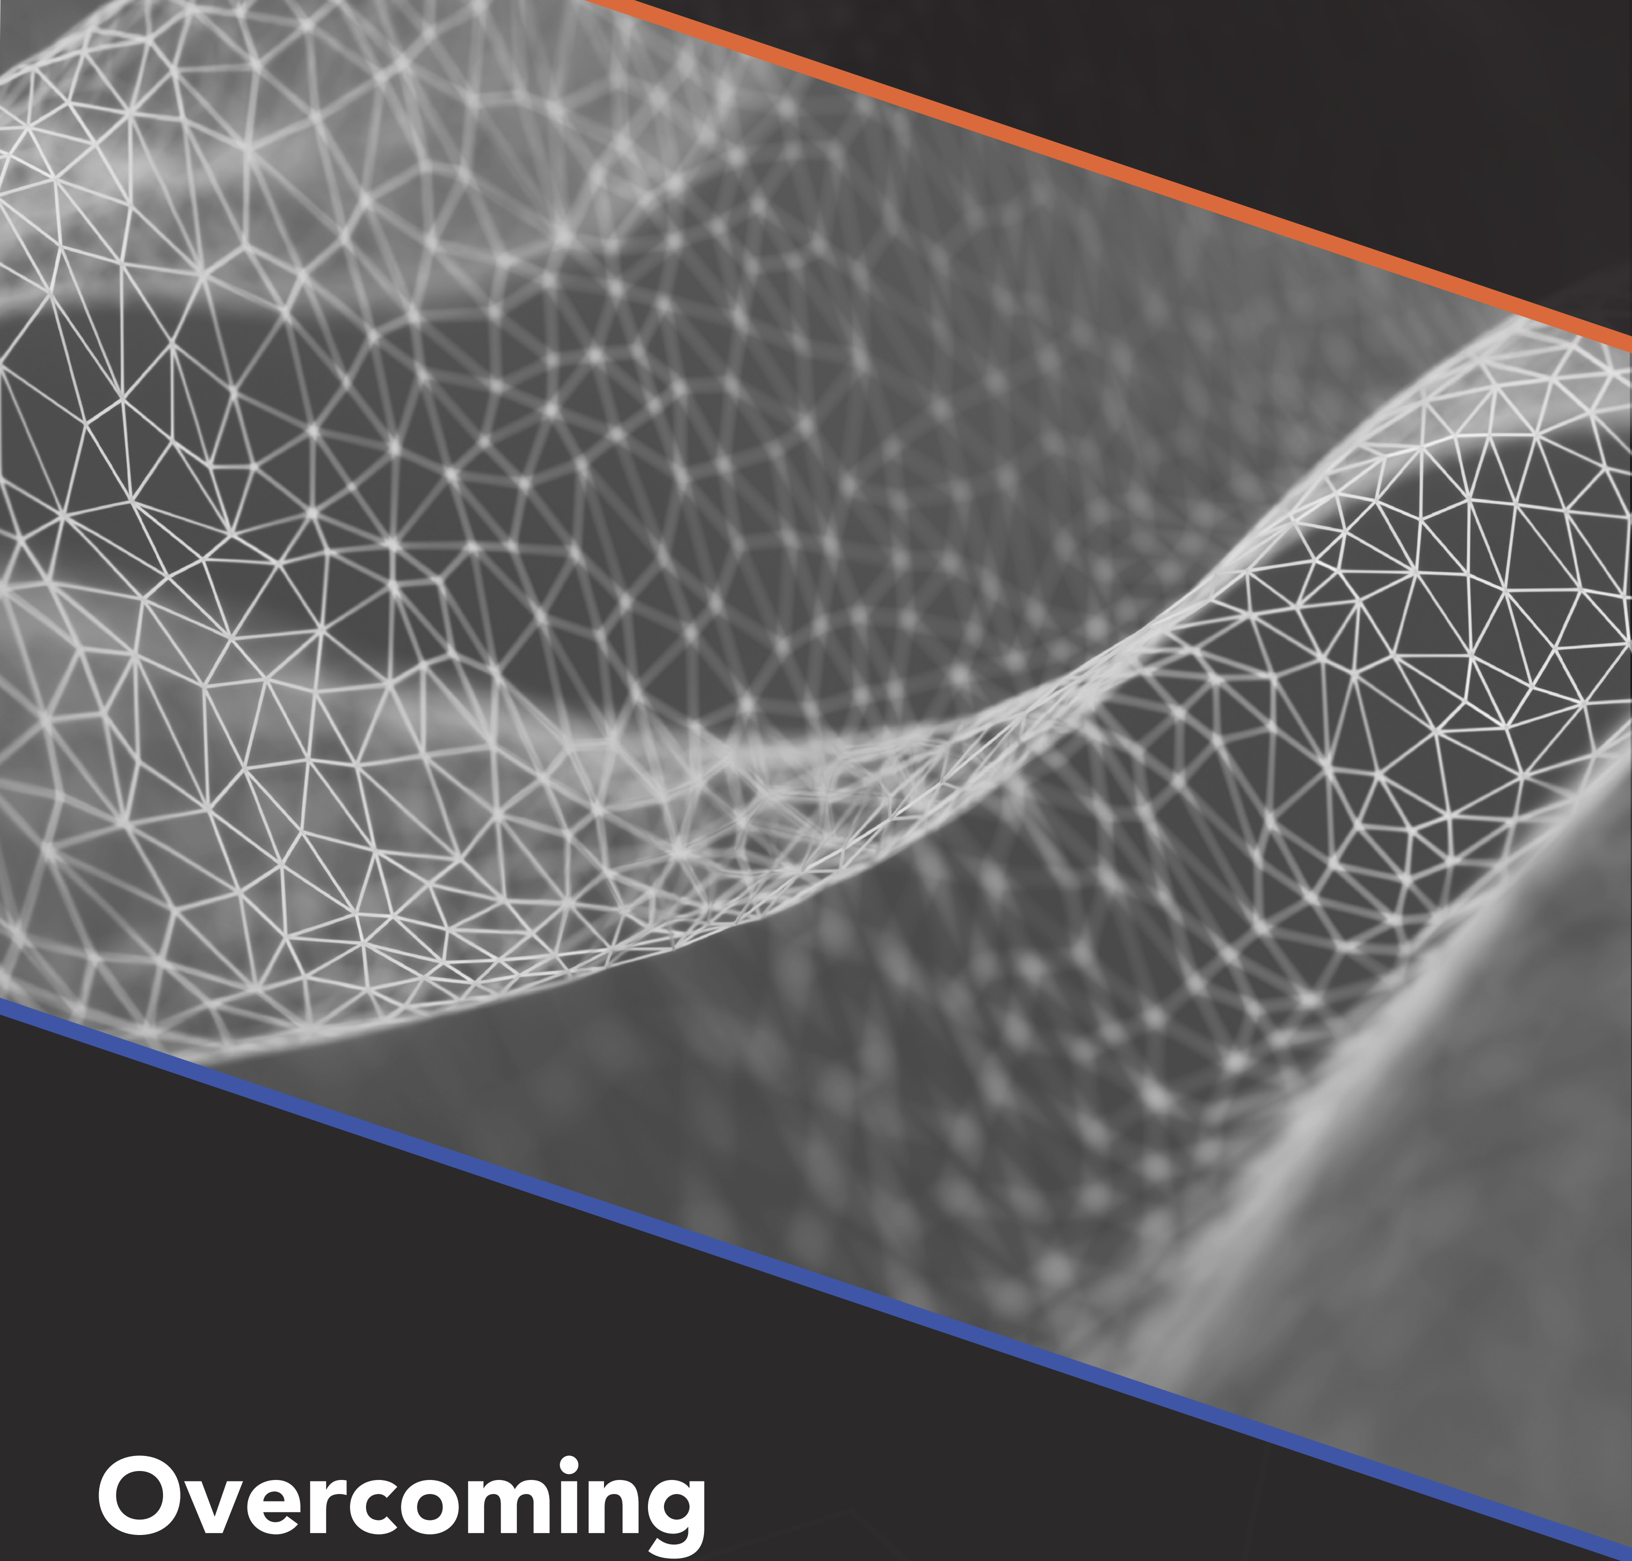

Overcoming Infrastructure Instability

verge.io

Legacy infrastructure is unstable, and the problems are accelerating. VMware raises license and support costs and pushes bundles. Nutanix adds per-node expense and strict hardware rules. Public cloud platforms such as AWS and Azure drive unpredictable recurring bills that grow with data egress and workload expansion. Hyper-V reaches its scale and performance limits under a mixed load. Scale Computing faces questions about its roadmap following the acquisition. Proxmox and OpenStack demand staff time and engineering depth that most enterprises cannot provide.

Budgets drift off plan, and refresh cycles are forced by infrastructure software, arriving years before the hardware reaches the end of its useful life. Hardware that is expected to serve five to seven years is retired by policy, not due to failure. Teams miss latency targets under growth. Outage recovery takes longer than the change windows allow. Security events drive demands for isolation and rapid rollback, which current stacks do not meet.

Incremental fixes do not stop this pattern. A patch within the same stack preserves the cost model and the hardware policy unchanged. A lift to the cloud alone trades one problem for another. A switch from one hypervisor to another often reveals that the destination platform is built the same way, with multiple software codebases loosely “integrated” behind a management GUI but lacking true cohesiveness.

