

VMware vCloud® Architecture Toolkit™
for Service Providers

Identifying and Calculating the Operational Savings of a VMware vSAN™ Solution

Version 2.9
January 2018

Martin Hosken





© 2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com



Contents

- Executive Overview..... 5
- Understanding the VMware vSAN Value Proposition 7
- Understanding the Changing Economics of Storage 10
- Identifying and Measuring the Operational Savings of vSAN 12
- Understanding Operational Transformation with vSAN 14
 - Scenario 1:** Initial Storage Platform Deployment 15
 - Scenario 2:** Modifying Service Levels / Workload Configuration 16
 - Scenario 3:** Increasing Capacity and Scaling..... 17
- Total Cost of Ownership Analysis 20
- Calculating the Cost of Operations 21
- Building a Business Case for vSAN 23
- How to Build a Cost Comparison for vSAN 27
- Understanding the Impact of Number of Failures to Tolerate 28
- Checklist for Decision Makers 30
- Summary 31
- Disclaimer 32
- References..... 33



List of Tables

- Table 1. Traditional Storage Array Management Lifecycle Tasks 13
- Table 2. Storage Activities Comparison Summary 18
- Table 3. Estimated Operational Costs (3 Scenarios)..... 21
- Table 4. Comparison of Total Cost of Three Scenarios Over Various Life Expectancies 22

List of Figures

- Figure 1. Storage Market Growth..... 5
- Figure 2. Storage Challenges and Issues 6
- Figure 3. Traditional Upfront Storage Provisioning Versus a Linear Scale-Out Model..... 8
- Figure 4. Hybrid Versus All-Flash Configuration..... 10
- Figure 5. Linear Scaling to Maximize Efficiency and Balance–Compute and Storage Resource 18
- Figure 6. Typical Storage Cost Factors 24
- Figure 7. OpEx – Over a Four-Year Period 24
- Figure 8. Simplified Annual Total Cost of Ownership 25
- Figure 9. vSAN Basic Design Factors..... 27
- Figure 10. Example Annual Storage Management OpEx and CapEx Cost Analysis 31

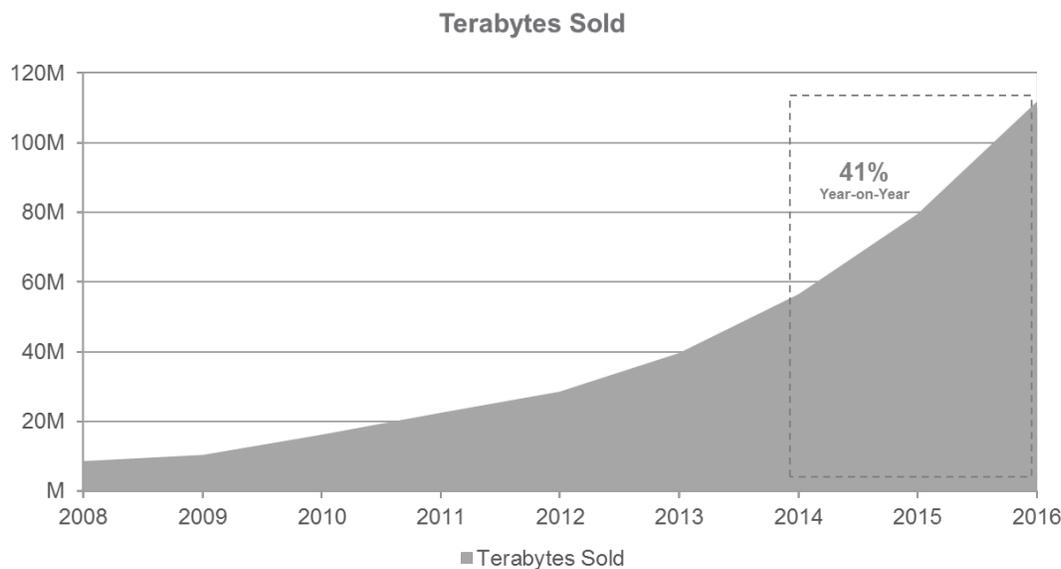


Executive Overview

VMware Cloud Providers™ offer a wide range of different service offerings to their customers, which themselves require a whole range of storage solutions. This extensive list of offerings is continuously being developed in order to differentiate providers from one another, to meet the ever-growing needs of modern applications, or to target specific industries and better serve their highly differentiated customer base. However, no matter what industries a service provider targets, whether computer gaming companies, financial institutions, or government customers, all sectors are seeing the same massive explosion in storage requirements, from both a capacity and a performance perspective.

The growth in data creation by human and automated systems that has been seen over the last decade has now reached an unprecedented rate. In addition, there are regulations and regulatory requirements that result in prolonged data retention. Because of this combination of factors, storage demands are exploding as never before, and putting significant pressure on service provider data centers, as shown in Figure 1.

Figure 1. Storage Market Growth

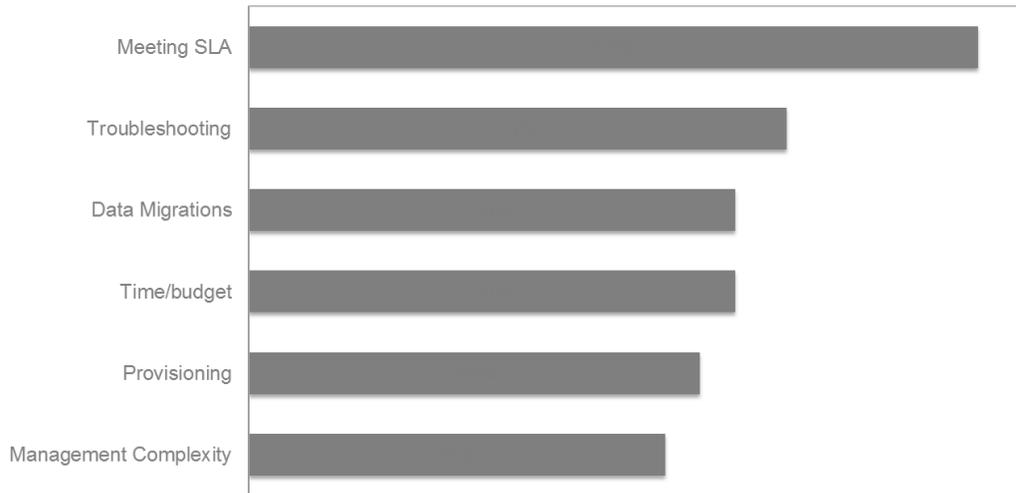


Source: IDC, Yezhkova, Worldwide Enterprise Storage Systems Forecast, November 2013, #244293; <http://www.idc.com/getdoc.jsp?containerId=WC20140109> (by subscription).

For service providers to meet the increasing storage challenges faced by their consumers, future expectations for delivering and simplifying storage resources must be addressed. VMware Cloud Providers are uniquely positioned with their customers to help address the growing and pressing storage challenges, and management and provisioning of storage are two of the top issues, as shown in the following figure.



Figure 2. Storage Challenges and Issues



Source: IDC, Storage Predictions 2014, January 2014, General Storage QuickPoll, #243511, n=307.

The growing requirements and expectations of storage systems, and the improvements and simplifications that a software-defined storage (SDS) system can offer, have progressed from being important to service providers to being imperative, not only in the implementation of these new technologies but also in their selection. Centralized storage and a storage-area network (SAN) or network-attached storage (NAS) have proven to be an excellent example of storage efficiency in the past. However, with limited potential for further improvement, service providers are evaluating and moving to SDS solutions, which aim to provide significant improvements in efficiency in the following four key areas:

- Scalability
- Price
- Performance
- Reliability



Understanding the VMware vSAN Value Proposition

VMware vSAN™ is a software-defined storage solution for VMware vSphere® environments. vSAN clusters employ the local servers HDDs and SSDs to create radically simple, high-performance, resilient, shared storage, designed for virtual machines and now with vSphere 6.5, offer block storage to external systems via the iSCSI storage protocol. Some of the key benefits of vSAN for service providers are:

- **Lower TCO:** Reduce TCO by up to 50 percent by leveraging server-side HDDs and SSDs to create a converged and resilient *shared* storage solution within the hypervisor.
- **Lower CapEx:** Server-side enterprise storage economics lowers capital expenses over array-based solutions.
- **Lower OpEx:** Automation, power and cooling, \$/TB, labor costs. With vSphere integration, current vSphere and storage skillsets can be used to manage the technology. In other words, “if you know vSphere, and a bit about storage, then you know vSAN.” In addition, management that is more efficient lowers operating expenses, which has the potential for significant overall OpEx savings.
- **Predictability:** No large upfront investments. vSAN allows you to scale granularly, and avoid overprovisioning and overpaying for future capacity and performance needs.

