White Paper: **January 2023**

Lower Infrastructure Costs with VMware vSAN 8 and vSAN Express Storage Architecture

vmware®

Table of contents

Executive Summary

VMware vSAN 8 can lower Hyper-converged Infrastructure (HCI) TCO by up to 40% through a variety of new capabilities, and this is achieved in no small part by increasing usable storage capacity by up to 70%. These improvements are delivered through a new single-tier architecture, RAID-6 Efficiency with RAID-1 performance, improved compression, and compute and performance optimization.

In this paper, we seek to show how vSAN 8 with Express Storage Architecture can deliver lower storage TCO in several common deployment types, focusing primarily on capital cost differences between the previous generation of HCI – the original storage architecture, or OSA – and the current generation. However, vSAN 8 can also deliver significant cost savings as compared to traditional external storage for both CAPEX and OPEX; we provide a high-level summary in this paper, but it is not the focus.

Introduction

Organizations of all sizes and across industries are becoming more digital, from restaurants with digital menus and online ordering to national governments with digital tools to better serve their citizens and residents. With the digitalization of business comes both an explosion of applications and the data that they rely on. Data has been growing at roughly 50% per year for many years, and organizations have seen their storage needs grow to accommodate the explosion of data.

Unfortunately, IT budgets haven't kept pace with data growth – many budgets are flat or in the single digits - leading IT leadership to analyze ways to keep their storage spend in check. While IT decision-makers value many capabilities in a storage system, including performance and resiliency, both efficiency and low total cost of ownership consistently rank highly among the most sought-after features.

Many organizations have been turning to hyperconverged infrastructure, or HCI, to lower storage spend through both lower capital and operational costs.

VMware recently introduced vSAN 8 with Express Storage Architecture, our next-generation storage platform that raises the bar on HCI performance, efficiency, and resiliency.

vSAN Express Storage Architecture

VMware vSAN Express Storage Architecture (ESA) is a powerful next-generation distributed storage platform that offers incredible improvements in the total cost of ownership (TCO). Hardware and software costs can be reduced by as much as 40% for business-critical databases. For critical edge / remote deployments, incredible resilience can be delivered with only 1/3 of the storage devices required before.

What is vSAN?

VMware vSAN™ is VMware's software-defined storage solution, built from the ground up for vSphere virtual machines. It abstracts and aggregates locally attached disks in a vSphere cluster to create a storage solution that can be provisioned and managed from vCenter. vSAN is built inside of the hypervisor. It integrates with the entire VMware stack, including features like vMotion, HA, DRS, etc. VM storage provisioning and day-to-day management of storage SLAs can all be controlled through VM-level policies that can be set and modified on-the-fly. vSAN delivers enterprise-class features, scale, and performance, making it the ideal storage platform for VMs.

What is vSAN Express Storage Architecture?

vSAN Express Storage Architecture in vSAN 8 is a single-tier architecture optimized for high-performance NVMe-based TLC flash devices for both on-premises environments, and for the public cloud hyper-scalers. vSAN ESA provides a new way for vSAN to process and store data efficiently. A new, fast, and efficient data path paired with an efficient and resilient data structure gives our users the ability to store and access data fast and efficiently. Express Storage Architecture also provides an all-new snapshot engine that meets high levels of performance and scalability while remaining fully compatible for use with our ISV partners who provide backup solutions. And finally, vSAN Express Storage Architecture makes administering an environment simpler by improving how devices contribute as a storage resource to the cluster.

Traditional Storage Costs Challenge IT Budgets

Storage infrastructure can consume a considerable amount of the annual IT budget. While demands for IT infrastructure continue to grow rapidly, data growth – roughly estimated at 50% annual growth - is fueling increased demand for storage. At the same time, data feeds innovation, enabling the development of new products and services, and is often seen as valuable as gold for organizations. New products and services, which generate new revenue streams for organizations, create their own data and storage requirements. These newer applications often require highly performant, resilient storage and are often seen as business-critical. IT simply can't reduce the amount of data businesses generate but must rather move to accommodate these new applications' rapidly growing storage requirements. Due to the pace of data exceeding IT budget growth, companies of all sizes are looking for ways to reduce IT spending. Storage costs broadly fall into two distinct categories, capital expenditures – large, upfront purchases – and operational expenses – ongoing expenses to run and maintain storage infrastructure.

