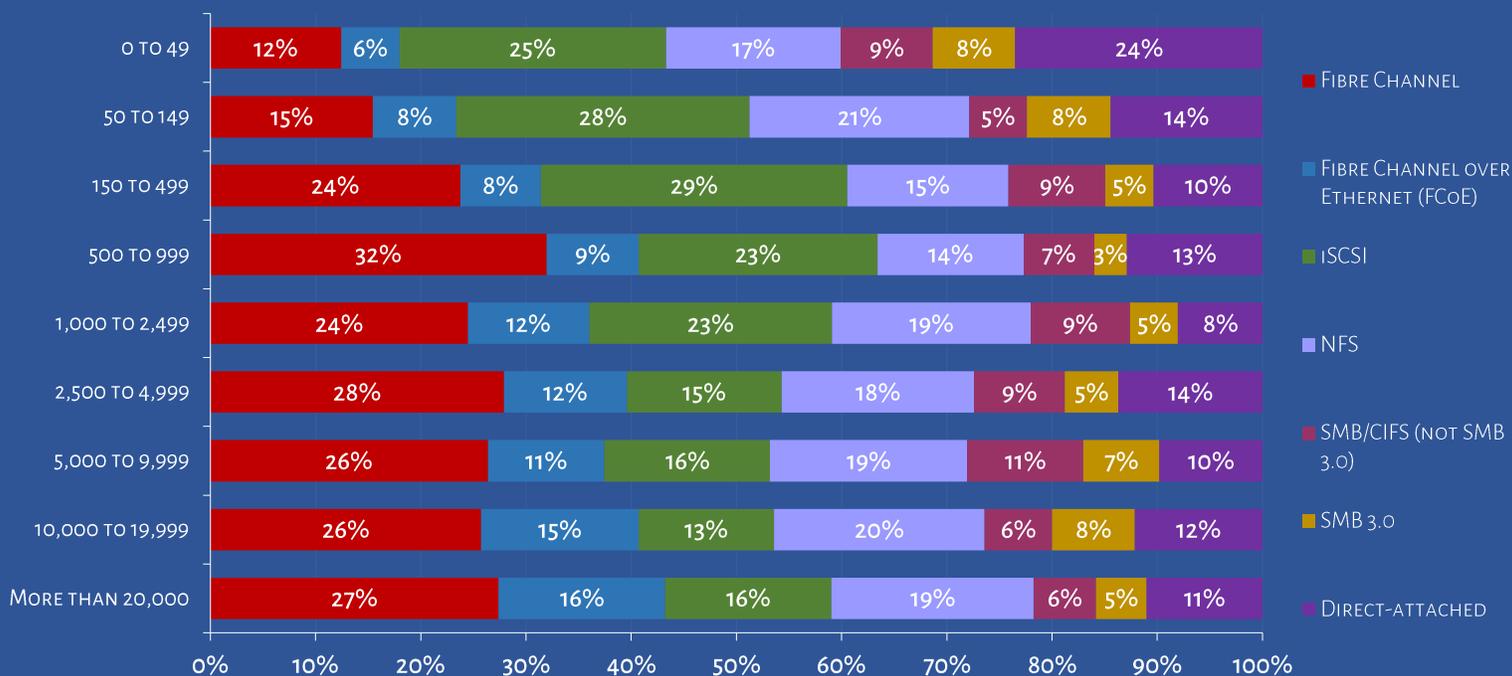


Figure 14: Protocol choice as a function of company size

Storage Capabilities

What used to be add-on features for extra cost are capabilities that are now built into many storage systems, although some vendors still charge for them. Storage is about much more than just holding on to your data in as efficient a way as possible. To an extent, storage is also about data protection and security. Replication is a practice of writing the same data to two separate locations (or two separate storage systems). Replication is often used as part of a disaster recovery process and is often used to copy data from one site to another. Different types of replication exist, including synchronous and asynchronous replication. Synchronous replication writes data to both Site A and Site B at the same time. Asynchronous replication writes data first to Site A before writing it to Site B. In asynchronous replication, time delay and frequency of writing data to Site B may be associated with an organization's disaster recovery (DR) policy that contains specific RTO (recovery time objective) and RPO (recovery point objective) goals.

With that in mind, Figure 15 gives you a look at the various storage features that are built into respondent's 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.

Does your existing storage provide any of the following features?

