Nutanix vs. VergeIO
Data Availability Assessment

verge.io

Comparing data availability between Nutanix and VergeIO is essential when evaluating VMware alternatives. Both platforms consolidate critical components, such as data, metadata, hypervisor configuration files, and virtual machine (VM) metadata, onto a unified server infrastructure. Nutanix employs a traditional hyperconverged infrastructure (HCI), while VergeIO introduces an advanced ultraconverged infrastructure (UCI), aligning all this into a single code base for efficiency and further centralizing resources.

This integrated design approach increases the responsibility placed on each platform to ensure robust resiliency against hardware failures. While both companies offer comprehensive protections against typical failure scenarios, provided their best practices are meticulously followed, the critical differentiator is the cost and complexity associated with these practices. If the recommended protections are too costly or complicated to implement practically, the resulting exposure could be as significant as having no protection.

Executive Summary

Ensuring robust data availability is crucial when evaluating VMware alternatives like Nutanix and VergeIO. Although both Nutanix's traditional hyperconverged infrastructure (HCI) and VergeIO's advanced ultraconverged infrastructure (UCI) aim to provide continuous access to data in the event of hardware failures, the underlying design differences yield significant variations in efficiency, complexity, and cost-effectiveness.

Nutanix relies on a combination of replication factor 3 (RF3) for redundancy and data locality to optimize performance. While these techniques provide effective resilience and performance, they have notable downsides, including high storage overhead, reduced resource efficiency due to reserved capacities, complex recovery processes, and performance degradation during node or drive failures.

In contrast, VergeIO's unified ultraconverged architecture integrates storage directly into the VergeOS environment through VergeFS, eliminating the need for data locality and significantly reducing operational complexity. Its ioGuardian technology further enhances availability by maintaining a highly efficient third data copy. This unique approach simplifies protection against multiple simultaneous failures without imposing heavy resource requirements, enabling even small organizations to achieve enterprise-grade data resilience at a lower overall cost.

VergeIO's tightly integrated, streamlined design ultimately offers superior efficiency, simplified operations, lower total cost, and consistently high performance under normal conditions and during recovery scenarios. Organizations prioritizing robust yet straightforward data availability will find VergeIO distinctly advantageous compared to Nutanix.

Data Availability vs. Data Protection

Understanding the difference between **data availability** and **data protection** is crucial when evaluating infrastructure solutions like Nutanix and VergeIO. Both concepts address distinct scenarios, each with specific methodologies and outcomes.

Data Availability refers to ensuring continuous, uninterrupted access to data, even during hardware or component failures, such as the loss of drives or entire server nodes. This capability is especially critical in converged infrastructures, where compute, storage, and network resources share the same physical hardware. In these environments, a single server outage can impact compute and storage availability simultaneously, potentially causing substantial disruption. Effective data availability solutions mitigate these risks through features such as distributed data replication, rapid failover mechanisms, and real-time redundancy. Ideally, users and applications experience minimal or no interruption during these failures, enabling business operations to continue seamlessly.

Data Protection, in contrast, provides a contingency to restore data access after exceeding the data availability limits. When multiple simultaneous failures or catastrophic events exceed the infrastructure's built-in redundancy, data protection strategies—including backups and disaster recovery (DR)—are employed. Unlike data availability, which maintains ongoing data accessibility, data protection usually involves an operational interruption, even if brief, while data is restored or recovered. Data protection addresses scenarios beyond hardware failures, such as accidental deletions, data corruption from software errors, ransomware attacks, and other logical issues requiring recovery from historical backups or snapshots.

A comprehensive infrastructure strategy requires robust data availability, minimizing downtime during hardware failures, and thorough data protection to reliably recover from catastrophic events and logical data-loss incidents. While this article focuses primarily on data availability, we will provide an in-depth comparison of the data protection capabilities between Nutanix and VergeIO in a future article.

Overview of Nutanix AOS and VergeIO VergeFS

While this document primarily focuses on data availability rather than a detailed storage architecture comparison, understanding the underlying storage implementations of each platform is essential due to their direct impact on data availability performance.

VergeIO integrates storage directly into its ultraconverged operating environment, VergeOS. VergeFS, integrated into VergeOS, is a highly efficient file system that is purpose-built for the integrated infrastructure stack. In contrast, Nutanix relies on its Acropolis Operating System (AOS), a storage software layer atop the hypervisor, to manage storage and provide distributed data services across nodes.