In addition to the above, vSAN is simple and automates time-consuming manual storage tasks. Not only is vSAN managed through the vSphere Web Client, but it also integrates with other VMware products, such as VMware vCloud Director® and VMware vRealize® Automation™. This integration into VMware's Cloud Management Platform solutions further simplifies the provisioning and management of storage in service provider data centers.

The cold hard truth is that storage infrastructure for service providers has become exceedingly complex, and often the “ad-hoc” solutions that are employed in order to try to achieve a *Cloud* delivery model have only added additional complexity. The fact is, most service provider problems can be traced back to poor infrastructure design choices, and rather than fix the root cause, service providers’ operational teams spend the vast majority of their time trying to remediate the environment. This leads to increased operational overhead through perpetual troubleshooting and recurring optimization activities. Therefore, the simplification offered through vSAN and a move towards a Hyper-Converged Infrastructure (HCI) offers service providers a break from their longstanding frustration with the complexity and high cost of existing storage systems, making them, in most cases, highly receptive to VMware’s vSAN “best of breed” storage offering.

In this vCAT-SP paper, we aim to address how the vSAN product affects operations and therefore operational costs. vSAN brings both a software-defined storage infrastructure and a VM-centric storage model together into a single solution, which can significantly reduce the TCO of storage, when compared to a traditional SAN and array architecture.

The vSAN hyper-converged storage model offers a software-defined storage approach that alters the acquisition cost of storage and fundamentally alters the ongoing operational costs in a number of different ways, including, but not limited to, the following:

Acquisition Costs (CapEx)

- Avoids proprietary storage hardware lock-in and associated costs
- No dedicated storage adapters (Fibre Channel HBAs) in hosts
- No requirement for fabric hardware, such as Fibre Channel director or Top of Rack switches
- Based on server-side economics – server components (hardware) competitively priced versus traditional external arrays (hardware + software)
- Lower upfront costs with granular scaling
- Leverage falling storage costs year-on-year by:
 - Spreading your purchases to benefit from declining prices



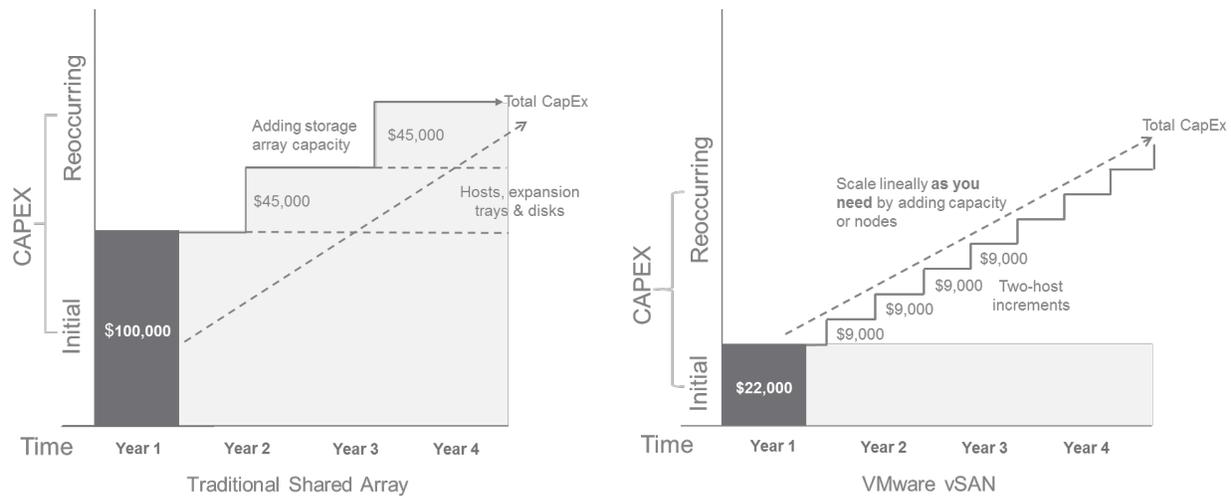
- Growing storage alongside compute (creating a building block architecture)
- Taking advantage of new hardware technologies coming to market
- Benefit from higher resource utilization through:
 - Eliminating overprovisioning of resources and SLAs
 - Avoiding low utilization due to LUN-VM mismatches
 - Simpler budgeting, forecasting, and capacity planning

Operational Costs (OpEx)

- Merge vSphere and storage administrative roles and operational costs
- Daily task efficiency through higher admin productivity and simple automation
- Deployment simplification and reduced footprint
- Simpler budgeting
- Lower the cost of specialized operational skillsets (labor costs)
- Simpler automation
- Power and cooling

As a result of this significant change in CapEx and OpEx models, the typical pattern of hardware acquisition seen in service providers using the vSAN model can be contrasted with the traditional model, as shown in Figure 3:

Figure 3. Traditional Upfront Storage Provisioning Versus a Linear Scale-Out Model



Source: Sample pricing based on vendors/resellers public website prices

As seen in Figure 3, another significant contributor to the lower TCO of a VMware vSAN solution is the lower cost of entry for hyper-converged server hardware, with no requirement for two or more dedicated storage fabrics in the architecture. Not only is the initial CapEx investment lower, because vSAN Ready Nodes or appliances typically cost much less than traditional storage. You can also scale them incrementally, by adding smaller numbers of HCI nodes or even by just adding storage capacity through additional disk groups, when required. This model of scaling an infrastructure compares very favorably with buying larger bundles of disk trays, disks, compute-only hosts, and fabric components separately.



Furthermore, if there is space available in the rack, you can simply configure a new vSAN HCI node into an existing cluster, and in a matter of minutes be ready to scale the storage and compute infrastructure. This building block approach saves a significant amount of time, versus racking and cabling new array disk trays, configuring new host servers, cabling the new Fibre Channel HBAs across multiple isolated fabrics, configuring new fabric zones on Fibre Channel switches, and so forth.

Though vSAN delivers a compelling story for both CapEx and OpEx, this white paper focuses on how software-defined storage can transform the long-term OpEx costs for service providers. I examine vSAN product operational management, and drill down into the detail of a number of key operational tasks associated with service provider storage provisioning and management and the associated vSphere infrastructure. We also address the key data points in a comparative examination, which aims to estimate the overall time and effort required for common OpEx tasks on VMware vSAN, when compared with traditional array and fabric SAN-based storage systems.



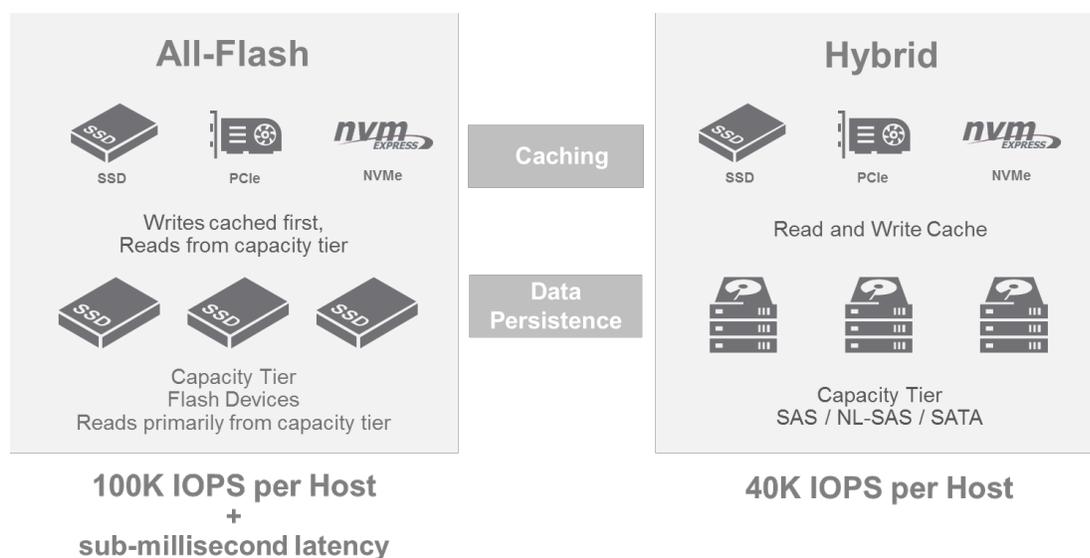
Understanding the Changing Economics of Storage

Before addressing the VMware vSAN total cost of ownership (TCO) question, or building a business case for vSAN, it is important that we cover how the storage market is in the midst of serious change and realignment.