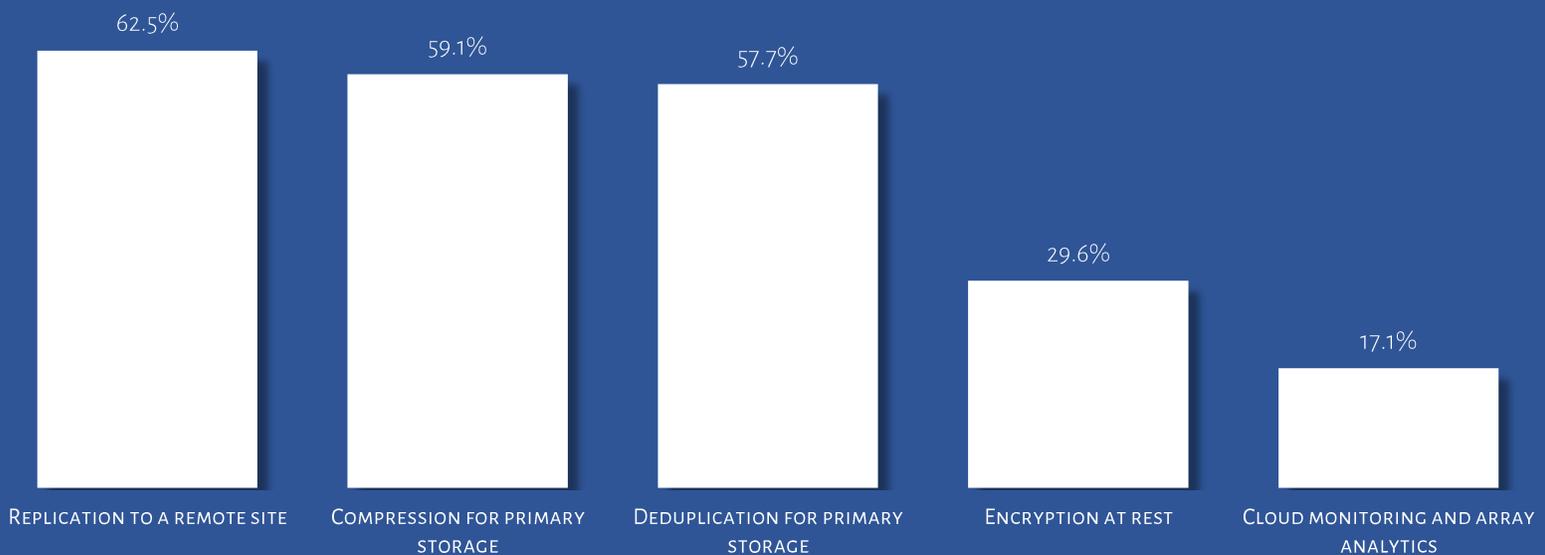


Figure 15: Capabilities inherent in existing storage systems

You may be wondering what is meant by “cloud monitoring and array analytics.” On a regular basis, arrays with these capabilities send to the corporate mother ship — the storage array vendor — a plethora of statistics around storage environment usage, performance, and health characteristics. In many cases, such systems are sending hundreds or thousands of data points to centralized vendor databases on an ongoing basis.

At first glance, these kinds of systems might appear to be ways to just streamline the support experience, but their use has far more benefits.

Customer Support and Proactive Issue Correction

Yes, these systems are absolutely used to make the customer support experience far better than the traditional support methodology that usually goes something like this eight step process:

1. Customer (or worse, an end user) notices a problem
2. Customer determines that the storage array has a problem or is running slowly
3. Customer calls vendor and provides them with case details
4. Customer hangs up and waits for a call back
5. Vendor calls back and requests log files to be sent and tells customer they will call back after analysis
6. Customer sends log files to vendor
7. Vendor analyzes log files
8. Vendor calls back to help customer correct whatever issue was being experienced

What if, instead, the support process looked like this for the majority of issues?

1. Vendor fixes problem or vendor notifies customer that storage capacity is running low and needs to be upgraded.

With automated support systems that report stats back to the vendor on a regular basis, this one step process is actually achievable. Of course, there will remain instances in which the first support paradigm still needs to be followed, but these can become the exception rather than the norm. When a vendor sees a problem even before the customer and can easily fix the problem, the customer has the ability to focus their efforts very differently . . . and in a good way.

This support improvement goes beyond just the vendor, though. It can be a critical enabler when it comes to improving the relationship between a VAR and a customer, as long as that VAR is able to access the statistics.

Issue Avoidance

Support is great, but it's a reactive service. Support happens when some event takes place that requires the support call. What if whatever issue that might result in a support call can actually be avoided? When an array vendor has the ability to match your array's performance, health, and application characteristics against thousands of other customers, there are analytics opportunities that make issue avoidance a reality. As

customers grow, they can be proactively notified about things like unacceptable cache miss ratios, which would imply that flash or cache storage needs to be increased. These systems can also begin to understand individual application performance characteristics in order to improve the experience for other customers that might be likely to encounter similar issues.

Product Development

Just as important, with billions of actionable data points at their disposal, vendors can use their massive customer-provided data statistics as a part of their ongoing product development plans. For example, for a hybrid vendor, are customer environments beginning to demonstrate that the time is right to introduce an all-flash system? While that's a very general example, it does demonstrate how data can be used for product development.

Or, for example, is the vendor seeing particular challenges for organizations running 2,500 Exchange mailboxes? They can take product development steps to address such challenges and further improve the ability to address the market they serve.

Storage Variety

The more of 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 16). 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 its 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=977)

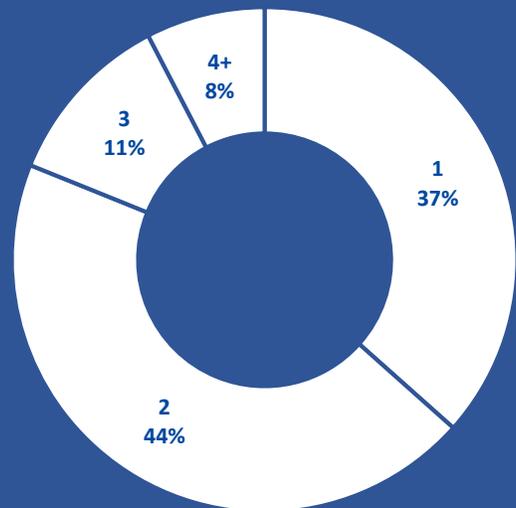


Figure 16: Number of storage vendors in place

VMware VVols

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.

VMware's VVols aims to change that situation. However, in order for any trend like this to be successful, end users need to understand the technology. There is a lot of work to do on this front, as evidenced by respondent's self-reporting their knowledge around VVols, shown in Figure 17. 55% of respondents indicate that they have little to no knowledge of VVols. Only a scant 5% consider themselves experts, with 40% indicating moderate understanding.

