Comparing vSAN Alternatives

verge.io

Broadcom's acquisition of VMware and increasing license costs are prompting many IT organizations to reevaluate their infrastructure choices. For those using VMware vSAN or considering a hyperconverged storage alternative, it is essential to understand how different vSAN architectures impact performance, scalability, data protection, and total cost of ownership.

This white paper compares VMware vSAN, Nutanix AOS Storage, and VergeIO VergeFS, focusing on their architectural differences and the resulting business implications.

Key Takeaways

1. vSAN Architectures: Application vs. Service

- VMware vSAN & Nutanix AOS: Operate as storage applications running inside VMs, adding latency, complexity, and overhead.
- VergeOS VergeFS: Runs as a service within the hypervisor, eliminating inefficiencies and delivering better performance on standard hardware.

2. Performance & Scalability Considerations

- VMware vSAN (ESA) mandates NVMe storage and 25Gbps networking to overcome architectural inefficiencies, yet struggles with latency under load.
- Nutanix AOS Storage uses a controller VM, which adds CPU/memory overhead and limits network efficiency.
- VergeOS VergeFS achieves sub-millisecond latency & 1M+ IOPS using realistic workloads, efficiently utilizing commodity hardware.

3. Data Protection & Recovery

- VMware & Nutanix rely on per-VM snapshots, limiting macro recovery and requiring manual VM-by-VM restoration after failures.
- VergeOS supports macro recovery, allowing entire data centers, tenants, or workloads to be restored instantly from one snapshot.
- Inline Repair Server (VergeIO only) ensures continued data access even during multiple drive failures, reducing reliance on backups.

4. Cost & Hardware Flexibility

- VMware vSAN requires vSAN Ready Nodes, limiting hardware flexibility and forcing expensive refresh cycles.
- Nutanix locks customers into a single hardware type, restricting scalability and choice.
- VergeOS supports a mix of legacy and modern hardware, reducing capital expenses and extending hardware lifespan.

5. Deduplication & Storage Efficiency

- VMware's ESA version lacks deduplication, forcing customers to choose between storage efficiency and feature access.
- Nutanix's deduplication negatively impacts write performance and is typically only enabled for VDI workloads.
- VergeOS' deduplication is global, inline, and has no measurable performance penalty, significantly improving storage efficiency and WAN-aware replication.

Conclusion: The Best vSAN Alternative?

Organizations seeking a high-performance, cost-efficient, and resilient vSAN alternative should consider VergeOS. Its integrated approach to storage, networking, and virtualization eliminates the inefficiencies of legacy vSAN architectures while offering superior performance, recovery, and flexibility at a lower cost.

For a deeper technical comparison, read the complete analysis below.

COMPARING VSAN ALTERNATIVES

Broadcom's purchase of VMware, combined with increasing license costs, is prompting numerous IT organizations to look for alternatives to VMware. Many of these organizations have either adopted VMware's vSAN technology or are contemplating transitioning to a converged architecture as they move away from their VMware configurations. This paper will examine the architectures of VMware vSAN and its alternatives, highlighting how these design choices affect storage performance, scalability, and cluster flexibility.

What is a vSAN

A vSAN, or virtual storage area network, is a broad term that applies to a category of storage solutions offered by multiple vendors, not just VMware's vSAN. This technology works by deploying software either as a virtual machine (VM) alongside the hypervisor or as a service within the hypervisor on each server (node) designated to contribute storage to the cluster. These VMs or services utilize the physical storage within their host servers, aggregating that storage into a shared pool accessible across the cluster. Other VMs then use this pooled storage for booting and data storage.

The Genesis of vSANs

Early vSAN technology evolved from software-defined storage (SDS) solutions that were modified to function in virtual environments. The second and later generations of these SDS solutions were purpose-built for virtualized environments from the outset. Examples include VMware's vSAN, Nutanix's AOS Storage, and VergeIO's VergeFS.

Operating in a virtualized environment presents unique challenges compared to the traditional dual-controller architecture used by most SDS and dedicated storage arrays. First, the storage software no longer has exclusive access to CPU, RAM, and network resources—it must efficiently share these resources with other application VMs. As a result, the software must be highly optimized to manage resources dynamically without impacting performance.

The second challenge is that the storage software and physical storage media are not dedicated to a single hardware platform. The first generation of vSAN solutions addressed data distribution complexities by limiting users to two nodes contributing storage to the cluster, effectively mimicking a dual-controller architecture within a virtualized environment. In this setup, data is mirrored between the two nodes. While simple to implement, this design introduces challenges during server or drive failures and imposes scalability limitations.