Capital Expenditures

Initial cost: The initial cost of traditional enterprise storage infrastructure can be quite high. This can be due to the need for purchasing hardware, software licenses, and storage fabric infrastructure. Procurement for storage often happens as a separate process that has a great deal of complexity in reviewing competition offerings, delaying the deployment.

Additionally, the cost of implementation, installation, and training can also add to the overall cost of the system. Storage often comes from a different vendor, requires specialty installers, and in some cases may require specialty rack or power infrastructure.

Expanding an array after initial procurement adds additional capital cost complexity. Needing to add capacity to a storage array with only one year of remaining life may prove to be a bad investment, forcing early product refreshes and new array procurement.

Operational Expenditures

Storage operations, including provisioning, performance monitoring, capacity planning, and recovery: All storage systems require common administrative tasks with varying degrees of frequency and complexity. Traditional storage systems typically require specialists to perform these routine operations; while the latest systems have become simpler to use, they still lack the automation and simplicity of other storage options on the market today.

Maintenance and Upkeep: Traditional enterprise storage infrastructure is often quite complex and requires regular maintenance and upkeep. This can include regular updates and upgrades, as well as periodic repairs and replacements. All these tasks can add to the overall cost of the system.

Power, Real Estate, and Cooling: Traditional storage systems have significant physical footprints and have higher power and cooling costs than servers.

Hyperconverged Infrastructure Lowers Storage Costs

Hyperconverged infrastructure, or simply HCI, delivers storage and compute for VMs and containers from the same x86 server platform running the hypervisor. Industry-standard components are used with direct-attached storage devices, supporting a wide variety of storage technologies, including the latest NVMe-based flash devices.

HCI provides both capital and operational savings over traditional storage systems. Capital savings include lower-cost access to storage from “server economics,” avoiding expensive, proprietary storage networking technology, granular scaling, and space efficiency features that reduce the upfront costs of storage versus traditional storage systems.

Capital Cost Savings

Cost-benefit of Server-based Storage

The server storage market is a highly competitive market that brings the newest, fastest, most performant drives to market many times years before array controllers can be adapted/updated/replaced/qualified to process them. “Server Economics” gives the customer the most choice and flexibility to control costs.

Networking Advantages

vSAN is powered by Ethernet rather than proprietary specialty networking or cluster interconnects. vSAN has evolved to take advantage of powerful

commodity 100Gbps ethernet ASICs and newer low latency RDMA transport options that are cost-effective and high-performance. vSphere Distributed Services Engine provides a complementary offload of networking and security services today to help make the most out of Ethernet.

Expanding capacity for vSAN clusters is easy. Adding drives to an existing host enables scaling up or adding hosts with more drives to a cluster enables scaling out. Performance scales linearly as clusters are expanded, and there is never a need for a forklift replacement. New hosts can replace old hosts in a cluster providing a "Ship of Theseus" upgrade path that is compatible with slower more gradual allocations of capital for upgrades.

Space Efficiency

Space efficiency technologies in enterprise storage play an important role in improving value and decreasing costs. VMware vSAN has several technologies in place to help improve storage efficiency.

Opportunistic. These space efficiency techniques are dependent on the conditions of the data and are not guaranteed to return a predetermined level of savings. vSAN offers several types of opportunistic space efficiency features such as compression, TRIM/UNMAP space reclamation, and thin provisioning.

Deterministic. These space efficiency techniques can be relied upon to deliver a guaranteed level of capacity savings. vSAN offers deterministic space efficiency capabilities through data placement schemes that are optimized for storing data in a resilient but efficient manner. This includes RAID-5/6 erasure codes.