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

(N=1049)

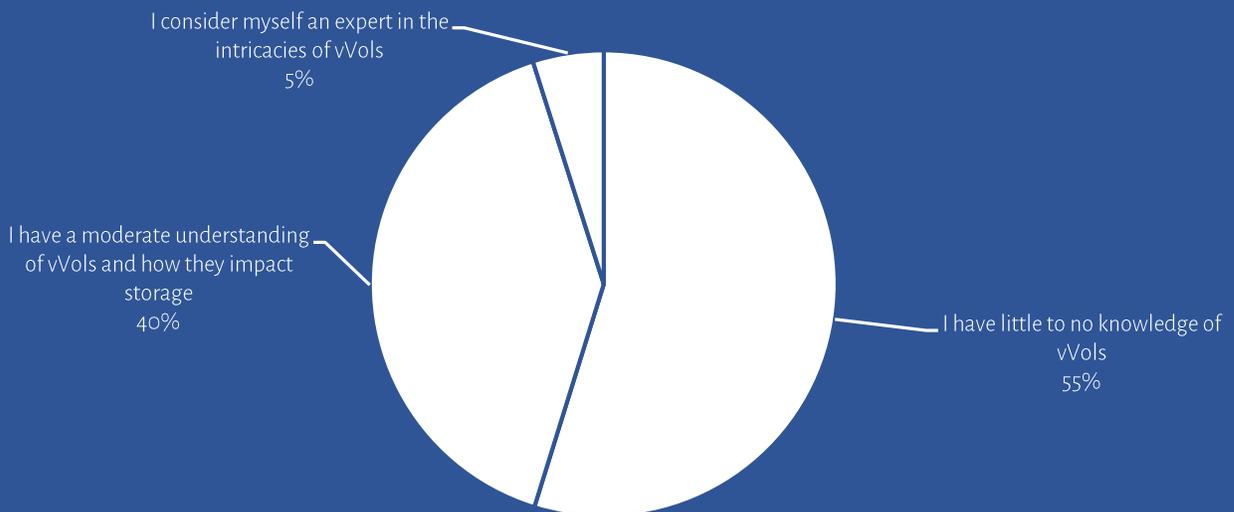


Figure 17: Respondent feedback regarding knowledge of VVols

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.

VVols are not a product you buy; they are a set of APIs that must be integrated by storage vendors. Once a vendor has added this API support, VMware vSphere can more directly integrate with the storage services. Policy-based administration improvements were cited by 41% of respondents as a positive outlook on how VVols can improve the storage environment (Figure 18).

In what ways do you believe that VVols can help to improve your storage environment?

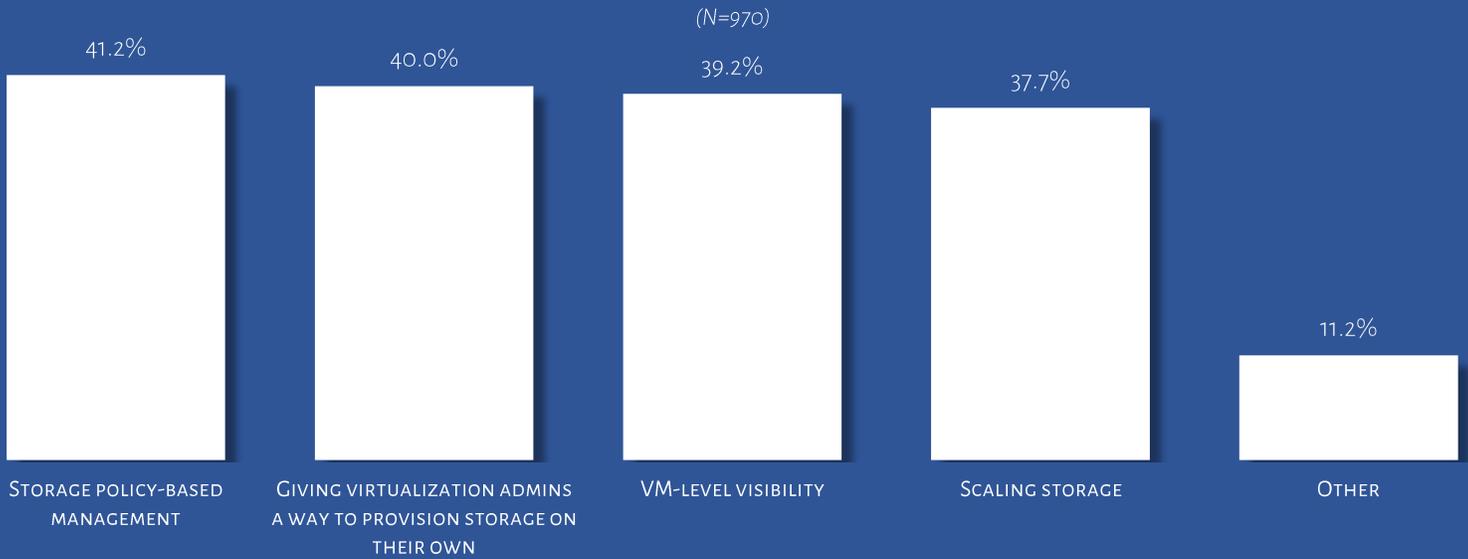


Figure 18: VVols improvements to storage environments

VVols can also help to eliminate silos within the IT department by enabling virtualization administrators to provision their own storage 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.

With VVols enabled, administrators are able to manage storage at a much more granular level as every virtual machine has a minimum of two VVols when powered off and three or more VVols when a virtual machine is powered on. With snapshots and many virtual disks, it's conceivable that a virtual machine could use 100+ virtual volumes. For this reason, administrators must be careful that their storage array not only supports VVols but that it can support the number of virtual volumes that their virtual machines will consume, over time.

