Enterprise Strategy Group™
by TechTarget

WHITE PAPER

Key Workloads and Use Cases for Virtualized Storage

How Virtualized Storage Solutions Are Enabling IT Modernization Across a Range of Critical Use Cases

By Simon Robinson, Principal Analyst
Enterprise Strategy Group

August 2024

Contents

Introduction

As organizations of all types and sizes seek to transform to take advantage of the digital revolution, many are facing the prospect of modernizing their underlying infrastructure. Software-defined approaches to compute—via server virtualization—are now standard practice, but recent advancements are making virtualized storage an increasingly compelling choice for IT organizations looking to modernize their storage environment.

But this isn't about modernizing for the sake of it; it's about establishing a standard operating model for IT that can be applied to a broad range of workloads and applications, which, in turn, can drive compelling gains in operational efficiency and provide the business with a platform for the workloads of today and tomorrow. Virtual storage is playing a key role in enabling this.

Why Does Storage Matter, and Why Is It So Challenging?

The storage layer has always been critical part of the infrastructure. The importance of its role in storing, managing, and protecting an organization's most important digital assets—its data—is not in dispute. But storage has always been a challenging part of the environment to manage. What's more, driven by several factors, many of these challenges are not going away:

- **Growing data volumes.** The digital revolution is being fueled by data, lots of it. And the volume of data that organizations have to store and manage—both structured and unstructured—continues to grow rapidly, presenting IT functions with an ongoing challenge. According to TechTarget's Enterprise Strategy Group research, 69% of organizations find data storage infrastructure demands are hard to predict.[1] The rapidly emerging AI era will likely see data volumes grow significantly as organizations begin to build and implement AI training and inference models across their environment.
- **Storage complexity.** With 9 in 10 organizations finding that their overall IT environment has become more complex over the last two years,[2] this is a major challenge overall. This complexity is also evident at the storage infrastructure level: 68% of respondents responsible for evaluating, purchasing, and managing data storage technologies agreed that IT infrastructure environment complexity is slowing down IT operations and digital initiatives. Enterprise Strategy Group research points to a lack of storage array integration with broader IT management and automation tools, time and effort required to provision capacity to consumers, and performance as key challenges related to on-premises, data-center-based storage.
- **Fragmented storage.** Not only is there more data, but there's more data in more locations: on premises, in the core data center, at edge locations, in colocation facilities, in the cloud, and so on. Enterprise Strategy Group found that 67% of organizations reported managing data at the edge as a significant burden. This is compounded by the proliferation of storage silos across storage area networks (SANs), network-attached storage, direct-attached storage, object storage, etc. Managing these different types of storage often requires specific expertise, which makes it difficult to drive consistency, as well as efficiency, across these locations and could also lead to greater risk in terms of data security and resiliency.
- **Changing application demands.** The application landscape is shifting rapidly as organizations embrace a range of cloud-native technologies, such as container technologies running Kubernetes orchestration, to build, manage, and deploy new apps in a highly dynamic and lightweight manner. As these apps begin to require more persistent storage, this places new demands on the storage environment, since it requires much higher levels of automation to support the ongoing creation and deletion of often thousands of containers. It also requires new levels of provisioning that enable developers to access storage capacity directly and instantly. This is quickly becoming a "must-have" capability: Enterprise Strategy Group research found that 80% of

[1] Source: Enterprise Strategy Group Research Report, *Navigating the Cloud and AI Revolution: The State of Enterprise Storage and HCI*, March 2024. All Enterprise Strategy Group research references are from this report, unless otherwise noted.
[2] Source: Enterprise Strategy Group Complete Survey Results, *2024 Technology Spending Intentions Survey*, February 2024.

respondents said their organization needs to find ways to increase or improve self-service storage provisioning for developers.

- **Cost.** All of this adds up to significant cost for IT organizations in terms of physical IT hardware capital cost and cost of installation, ongoing management, as well as data center space and power. This is a major motivator for organizations to simplify their storage infrastructure.