Granular Scaling

HCI scales at a much more granular level than traditional storage, which helps avoid large opportunity costs. Traditional storage typically "scales up," with large storage purchases upfront and additional capacity added in large increments. A scale-up model encourages customers to procure large amounts of storage upfront, usually multiple years' worth of capacity at the outset. The result is thousands of dollars parked in a data center that may not be used for months or years – money that could have been used in other ways. HCI takes a different approach, allowing capacity to be added on an as-needed basis.

Scale-Out. HCI systems usually scale out, with capacity added as servers are added to a cluster. Clusters can scale as little as one node at a time, and many customers add capacity as needed every quarter.

Scale Up. HCI systems can also scale up by adding storage to existing nodes. This is a great option for workloads that tend to require more storage capacity than compute resources.

Scale Compute/Storage Independently (Disaggregation). Some HCI software, including VMware vSAN, allows the creation of storage-only or storage-dense clusters that can serve as storage to compute-only (vSphere only) clusters, essentially serving as external storage for those servers. There are several advantages to this model, as it allows storage and compute to scale

independently, so there are minimal "stranded resources," storage can be tiered, offering different levels of resiliency and data services, and serves as storage to servers that may not be able to handle much capacity while retaining the operational simplicity of HCI.

Operational Benefits of HCI

vSAN has several technologies to make rapid deployment easy. Cluster QuickStart and other guided wizards make standing up a cluster in a predictable, reliable, and consistent manner just a couple of clicks. In addition, vSAN can be automated as part of larger cloud deployments in the public cloud with VMware Cloud on AWS™ and several of VMware's hyper-scale partners, as well as in private clouds using VMware Cloud Foundation™.

Storage provisioning in vSAN is a simple act of creating virtual machines and assigning storage policies. There is no need to create LUNs, manage multiple VMFS volumes, and balance capacity and performance across data stores.

vSAN has advanced capacity management capabilities built in. Native reservations for rebuild and operational reserve can be enabled with just a few clicks. Capacity alarms are enabled and tunable giving warning when the cluster needs additional capacity. Capacity expansion can be as simple as putting new drives into the cluster and clicking to adapt them into a single pool of storage which is the vSAN datastore. vSAN works seamlessly with VMware Aria Operations and other enterprise infrastructure management solutions.

The native vSAN performance service is a powerful, distributed performance tracking service that provides end-to-end performance visibility. This is supplemented with vSAN IO Insight which allows a customer to integrate and profile workloads, and IO Trip Analyzer which allows a customer to quickly break down the latency of every layer end to end of the vSAN IO path. Additional APIs are available to expose vSAN performance to VMware Aria Operations™, Grafana, and other customer performance monitoring tooling.

Consolidating storage inside the servers has several power efficiencies, rack space, and cooling improvements. Power-hungry dedicated storage arrays, trays, and storage fabric switches are not needed for a VMware vSAN deployment.

vSAN 8 ESA Takes HCI to the Next Level

Superior Performance at a Lower Cost

In many ways, the Express Storage Architecture in vSAN 8 is so much more than just discrete enhancements introduced in a layer of vSAN. It allowed VMware to look at, and rethink how data could be processed more efficiently in the stack. That is exactly what happened. We've moved some data services to the very top of the vSAN stack to reduce the resources and context switching needed to perform these tasks. For example, data is compressed and encrypted at the highest layer to minimize process amplification. Checksums are reused in an innovative way to minimize duplicate efforts. And finally, the

vSAN ESA takes advantage of our new patented log-structured filesystem to write data efficiently by ingesting a lot of small I/O and writing it as one large I/O. It is the combination of these capabilities that make the vSAN Express Storage Architecture greater than the sum of its parts.