In order to support VVols, arrays must be a VASA (vSphere API Storage Awareness) 2.0 provider, must have SAN firmware that supports VVols, and must support vSphere 6. Of course administrators should also make sure that their array is listed on the vSphere 6 compatibility list). VMware's vCenter will communicate with the array via the array's VASA 2.0 provider. Migration to VVols is usually done by moving a virtual machine to the array using storage vMotion (svMotion).

Once VVols are in place, administrators configure per-virtual machine storage policies and the array will automatically be able to adapt to those policies. For example, an administrator can specify that a virtual machine needs additional storage resiliency or performance, the array will adapt to that on the fly with no underlying change to the storage and no downtime. One of the most obvious benefits with VVols is that virtual machine snapshots can go from 10+ minutes down to seconds because the snapshot functionality is offloaded to the storage. Organizations also receive an improvement in storage space utilization because they no longer have multiple over-provisioned LUNs. Finally, VVols makes for simpler management of the storage resources.

Understanding performance characteristics in a granular way is the only way to ensure consistent, ongoing performance. VVols helps to bring improved storage performance characteristics to the hypervisor; this information is used for rebalancing and enables administrators to make data-driven decisions about the environment.

We asked respondents about their need for a storage solution that supports VVols. Only 8% (Figure 19) indicated that they would refuse to buy a storage device that doesn't support this technology. The rest, at 92%, indicated that it's either nice to have (65%) or not important at all (27%).

How important is it to you that a storage solution support VVols?

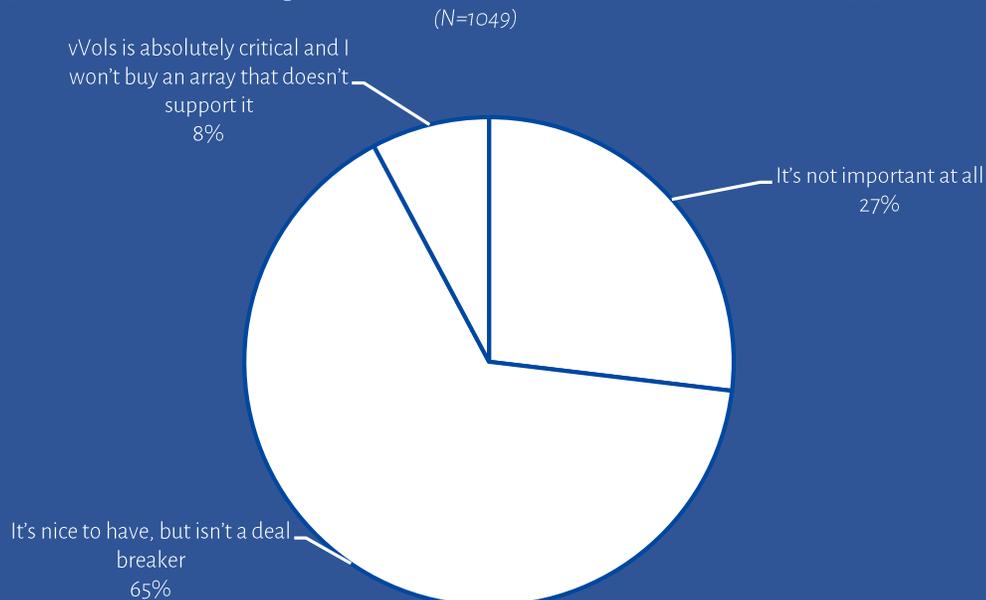


Figure 19: Criticality of VVols support in a new array purchase

At first, you may assume that no one thinks VVols is all that important. However, when the respondent data regarding VVols knowledge is combined with the importance of VVols, an interesting story emerges. As you can see in Figure 20, as people's understanding of VVols increases, so do people's feelings regarding the importance of VVols inclusion in an array. Figure 20 shows you in absolute terms that for people who consider themselves experts, only 2% feel it's not important at all, but that number grows to about 15% of those same people feeling the VVols is absolutely critical.

How important is it to you that a storage solution support vVols? What is your knowledge or understanding of VMware's VVols?

(N=1049)