A critical advantage of VergeFS is its deep integration within the VergeOS operating system. VergeFS interacts directly within VergeOS, significantly reducing software overhead and eliminating

the need for additional layers of abstraction between storage and hypervisor operations. This streamlined integration results in fewer processing steps for each data operation, minimizing latency and write penalties.

In contrast, Nutanix AOS operates as a software-defined storage layer separate from the hypervisor kernel. It actually installs as a VM. Consequently, each I/O operation involves traversing multiple software layers, which introduces additional write overhead. In scenarios requiring extensive data availability protections, such as real-time replication and redundancy, these extra layers magnify the performance penalty due to increased write amplification.

This efficiency difference is particularly significant in data availability scenarios, as ensuring continuous data access inherently involves additional writes for replication, parity checks, or metadata updates. The more write overhead an architecture incurs, the greater the performance degradation experienced by workloads, especially under heavy load or during recovery operations. VergeFS's direct integration ensures that these essential data availability processes occur with minimal impact on application performance, maintaining responsiveness and stability even during high-intensity operations.

Therefore, from a data availability perspective, the seamless, tightly integrated storage architecture of VergeFS within VergeOS provides measurable performance advantages over Nutanix AOS, where additional software abstraction layers can contribute to increased latency and reduced overall efficiency.

Data Availability vs. Data Locality

The data locality feature of Nutanix AOS is another capability to address before discussing data availability, because it significantly impacts how the data availability features perform and react to failure. Nutanix claims that AOS's data locality feature provides a performance advantage since it ensures that the primary copy of the VM's data is on the same node as the VM itself.

Data locality only assists in read performance since writes must still be synchronously written to an alternative node. It causes a performance bottleneck when VMs are moved since all data must be fetched remotely from other nodes. It will also cause congestion as it executes its background localization process, which moves a VM's frequently accessed data from a remote replica to the VM. This is especially evident during a node failure, where dozens of VMs will no longer have access to the remote, and an extensive background localization process must occur.

VergeFS intentionally does not implement data locality in order to avoid the complexities associated with migrating or restarting VMs following a node failure. Its deep integration into VergeOS, combined with VergeFS's optimized internode communication protocol and the ability for multiple nodes to respond to read requests simultaneously, enhances data access performance. Additionally, VergeOS extensively leverages RAM for caching, which is even further optimized by integrated global deduplication. As a result, VMs typically experience faster read and write performance over the VergeOS internal network than they would from accessing data stored locally on their node. Also, the VergeFS data placement design simplifies the movement of VMs and their responsiveness during a node failure.

Data locality offered benefits when slow hard drives, costly flash storage, and lower-speed networks dominated storage. Today's widespread use of high-speed NVMe storage, affordable RAM caching, and pervasive 10GbE (or faster) networking has largely erased those advantages.

Additionally, leveraging ample, cost-effective RAM caches—especially with global inline deduplication—makes remote data access extremely efficient and responsive. Deduplication maximizes RAM's effectiveness, enabling more unique data blocks to reside in memory, significantly accelerating read performance.

Maintaining data locality introduces unnecessary complexity and overhead in this modern context, such as continuous data placement tracking, balancing, and migration tasks. Instead, contemporary architectures like VergeOS utilize storage protocols optimized for efficient data movement across fast active-active networks, paired with intelligent RAM caching and deduplication, providing superior performance without the overhead associated with data locality.

Data Availability Baseline - Distributed Mirroring

Nutanix AOS and VergeIO VergeFS utilize distributed mirroring technology as the foundation for data availability. Distributed mirroring ensures resiliency by automatically creating multiple copies of data across distinct physical components within the infrastructure. When new data is written to either environment, the platform generates two mirrored copies, distributing these across different drives and separate servers or nodes. The replica copies of each VM are spread out across multiple nodes and disks within the environment. Each VM's replica is distributed across various nodes and drives using sub-VM granularity to distribute data.

This intelligent data distribution significantly mitigates risks associated with single points of failure. If a drive experiences a failure, the virtual machines (VMs) relying on primary data from that compromised drive instantly and transparently switch their data access to the replica copy located on another node. As a result, users and applications experience no disruption or downtime during these hardware failures, maintaining operational continuity.