The enterprise needs a clean break, and that takes two parts that work together to minimize cutover time.

The first part is a defined path. It must synchronize while production runs, protect application I/O, and give clear approvals and reports. It must accept sources across VMware, Nutanix, Hyper-V, Scale, Proxmox, OpenStack, and various cloud IaaS platforms. It must cut over in planned windows measured in seconds or minutes, not hours.

The second part is a stable destination. It must run virtualization, storage, and networking within a single operating model. It must extend hardware life and accept standard x86. It must deliver tenant isolation for projects and departments and operate consistently across core, edge, and ROBO. It must lower operating costs by removing stacked products and tool sprawl

This paper names the causes of instability and shows where each vendor fits. It then defines the path and the destination in measurable terms. It includes a three-year cost framework, a risk register, and a rollout plan that begins with a small pilot and scales to the entire estate. It concludes with Cirrus Data and VergeIO as examples of a path and destination working together to deliver an infrastructure that can solve today’s challenges and is prepared for the future.

THE INFRASTRUCTURE INSTABILITY CRISIS

Enterprises need infrastructure to provide stable cost models, predictable refresh cycles, and dependable recovery. Infrastructure software vendors are delivering the opposite. Costs rise faster than budgets, hardware is retired long before it fails, and outages expose weak recovery. Security incidents grow in frequency and impact. The result is an unstable, expensive, and difficult-to-manage infrastructure environment that is now sprawling across multiple hypervisor software stacks, various dedicated hardware products, and half-steps toward the cloud.

Rising costs are the most visible sign. Licensing fees rise with little connection to the delivered value, and forced bundles include products that many organizations do not want. Public cloud bills from AWS and Azure exhibit the same trend, with recurring costs that increase each quarter and spike with egress charges or the introduction of new workloads. But worse, they run on hardware the enterprise will never own. Organizations pay the same monthly cost for hardware that is declining in value, with no asset left at the end of the term.

Hardware churn compounds the problem. Vendors enforce strict compatibility lists that cut short the useful life of servers and storage. Equipment that should remain in service for five to seven years is pushed out after three or four. Capital is wasted, refresh cycles accelerate, and planning loses credibility.

Performance and scalability gaps are another driver. Hypervisors meet initial needs but struggle as estates grow. The cloud may meet those needs, but at an added monthly cost. Latency targets are missed, mixed workloads stall, and service levels fall. Complexity rises as teams layer tools and separate infrastructure to fill gaps.