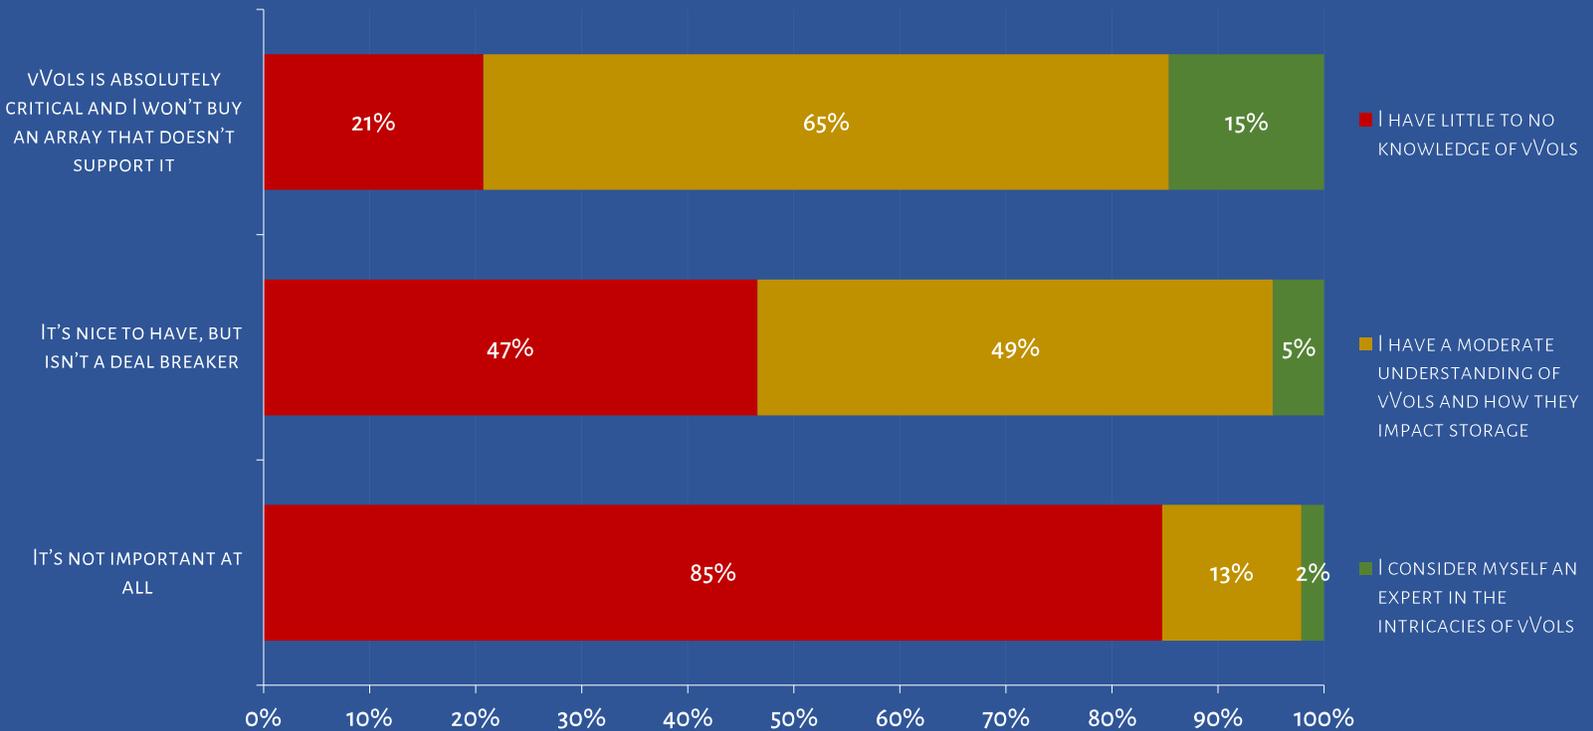


Figure 20: VVols criticality as a function of knowledge around VVols

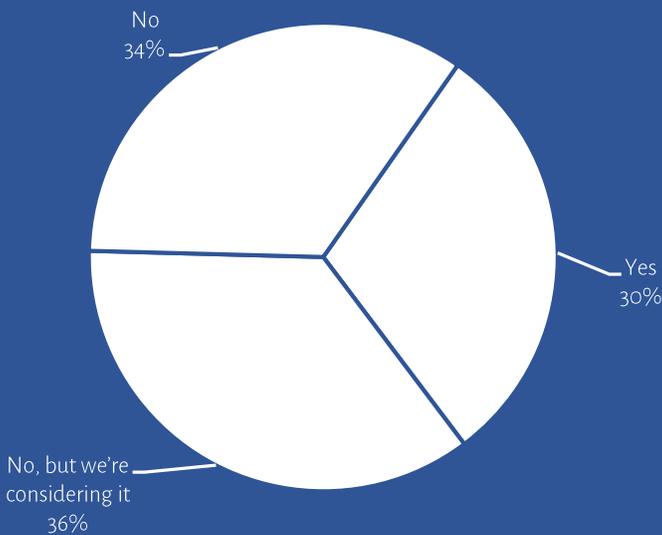
Cloud-based Storage

Of course, on-premises storage isn't the only game in town. Plenty of organizations —30% according to our survey (Figure 21, left side) — have adopted cloud storage in some way. An additional 36% are considering some kind of adoption with 34% indicating that they do not use and do not plan to use cloud storage at present.

In the right-hand graph in Figure 21, for those that indicated that they are using cloud storage, you can see the various ways that these respondents use the technology: 12% have all (or most) of their data centers in the cloud; 42% use cloud as a storage tier, while 46% are using cloud for backup and recovery.

Are you using cloud storage at all in your environment?

(N=1049)



In what ways are you leveraging cloud storage?

(N=314)

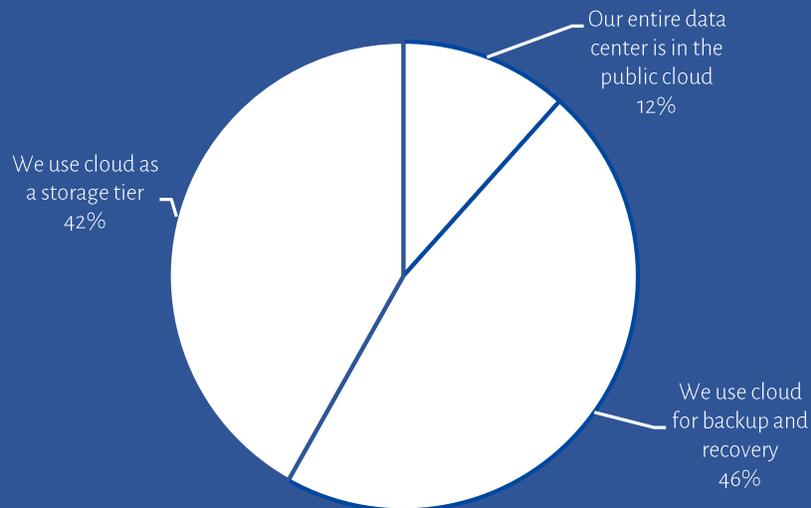


Figure 21: Use of cloud storage

There are any number of reasons that businesses choose to use cloud for storage. As is evidenced in Figure 22, the most popular reason is economics. CIOs and CFOs really like things that are “pay as you grow” and this has been a popular reason for cloud adoption.