Twenty years ago, servers were installed with local attached disks, which resulted in massive islands of storage, poor utilization of resources, and complex operational management. As a result, the industry gradually moved to a shared storage model with SAN and NAS arrays, which was largely accelerated by the increased adoption of virtualization. Though this significantly improved the TCO, these traditional storage arrays have become a large part of IT infrastructure costs and have struggled, in many cases, to continue to meet the performance and feature requirements of the modern virtualized application.

vSAN enables both hybrid and all-flash architectures. Therefore, one of the key design and economic factors in adopting vSAN is hybrid configuration versus all-flash architecture, shown in Figure 4.

Figure 4. Hybrid Versus All-Flash Configuration



As shown in Figure 4, irrespective of the architecture choice, there is always a flash-based caching tier, which can be configured from flash devices such as SSDs, PCIe cards, NVMe, or Ultra DIMMs. This flash caching tier acts as the read cache/write buffer in a hybrid configuration, which dramatically improves the performance of storage operations. In an All-Flash configuration, this caching tier acts as a write buffer only, to provide a layer of endurance flash, with all reads typically being sourced directly from the persistent storage flash layer.

In the hybrid architecture, server-attached magnetic disks are pooled to create a distributed shared datastore, which persists the data. In this type of architecture, you can typically get up to 40K IOPS per server host. However, in an All-Flash architecture, the flash-based caching tier is intelligently used as a write-buffer only, while another set of SSDs forms the persistence tier to store data. Since this architecture utilizes only flash devices, it can deliver extremely high IOPS of up to 100K per host, with predictable and low latency across the vSAN datastore.

A very common question from service providers is “Have we reached that cost tipping point for SSDs versus HDDs for persistent storage yet?” At the time of writing, for most use cases the answer is yes, particularly when you factor in the impact of space efficiency technologies, such as erasure coding and deduplication and compression, which are only available in an all-flash configuration. With the falling prices of flash and the space efficiency technologies accounted for, the All-Flash value proposition for persistent storage is often far more compelling than hybrid. However, that does not mean that All-Flash is



always less expensive. A hybrid configuration still maintains a place for some vSAN storage use cases. The key cost factors that must be considered include:

- On a \$/GB basis, SSDs are more expensive
- On a \$/IOPS basis, HDDs are more expensive
- When building a system, you must consider:
 - Usable versus raw capacity (the impact of erasure coding/RAID)
 - Space efficiency from dedupe and compression
 - Pricing differences between disk vendors
 - Power savings
 - Reliability and availability (annual failure rates)

In summary, key factors that are pertinent to the TCO of storage and to building a business case for a vSAN hyper-converged model include:

- We are seeing server flash become far more viable in terms of cost and cost per GB/IO
- The cost of rotating magnetic disks continues to fall
- The capacity of flash and magnetic disks continues to increase
- We find we have far more CPU cycles in servers, and as a result we can implement a much higher level of functionality in software instead of hardware

Finally, we are also seeing object storage and cloud economics significantly influence the overall cost of storage.



Identifying and Measuring the Operational Savings of vSAN

Storage has traditionally been one of the most expensive and complex infrastructure components within the data center. For most of the history of virtualization, however, it only appeared to become more and more expensive, largely due to the explosion in the growth of data and ever-increasing demands for higher storage performance by a similarly increasing number of applications. A primary aim of vSAN is to give service providers a storage solution that can contribute significantly to controlling these growing costs.

vSAN also fundamentally changes the management of storage, so that it becomes far more VM-centric. As a result, this helps to yield a reduction in OpEx cost impact, versus the upfront storage acquisition model. We estimate the reduction, based on our service provider industry experience, to be \$3-\$5 per gigabyte, with the ongoing costs of operating and maintaining a typical mid-range storage system at a run rate of between \$5 and \$8 per gigabyte, annually. These estimated operational costs consist of a number of elements, including:

- Data Center / Facilities
- Power
- Cooling
- Floor Space
- Physical Maintenance / Data Center Operations
- Vendor Support for Storage Systems
- Software Licensing

The vast majority of these costs, typically estimated at around three quarters or as high as \$6 per gigabyte, revolve around the ongoing management of storage.

With vSAN, VMware integrates the storage platform into the same core vSphere compute infrastructure. Therefore, storage and compute can be deployed and scaled as required, in a far more iterative and operationally simplified manner. On top of this single datastore storage pool, provisioning and management tasks occur at the VM-level, with each VM's storage characteristics, including availability, being maintained at the VM-level. This VM-centric approach enables vSphere operational teams to deal with storage within the same skillset and in the same operational procedures as other vSphere administrative tasks.

As a result, vSAN has a potentially significant impact on storage acquisition and operations, which must be addressed by service providers who are considering making the move towards operating a software-defined storage model. For this paper, the aim is to focus on providing a hands-on analysis of the most time-intensive storage tasks within service provider data centers, where vSAN is having a significant impact.

There are a number of operational tasks that are not being assessed as part of this example. For instance, the routine replacement of failed drives, storage fabric operations, and off-site backup procedures. Even though these operations are not being addressed here, they do make up a significant portion of the overall operational cost of maintaining a storage platform.

The following is a summary of the estimates on which the TCO example is based and defines what portion of the annual costs of ownership for storage are made up by the time and effort involved in specific tasks. This provides a clear business case for the VMware hyper-converged model, making it a clear choice for adoption by VMware Cloud Providers.

The following table outlines the frequency of occurrence of the operational tasks we are measuring as part of the example TCO assessment. The aim here is to demonstrate how these activities contribute to the ongoing time and effort, and therefore the operational costs of storage.



Table 1. Traditional Storage Array Management Lifecycle Tasks

Traditional Storage Array Management Lifecycle Tasks	Frequency of Activity	Impact on Existing Environment	Total Number of Annual Events Calculated
Scenario 1: Initial Platform Deployment	Twice annually	Low	2
Scenario 2: Modifying Service Levels / Workload Configuration	Twice weekly per 500 VMs	Large	104
Scenario 3: Increasing Capacity and Scaling	Monthly	Medium	12
Other Traditional Shared Storage Related Activities	Daily	Various	N/A
Other storage management tasks, such as replacement of failed disks, storage networking, and off-site backups, typically required by traditional storage arrays, as well as by vSAN	Weekly	Large	N/A

Table 1: Estimated operational overhead for scenarios when using vSAN versus a medium-sized traditional array-based storage system with 1000 VMs.

As addressed later in this paper, a deeper analysis of each of the three common activities being highlighted above demonstrates the total impact of vSAN on these specific operations leads to a 33 percent reduction in the total storage management OpEx. However, these are examples only, which differ significantly across environments, and there are other variables, such as administrator skill level, hardware, software tools, or automation that might already be in place with the service provider. All of these factors have a measurable impact on a total cost of ownership analysis, and must be considered.



Understanding Operational Transformation with vSAN

Adopting a Hyper-Converged Infrastructure (HCI) and vSAN approach can present a huge shift in how storage is deployed, operated, and maintained by a service provider when compared to traditional SAN storage arrays. For instance, at a high level, vSAN adoption is likely to impact tasks such as:

- Storage planning
- Design and deployment
- Growth and upgrades
- Maintenance and configuration management
- Service level assurance (monitoring, troubleshooting, and remediation)
- Optimization and performance tuning
- Capacity management
- Data management (data protection, lifecycle policies)
- Data validation (backup recovery, data integrity)

The reason vSAN adoption has such a significant impact on so many areas is that it leverages a converged platform for both compute and storage in a VM-centric mechanism, to transform complex and often manual tasks into a simple and clean single point of management with policy based automation.