Traditional three-tier storage architectures—i.e., those comprising physical compute hosts, a dedicated storage fabric providing connectivity to centralized shared storage, and shared central storage array—have been the predominant approach for many organizations over the last several decades. However, three-tier storage is expensive to buy, run, and manage, and it requires specific expertise to manage at a time when the IT workforce is becoming more generalist. It can also be difficult to scale, for both performance and capacity, meaning organizations typically buy more capacity than they require, which is inefficient. The result for many organizations is a suboptimal collection of disparate storage silos that are difficult and expensive to manage.

It's time for a new approach that makes flexible use of infrastructure for all applications, across all locations, on a consistent basis.

Enabling Critical Use Cases Through Storage Virtualization

What Is Virtualized Storage?

The last decade has seen the emergence of virtualized storage solutions as an alternative to address the shortcomings of three-tier storage architectures. Virtualized storage solutions such as VMware vSAN eliminate the need to deploy three-tier storage by aggregating the storage and compute of commodity server hardware into a shared resource pool. Distributing resources this way across server clusters provides resiliency and performance while enabling simplified, incremental scalability since additional servers can be added into the cluster as growth occurs.

This new approach to enterprise storage is proving extremely popular in a range of IT organizations. Today, virtualized storage solution approaches, which also form the basis of hyperconverged infrastructure (HCI), are deployed across a vast range of enterprise use cases and are now regarded as a cornerstone of data center modernization. According to Enterprise Strategy Group research, 82% of organizations deploying HCI said it is core to their data center modernization plans, while 9 in 10 said HCI is suitable for mission-critical workloads. 84% noted that deploying HCI has helped them become more agile and agreed that HCI is a core part of their edge modernization plans.

Virtual Storage – Driving a Range of Benefits

Diving in a layer, the specific benefits that organizations deploying virtualized storage and HCI are experiencing are multiple, as detailed in Figure 1.