Disaster recovery is another popular use case, with 40% of those using cloud saying that this critical insurance was one of their key cloud drivers. It prevents the need to build a second data center. The next two drivers, a need to react to peaks and valleys in demand and the need for immediate storage, are popular for similar reasons. Cloud allows instant deployment and can be easily turned up and turned down.

The last option, on-premises storage being too complex, can't be overstated. And, before you think that this is only an issue in small organizations, consider the chart shown in Figure 23 (next page). It is readily apparent that this is not just a small organization issue. Even the largest organizations face ongoing storage complexity. Bear in mind that the sample size for this particular data point is relatively small, but it is interesting nonetheless.

The message for storage vendors here is clear. If you can provide customers with a solution that addresses the complexity issue, you will do well.

What drivers drove you to use or consider cloud for storage?

(N=314)

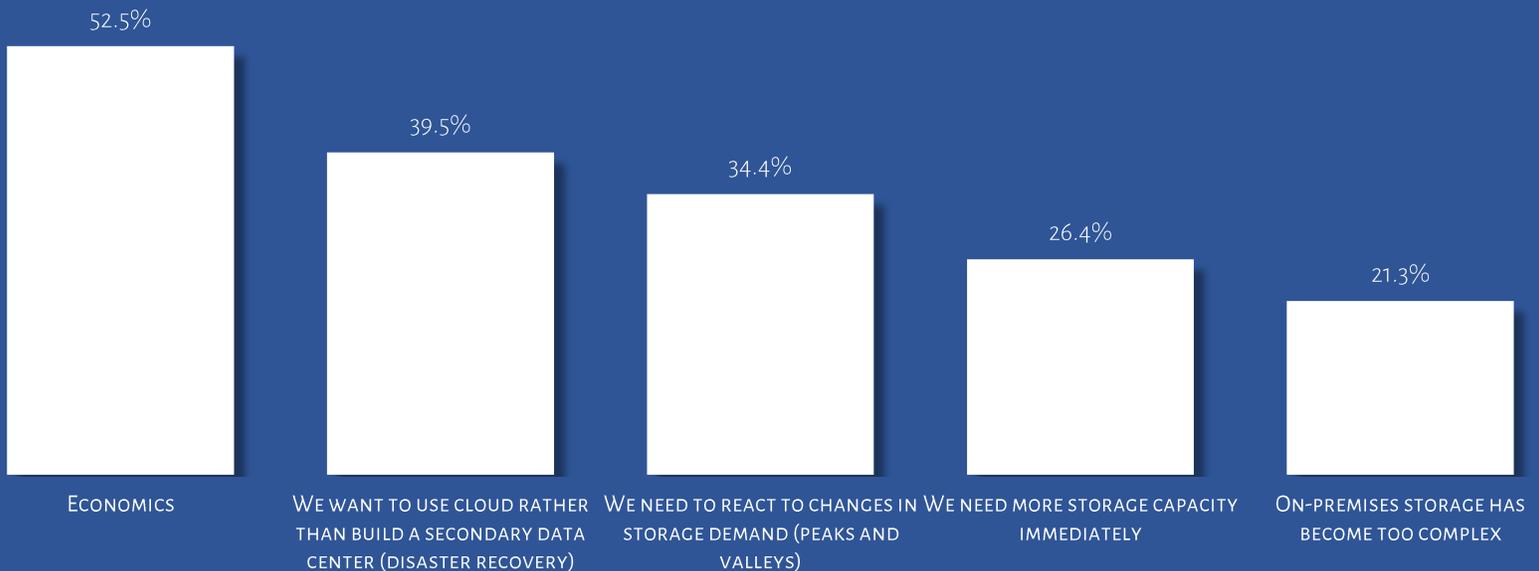


Figure 22: Cloud implementation drivers

How many people work in your company? (Those that chose 'On-Premises storage has become too complex')

(N=67)