Subsequent generations of vSAN solutions distribute both the software and storage resources across multiple nodes within the virtualized environment. The vSAN software running on each node must coordinate efficiently, combining the physical capacity into a single virtual storage pool that is universally accessible.

This distributed approach means that storage operations—such as reading and writing data—must occur across multiple servers within the cluster. Ensuring an even distribution of data writes requires additional intelligence within the storage software, which can introduce performance inefficiencies if not properly optimized. However, when executed correctly, this approach enables high-performance, scalable read and write operations across the entire cluster.

COMPARING VSAN ARCHITECTURES

vSAN as an Application

Most vSAN technology is an add-on to the hypervisor rather than an integrated component. As a result, many vSAN solutions operate as virtual machines (VMs) running on the hypervisor. In this architecture, storage functions as a separate application within a VM, which runs on each node that contributes storage. Prominent examples of this second-generation vSAN technology include VMware vSAN and Nutanix AOS Storage.

While implementing vSAN as an application simplifies development for vendors, it introduces performance challenges. In this model, each storage request initiated by an application VM must pass through multiple layers before completion:

1. The application VM sends an I/O request to the hypervisor.
2. The hypervisor forwards the request to the vSAN storage VM.
3. The storage VM processes the request and sends a response back to the hypervisor.
4. The hypervisor delivers the response to the application VM.

This multi-step process introduces additional latency and creates a dependency on the hypervisor for all storage operations, which negatively impacts performance and efficiency. Furthermore, the storage VM must coordinate metadata updates with other storage VMs in the cluster, increasing network traffic and further straining resources. While vSAN solutions recommend dedicated networking connections for internode communication, the additional burden of metadata management and data resiliency operations compounds the inefficiencies of this architecture.

As detailed in the Architecture Impact on Hardware Selection section below, this design choice also restricts the hardware options available to customers. Vendors require the use of "vSAN-Ready" nodes—pre-configured servers with high-performance processors, large amounts of RAM, and premium flash storage—to mitigate the inefficiencies of this architecture. This requirement ultimately raises hardware costs for customers.

vSAN as a Service

Some vendors have taken a different approach by integrating vSAN functionality directly into the core operating environment. In this model, vSAN operates as a service rather than as a separate VM. Running vSAN as a service places the storage software on equal footing with the hypervisor and, in some cases, networking. This eliminates the need for dedicated storage VMs on each node, reducing complexity and improving efficiency.

With vSAN as a service, every node within the cluster has the necessary software to contribute storage to the global volume(s), if the customer chooses to enable this feature. More importantly, this approach eliminates the overhead and inefficiencies inherent in vSAN as an application.

This is the technique VergeIO uses in its data center operating platform, VergeOS. Within this platform, the storage software (VergeFS), hypervisor (VergeHV), and networking (VergeFabric) all function as integrated services. These services operate without requiring secondary communication paths, ensuring a streamlined, high-performance infrastructure. By integrating storage directly into the hypervisor, VergeOS eliminates the inefficiencies found in traditional vSAN architectures, resulting in superior performance and resource utilization.

Which vSAN Model is Best?

A vSAN as a service implementation model delivers higher performance and efficiency compared to the vSAN as an application approach. By reducing overhead, this model enables better storage performance on less powerful hardware, lowering costs for customers.

Additionally, some VMware alternative solutions support repurposing existing VMware hardware. Customers transitioning to a vSAN as a service solution can expect a measurable performance improvement, extending the useful life of their existing hardware while optimizing overall infrastructure efficiency.

ARCHITECTURE IMPACT ON HARDWARE SELECTION

VMware's vSAN is restricted to using "vSAN Ready Nodes," which are pre-configured to meet vSAN's requirements and attempt to compensate for some of its inherent inefficiencies. Additionally, vSAN has a history of deprecating hardware. For example, vSAN 8.0 mandates support for the "Express Storage Architecture (ESA)," which typically requires high-endurance NVMe SSDs and a dedicated 25Gbps network for optimal performance. Furthermore, vSAN 8.0 exclusively supports flash drives, preventing customers from utilizing hard disk drives for archiving or staging purposes.