Enterprise AI demand adds even more pressure. Organizations have valuable data that they cannot place in the cloud, and they do not want to be dependent on cloud platforms for AI training or inference. Private AI is the clear path, but the inflexibility of today's infrastructure software forces many enterprises to deploy new, separate platforms dedicated to supporting AI workloads. This duplication adds more cost and complexity to an already unstable environment.

Taken together, these factors create a crisis of instability. Enterprises find themselves paying more for less, replacing hardware prematurely, and operating platforms that do not scale or recover well. The pattern is consistent across VMware, Nutanix, Hyper-V, Scale Computing, Proxmox, OpenStack, and public cloud IaaS like AWS and Azure. Enterprises need an exit from this cycle and a model that brings stability back to cost, hardware life, performance, and recovery.

WHY LEGACY PLATFORMS ARE FAILING

The instability crisis is not abstract. It comes from specific, recurring weaknesses across the leading infrastructure platforms. These weaknesses fall into seven categories:

- Rising prices
- Inflexibility and hardware lock-in
- Performance and scalability gaps
- Roadmap uncertainty
- Future readiness
- Security and compliance gaps
- Lack of interplatform movement

Each category creates pressure on budgets, operations, and long-term planning. Together, they explain why enterprises are struggling to maintain stability.

Rising Prices

Many vendors are increasing licensing and support costs. Vendors now generate more revenue through forced bundles and complex terms than through the value they deliver to customers. VMware has sharply increased prices under new ownership. Nutanix, as it switched to a subscription model, increased pricing sharply, and its strict server requirements mean that hardware is priced at a premium. Public cloud IaaS providers, such as AWS and Azure, issue recurring monthly bills that increase each quarter, sometimes accompanied by egress charges or new workload growth. As a result, IT budgets quickly slip off track, and long-term cost models lose their reliability as a guide for decision-making.

Inflexibility and Hardware Lock-In

Infrastructure software should extend hardware life, but vendors are forcing the opposite. Strict compatibility lists shorten useful service life and force refreshes years ahead of schedule. Nutanix limits support to a narrow hardware list, pushing customers into new nodes before existing systems reach the end of serviceable life. VMware and Hyper-V deprecate servers through software support policies, even when those servers remain technically sound. Proxmox, while popular in smaller deployments, lacks the enterprise-grade features required to run mixed hardware and multi-tenant isolation at scale. The net effect is wasted capital and accelerated refresh cycles.

Performance and Scalability Gaps

Many platforms work well at small scale but fail as estates grow or workloads diversify. Hyper-V struggles to maintain performance at higher node counts and under mixed workload intensity. OpenStack deployments introduce complexity, making upgrades and scaling unpredictable without deep engineering resources. Scale Computing has a hard limit of eight nodes and limited performance headroom, leaving enterprises unable to meet latency or throughput requirements.

Other platforms show the opposite problem: they cannot scale down to two or three nodes, which is required for edge sites, remote offices, and venues. For these organizations, the gap forces IT to deploy one platform at the core and another at the site, creating yet another form of infrastructure sprawl with multiple systems to license, manage, and support.

Roadmap Uncertainty and Support Risk

Vendor stability is as important as product features. VMware faces uncertainty as ownership changes reshape licensing, support, and product direction. Customers question whether the long-term roadmap will favor enterprise needs or short-term financial goals.

Microsoft has demonstrated an inconsistent commitment to Hyper-V, shifting its focus toward Azure and leaving on-premises customers uncertain about the continued investment. Scale Computing faces uncertainty following its reverse merger with Acumera, leaving customers unsure about future development and support models. OpenStack suffers from uneven quality across modules and a fragmented roadmap. Proxmox relies heavily on a community-driven model, creating gaps in long-term enterprise support and integration. Enterprises cannot make multi-year infrastructure plans on platforms where the future is unclear.