Figure 1. Benefits of HCI Deployment

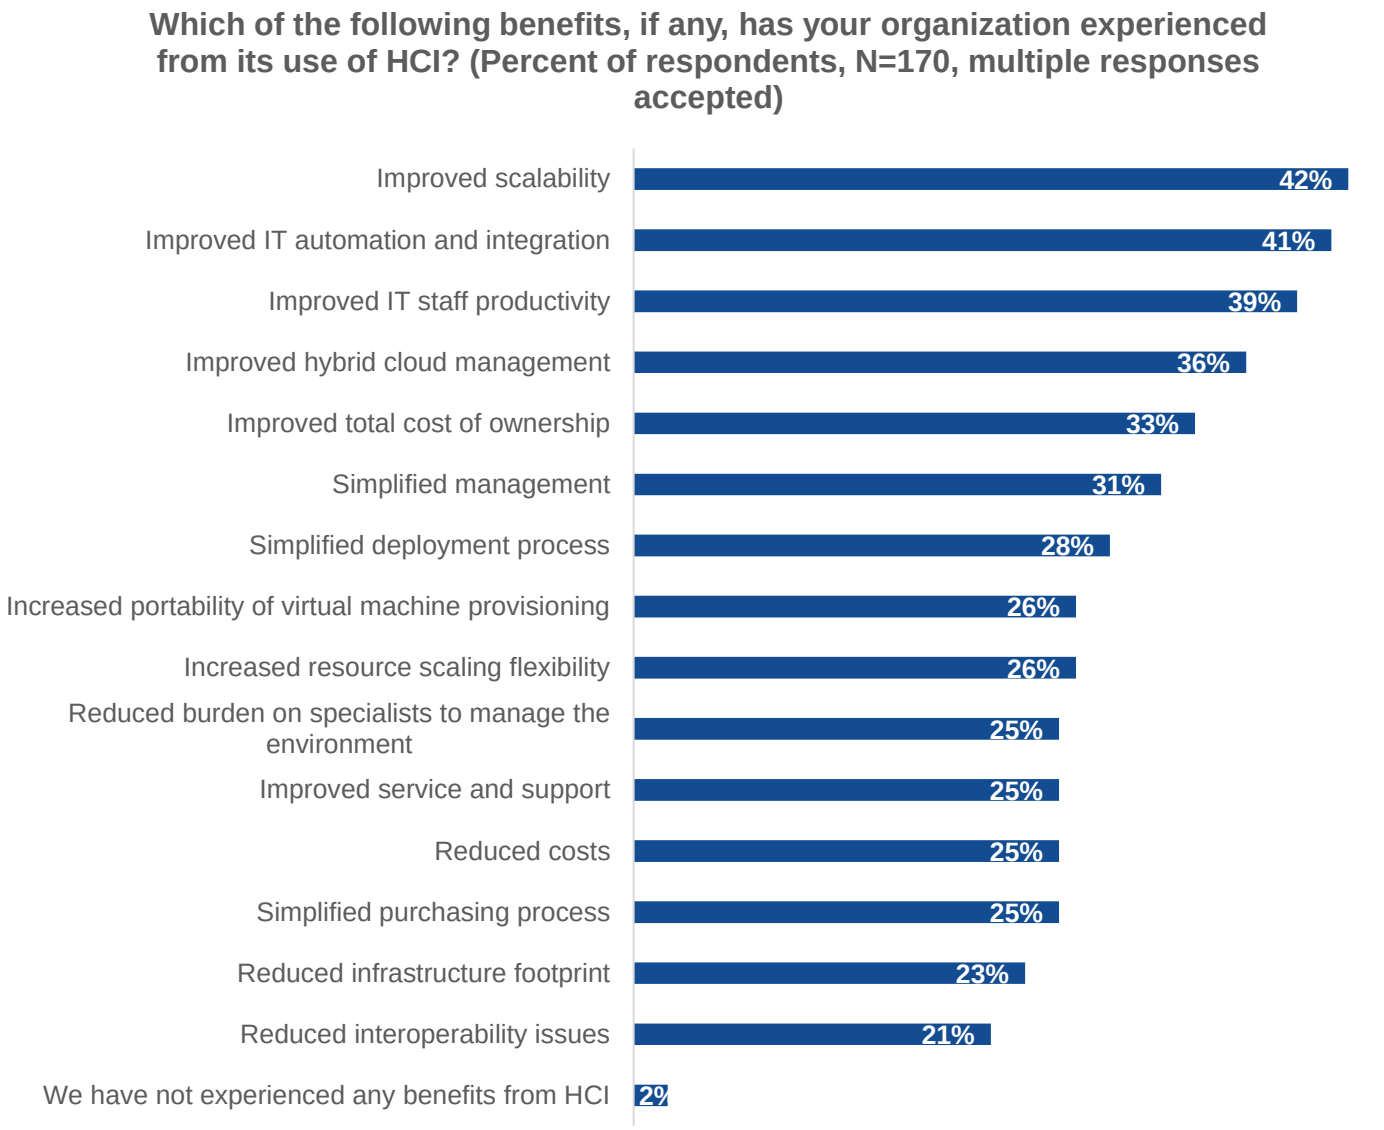

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

Notably, these benefits include hard, quantifiable outcomes, such as improved total cost of ownership (TCO), reduced costs, and reduced infrastructure footprint. However, organizations are clearly also experiencing benefits that might be harder to quantify but are no less tangible. In particular, rapid scalability was the most frequently cited benefit, selected by 42% of respondents. In an era of rapid data growth but also one of uncertainty as to the impact of new applications (such as AI), the ability for IT organizations to quickly and cost-effectively scale their IT environments is paramount.

"Must-have" Capabilities for HCI

- Independent scaling of storage and compute (cited by 85%)
- Support for container-based workloads (cited by 88%)

Storage virtualization and HCI software are particularly well suited to this since capacity can be added to increase the size of a cluster relatively easily with these technologies. Recent innovation in virtualized storage solutions has enabled high levels of scale; for example, VMware's vSAN can scale up to 64 nodes or be deployed as a stretched cluster, delivering high scalability and flexibility. It can also deliver high levels of capacity: up to 70% more capacity with the same number of servers since every virtual device contributes.

Note that this may not always require scaling compute and storage at the same rate. Some organizations want the option to add storage capacity without adding compute if data growth is outpacing compute growth; alternatively, an organization might want to add compute-centric nodes if more processing is required on the same data set. Similarly, the ability to scale down, as well as scale up, might be important to organizations that, for example, are moving the location of a particular application or that wish to retire or downsize a particular application or workload. Traditionally, scaling down three-tier storage is all but impossible.