Both platforms continuously manage and monitor the placement of these mirrored copies, actively balancing data distribution across nodes and storage resources. This proactive management ensures optimal storage performance and evenly distributes workload and wear, which helps prevent performance bottlenecks or premature hardware degradation. Additionally, distributing mirrored data across multiple nodes enhances overall resilience, safeguarding against broader system impacts from isolated component failures.

While distributed mirroring forms a fundamental baseline for data availability in both solutions, it is essential to recognize that mirroring inherently doubles the required storage capacity, as each write generates two distinct data copies. This overhead makes the efficiency of the underlying storage system especially critical. As discussed previously, the tighter integration of VergeFS within VergeOS reduces write overhead, providing an advantage in performance and efficiency under the increased write workloads associated with continuous data replication and mirroring.

Distributed mirroring is an essential and effective baseline for protecting against common storage failures. Both Nutanix and VergeIO leverage this technique to maintain seamless availability, but the underlying efficiency of the storage platform can significantly influence performance and operational cost.

Single Drive Failure Protection

Single-drive failure is one of the most critical—and common—scenarios for data availability protection in converged infrastructures. Considering converged and hyperconverged environments often contain dozens or even hundreds of drives, the likelihood of encountering a single-drive failure is substantial. Proper handling of these failures involves three key steps: proactive detection, transparent data access (failover), and efficient healing or recovery.

Both Nutanix AOS and VergeIO VergeFS actively monitor the health of each drive, proactively identifying potential failures before they become catastrophic. Both platforms can alert administrators of impending drive issues by continuously assessing drive metrics and performance indicators, allowing for preemptive action and minimizing the risk of sudden operational disruption.

Once a drive fails, both AOS and VergeFS transparently ensure continuous VM operation by retrieving any missing data segments from replicated copies stored on other nodes. However, the performance impact of this network-based data retrieval differs significantly between the two architectures.

For Nutanix AOS, accessing primary data remotely across the network represents an abnormal operational state, as VMs typically rely on local data access (data locality). Consequently, the affected virtual machines may experience noticeable performance degradation during the failure state due to increased network traffic and higher latency. Moreover, the additional network load can negatively impact other VMs within the same environment, potentially causing wider operational disruption.

In contrast, VergeOS and its VergeFS storage architecture inherently distribute data access across nodes as part of normal operations. As a result, accessing data over the network following a drive failure is not a deviation from regular operation for VergeFS-managed VMs. The architecture's optimized internode communication, extensive caching mechanisms, and integrated global deduplication ensure that affected VMs—and others sharing the environment—experience minimal or no performance impact during drive failures.

Healing and recovery approaches also vary significantly between AOS and VergeFS. Nutanix AOS initiates an automatic "self-healing" process immediately upon detecting a drive failure. This process redistributes and replicates the data segments from the failed drive onto healthy drives within the cluster. However, this operation requires substantial free storage capacity—typically

equivalent to at least one full node's capacity (N-1 nodes' worth of storage in an RF2 cluster)—to redistribute data safely. This capacity requirement can constrain environments with limited available space, complicating storage planning and potentially requiring costly reserve capacities.

VergeFS, on the other hand, initially takes no immediate action beyond alerting administrators of the drive failure. Instead, it leverages its inherent redundancy and real-time data recovery provided by ioGuardian technology (discussed later), ensuring continuous data availability without immediate rebalancing or redistribution. VergeIO recommends maintaining a spare commodity SSD for quick physical replacement, an economical approach made possible by VergeOS's compatibility with off-the-shelf drives, regardless of brand or model.

Administrators commonly find that ejecting and reinserting the failed drive can temporarily bring it back online; however, replacing the failing drive promptly remains a best practice. Once a failed drive is physically replaced, VergeOS automatically detects the new drive, formats and initializes it, and seamlessly begins the recovery process. Data is efficiently redistributed and rebuilt onto the replacement drive, while the environment remains fully operational. Throughout this simple, automated procedure, administrators can track repair progress and estimated completion time via the intuitive management interface, ensuring minimal operational overhead and clear visibility into system status. If, for some reason, the drive can't be replaced, the administrator can issue a command that redistributes that data across the remaining drives and nodes in the environment.

VergeFS's approach delivers robust single-drive failure protection while maintaining optimal resource efficiency, reducing complexity, and avoiding significant performance penalties associated with traditional recovery processes.