Nutanix initially launched with a turnkey hardware-software bundle. While Nutanix still offers its own hardware solutions, it has expanded its ecosystem to include OEM platforms and recently announced support for vSAN-Ready nodes. However, the Nutanix website states that different types of nodes cannot be mixed. Once a customer commits to a specific server strategy, they must continue with it throughout the lifecycle of their deployment.

Nutanix Hardware Platforms: Specs for Nutanix, OEM & Partner Platforms.

VergeOS provides far greater flexibility by supporting a diverse range of node types within the environment. Hardware nodes are grouped logically, but VMs can access resources from any of those groups. Unlike VMware, VergeIO has never deprecated hardware. Customers can seamlessly run servers that have been in production for six or more years alongside newer hardware deployed within the last six months.

VergeOS's adaptability is driven by its integration of core infrastructure services—virtualization, storage, and networking—into a unified platform. Additionally, VergeOS leverages narrow AI to optimize resource allocation and provide a high level of abstraction from the underlying hardware. With VergeOS ioOptimize, customers can dynamically scale resources up or down as needed, enhancing overall efficiency.

VergeOS is the superior choice for customers who require the flexibility to deploy storage nodes of various media types and capacities and the ability to use dedicated compute and storage nodes. Its relatively lenient minimum requirements allow VergeIO customers to integrate a wide range of server hardware and storage solutions into a single VergeOS instance. Using its Virtual Data Center (VDC) technology, these resources can be made globally accessible or assigned to specific workloads, providing a level of adaptability unmatched by VMware or Nutanix.

DRIVE FAILURE PROTECTION

Data protection is the most critical aspect of any vSAN or storage system. Vendors' methods of protecting against failures significantly impact performance and architectural flexibility. Data protection encompasses safeguards against hardware failures, as well as protection from data corruption or accidental deletion.

vSANs present a unique data protection challenge: Storage media is distributed across multiple servers rather than housed within a single storage array. Generally, there are three primary methods for protecting data distributed across nodes: synchronous mirroring, synchronous replication, and erasure coding.

Synchronous Mirroring

When a new data block is written to a vSAN using synchronous mirroring, it is simultaneously written to two designated nodes. The application VM does not receive an acknowledgment of a completed write until both copies have been successfully written. This technique is straightforward and similar to how external dedicated storage arrays with dual controllers operate. It works well for small, remote environments requiring only two nodes. Still, it does not scale effectively for enterprise deployments, as only those two specific nodes are responsible for serving storage to the larger environment.

Some vendors attempt to address this scalability issue by enabling multiple two-node vSAN deployments within the same environment. However, this approach is functionally equivalent to deploying multiple standalone external SANs. Storage remains siloed within each pair of nodes, and moving a VM's data between them requires live storage migration capabilities, which many vSAN vendors lack.

From a resiliency standpoint, if one of the storage nodes or drives fails, the customer is left vulnerable, one failure away from total data loss. At that point, recovery must rely entirely on backups, which may lead to downtime and data loss.
Distributed Synchronous Replication
Distributed synchronous replication operates similarly to synchronous mirroring in that it maintains two copies of data, but it distributes the redundant copies across multiple nodes within the environment. This method ensures that redundant blocks do not reside on the same server, thereby enhancing fault tolerance and eliminating the scalability limitations of synchronous mirroring.

Because multiple nodes participate in storage operations, distributed synchronous replication generally delivers better read and write performance than synchronous mirroring. Additionally, while drive or server failures still pose a risk, the failure of an individual drive or server does not necessarily lead to data loss, making this approach more resilient.

Erasure Coding

While all drive failure protection methods introduce some overhead, one concern with both replication models is that they create a full duplicate of the data, doubling storage requirements. To reduce this overhead, some vSAN vendors implement erasure coding, a technique similar to RAID 5 or 6, commonly used in traditional storage arrays. Erasure coding uses a parity-based protection scheme to ensure continued access to data in the event of drive failure.

Like RAID, erasure coding allows users to configure how many drive failures the system can tolerate. However, increasing fault tolerance also increases the associated data overhead. Most organizations opt for an "n+2" protection level, which results in approximately 35% to 40% storage overhead, compared to the 100% overhead of synchronous mirroring or replication.

The primary drawback of erasure coding in a vSAN architecture is the significant CPU and memory overhead required to manage parity calculations. This computational burden can severely impact performance, particularly in environments where CPU and RAM resources must also be allocated to application VMs. To mitigate this, vSAN vendors enforce strict guidelines for erasure coding deployments. In most cases, these nodes require more powerful processors and increased memory compared to other nodes in the cluster. Additionally, many vendors recommend disabling deduplication using erasure coding, negating much of the anticipated storage efficiency gains.