Another notable cited benefit is the improvement organizations are seeing around IT management, automation, and integration. This is especially notable, given that this is a top-cited challenge with traditional external storage approaches. The tight coupling that some providers offer between virtualized storage and the other foundational infrastructure components—virtualized servers and virtualized networking, along with an overall consolidated management plane—further promotes this benefit, in particular by offering simplified storage management that enables admins to provision storage in minutes and modify storage on the fly to deal with changing requirements. This integration also extends to public cloud environments: A key goal for IT organizations is to be "multi-cloud-ready" and to adopt on-premises infrastructure that is consistent, enabling them to migrate their applications across clouds without refactoring or re-platforming.

Meanwhile, the fact that almost 40% of respondents cited improved IT staff productivity as a top benefit should also not be overlooked. Freeing up IT staff from basic provisioning, management, and other "keeping-the-lights-on" tasks to focus on more value-added activities both drives greater efficiency and can often lead to a happier, more satisfied IT workforce.

Finally, a key advantage of virtualized storage solutions displacing traditional storage approaches is that they can run on a variety of hardware platforms, enabling customers to choose the hardware design profile that best meets their needs in terms of cost and performance.

Virtualized Storage – Key Use Cases

Virtualized storage is proving increasingly popular among IT organizations looking to simplify and more cost-effectively manage their IT environment. Though HCI can be deployed for a wide range of use cases, those getting started may wish to consider the following use cases, for which virtualized storage is especially well suited.

Virtual Desktop Infrastructure

Managing end-user devices is an increasing challenge for IT organizations, owing to a combination of increased security threats, the cost of maintaining multiple devices, and an increasingly distributed workforce after the COVID-19 crisis fundamentally changed working patterns and locations for many employees. The response for many IT organizations is to deploy virtual desktop infrastructure (VDI), which hosts virtual desktops on a centralized server and is deployed to end users on request.

However, deploying infrastructure for VDI—particularly at the storage level—can prove challenging. Dedicated external storage is expensive to buy and manage, requiring proprietary hardware such as Fibre Channel networking and, potentially, extensive configuration to optimize the VDI environment. One challenge, in particular, is the "boot storms" that occur when hundreds or even thousands of workers log on simultaneously each day. These traditional approaches, utilizing fixed I/O controllers, are often rigid and inflexible, making it difficult and expensive to

scale as the company changes in size. A simple but expensive and inefficient response is to massively overprovision storage to account for future growth.

By contrast, virtualized storage and HCI offers a simple, cost-effective, and scalable approach to infrastructure for VDI. A key advantage is the linear scaling model: IT only needs to buy the capacity they need today and can then scale incrementally by adding capacity as requirements grow. This can be scaled "out" by adding a single node at a time or scaled "up" by adding more storage capacity. In addition, space-saving features such as deduplication and compression reduce the total amount of capacity required. For these reasons, HCI is proving to be an extremely popular choice for running VDI applications.

Business-critical Database Applications

Database applications such as Oracle, SQL Server, MySQL, and SAP HANA, as well as vertically oriented applications like Epic and Meditech, continue to play a critical role in the modern data center, supporting key transactional systems that the business relies on and forming a foundation for advanced data analytics. These applications are demanding, especially at the data and storage level, and require high-performance storage as well as on-demand scalability, optimal efficiency, and simplified management.

Virtualized storage offers a foundation for running business-critical applications with support for the latest storage technologies, including all-flash and NVMe-based triple-level-cell (TLC) drives, to cost-effectively boost transaction speed in a consolidated footprint. It also offers a simplified approach to scaling that can grow rapidly, predictably, and incrementally to linearly and consistently scale database performance in terms of transactions per second. In solutions such as VMware vSAN, this capability is delivered in an efficient manner through data storage technologies that offer RAID 1 performance with the storage efficiency of RAID 6 and up to 50% less CPU overhead.

One important architectural requirement here is the notion of disaggregation. Bottlenecks in database environments that occur as data stores grow in size tend to be at the storage layer, rather than the compute layer. Hence, the ability for a virtual storage or HCI solution to expand high-performance storage resources separately from compute becomes an important way to cost-effectively scale. Scaling in this disaggregated manner also helps optimize database licensing spending since licenses can be consolidated by provisioning tier 1 databases on smaller vSphere clusters.