Multi-Drive Failure Within a Single Server

In a converged environment, multiple drive failures can occur in three scenarios. First, when multiple drives within a single node fail, second, when multiple drives across multiple nodes fail, or third, when an entire server goes down, all the drives within that server have essentially failed.

If multiple drives fail within a single server, the standard replication policy, be it one or two replicas, should provide adequate protection since both Nutanix and VergeIO ensure that replicas do not occur within the same node. The only difference between a single drive failing and multiple drives within the same node is the number of drives that must be repopulated once a failure occurs.

Single Node Failure

From a data access perspective, a single node failure is similar to multiple drives failing within a single server; the standard replication policy should maintain data access. The big difference is that the server is not able to provide compute resources, and the VMs on that server must be repositioned via a live migration function, which again, both Nutanix AOSAHV and VergeOS provide. This situation, however, is where Nutanix's Data Locality feature adds significant complexity, and understanding how the virtual machines are distributed is a critical differentiator between the two products.

Virtual Machine Distribution on Node Failure

After a node failure, Nutanix and VergeIO can automatically restart virtual machines (VMs) on alternative nodes. However, the specific strategies employed by each platform to determine VM placement post-failure differ significantly, resulting in notable impacts on cost, resource efficiency, and overall operational complexity.

Nutanix AOSAHV relies on predefined policies to govern VM redistribution upon node failure. The original method, still available as a legacy option, follows a "best-effort" policy. Under this approach, AHV attempts to restart affected VMs on any remaining nodes with available resources. While simple, this method risks performance degradation or even failed VM restarts if adequate resources are not immediately available on surviving nodes.

From AHV version 5.0, Nutanix introduced the **"Reserved Segments"** approach as the default VM recovery policy. Reserved Segments proactively reserve a predetermined percentage of each node's resources, effectively creating an overhead to handle potential failures. This reservation proportionally reduces available computing resources across the cluster, leading to substantial resource utilization inefficiencies, particularly in smaller clusters:

Cluster Size (Nodes)	Per-Node Reserved Percentage	Total Reserved Capacity (Equivalent Nodes)
3	33.33%	1 node
4	25%	1 node
8	12.5%	1 node
12	8.33%	1 node

It is also important to note that Nutanix applies this reservation at the cluster level rather than the broader environment. Nutanix customers frequently deploy multiple clusters within the data center to achieve specific workload isolation, compliance adherence, or optimized resource allocation. While beneficial from a management standpoint—allowing isolation of production from development workloads, or regulatory compliance-sensitive applications from general-purpose workloads—this strategy leads to **"cluster-sprawl"**. It multiplies the reservation overhead, significantly increasing overall infrastructure costs and operational complexity.

In contrast, VergeIO's VergeOS addresses VM placement after node failure using its built-in intelligence, ioOptimize. ioOptimize employs narrow AI algorithms that dynamically assess the optimal node for restarting each VM in real time when a failure occurs. Rather than relying on fixed reservations or predetermined failover nodes, ioOptimize evaluates resource availability, current node utilization, and VM workload characteristics to instantly identify the most suitable surviving node.

This predictive, dynamic placement ensures that resource efficiency remains high under normal operations, where resource availability and performance matter most, and during recovery scenarios. With ioOptimize, VergeOS eliminates the need for pre-reserving resources or designating specific failover nodes, allowing organizations to achieve higher resource utilization without the overhead penalties of traditional reservation-based models.

Additionally, the automated nature of ioOptimize significantly reduces operational overhead. IT teams do not need to manually configure or periodically reassess failover settings as nodes are added, removed, or updated. Instead, VergeOS transparently manages these adjustments, optimizing resource distribution across the available nodes. This adaptive capability becomes increasingly valuable in dynamic environments, providing seamless scalability and operational simplicity.

VergeOS also eliminates "cluster sprawl"—the practice of deploying multiple smaller clusters to achieve workload isolation, regulatory compliance, or performance optimization. Its Virtual Data Center (VDC) capabilities enable IT administrators to securely segment workloads and allocate distinct resources within a cohesive infrastructure. Each Virtual Data Center supports unique configurations, dedicated resource pools, and customized security policies, delivering all the isolation benefits traditionally achieved through multiple clusters, without the associated complexity or additional cost. VergeOS significantly simplifies management, reduces operational overhead, and lowers infrastructure costs compared to managing separate Nutanix clusters by consolidating numerous workloads into a unified operating environment. In VergeOS environments, additional clusters are only necessary when isolating fundamentally dissimilar hardware platforms.