Performance Without Trade-off

Simplifying administration was an important goal for VMware in the development of the Express Storage Architecture. We wanted to ensure that customers could have optimal levels of resilience that are space efficient and offers the highest levels of performance. Thanks to the assortment of architectural changes the ESA introduces, this is precisely what the ESA in vSAN delivers. Going forward, administrators can now simply go with RAID-6 to gain enhanced levels of resilience, and predictable levels of space efficiency, all without compromising performance. When using the vSAN ESA, the performance of vSAN's RAID-5/6 erasure coding will be equal to that of a RAID-1 mirror. How this is achieved is in part, through the vSAN ESA's log-structured filesystem (LFS) and data structure that makes sure that a parity stripe is always fully updated. This removes the need for traditional read/write modification for a small block write.

Enhanced Compression

The Express Storage Architecture in vSAN 8 allowed us to redesign how compression is implemented. First, new techniques are in place to improve the potential data reduction of the compression feature to as high as an 8:1 compression ratio for every 4KB block written, which is a 4x improvement from the original storage architecture. The new design of the compression feature in vSAN occurs in the upper layers of vSAN as it receives the incoming writes. We can not only reduce the amount of CPU effort needed to compress the data, but this also reduces the amount of network traffic, as any replica traffic is always compressed. This highly efficient design allows for it to be enabled by default, but for workloads already performing application-level compression, it can be turned off on a per-VM basis through storage policies.

Every Device Contributes Capacity

Since the Express Storage Architecture found in vSAN 8 is an architecture that is optimized for today's hardware, it gives us the ability to build a high-performing storage pool using a single tier of storage devices. This means that all devices selected for use with vSAN will be contributing to the capacity and performance of a vSAN cluster. A user simply claims the desired devices on a host to create a single storage pool per host, which contributes its resources to the single data store for the cluster. The result is a simple user experience, better serviceability, and availability of the storage devices that power a vSAN cluster.

New RAID-5 Schemes in ESA

Storage policies have always allowed vSAN administrators to prescribe specific levels of resilience and space efficiency for VMs powered by vSAN. With the Express Storage Architecture in vSAN 8, we've introduced new levels of

intelligence into our storage policy engine to make the data more space efficient, and resilient, all while being easier to manage. When customers use a RAID-5 erasure code policy rule, ESA uses the host count of the cluster to automatically determine which of the two possible RAID-5 data placement schemes should be used. A new 4+1 RAID-5 data placement scheme spreads data with parity across 5 hosts and is used in clusters with 6 or more hosts. A new 2+1 RAID-5 data placement scheme spreads data with parity across 3 hosts and is used when a cluster consists of 3-5 hosts. If the cluster host count is reduced below 6 hosts for a sustained period (24 hours or more), ESA will automatically adjust itself to the 2+1 RAID-5 data placement scheme so that the cluster maintains at least one spare host for rebuilding data in the event of subsequent failures. Since the vSAN ESA delivers the same levels of performance using erasure coding as it does with RAID-1 mirroring policies, this means that customers will by default, be able to store data at a more space-efficient level, with a higher level of resilience that adapts to the size of the cluster.

Superior Economics of vSAN 8 with ESA

Example 1 - Resilient Edge Cluster

A 2-Node vSAN cluster needs to be configured with the ability to tolerate not only the failure of a host but also subsequent device failures using the nested fault domain capability, as well as heal from subsequent loss of a disk group within a host.

vSAN Original Storage Architecture (OSA) requires a minimum of 3 disk groups to configure a RAID-1 configuration within a host. 4 disk groups (3 + 1 extra for rebuilding) would be needed for self-healing within a host while using RAID-1. Using recommended practices of 2 capacity devices per disk group this means 4 cache devices and 8 capacity devices (24 devices in total for both hosts) would be required in each host to meet the design requirements.