Traditionally, storage is used by the vSphere platform to provide the required disk components of the environment, meaning that at least two different means of operations and expertise are required, with storage fabric networking often considered yet another area of expertise and operations. When employing vSAN, there will be only one single point of management to provision and manage the vSphere infrastructure, which includes all the necessary IT resources required for day-to-day operations, such as the provisioning and monitoring of virtual servers, storage, and networks.

However, like any other storage platform, vSAN storage architects must work with application owners and other IT teams to plan an appropriate storage infrastructure, and provision a suitable platform to meet the requirements of the service providers Service Level Agreements (SLAs). In most service provider architectures, deployment issues in this phase include providing sufficient:

- Capacity
- Performance
- Availability
- Continuity
- Data Protection
- Disaster Recovery

The very nature of traditional shared storage systems means that the many-to-one relationship of VMs to physical hypervisor requires multiple VMs to be deployed on a single class of storage, or to suffer the massively increased complexity and management overhead for having hundreds of LUNs presented to vSphere clusters. The vSAN operational model is significantly different from this traditional storage design and administration, which typically requires a multi-disciplinary approach to these skills.

As a part of the storage lifecycle, the service provider's storage administrators must operate and maintain the day-to-day delivery of storage services, which must meet SLAs with tenants. Typically, as part of these contracts, data must be protected and secured at all times. One of the main advantages of operating storage at the VM level with vSAN is that it provides and enforces a per-VM/VM-disk storage policy through the Storage-Policy Based Management (SPBM) mechanism. With SPBM, each VM is



provisioned with the required set of availability, performance, and storage capacity attributes required to meet its SLA. The SPBM mechanism then enforces policies automatically throughout the VM lifecycle.

However, when designing a traditional array-based storage system, storage is initially provisioned according to various policies required to meet the SLAs. These service levels must then be maintained through various management activities, including monitoring, troubleshooting, and remediation. This requires a team of storage administrators working behind the scenes preserving capacity, expanding as necessary, monitoring, and maintaining utilization, managing DR processes and performing regular testing of backups, and confirming their recoverability is validated.

This is one of the key operational advantages of vSAN, and where operational teams can see the most significant improvement over that traditional inflexible storage architecture. vSAN offers a per-VM capability to deliver to each workload its specific requirements, as opposed to delivering tiered services. The result in operational efficiency can release a significant saving in both upfront costs and operational costs over time, especially in environments where workloads have shifting requirements.

This document focuses on providing a detailed assessment of three key storage operational tasks. These tasks have been determined to have the highest impact in terms of both time and operational effort, from either a single task, such as initial deployment, or repeated tasks, which can be measured over the lifecycle of a traditional mid-range Fibre Channel storage array. The aim of this exercise is to estimate the upfront and ongoing savings in terms of both operational effort and monetary cost. This breakdown will then be compared with similar vSAN operations to help build a total cost of ownership analysis, which in turn forms part of a business case for decision makers for a VMware-based hyper-converged transformation project.

Scenario 1: Initial Storage Platform Deployment

In this scenario, we are comparing a simple deployment of a traditional medium size storage array system employing a Fibre Channel fabric, and an 8-host VMware ESXi™ cluster. In this example the Fibre Channel fabric is already in place, the ESXi hosts have host bus adaptors (HBAs) installed and cabled, and all physical equipment is racked, cabled and, ready for provisioning.

The comparative savings between vSAN and this traditional storage mechanism are generated by a reduction in storage array configuration tasks, which is likely to include a number of major elements, including:

1. The initial setup of system information such as interface settings, network settings, time settings, user and administrative access settings, and general Fibre Channel fabric configuration. In a vSAN configuration, this is pre-configured with the vSphere cluster settings. This activity has been estimated to be 40 minutes.
2. The initial configuration of disk and volume parameters on the storage array, including configuring items such as disk groups, disk group members, parity protection, storage tiers, and any tiering and caching policies required by the system. This configuration requires relatively infrequent activities. Therefore, such tasks are often slowed down by learning/remembering the activities to complete the task. This overhead has been included in this estimation of time. Based on the information gathered, this has been estimated to be 110 minutes, which varies depending on the number of disks, disk types, and the skill level and experience of the engineer performing the tasks.
3. The provisioning of initial volumes to ESXi hosts, including the configuring of multiple volumes to serve different requirements or SLAs, the masking of volumes to initiators, LUN masking, attachment of volumes to ESXi host, and verification/rework of configurations and host scanning. These activities have been estimated at 90 minutes, which varies depending on the number of physical host changes required.

For this scenario, we are estimating the potential for up to a 20x improvement in operational efficiency, assuming the deployment of an environment with sufficient network connectivity, bandwidth, and pre-installed disks within the vSAN cluster, as shown in Table 2.



Scenario 2: Modifying Service Levels / Workload Configuration

Scenario 2 demonstrates a particularly powerful feature of vSAN: providing a simple mechanism that allows the modification and reconfiguration of differentiated storage service classes over time. For instance, modifying the storage's performance or availability configuration, in order to tune for different application workloads. For vSAN, this is as simple as modifying the storage policy associated with the virtual machine or individual virtual disk, as compared with the traditional Fibre Channel SAN approach.

vSAN provides a way for service providers to easily and rigorously define and automate service-class based provisioning. Policies for service classes can be assigned based on business SLAs or required application service levels. vSAN policies are easy to dynamically change for any given workload, without migration or reconfiguration overhead. This approach through SPBM produces significant efficiencies, and enables service providers to benefit from establishing multi-tiered service levels.

SPBM is an intrinsic part of vSAN, which can help to manage, maintain, and automate provisioning to the point that the manual processes are significantly decreased. In addition, SPBM with vSAN lets service providers identify and immediately modify any initial configuration that fails to meet SLA requirements. This is a process that traditionally involved operational effort from multiple IT silos or administrative teams.

This simplified reconfiguration capability means storage resources can now be dynamically realigned and optimized to serve specific application requirements, across the entire application's lifecycle. For instance, any changes to the application requirements requires significantly less time and effort to accommodate. Additionally, when tenants wish to differentiate their storage or move applications between test/dev and production, or when an application no longer needs the highest levels of availability, traditionally this could be quite challenging, requiring the movement or migration of physically stored data from one volume to another. In contrast, vSAN allows tenants or service providers to immediately, non-disruptively, and almost effortlessly perform such reconfigurations.

As a result, the findings are likely to be significant when examining the process of reconfiguring a single application on vSAN, when compared to a traditional Fibre Channel SAN,. The reason for this is that with traditional array-based storage, the process of actually adjusting workload storage tiers or configurations can be a time consuming task. In the vSphere virtual infrastructure, such changes typically occur with VMware vSphere Storage vMotion[®], which is far superior to the capabilities of the physical infrastructure, but still requires the movement of data from one volume to another. This process can be slow, have a significant overhead on the array, and can also be slowed down by mismatches in datastore block size. vSphere Storage vMotion is required because it is not typically the case that a storage array can change the service level of a volume without serious reconfiguration.

To provide a comparison for this scenario, we have created an estimate based on the vSphere Storage vMotion time for 4x 25 GB virtual machines across two different datastores, using different block sizes. In this scenario on traditional storage, while the storage stays online during this activity due to the significant data movement required, it is far more disruptive than the transparent nature of a similar event through the modification of a VM-level storage policy performed on a vSAN datastore through the SPBM mechanism.

The vSphere Storage vMotion activities, outlined above, have been estimated at 120 minutes, which varies significantly depending on the array hardware, SAN design and configuration, VAAI support, and disk types, among other things. It might also be possible that during these operations the administrator could perform other tasks and duties. However, in order to provide a simple comparison, in this sample scenario it is assumed that the administrator is monitoring the operations throughout, and is therefore unable to perform secondary tasks.



Scenario 3: Increasing Capacity and Scaling

Increasing storage system capacity to meet the needs of ongoing growth is a frequent task for traditional storage array solutions, as capacity among shared storage systems is often allocated gradually over time because there is typically a need to conserve expensive disk resources.