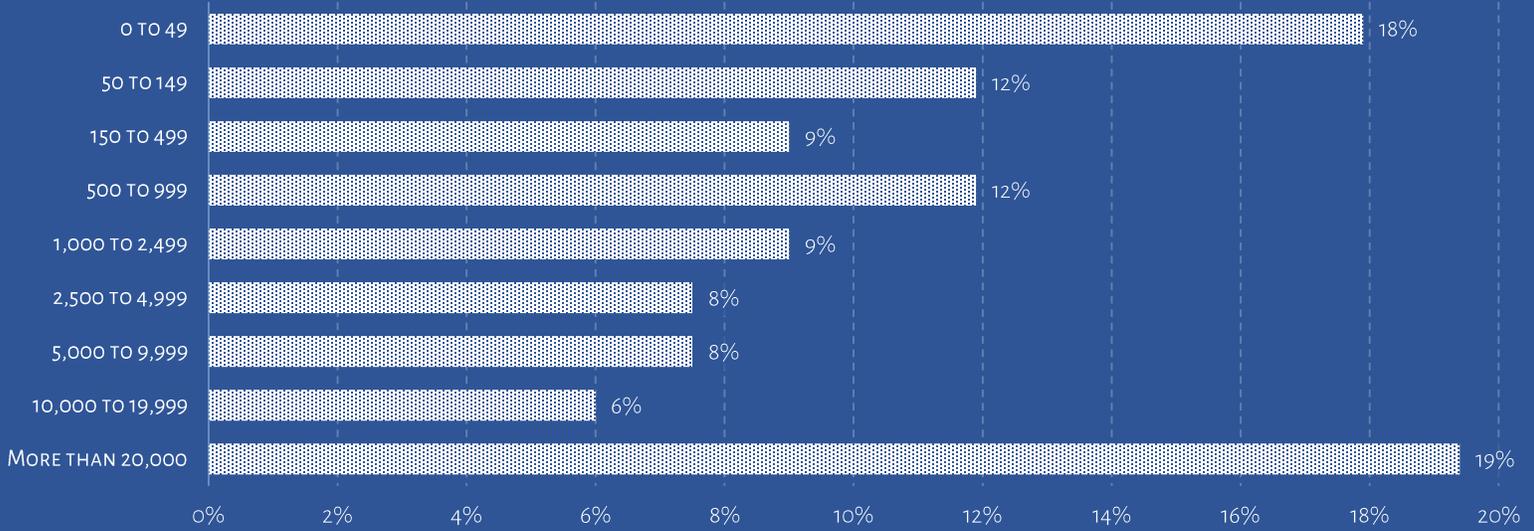


Figure 23: The number of people indicating storage complexity-- by company size

Respondent Demographics

The following figures provide you with detail regarding the makeup of the survey respondents

WHAT IS YOUR INDUSTRY VERTICAL? (N=1049)