Another primary concern is the impact on drive rebuild times. As flash drive capacities grow, rebuild times increase exponentially, often taking multiple hours or even days to complete. During this period, the environment remains degraded, leaving data at heightened risk. Finally, erasure coding requires a larger minimum cluster size, with Nutanix recommending at least five nodes for an erasure coding configuration.

Inline Repair Server

Each of the aforementioned protection methods has a critical limitation: They cannot protect against a total storage system failure. A storage system failure occurs when enough drives fail to exceed the system's ability to maintain data integrity or when network, controller, or software update failures render the storage system inaccessible.

In these cases, organizations must depend on backup software for recovery. However, recovering from a failed storage system can be quite challenging. Many organizations conduct full backups only once a day, leaving them at risk of losing up to 24 hours of data. Furthermore, restoring from backups requires available storage capacity, which may not be readily accessible in the event of a total system failure. This means customers must wait for the storage system to be repaired—potentially needing to order replacement drives or components—before starting the lengthy process of restoring data.

To tackle this challenge, some vSAN vendors deploy an inline repair server, which acts as a parity node containing an extra copy of the data. Unlike traditional backup appliances, this parity node is fully integrated into the vSAN environment and recognized as part of the storage system. When multiple simultaneous drive or server failures occur, the vSAN can recover lost data segments in real-time from the repair server, enabling applications to keep running without interruption.

Which Data Protection Method is Best?

Erasure coding appears appealing due to its theoretical 60% to 70% storage efficiency compared to synchronous mirroring and replication. However, its stringent hardware requirements, performance trade-offs, and limitations—such as the inability to leverage deduplication—eliminate much of its theoretical advantage. Additionally, the increasing affordability of high-performance NVMe flash drives (now under $150 per terabyte) diminishes the importance of the storage savings offered by erasure coding.

Of the two replication-based protection methods, synchronous replication provides the best balance of performance and scalability. Its simplicity ensures that more CPU and RAM resources remain available for application workloads, and it provides faster recovery times than erasure coding in the event of a drive failure. Furthermore, it is better suited to work alongside deduplication and snapshots without negatively impacting performance.

VMware, Nutanix, and VergeIO all support distributed replication models, except in two-node configurations. VMware and Nutanix offer erasure coding, but both solutions have significant limitations, particularly in the area of hardware requirements and performance impact. Neither VMware nor Nutanix supports an inline repair server, leaving customers vulnerable to data loss if drive failures exceed the protection level—especially in common two-node deployments.

VergeOS takes a different approach by forgoing erasure coding in favor of an inline repair server. VergeOS's ioGuardian integrates repair server functionality directly into VergeFS, and it is included with the software at no additional cost. With ioGuardian, VergeIO customers can continue accessing data and running applications even in the event of multiple simultaneous drive or server failures. By combining high-performance synchronous replication with the ability to withstand storage failures beyond the protected state, VergeOS offers the most resilient and efficient data protection strategy of the three solutions.

DATA FAILURE PROTECTION

Data failures can result from software or hardware configuration issues or accidental deletions, often caused by users saving an older file with the same name. While most vSAN solutions offer snapshot technology to help administrators recover from such disasters, many impose significant limitations on the number of snapshots allowed and the granularity of recovery.

Snapshot Depth

Because snapshots execute quickly, they can be taken frequently and serve as the fastest recovery point of any data protection method. Many organizations would prefer to rely solely on snapshots for data failure recovery. However, even a modest snapshot schedule—taking a snapshot every 30 minutes with one week of retention—requires the vSAN solution to manage 336 snapshots, which most are not designed to handle efficiently.

While taking snapshots every 15 to 30 minutes may seem excessive, it has become a best practice in the era of cyberattacks and ransomware. Ransomware can corrupt an entire environment in as little as 15 to 20 minutes. Recovering from a backup taken the previous night can result in unacceptable data loss, and snapshots taken too infrequently may still contain infected data—the ability to granularly "dial back" to a moment just before the attack is critical. With frequent and rich snapshot history, organizations can recover from a ransomware attack with minimal data loss.