Additionally, moving to such an approach enables IT administrators to abstract the management from underlying hardware, enabling policy-based storage management that can evolve as requirements change. And modern virtualized storage solutions offer a rich range of data protection and resiliency features such as high-performance, high-capacity snapshots, as well as replication and high availability features such as stretch clustering, which ensure protection and resiliency for mission-critical database applications. Data reduction features such as compression (rather than deduplication, which is less effective for databases) offer more cost-effective use of storage capacity.

Mainstream Applications – Web Servers, Data Protection, and DR

Virtualized storage solutions are an excellent fit for many mainstream enterprise applications that are already running in a virtualized server environment. For example, web HTTP server applications (e.g., Apache, Nginx, Microsoft Internet Information Services, etc.), as well as backup applications such as Veeam, Veritas NetBackup, CommVault, and VMware Site Recovery Manager (SRM), are well-understood and mature "workhorse" applications, and, in many cases, they need storage that "just works" with low running costs and simple management. Virtualized storage offers both types of applications a lower TCO compared with dedicated external storage that is often overengineered and costly for these use cases. Such applications often grow with the business as well, so they need to run on a platform that can scale both capacity and I/O as requirements change. Meanwhile,

support for a range of storage types, such as native block and file storage, as well as for integration with object stores, becomes important as different applications require access to storage via multiple protocols.

Additionally, in an increasingly uncertain world fraught with risk, the role of disaster recovery (DR) has never been more relevant. Yet, historically, implementing DR solutions was a complex and expensive business, requiring identical hardware in the DR site as in the production site, not to mention replication software and the cost and resources of managing a second site.

Virtualization removed the need to run the same hardware by abstracting the physical hardware from the applications. This made it much easier to copy a virtual machine from one site to another since VM configurations, applications, and data are stored as files. However, storage is still required at the second site, and this remains a challenge because VM files are large and can affect costs, performance, and, ultimately, an organization's ability to quickly recover from an outage or disaster.

Virtualized storage solutions such as VMware vSAN are an ideal use case here because they enable inexpensive, industry-standard hardware to be deployed, significantly reducing the upfront investment and removing the dependency on external shared storage hardware, as well as the need for separate administration and management tools and connections. This vastly simplifies the operation of a DR location, which might be essential if there are little to no local IT staff. In addition, virtualized storage solutions enable protection policies to be set on a granular, per-VM basis, offering precise control over which, and how, different workloads are protected.

For example, tier 1 workloads could be protected using replication for fast recovery or stretch clustering across three sites for high availability, while tier 2 workloads could be protected through a local backup, which can then be copied to the DR site. Integration with tools such as SRM offers automation of the recovery process, nondisruptive DR testing, and precise control over the startup order of VMs. In addition to offering low recovery point objectives, these approaches further save costs because they avoid providing excess capacity at a DR site to accommodate "all-or-nothing" replication alternatives.

Edge Applications

Organizations of all types are increasingly adding all manner of new digital capabilities locally—including IoT, point of sale, security, and surveillance, to name just a few—to ensure high performance and low latency. But doing so effectively and efficiently is not without its challenges. These can include a lack of IT skills, growing complexity, space and power issues, and scaling limitations, among others. For such environments, deploying traditional external storage solutions is a complex and unwieldy option, and, as such, it is often a nonstarter.

A further challenge for edge environments, in particular, is that many emerging edge-based applications are deployed on container-based technology such as Kubernetes, owing to their lightweight nature. This creates another issue since infrastructure needs to manage these cloud-native architectures as well as existing virtualized applications.

Virtualized storage solutions offer a path forward by providing scalability, flexibility, and enhanced performance at the edge at a lower cost, in a smaller footprint, and, crucially, in a simplified form factor. Key capabilities for VMware vSAN include the ability to scale cost-effectively—for example, supporting up to 64 remote clusters with a single remote witness—as well as support for native file and block storage to reduce footprint; high performance via support for NVMe-based TLC NAND flash; centralized management that enables management of all remote sites from a single tool (with no limit to the number of hosts or VMs); and, crucially, cloud-native support that enables both containers and VMs to be managed with a single, general-purpose infrastructure. Additionally, a range of availability options enables centralized IT to effectively protect data with high-performance native snapshots and integration with third-party backup solutions.