In terms of operational management, we estimate that vSAN has a significant impact on the ongoing expansion and modification of array-based storage, because the demand for additional storage performance and capacity can change over time.

The reason for this is that vSAN brings about a significant change in the processes required to deploy additional capacity and performance within the storage platform. This is a very common task in larger environments, and more so when a traditional shared storage array is employed by both virtual hosts and non-virtual systems. In environments such as these, the typical behavior of storage administrators is to maintain the environment by a practice of carving out relatively small chunks of capacity on an ongoing basis, and expanding storage volumes as and when actual utilization demands it. In addition to this obvious operational overhead, there is also typically an impact on agility as well, when multiple operational teams must negotiate over the “how much”, “what”, “where”, and “when” of provisioning additional storage capacity. However, in order to break this scenario down into specific tasks, when the actual provisioning activities occur, storage administrators typically:

- Provide physical installation of disk hardware, typically scheduled during a maintenance window, to avoid disruption and potential downtime. Estimated as 45 minutes.
- Initialize disks, create disk groups, and add the new capacity to existing volumes or LUNs, or configure new volumes/LUNs with the new capacity. Estimated at 45 minutes.

Some service providers might avoid the full operational overhead of this problem by acquiring arrays configured at their full raw capacity upfront, and only carrying out allocation tasks (45 minutes). However, the disadvantage of this model of operations is both in the initial CapEx and in the OpEx cost of maintaining a significant proportion of unused capacity across the storage array.

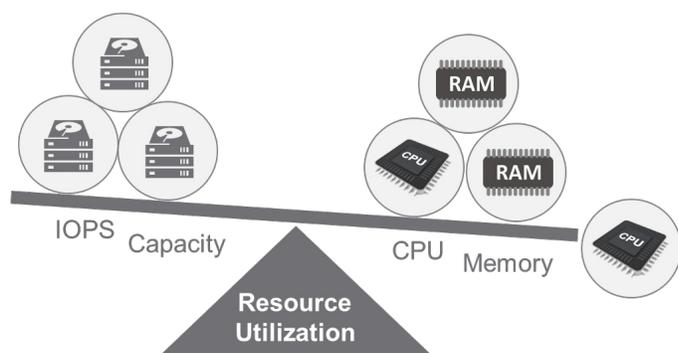
The VMware hyper-converged infrastructure model significantly simplifies these activities, when service providers are deploying standardized servers with consistent internal disk capacity in a building-block architecture. As hosts are typically more frequently provisioned than storage, the natural expansion of the physical server resources is likely to bring about more than sufficient storage capacity and performance, in a newer and more operationally efficient way. That is not to say that this new smoother and more gradual means of provisioning infrastructure does not require appropriate planning. It is also noteworthy that when a service provider does not wish to expand their storage, new hosts can still operate in a vSphere cluster without employing their local storage resources. The datastore can be expanded later by selecting the disks you want to use with vSAN. When adding additional hosts to an existing vSAN cluster, this newly added capacity is made available immediately on the vSAN datastore, allowing the provisioning of new workloads within minutes.

A final consideration for increasing capacity and scaling is the balancing and maximizing of compute and storage resource utilization. One of the key design factors associated with a hyper-converged architecture is the balancing of compute and storage resources – so that disk resources, CPU, and memory resources are consumed and exhausted in parallel, as far as possible.

For instance, the total cost of ownership of any hyper-converged infrastructure would be significantly increased if 95 percent of usable storage resources have been exhausted on hosts when only 35 percent of compute (CPU and memory) is being consumed. This can then be exacerbated significantly when a large-scale or web-scale platform is being designed, and a disproportionate amount of resources are not being effectively utilized across hundreds, or even thousands, of hosts.



Figure 5. Linear Scaling to Maximize Efficiency and Balance—Compute and Storage Resource



The following table provides a summary of the estimations and assumptions of time and effort outlined in the sample scenarios above.

Table 2. Storage Activities Comparison Summary

Administrator Activity	vSAN / Estimated Per Operation	Traditional FC SAN / Estimated Per Operation	Estimated Improvement Realized Through vSAN
Scenario 1: Initial Platform Deployment	12 minutes	240 minutes	20x
Scenario 2: Modifying Service Levels / Workload Configuration	3.5 minutes	120 minutes	34x
Scenario 3: Increasing Capacity and Scaling	< 2 minutes (1m,26s)	90 minutes	45x
Other Traditional Shared Storage-Related Activities	N/A	N/A	N/A
Other storage management tasks, such as replacement of failed disks, storage networking, and off-site backups, typically required by traditional storage arrays, as well as vSAN	N/A	N/A	N/A



These are just a few examples of how vSAN can reduce the sometimes complex tasks associated with traditional storage administration. In these three examples, we have calculated an overall mean improvement in operational management efficiency of 33x.

For the operational tasks outlined in these scenarios, specifically tasks that involve provisioning storage at an explicit service level, the service provider with traditional array-based storage might often spend a significant amount of time and effort in planning the multi-disciplinary activities. These activities, such as determining requirements and modeling for capacity and storage capability between application owners and infrastructure administrators, might with some service providers be more rigorous, predetermined, and automated. However, the reality is that this is rarely the case, and such determinations typically occur over time, as applications are provisioned, or as business requirements change.

Finally, many storage tasks remain unchanged by the adoption of vSAN. For instance, storage management tasks, such as the replacement of failed disks, storage networking, monitoring, and off-site backups, account for the remaining total operational time associated with storage administration. As you would expect, many of these activities are typically required by vSAN in an operational model similar to that required for traditional storage arrays.



Total Cost of Ownership Analysis

The total cost of ownership (TCO) is the measure we use to compare storage infrastructure platforms. The TCO incorporates the purchase and support costs of the storage solution, along with ongoing operational and management expenses. The operational efficiency built into the storage stack can significantly affect the bottom line. Once you have procured and implemented the platform, the operational costs for administration and maintenance can easily balloon to between x2 and x4 of the CapEx. A solution that streamlines and automates routine maintenance tasks can increase uptime and save the service provider significant time and money.

In this limited study, we use the results from the previous operational scenarios and examine the operational expense of the two platform options, using the three scenarios to provide a hypothetical TCO comparison. In addition to operational costs, TCO must also be broken down into acquisition costs (CapEx), which is address later in this document.

Most service providers require the TCO analysis, alongside a return on investment (ROI) investigation, to justify their business case to move from a traditional storage model to a VMware HCI solution. However, understanding how vSAN differentiates itself from traditional storage solutions is key to understanding the value proposition, and therefore the business case.

When choosing a storage platform, considering all costs, both acquisition-related and operational, is essential. The service provider must account for not only the cost of acquiring the hardware, licenses, and software, but also the costs associated with the time a system administrator devotes to maintenance and management tasks within each environment. Since system administrator time can be more valuable when used on strategic IT initiatives that deliver a competitive edge for their organization, instead of routine maintenance, it is always beneficial to reduce the operational overhead of administration and its associated cost. Because these operational costs add up over time, they can become a significant proportion of the overall costs for a data center. As a result, when comparing the TCO of vSAN with traditional storage arrays, both CapEx and OpEx are impacted and must be accounted for as part of the business case.



Calculating the Cost of Operations

To illustrate how these time savings can affect a service provider's bottom line, we have calculated the operational cost savings for a service provider that chooses vSAN over a traditional array-based storage solution. These calculations take account of tasks that must be repeated many times throughout a typical four-year period, considered to be a normal hardware and software refresh lifecycle. These calculations have assumed the tasks will be completed by a senior storage administrator, and we have calculated costs based on that individual's salary, plus benefits.

The average national base salary for a senior storage administrator was found to be \$88,599, and total compensation estimated to be \$126,662, based on information gathered from salary.com. In this example, the total compensation includes base salary, employer contributions for bonuses, social security, 401k and 401b, disability, healthcare, pension, and paid vacation time. We calculated the average cost per minute for a senior storage administrator's time at \$1.02, based on 52 40-hour weeks.

We then estimated the number of times the storage administrator would need to carry out these tasks per four-year period, for each of the three scenarios, based on industry experience, to come up with reasonable estimates of maintenance events, storage additions, deployments, and so on. In Table 3, we present the assumptions we used to calculate the number of tasks for the cost comparisons.