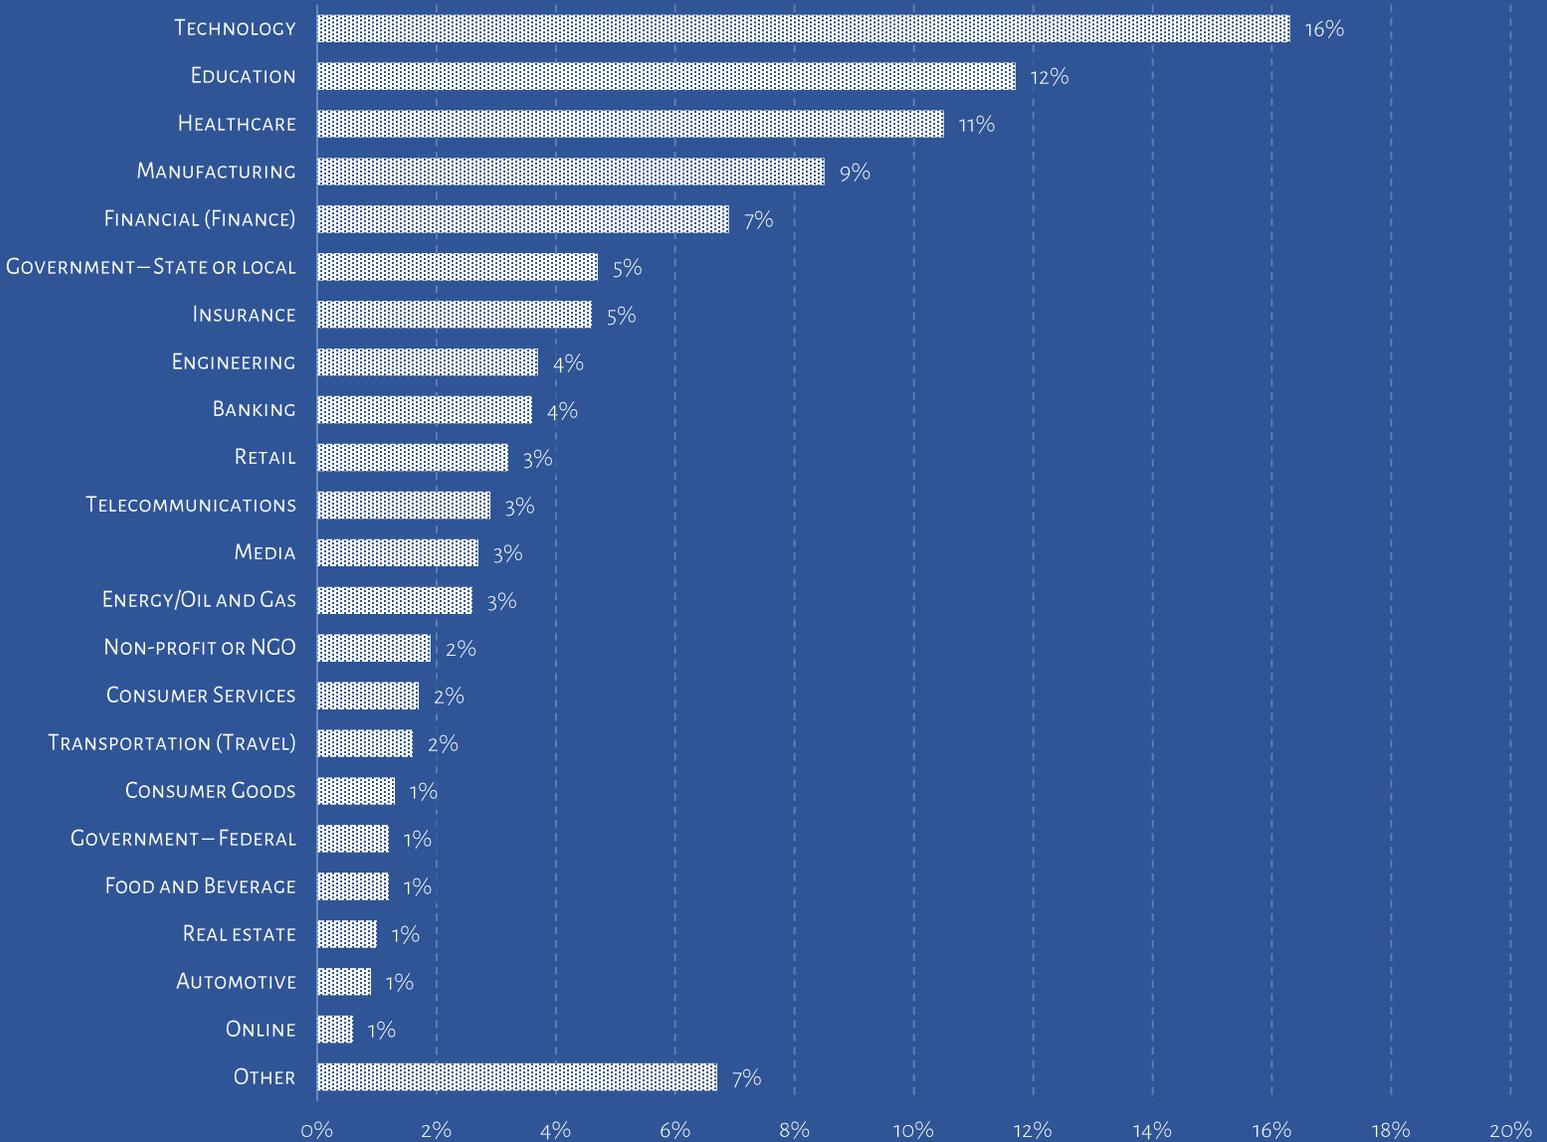


Figure 24: Respondent vertical

Which phrase best describes your title?

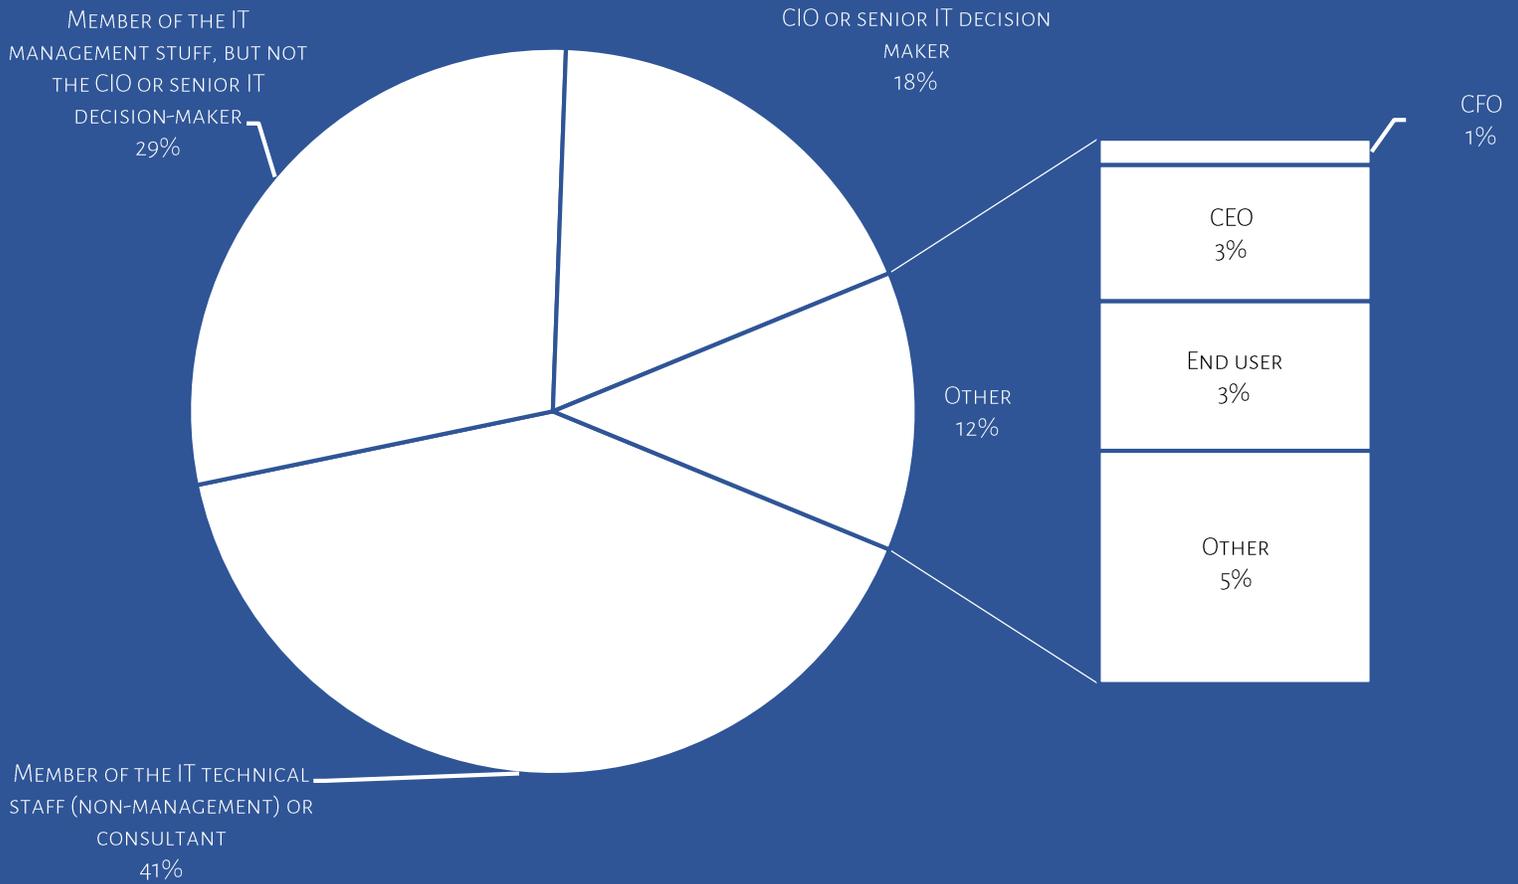


Figure 25: Respondent title category