Most vSAN solutions limit the number of active snapshots they can maintain due to performance concerns; each additional snapshot introduces overhead, slowing down the entire storage infrastructure. Historically, VMware vSAN was limited to 32 snapshots per VM. With version 8 update three, it now supports 200 snapshots per VM, which requires significant hardware investment. Additionally, as the number of snapshots approaches 200, production performance continues to degrade. Nutanix still limits snapshots to 35 per VM, an increase from its previous limit of 25.

As a result, VMware vSAN and Nutanix customers primarily use snapshots to support their backup process rather than as a replacement for traditional backups. Both solutions offer immutable or read-only snapshots, but this feature can be misleading. While these snapshots prevent ransomware from modifying the data within them, they do not stop the snapshots from capturing newly corrupted data. Additionally, the limited depth of snapshots makes them ineffective for full ransomware recovery.

Macro Recovery

Beyond snapshot depth, organizations looking to enhance or replace traditional backups with snapshots need macro-level recovery capabilities. Ideally, a vSAN should be able to restore an entire environment from a single snapshot. Organizations should also be able to restore data at the business unit, individual VM levels, and even down to a single file.

VMware ESXi is deployed in clusters, and its add-on solution, Cloud Director, allows customers to separate workloads into tenants, which VMware calls Organizations. Nutanix also supports multi-tenancy, referring to its tenants as Virtual Private Clouds. VergeIO also offers multi-tenancy, with its tenants known as Virtual Data Centers.

However, VMware vSAN and Nutanix AOS Storage can only take snapshots at the VM level. This lack of macro recovery capabilities creates challenges for performance and recovery speed.

For example, if a customer has three tenants, each with 300 VMs, VMware and Nutanix can use VM tags to group the VMs into several snapshot schedules. However, each schedule must manage 300 separate snapshot instances, contributing to the overall snapshot limitation. From a recovery perspective, returning one of these tenants to a specific point in time is cumbersome. Although administrators can filter snapshots using tags, VMware vSAN and Nutanix AOS Storage must still execute, in this scenario, 100 simultaneous VM recoveries.

In contrast, VergeOS has a distinct advantage. VergeFS, like VMware and Nutanix, can protect individual VMs and entire instances and tenants (Virtual Data Centers). Administrators can drill into these macro-level snapshots to recover specific VMs. In practice, VergeOS manages snapshots at the Virtual Data Center level by default, with VM-level snapshots being the exception rather than the rule. This approach enables VergeOS to handle significantly fewer snapshots than competing solutions while providing deeper recovery options.

The three vendors have fundamentally different philosophies regarding multi-tenancy. VMware and Nutanix treat multi-tenancy as a way to group VMs, but their snapshot and recovery models still operate at the individual VM level. VergeIO, however, treats multi-tenancy as a level of encapsulation through its Virtual Data Center technology, allowing for greater snapshot efficiency and streamlined recovery operations.

Granular Recovery

Granular recovery remains an essential feature of any storage solution. While macro recovery is critical for large-scale disaster recovery scenarios, granular recovery is used daily to help users recover lost files or roll back small but impactful changes. Granular recovery includes restoring entire VMs or recovering specific files within a VM.

All three solutions—VMware vSAN, Nutanix, and VergeOS—support full VM recovery. However, VMware vSAN does not allow administrators to recover individual files from a snapshot. Instead, they must restore an entire VM to a different location, ensure that the networking does not conflict with the production environment, start the VM, manually extract the required file, and then transfer it to the original VM. This process is cumbersome and time-consuming.

Nutanix and VergeIO both support single-file recovery by allowing administrators to mount a VM snapshot as a drive on the production version of that VM. This approach enables users to easily copy the required file between the two, eliminating the challenges associated with networking conflicts and data transfer times. However, Nutanix's implementation requires the Nutanix Guest Tools, and due to its snapshot depth limitations, its single-file recovery capabilities are less functional in practice.

In many cases, VMware vSAN's lack of a single-file recovery method and Nutanix's snapshot limitations force organizations to rely on backup and recovery software. While modern backup software can handle this task efficiently through intuitive graphical interfaces, it still has its challenges. The most notable limitation is recovery frequency—most backup solutions run on a scheduled basis, often once per day. This means it may be permanently lost if data is created or modified between backup intervals. Additionally, extracting a single file from a backup set requires time, as the backup software must retrieve and copy it back to the VM.

VergeOS stands alone in its ability to execute snapshots frequently without impacting performance. This enables administrators to extract data at any level of granularity, making VergeOS a superior choice for both macro and granular recovery.