Cloud-native Applications

Cloud-native approaches, such as container-based technologies running Kubernetes, are transforming applications and application development. The emergence of modern container-based applications, including AI, will further accelerate the adoption of hybrid cloud application deployment. Moreover, as cloud-native applications become increasingly stateful, running data applications such as MySQL, Redis, MongoDB, Kafka, and Cassandra, they will require underlying storage infrastructure with rich data services and high levels of automation, plus the scale and agility to match developers' requirements.

This is challenging from an IT infrastructure perspective for a couple of reasons. First, from a storage perspective, containers are highly dynamic, high-churn environments, with hundreds or even thousands of containers being created and deleted every hour. Managing storage manually in such an environment is impossible.

Second, IT doesn't want to manage containers in an entirely separate environment—that would be inefficient and add complexity. Instead, they need to manage this emerging class of applications alongside existing mission- and business-critical applications. However, this further cranks up the pressure on the underlying infrastructure to be highly performant, scalable, and agile. Rigid infrastructure stacks no longer cut it, and IT organizations can ill afford to deploy even more silos.

Virtual storage solutions are, again, an ideal approach, as they can offer cloud-native storage support to automate storage provisioning for container-based workloads. This enables developers to get the storage they need, when they need it, while at the same time enabling IT admins to run, monitor, and manage storage for both containers and virtual machines on the same platform.

Specifically, virtual storage solutions such as VMware vSAN enable developer self-service and scaling through the automatic provisioning of persistent volumes that automatically scale as applications scale, without IT admin involvement and with access to multiple storage classes that fit application needs. This can be managed via a Kubernetes API and console so developers don't have to understand the specifics of storage to consume it. It's also important that this is done in a secure manner that stays inside enterprise governance guidelines. Here, the same storage classes and policies that organizations use for their VMs can also be applied to their container environments. This extends to workload-centric storage management so developers can self-serve to provision storage but stay within defined boundaries. Finally, providing a consistent operating model for cloud-native storage across any vanilla Kubernetes distribution enables better traceability and faster troubleshooting, establishing a common language between developers and IT operations teams. For VMware vSAN, this model can extend across multiple clouds, including AWS and Microsoft Azure.

Conclusion

vSAN is VMware's market-leading modern storage software for VMware Cloud Foundation (VCF) and VMware vSphere Foundation (VVF), enabling management of compute and storage with a single, integrated platform. It offers the agility, scalability, and performance needed to meet the evolving demands of modern applications and workloads. Its direct integration into VCF and VVF ensures high levels of storage efficiency, performance, and operational simplicity that cannot be achieved with traditional storage.

VMware vSAN is a critical component of VCF, playing a pivotal role in helping customers achieve a cloud operational model in their data centers. Alongside networking, management, automation, and security features, VCF extends vSAN's capabilities to provide the agility, scalability, and performance necessary for both traditional and modern applications. This enables organizations to build, manage, and optimize private clouds using modern storage building blocks, with the added flexibility to integrate seamlessly with public and hybrid cloud environments.

VCF provides a flexible and simplified private cloud platform with public cloud extensibility that integrates leading components, including vSphere (compute), vSAN (storage), NSX (networking), and management, into a single solution.

Over the last decade, the emergence of HCI and vSAN has effectively laid the groundwork for a completely new storage operating model that is simple, performant, resilient, and compatible with modern applications. With the availability of options such as Express Storage Architecture and vSAN Max, VMware is enabling customers to take the next step with a disaggregated approach to HCI that enables them to supercharge their storage environments in step with their evolving digital strategies. Crucially, this now extends to a broad range of enterprise use cases—from the core to the edge—that provides a basis for a simplified IT operating model across a large part of the IT application estate, lowering acquisition and operating costs, improving scalability, and future-proofing the environment for tomorrow workloads.

For more details about VMware vSAN, visit https://www.vmware.com/products/vsan.html.

About Enterprise Strategy Group
TechTarget's Enterprise Strategy Group provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

contact@esg-global.com
www.esg-global.com