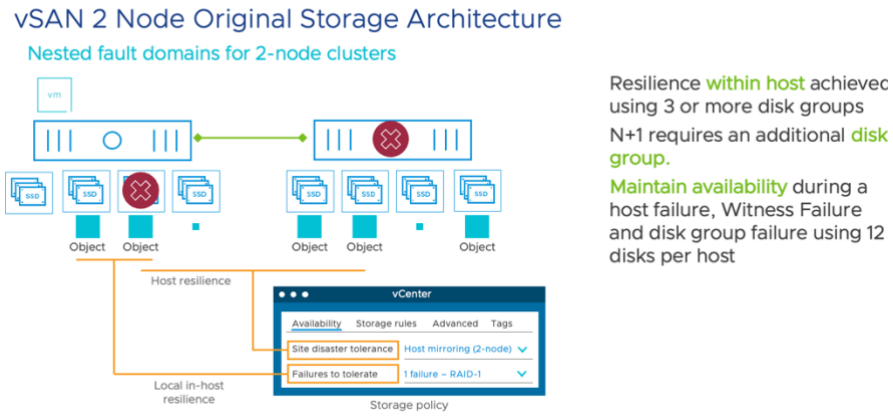

vSAN Express Storage Architecture (ESA) can perform RAID-5 with as few as 3 devices (2+1) within a host and could achieve self-healing within a host with as few as 4 devices. In total, this requires 8 drives total.

The vSAN ESA Cluster could meet the design requirements with 1/3 the number of devices as the vSAN OSA cluster.

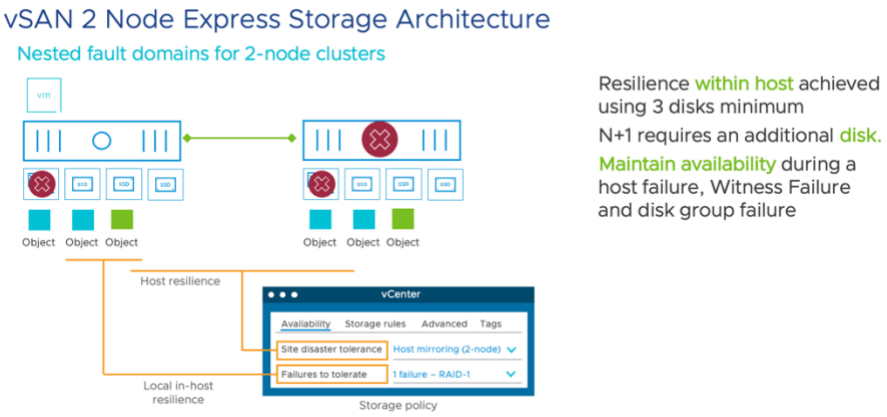

Example 2 – Core Datacenter OLTP Database Cluster

Several database clusters running a highly random write workload will be consolidated in the core data center. Due to high per-core application licensing, host consolidation is important to control ancillary costs. A requirement for the ability to tolerate the failure of 2 hosts is required to meet the durability SLA of the application owners.

Workload:

SQL – 100VMs, 50/50 read/write 8KB, 3.5TB each, 9000 IOPS per VM.

Sizing Requirements

Requirements	
Hosts Failures to Tolerate	2
Virtual Machines	100
Read/Write Ratio	50/50
I/O Size	8KB

Host Configuration

For the purposes of comparing VSAN OSA to ESA architecture similar configurations were chosen. It is understood that further optimization (such as larger flash devices on ESA) may be chosen to ultimately tighten TCO in the design.

Host Configuration	
CPU	2 sockets, 32 Cores, 3Ghz
Disks	24 x 3.84TB Mixed use NVMe (3 Disk groups for OSA)
Networking	100Gbps

Other Assumptions:

While CPU load was added to the sizing, this was not a CPU-bound working set.

A compression Ratio of 1.7 was chosen for OSA and 2 for ESA. This is conservative as ESA Compression is up to 4x more effective than OSA compression.

Designing for Failure – RAID level choice

For vSAN OSA Sizing calculations were tested with both RAID-6 and RAID-1 (mirroring). It was found that RAID-6 would require more hosts due to the write performance requirements.

vSAN ESA can leverage RAID 5/6 without any performance penalty for writes, and RAID-6 was found to require fewer hosts.