Table 3. Estimated Operational Costs (3 Scenarios)

Scenario / Activities	Total Events Per Year / FC SAN & vSAN	Cost Per Event / FC SAN	Total Cost Per Year / FC SAN	Cost Per Event / vSAN	Total Cost Per Year / vSAN
Scenario 1: Initial Platform Deployment	2	\$244.80	\$489.60	\$12.24	\$24.48
Scenario 2: Modifying Service Levels / Workload Configuration	104	\$122.40	\$12,729.60	\$3.57	\$371.28
Scenario 3: Increasing Capacity and Scaling	12	\$91.80	\$1,101.60	\$2.04	\$24.48

Table 3: Estimated operational cost savings based on these scenarios when using vSAN versus a medium-sized traditional array-based storage system with 1000 VMs.

In the following table, we have calculated the two-year, three-year, and four-year total cost of management operations, as the sum of the platform's operational cost, of the three scenarios we evaluated for the hypothetical 1000-VM data center.

**Table 4. Comparison of Total Cost of Three Scenarios Over Various Life Expectancies**

	Array-Based Solution in \$	vSAN Solution in \$	vSAN Management Operational Saving in \$
2-year CapEx (from 3 evaluated scenarios)	\$28,641.60	\$840.48	\$27,801.12
3-year operating expenses (from 3 evaluated scenarios)	\$42,962.40	\$1,260.72	\$41,701.68
4-year operating expenses (from 3 evaluated scenarios)	\$57,283.20	\$ 1,680.96	\$55,602.24

As shown in Table 4, the results show that when compared to a traditional shared storage array, the VMware vSAN platform provides a significantly lower operational management cost, which in turn can result in a lower TCO for the VMware solution, when considering the three scenarios we have outlined.

These three scenarios demonstrate only a small subset of the typical storage operational requirements of a service provider, but other studies of cross-industry IT spending show that annual operational expenses are typically between two and four times capital expenses. This means the impact of operational cost savings for vSAN might be multiplied well beyond the totals for the three tasks we have included in this analysis. As a result, VMware Cloud Providers might find that additional features of vSAN, such as iSCSI block storage or tight integration with VMware vRealize Operations™ and VMware vRealize Log Insight™, could lead to even greater management operational savings.



Building a Business Case for vSAN

Building a business case for a change of storage platform strategy is based on selecting the correct storage infrastructure for the tenants' applications. Understanding which technology works with their applications, providing for the consumer's availability and performance requirements, and meeting the data storage needs that the tenants have today and are going to have going forward, is key to the decision making process.

The challenge for service providers is to see beyond the short term, and *what's new and trending* in the storage market, and to arrive at a strategy to achieve a resilient and high performance storage infrastructure for their tenants in the long term at a cost acceptable to both provider and consumer. This obviously requires an objective approach, with a clear assessment of the components that make up a storage infrastructure's cost, and a clear-headed assessment of the alternative ways in which storage can be delivered to their consumers, storage that provides the storage functionality that the market requires, but that also optimizes both CapEx and OpEx.

Therefore, this requires that several questions be considered as part of building the business case. These might include:

- Which technology enhances application performance?
- Which technology provides greater data availability?
- Which technology can be deployed, configured, and managed quickly and effectively using available, on-staff skills?
- Which technology provides greater, if not optimal, storage capacity efficiency?
- Which technology enables storage flexibility; that is, ability to add capacity or performance in the future without affecting the applications?

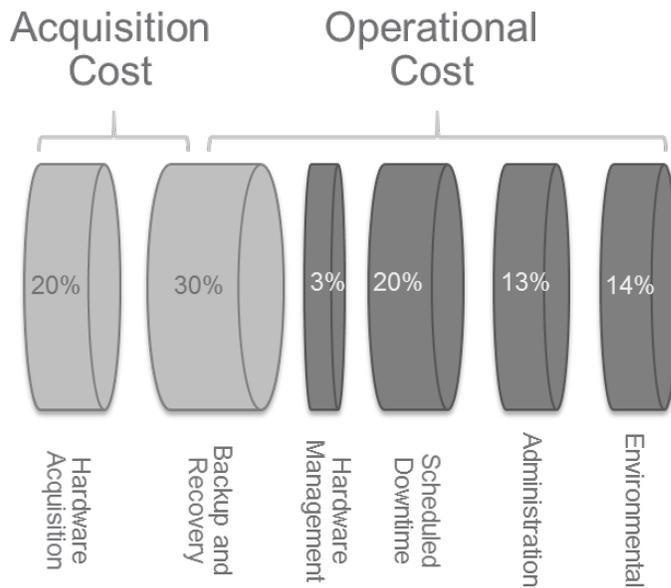
To these questions, service provider architects and decision makers also add another: which technology fits with available budget?

Ultimately, hardware costs and software licenses, plus leasing costs and a number of other factors, all contribute to the acquisition expense of storage infrastructure, which service providers typically seek to spread over the increasingly lengthy lifecycle of the hardware and software license agreements. It is noteworthy that most storage hardware ships with a three-year warranty and maintenance package, and increasing that agreement when it ends typically costs as much as purchasing a new array.

In addition, as outlined previously, hardware and software acquisition expense typically only accounts for a relatively small amount of the estimated annual cost of ownership of the storage infrastructure. According to leading analysts, the cost to acquire (CapEx) is overshadowed by the cost to own, operate, and manage (OpEx), as shown in the following figure.



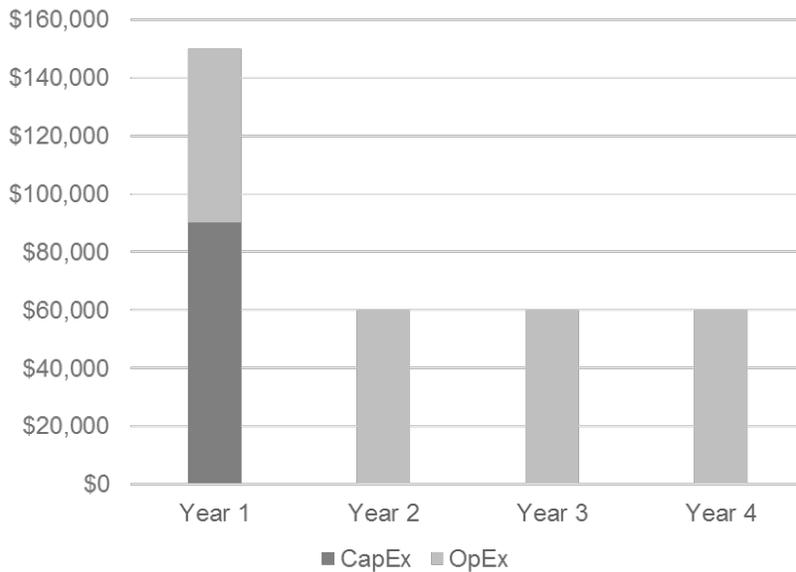
Figure 6. Typical Storage Cost Factors



Source: Multiple storage cost of ownership studies from Gartner, Forrester, and Clipper Group.

The figure above illustrates how the storage budget is distributed, and the following figure looks at the OpEx costs over the expected lifecycle of the solution. As shown in Figure 7, OpEx is measured at approximately triple that of the initial CapEx investment over the four-year period. Therefore, ignoring OpEx is a big mistake.

Figure 7. OpEx – Over a Four-Year Period



Source: Estimated using Gartner CP Profiles Database and IT Key Metrics Data 2014 Report.



As clearly visible from Figure 6, a storage system total cost of ownership includes far more outlays than the initial obvious costs on hardware, software and disk capacity. Taking this holistic view of these various factors affecting the service provider's budget over the storage infrastructure's lifecycle facilitates a clear assessment of the TCO, and ultimately a more informed purchasing decision. These factors include, but are not limited to the following elements:

- Cost to buy, implement, run, cool, and expand the storage system
- Manage, integrate, and test the storage system
- Confirm storage system reliability and availability
- Downtime incurred, caused by routine maintenance or unexpected equipment failure
- Vendor lock-ins associated with storage systems that use proprietary components
- Licensing software at time of purchase and upon each capacity increase

As shown in the preceding two figures, Gartner and other analysts suggest that management and operational costs are the real driver behind the total cost of ownership associated with storage infrastructure. As shown in the following figure, storage infrastructure TCO is calculated by combining both CapEx and OpEx costs and dividing that figure by the product life expectancy.