While Nutanix's Reserved Segments approach ensures predictable failover capacity at the cost of significant resource reservation overhead, VergeIO's ioOptimize dynamically allocates resources at the point of failure, providing superior efficiency, lower operational complexity, and reduced infrastructure costs.

The Impact of Data Locality on Unplanned VM Migrations

Systems relying on data locality, such as Nutanix, face significant performance and operational challenges during unplanned VM migrations caused by node failures. Although Nutanix attempts to mitigate complexity through incremental, 1MB extent-based data relocalization and background optimization, substantial issues persist. Immediately following migration, virtual machines experience noticeable performance degradation due to latency introduced by remote data access. This latency issue is exacerbated by inefficiencies inherent in Nutanix's internode communication protocol, severely affecting workloads characterized by intensive, random I/O patterns, such as transactional databases.

Moreover, Nutanix's incremental data-relocation strategy results in extended network strain, as data is transferred gradually rather than in bulk. This continuous network utilization can degrade cluster performance and negatively impact other workloads. Resource contention on the destination node may also complicate or even halt migrations, especially when CPU, memory, or storage resources are insufficient. **Furthermore, temporary redundancy overhead occurs as newly localized replicas coexist with older remote copies, temporarily increasing storage utilization until full cluster rebalancing completes.**

Even after resolving the initial failure, data locality once again complicates matters when the original node returns online. Nutanix attempts to reduce the impact through intelligent handling, migrating only data blocks modified during the node's downtime rather than the full VM dataset. Despite these optimizations, the system initiates incremental data transfers at a granular (1MB) level, triggering additional network overhead and performance impact. This second data transfer cycle further extends the network strain, prolongs latency effects, and adds complexity as the system works to rebalance replicas and maintain redundancy.

In contrast, VergeIO's architecture intentionally eliminates the complexities associated with data locality. VergeFS uniformly distributes data across nodes, leveraging high-speed networks combined with an optimized, active-active storage protocol, efficient RAM caching, and global inline deduplication. This design ensures consistently high performance under normal operations, during node failures, and when failed nodes return online. Unlike Nutanix, VergeOS experiences no noticeable performance differences between active, failed, or recovery states. By avoiding unnecessary data movement altogether, VergeIO simplifies migration processes, ensuring rapid recovery and maintaining stable, predictable application performance throughout the entire node failure and recovery cycle.

Multiple Drive Failures Across Multiple Nodes

Multiple drive or node failures present significant risks in environments employing distributed mirroring. If two or more drives fail across different nodes, particularly drives containing related or mirrored data, the risk of data loss increases substantially. Data loss doesn't require simultaneous failures; the risk window begins with the initial drive failure and persists until the data is completely rebuilt or redistributed across remaining nodes. Suppose an additional drive failure occurs within this rebuild timeframe. In that case, the situation becomes equivalent to multiple simultaneous drive failures, potentially forcing administrators to resort to external backup solutions and endure costly downtime.

Nutanix and VergeOS address this risk using distinctly different approaches. Nutanix relies primarily on a Replication Factor of 3 (RF3), creating three complete copies of data across separate nodes. RF3 can sustain simultaneous failures of two drives or nodes holding overlapping data sets. Although robust, RF3 introduces substantial storage overhead, reducing usable capacity by approximately 66% compared to 50% for the more common RF2 configuration. Additionally, Nutanix's RF3 requires a minimum of five nodes to maintain redundancy and metadata quorum, significantly increasing infrastructure complexity and cost.

Once RF3 is activated, its metadata configuration (five copies) permanently persists—even if containers are later downgraded to RF2. Reverting entirely to RF2 thus demands a disruptive, cluster-wide reimaging. This inflexibility often compels organizations to manage complex dual-container strategies, placing critical workloads into RF3-protected containers and relegating less-critical applications to RF2. Consequently, IT administrators face challenging performance, availability, and cost trade-offs. Even within RF3 environments, a third failure occurring before data rebuild completion may still result in data loss, emphasizing the critical importance of rapid hardware replacement and diligent infrastructure management.