Results

vSAN Original Storage Architecture requires significantly more hosts to deliver the performance requirements. In total 18 hosts were required to meet the workload requirements with the hardware bill of materials provided.

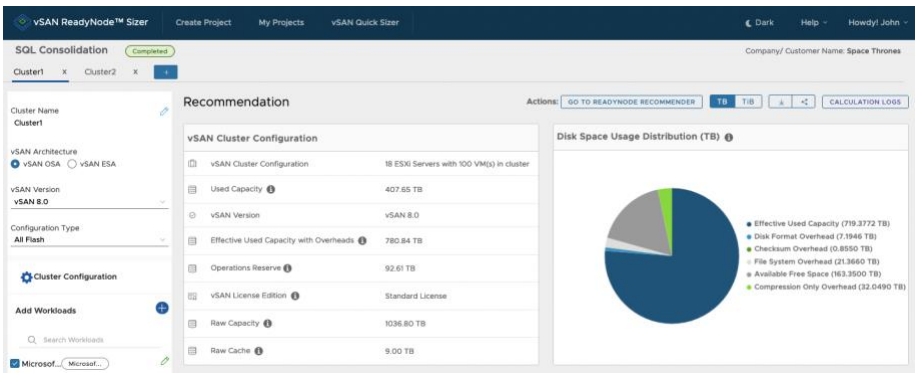

vSAN Express Storage Architecture was able to deliver the solution with 8 fewer hosts. In total, only 11 hosts were needed to meet the solution requirements.

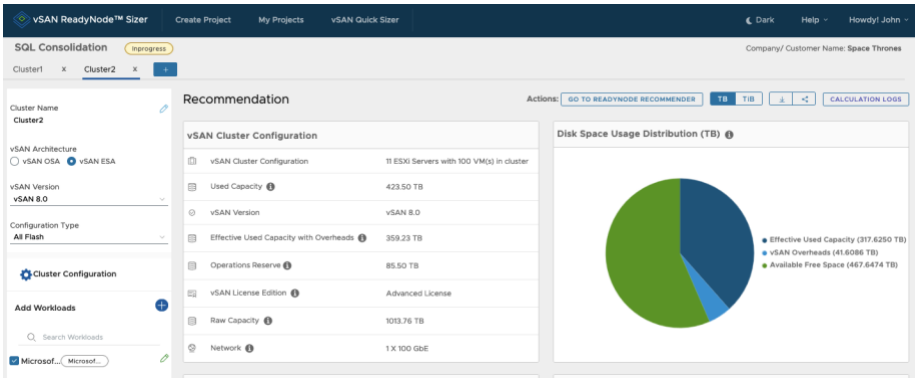

TCO comparison

7 fewer hosts demonstrate almost a 40% reduction in hardware requirements. In addition to fewer hosts, the software licensing requirement savings for database licensing is significant. Quantity

Product	Unit Cost	OSA Quantity	ESA Quantity	Hardware Savings %	Total Hardware Savings
Servers	$49,296	18	11	~40%	145,072

7 fewer hosts reflect 448 additional CPUs that will need an operating system and database licensing. Using Microsoft SQL 2019 list pricing these numbers can be calculated.

Product	Unit Cost	ESA	vSAN OSA	Savings
SQL Server Standard	$3586 - 2 core license	$1,262,272	$2,065,536	$803,264
Server Hardware	$49,296	$542,256	$887,328	$345,072
Total Savings				$1,148,336 40%

Conclusion

vSAN 8 with Express Storage Architecture, our next-generation storage platform raises the bar on HCI performance, efficiency, and resiliency. VMware vSAN 8 can lower HCI TCO by up to 40% through a variety of new capabilities, and this is achieved in no small part by increasing usable storage capacity by up to 70%. When looking to deploy your next VMware software-defined data center be sure to reach out to your account team and discuss how vSAN ESA can be used to run more workloads with fewer costs.

Item No: vmw-wp-temp-uslet-word-2021 8/21