From a simplified standpoint, the total annual cost of ownership of storage can be calculated with the following formula.

Figure 8. Simplified Annual Total Cost of Ownership

$$\frac{\text{Cost of Device} + \text{Warranty \& Maintenance Agreement} + \text{Software License Agreement}}{\text{Years of Service}} = \text{CapEx Cost}$$
$$+ \text{Annual OpEx Cost (Typically 2 to 4 x CapEx)} = \text{Total Cost of Ownership (TCO)}$$

In many cases it is the mixed or heterogeneous nature of storage infrastructure that contributes to, and increases, the difficulties associated with the unified management of a storage platform. While this is true for a number of service providers, the challenges associated with the heterogenous nature of many storage infrastructures do not justify the cost of replacement with a homogeneous suite of hardware. This heterogeneous infrastructure has emerged in most cases as service providers have tried to deliver choice within their data centers, or in some cases as a result of continually trying to leverage a new, best-of-breed technology each year, in order to attract forward-thinking consumers. This situation has been compounded by storage hardware vendors who have failed to diversify their product offerings enough to meet the varying storage needs of different workload types, data, or storage delivery protocols. Even those vendors who do offer a full range of storage hardware solutions typically do not offer any common management tools across all solutions, especially when some products have been acquired as a result of the vendor's acquisitions.

Two other factors to take into account are *capacity allocation efficiency* and *capacity utilization efficiency*, both of which can have a significant impact on the bottom line when it comes to storage infrastructure and its cost of ownership. Allocation efficiency is the measure of how efficiently storage capacity is allocated. Utilization efficiency refers to the placement of the right data on the right storage, based on factors such as frequency of use, consumer performance expectations, or service-level driven availability requirements.



In the past, with a traditional shared storage model, storage administrators tended to deploy capacity in a sub-optimal manner, often purchasing Tier 1 storage to host customer data from applications that did not require the expensive and high performance attributes of the hardware provided.

However, the movement to flatten storage infrastructure, as seen by applications such as Hadoop and also the model employed by many of the hyper-converged software-based storage models, including vSAN, is eliminating many of the benefits of tiered storage altogether, alongside the negative impact of capacity allocation efficiency and capacity utilization efficiency. Tiered storage, when employed correctly, was supposed to enable the placement of data on the performance, capacity, or cost-appropriate storage, and as such, reduce the overall cost of storage ownership. A flattened storage model simplifies this approach through both cost of acquisition and operational management, so do these benefits outweigh the value of tiered storage, and consign the tiered model to the history books?



How to Build a Cost Comparison for vSAN

The first key component to building a cost comparison is understanding the requirements, which at a minimum include the design factors listed in the following figure.

Figure 9. vSAN Basic Design Factors

General Storage Requirements	VMware vSAN System Technical Considerations	VMware vSAN Hardware Decisions That Impact Cost
<ul style="list-style-type: none"> • Capacity • Performance • Availability 	<ul style="list-style-type: none"> • Number of Failures to Tolerate • Number of Disk Groups per Host • Flash Capacity Sizing • Memory & CPU • Network 	<ul style="list-style-type: none"> • Number of HDDs / SSDs Required • Type of HDDs / SSDs • Type of Controller • Type of Networking

For more information on designing a vSAN environment, refer to the following vCAT-SP documentation: <http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vcat/architecting-vmware-virtual-san-62-for-vmware-vcloud-air-network.pdf>

When building a cost comparison, always compare *like* storage systems. Comparing for capacity, performance, and availability is key for demonstrating a like-for-like comparison. In addition, take the following factors into consideration:

- Employ only supported hardware components, based on the latest vSAN Compatibility Guide (VCG)
- Factor in operational savings
- Assume conservative consolidation ratios when calculating capacity and IO requirements based on assumptions, as opposed to empirical data
- Include savings on Fibre Channel components where applicable
- Factor in hardware refresh cycles into cost comparison
- Allow for support, warranty, licenses, and maintenance costs as part of the cost comparison
- vSAN policies are assigned on a per-VM or per-VM-disk basis. For instance, one VM with FTT=2 does not require the entire cluster to be FTT=2 (see the following section)
- Calculate the cost of power, cooling, and floor space as part of the comparison
- Consider \$/IOPS in addition to \$/GB, particularly when comparing to a low-end storage solution with no flash
- The percentage of storage utilization must be factored into effective \$/GB paid for storage



Understanding the Impact of Number of Failures to Tolerate

As highlighted previously in this paper, vSAN utilizes Storage Policy-Based Management (SPBM), on a per-VM or per VM-disk basis, to assign attributes to workloads. There are a number of capabilities employed by vSAN that can be assigned to policies, linking their attributes to a VM or VM-disk, including:

- Number of disk stripes per object
- Object space reservation
- Flash read cache reservation
- Force provisioning

However, one specific capability, the *Number of Failures to Tolerate (FTT)*, has a significant effect on storage capacity and therefore cost, so it is addressed here. For more information on assigning storage policies in vSAN refer to:

<http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vcat/architecting-vmware-virtual-san-62-for-vmware-vcloud-air-network.pdf>

The Number of Failures to Tolerate capability addresses the key customer and design requirement of availability. With FTT, availability is provided by maintaining replica copies of data, to mitigate the risk of a host failure resulting in lost connectivity to data or potential data loss. The FTT policy works in conjunction with VMware vSphere High Availability to maintain availability and provide consistent and near continuous uptime to workloads.

In its most basic form, the FTT policy provides mirrored (RAID1) availability. For instance, a FTT=1 supports $n+1$ availability by providing a second copy of data on a separate host in the cluster. However, as you would expect, the resulting impact on capacity is that it is doubled. For example, if you deploy a virtual machine with a 100 GB disk, and copy 20 GB of data to the VM, without any further configuration of space efficiency technologies, the following impact on sizing would result:

- FTT=0 Results in 20GB of used capacity (not recommended)
- FTT=1 Results in 40GB of used capacity ($n+1$)
- FTT=2 Results in 60GB of used capacity ($n+2$)
- FTT=3 Results in 80GB of used capacity ($n+3$)

However, as mentioned, this is FTT in its most basic form. If you employ space efficiency technologies, such as erasure coding and deduplication and compressions in an all-flash configuration, erasure coding will yield at least a 33 percent space saving, and deduplication and compression, which is more difficult to calculate because it is dependent on workload type, could yield another 30-40 percent saving. The use of these space efficiency technologies in an all-flash configuration is one of the reasons why the all-flash value proposition is so compelling. Also, consider that all VMDK objects are by default sparse (thin), with no performance benefit of pre-allocating space.

vSAN license metering for VMware Cloud Providers is calculated based on consumed capacity of the vSAN datastores, irrespective of all other factors, such as the number of hosts in the cluster, number of VMs, or type of data stored. The metered value on which the cost is calculated is based on consumed datastore capacity, which is derived by subtracting **freeCapacityB** from **totalCapacityB**, leaving us with **Used Capacity**. As a result, the FTT policy not only has a significant impact on capacity sizing, but also has an impact on OpEx-based VMware Cloud Provider Program licensing costs.



Does this mean that VMware Cloud Providers pay for replica copies of data provided for redundancy? Yes, although it is not quite as straightforward as that from a cost comparison perspective with other solutions, because the following storage capacity efficiency technologies have a big impact on the **Used Capacity** being reported:

- Deduplication and Compression (All-Flash Only)
- Erasure Coding (All-Flash Only)
- Default Sparse VMDK Objects

Therefore, while the Number of Failures to Tolerate policy must be considered when designing and preparing a business case and total cost of ownership analysis for vSAN, a VMware Cloud Provider will *not* be doubling or tripling the storage costs for replica data when these space efficiency technologies are taken into account.