In contrast, VergeOS provides a streamlined and highly effective alternative through its advanced data availability technology: **ioGuardian.** While VergeOS supports traditional three-way mirroring, most customers prefer ioGuardian because it delivers enhanced protection without the associated overhead. ioGuardian maintains an independent third copy of the data, stored separately from the primary mirrored dataset. Rather than continuously replicating data three times, ioGuardian is a highly available backup, providing superior resilience at significantly lower resource costs.

A distinctive advantage of ioGuardian is its exceptional efficiency. Leveraging VergeOS's global inline deduplication, ioGuardian substantially reduces storage requirements compared to traditional three-way replication strategies. This efficiency allows organizations to deploy high-density, cost-effective storage, such as QLC SSDs, on the ioGuardian server, further decreasing infrastructure costs. The ioGuardian server requires minimal compute resources and does not host active virtual machines. It exists purely to store and efficiently serve data.

Moreover, ioGuardian imposes no additional burdens or limitations on the size or complexity of the production cluster. In fact, many smaller VergeIO deployments—such as two—or three-node clusters—effectively utilize an ioGuardian server, ensuring robust data protection without needing to scale up cluster size. This flexibility enables small and medium-sized businesses to achieve enterprise-class data availability with minimal incremental cost and complexity.

If multiple drives or nodes fail, ioGuardian transparently provides immediate data access without requiring a traditional recovery process. Affected virtual machines instantly retrieve data directly from the ioGuardian storage, maintaining uninterrupted operations without administrator intervention or complex recovery workflows. This instantaneous recovery capability surpasses the limitations of a conventional three-way mirror, enabling organizations to withstand extreme multi-node failures without downtime or data loss.

VergeOS also simplifies operational management by eliminating the need to manage multiple replication tiers or make complex workload placement decisions. Instead, administrators can uniformly apply ioGuardian's robust protection across all workloads, significantly reducing complexity and ensuring consistent, comprehensive data protection.

Furthermore, ioGuardian seamlessly integrates backup-like functionality into the primary data availability strategy. Since most organizations already maintain backup copies, **ioGuardian transforms a passive backup into an active component of continuous availability,** providing substantial additional resilience without added complexity or costs.

VergeOS's ioGuardian delivers a superior approach for managing multiple drive or node failures compared to Nutanix's RF3 solution. By ensuring uninterrupted data access, simplifying management, minimizing storage overhead, and effectively integrating backup capabilities, ioGuardian significantly improves availability, reduces downtime risks, and enhances operational simplicity—even during severe, multi-failure scenarios.

Conclusion

When evaluating alternatives to VMware, it is essential to select an infrastructure solution that provides robust data availability. Nutanix and VergeIO each deliver solutions designed to ensure continuous data access despite hardware failures. However, the architectural differences between the two platforms lead to substantial efficiency, complexity, and cost-effectiveness disparities.

Nutanix's hyperconverged architecture relies heavily on features such as data locality and RF3 to deliver availability. While effective in many situations, these methods incur significant operational complexity, substantial storage overhead, and considerable performance degradation during node or drive failures. Furthermore, Nutanix's requirement for reserved resources and additional storage for RF3 significantly increases overall infrastructure costs and management complexity.

VergeIO, leveraging an advanced ultraconverged infrastructure with its integrated VergeFS file system and innovative ioGuardian technology, presents a simpler, more efficient approach. VergeOS eliminates the need for data locality, instead using optimized network protocols, extensive RAM caching, and global inline deduplication to deliver consistently high performance, even during failure scenarios. Its intelligent ioOptimize feature dynamically allocates resources in real-time during failures, avoiding wasteful resource reservations and complex configurations.

Most importantly, VergeIO's ioGuardian significantly simplifies protection against multiple simultaneous drive or node failures. By maintaining an efficient independent data copy, ioGuardian seamlessly integrates active data protection with continuous availability, significantly reducing storage overhead and operational complexity. This enables organizations—even those operating small clusters—to achieve enterprise-level availability without additional cost or complexity.

While both solutions offer robust data availability, VergeIO distinguishes itself through superior efficiency, lower complexity, reduced storage requirements, and seamless operational simplicity. Organizations seeking dependable, cost-effective, and easily manageable infrastructure resilience will find VergeIO uniquely positioned to meet and exceed these critical needs.