DEDUPLICATION

Because of the potential for redundant data, deduplication is a necessary feature for any vSAN or storage system supporting a virtualized environment. However, vSANs face the challenge of executing a complex deduplication algorithm in an environment where computing and memory resources must be shared with application VMs.

VMware does not support deduplication in its new vSAN storage architecture, Express Storage Architecture (ESA); it is only available in the Original Storage Architecture (OSA). As a result, companies must choose between enabling deduplication or taking advantage of ESA's advanced capabilities, such as support for more than 32 snapshots and superior scalability. Under OSA, VMware claims that deduplication is inline, but this claim is subject to interpretation. OSA uses a two-tier architecture in which the first tier, known as the cache tier, is not deduplicated. Instead, it functions primarily as a large write buffer and read cache.

This design has significant ramifications. All data is written at least twice—first to the cache tier and then to the capacity tier. Deduplication is applied "inline" only when the data is transferred to the capacity tier. Frequently written but similar data, such as operating system images, may be written hundreds or even thousands of times to the cache tier. Another concern is that VMware's deduplication significantly increases metadata demands, and because metadata is stored on the cache tier, it reduces the available capacity for write buffering and read caching. This inefficiency impacts performance, forces companies to invest in larger, more expensive high-performance flash drives for the cache tier, and can accelerate flash drive wear.

Nutanix offers deduplication, but the company cautions that enabling deduplication alongside compression can negatively impact performance, particularly during write operations. Many Nutanix users follow a general best practice of enabling deduplication only for virtual desktop infrastructure (VDI) workloads, where the storage capacity savings outweigh the potential performance loss.

Of the three solutions, VergeOS has the most integrated approach to deduplication, having supported it since its inception. Like VergeOS' virtualization, storage, and networking, deduplication is built directly into the core OS. VergeOS' deduplication is global and inline, meaning all aspects of the environment, including compute and networking, benefit from its efficiency. There is no measurable performance penalty for using deduplication, and no specific workloads have been shown to degrade performance when deduplication is enabled. While VergeFS does incorporate a high-speed storage tier, this tier is used exclusively for metadata.

VergeFS also uses RAM for read caching. Because deduplication is integrated at the OS level, the cache benefits from the reduction in redundant data and can store three times or more data than competing solutions. As a result, the VergeFS read cache, despite being RAM-based, can retain recently accessed data for extended periods—ranging from double-digit minutes to even hours.

Another key advantage of VergeOS' deduplication is its WAN-awareness. Customers can replicate dozens or even hundreds of remote sites to a central location, transmitting only the unique data from each site. This capability significantly reduces bandwidth usage and replication times while ensuring efficient multi-site data management.

TIERING

Most vSAN solutions do not support multiple drive tiers. Supporting different drive tiers was once common in the early days of enterprise flash storage. Due to its high cost, organizations implemented small flash tiers and relied on tiering to migrate less performance-sensitive workloads to secondary hard disk tiers. However, inline deduplication for production storage and the rapid decline in flash drive costs have made single-tier storage more practical in most scenarios.

Recently, however, the industry has seen the emergence of high-density quad-level cell (QLC) flash drives, with companies like Solidigm delivering capacities as high as 122TB per drive. While QLC drives offer excellent read performance and acceptable write performance, they are not designed for sustained high write volumes, which can lead to premature wear. If the storage software supports intelligent tiering, these high-capacity QLC drives can be safely leveraged by automatically migrating workloads between performance and capacity tiers as data access patterns change.

VMware vSAN, Nutanix AOS Storage, and VergeOS all support live migration of VMs between multiple storage tiers.

COMPARING VSAN PERFORMANCE

Performance is one of the most critical factors when evaluating vSAN alternatives, as storage latency, throughput, and efficiency directly impact application responsiveness and overall infrastructure scalability. VMware vSAN, Nutanix AOS Storage, and VergeIO VergeFS take significantly different approaches to performance optimization, and these architectural differences lead to measurable variations in real-world results.

VMware vSAN, particularly in its Express Storage Architecture (ESA), attempts to improve upon the Original Storage Architecture (OSA) by requiring high-performance NVMe drives and a 25Gbps network to compensate for inefficiencies inherent in its design, where storage operates as a separate entity within the hypervisor. Even with these hardware improvements, vSAN struggles with latency under load, frequently exceeding five milliseconds. Additionally, vSAN's reliance on cache tiers means that all writes occur twice—first to the cache tier and then to the capacity tier—introducing significant overhead and reducing the lifespan of storage media. VMware's erasure coding implementation also imposes a heavy computational burden, further degrading write performance.