For more information on storage policies, including the number of failures to tolerate, erasure coding and deduplication and compression, please refer to:
<http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/vcat/architecting-vmware-virtual-san-62-for-vmware-vcloud-air-network.pdf>



Checklist for Decision Makers

For service providers looking to build a business case around the vSAN storage model, a clear assessment of available options is required to compare the total cost of ownership of the solutions available. This can then be used to calculate the improved return on investment on storage hardware, software, and operating costs.

At a high level, CapEx cost reductions derive from the removal of proprietary hardware and from the ability to employ commodity-type server components (within the constraints of the VMware Compatibility Guide), which use local server disks intelligently to maximize performance, efficiency, and availability of customers' applications.

Additionally, OpEx cost reductions associated with vSAN are derived from the management of the allocation of storage capacity in a more agile manner, the management of utilization of capacity with automation, and the management of multiple infrastructure components in a unified and simplified way. With a vSAN software-defined storage solution, applying the right software services to storage infrastructure is a far simpler process. In short, the implementation of vSAN technology enables the unified management of a large number of moving parts that go to make up the combined vSphere platform and storage infrastructure.

The following is a list of additional items that might find their way onto a service provider's checklist, when a decision maker is evaluating storage offerings. There is, as always, a keen eye on reducing storage costs:

- Support with currently deployed infrastructure
- Support for planned hardware / software technology implementations
- Ease of installation
- Ease of storage capacity allocation
- Ease of storage service allocation
- Effective capacity efficiency (for instance, compression, deduplication, and erasure coding/RAID) capabilities
- Mirroring/replication service without hardware constraints
- Data migration between different virtual storage pools, without hardware constraints
- Common management of centralized and distributed storage components
- Failover/failback capabilities



Summary

As outlined in this white paper, vSAN is gaining momentum with service providers, in part because of its ability to simplify and automate time-consuming manual storage tasks and the resulting impact on storage operating costs.

However, VMware Cloud Providers will find that their gains in efficiency and operational overhead vary significantly, based on a number of factors, such as size and complexity of the provider, legacy hardware footprint, scale, cost of resources, and so forth. In this white paper, we have provided examples of how these operational efficiencies can be realized in common practice.

While vSAN reductions in CapEx costs, in most cases, stand to transform traditional storage procurement, we believe that the operational transformations that come about as a result of vSAN have the potential to have a much larger impact with service providers, through reductions in operational or ongoing running costs.

Figure 10. Example Annual Storage Management OpEx and CapEx Cost Analysis

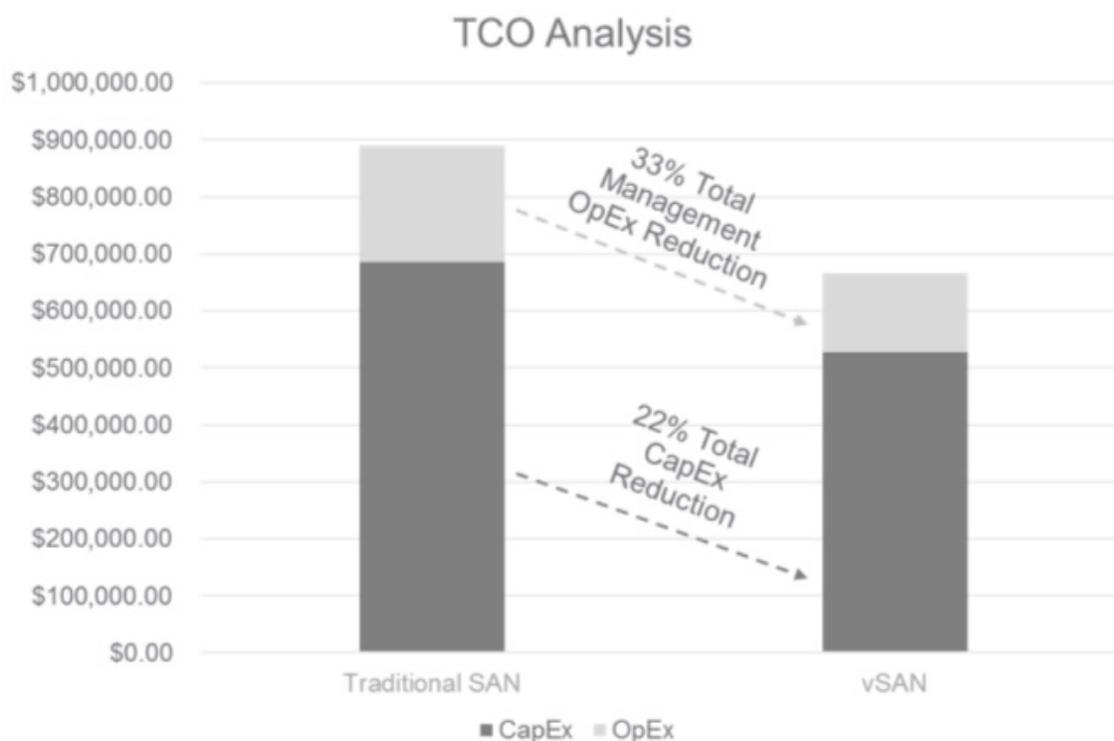


Figure 10 shows the estimated vSAN reduction in the tasks assessed, and the time, effort, and cost of the three scenarios provided in this paper. As shown previously, this leads to a one-third or 33 percent reduction in the annual management (labor) operational costs for the storage infrastructure, based on the sample scenarios.

The reason behind this is that vSAN delivers unmatched simplicity when compared to other traditional mid-range arrays or HCI storage offerings available today. It does this by enabling application workloads to be more efficiently managed, with changes rapidly implemented to meet the requirements of the service provider's consumers. Capacity, performance, or protection levels are all modified on the fly, with just a few clicks. It is this software-defined storage model, which provides a new approach to making better use of storage resources in the virtual environment. This is particularly the case when it comes to VM-centricity combined with an automated approach to policy management. vSphere operational teams gain the ability to handle storage with the same mind-set and in the same way as other vSphere-related tasks.



Based on the time and effort analysis provided in this white paper, vSAN operational costs are significantly less in each of the three scenarios outlined. As a result, this product represents one of the most significant cost saving opportunities for service provider data centers we have seen since the introduction of server virtualization by VMware 15 years ago. The growing adoption of vSAN by service providers signifies that VMware is delivering a competitive storage solution, which can scale, provide the high levels of efficiency and availability demanded by consumers, and not only simplify operations, but also dramatically reduce the time and effort associated with hardware acquisition.

Using the operational time and effort estimates of the three scenarios provided in this paper, we saw a positive task-level impact that ranged from 20x to 45x reductions in operational time and effort. This allows service providers to allocate resources to different task areas, and provide a far greater potential for cross-training of personnel.

We then converted the time and effort metrics into monetary values, and provided a cost comparison between vSAN and a traditional shared storage array. The comparison reflected the potential magnitude of the operational cost savings of vSAN, and its overall impact on storage OpEx, achieved through software-defined storage transformation.

As shown in Figure 10, our assessment suggests an average of 33 percent reduction in cost in monetary terms, thanks to the lower costs associated with reduced time and effort. In other words, vSAN requires only one-third of the management OpEx of traditional storage, based on the scenarios tested. However, this is not just an economic decision, but is also based on the product's inherent alignment with virtualized and container workloads, and the overall software-defined data center (SDDC) vision.

Disclaimer

The information provided in this white paper is based upon public information and the various sources listed in the following section. The paper also includes personal opinions of the author, all of which we believe to be accurate and reliable. However, as market conditions change, the information and recommendations are made without warranty of any kind.

Metrics and dollar values contained within this document are for guidance only. Actual savings vary based on a range of factors.



References

Additional information pertinent to this document and its topics:

Document Title	Link or URL
Software-Defined Storage and VMware’s Virtual SAN Redefining Storage Operations – Solution Brief by Taneja Group	http://www.vmware.com/content/dam/digital/marketing/vmware/en/pdf/solutions/vmware-taneja-group-solution-brief.pdf
VMware Software-Defined Storage: A Design Guide to the Policy-Driven, Software-Defined Storage Era	Martin Hosken (Published by Sybex) UK: https://www.amazon.co.uk/dp/1119292778 USA: https://www.amazon.com/dp/1119292778
Other	Various IDC and Gartner sources – as identified in the paper
Other	Various SABU customer-facing sales presentations