Nutanix AOS Storage faces similar challenges, as it, too, operates as a separate VM within the hypervisor. The Nutanix Controller VM (CVM) must process all storage requests, adding latency and increasing CPU and memory demands. Nutanix offers deduplication and compression but warns that enabling these features can negatively impact write performance. Many Nutanix customers limit deduplication to VDI workloads, where the benefits outweigh the performance trade-offs. While Nutanix has made efforts to improve storage efficiency, real-world testing indicates that it struggles to fully utilize available network bandwidth, often achieving less than 30 percent efficiency.

VergeIO VergeFS takes a fundamentally different approach by integrating storage as a native hypervisor service rather than running it as a VM. This design eliminates the unnecessary CPU and memory overhead found in VMware and Nutanix architectures. Benchmark tests of VergeOS 4.13 demonstrate its ability to achieve over 1 million random read IOPS using realistic 64K block sizes while maintaining sub-millisecond latency. Unlike VMware vSAN and Nutanix AOS Storage, VergeOS efficiently utilizes network bandwidth, achieving full throughput without requiring specialized hardware. Additionally, its software-embedded error-correcting code (ECC) and optimized write distribution allow organizations to use cost-effective NVMe storage without the risk of premature wear.

It is notable that neither VMware nor Nutanix has publicly released benchmark results demonstrating their storage performance capabilities. While both companies make claims about performance, there is a conspicuous absence of independent, vendor-published test results showcasing their solutions at scale. In contrast, VergeIO has conducted extensive performance testing, both internally and with third parties, demonstrating its ability to meet and exceed the demands of modern workloads. The lack of transparent performance data from VMware and Nutanix raises questions about their ability to deliver the levels of efficiency and scalability that modern IT environments require.

CONCLUSION

Selecting the right vSAN solution is critical for organizations seeking high-performance, scalable, and cost-efficient infrastructure. While VMware vSAN and Nutanix AOS Storage have been widely adopted, their architectures present inherent inefficiencies that impact performance, scalability, and flexibility.

VMware vSAN, particularly in its newer Express Storage Architecture (ESA), attempts to overcome performance bottlenecks by mandating high-performance NVMe storage and 25Gbps networking. However, its reliance on a two-tier storage model introduces excessive write amplification, increasing latency and wear on storage media. Additionally, vSAN's licensing model and hardware restrictions limit customer flexibility, forcing organizations into expensive infrastructure refresh cycles.

Nutanix AOS Storage faces similar challenges, operating as a separate VM within the hypervisor, adding unnecessary CPU and memory overhead to storage operations. While Nutanix offers features such as deduplication and compression, their performance trade-offs discourage widespread use outside of specific workloads like VDI. Furthermore, Nutanix's hardware compatibility constraints prevent customers from mixing different node types within a cluster, limiting long-term scalability.

In contrast, VergeIO's VergeFS adopts a fundamentally different approach by integrating storage directly into the hypervisor as a service. This design eliminates the inefficiencies found in VMware and Nutanix, reducing latency, improving resource utilization, and enabling greater hardware flexibility. Benchmark tests using commodity hardware show that VergeOS achieves sub-millisecond latency and over 1 million IOPS with standard 64K block sizes, outperforming traditional HCI storage architectures live VMware vSAN and Nutanix AOS Storage. Its global inline deduplication, WAN awareness, and advanced data protection features—such as the Inline Repair Server—offer a more resilient and efficient storage environment.

Beyond performance and efficiency, VergeOS offers unparalleled flexibility. Unlike VMware and Nutanix, which enforce rigid hardware compatibility requirements, VergeOS allows organizations to repurpose existing hardware, mix node types, and extend hardware lifecycles. Its Virtual Data Center (VDC) technology provides an additional multi-tenancy and data encapsulation layer, simplifying workload management and improving disaster recovery capabilities.

Ultimately, the choice of vSAN solution should align with an organization's performance, scalability, and budgetary requirements. While VMware and Nutanix require expensive hardware investments to mitigate their inefficiencies, VergeIO's fully integrated, ultra-efficient architecture enables high performance on commodity hardware, reducing capital and operational costs. Organizations looking for a future-proofed, high-performance storage solution that scales with their needs should consider VergeOS as the superior alternative.