Future Readiness

Enterprises must prepare their infrastructure to meet demands that exceed traditional workloads. Private AI is a leading example. Organizations hold sensitive data that cannot move to the public cloud, yet they want to run training and inference close to where that data resides. They want to avoid the expense and dependency that come with tying AI initiatives to AWS or Azure. Beyond cost, there is a growing unease about a few major cloud providers owning the majority of AI capabilities, leaving enterprises dependent on their roadmaps and pricing models.

Current on-premises infrastructure software solutions do not provide the capabilities enterprises need for AI. VMware, Nutanix, Hyper-V, and others require separate deployments to support GPUs and AI workflows. Companies, such as Nutanix, have attempted to address the problem with offerings like GPT-in-a-Box; however, this approach follows the same design strategy as other parts of the product: it is a separate module that requires a dedicated cluster. This lack of flexibility forces IT to set up new platforms to pursue AI, creating additional costs, increased fragmentation, and another form of infrastructure sprawl.

Security and Compliance Gaps

Enterprises face increasingly stringent regulatory requirements and a constant rise in cyber threats. Current infrastructure platforms are struggling to keep pace—other vendors bolt on security features instead of integrating them into the core platform.

VMware, for example, requires separate products or add-ons for microsegmentation, backup, and recovery, which increases complexity and cost. Nutanix relies on a mix of built-in controls and third-party integrations, creating gaps between compliance and daily operations. Hyper-V depends on the broader Microsoft ecosystem, forcing customers to layer Windows Server and Azure tools to meet security standards. Proxmox and OpenStack leave much of the responsibility to the enterprise, requiring deep in-house expertise to configure and maintain consistent controls. Public cloud providers, such as AWS and Azure, offer robust security frameworks. Those frameworks must be learned, and they create sovereignty and compliance challenges when data must remain on-premises.

The result is fragmentation in an area where you want cohesiveness, and security. This fragmentation forces organizations to manage multiple consoles, inconsistent policies, and overlapping tools. There are also more areas of exposure. Compliance audits take longer and cost more, while response to incidents is slower. Enterprises need infrastructure that embeds security and compliance as part of the design, not as afterthought or external modules.

Lack of Interplatform Movement

Enterprises learn that the first six weaknesses drive platform sprawl. Mergers, acquisitions, and tactical projects leave IT running VMware in one part of the business, Hyper-V in another, and Nutanix or Scale Computing at the edge. An initial set of workloads shifts into AWS or Azure, yet a complete migration never materializes. Proxmox and OpenStack can appear in smaller pockets, adding even more variation. Consolidation should reduce cost and complexity, but the inability to move workloads freely between these platforms makes it nearly impossible.

A few vendors provide migration tools, but they are narrowly scoped. Competitors may offer import utilities focused on VMware as the source, but do little to help organizations running a mix of hypervisors and cloud instances. They do not address the broader requirement of migrating workloads from any platform to a common target.
The result is that consolidation stalls, or more likely, is never undertaken. Enterprises are forced to maintain multiple platforms in parallel, each with its own licensing, support contracts, and operational overhead. Migration projects span years, and IT never achieves a single, consistent infrastructure model.

Hardware Flexibility	Scale & Performance	Roadmap Confidence	AI Readiness
Strict HCLs, early refreshes	Works at scale but costly	Ownership changes raise doubt	Add-ons only
Narrow HCL, premium hardware	Scales, but cost rises quickly	Slower growth trajectory	GPT-in-a-Box requires separate cluster
Hardware tied to Windows lifecycle	Degrades at high node counts	Microsoft shifting to Azure	No native AI
Vendor-bound hardware	Hard cap of 8 nodes	Reverse merger with Acumera raises concerns	No AI strategy
Basic multi-tenant features only	Limited enterprise scaling	Community-driven, no clear roadmap	No enterprise AI support
Hardware flexible but complex	Difficult upgrades at scale	Fragmented modules	No integrated AI
No asset ownership	Nearly unlimited but costly	Roadmap controlled by vendor	Advanced AI, but dependency risk

For an infrastructure consolidation strategy to achieve its objective, it must start with a universal migration path and end with a stable destination. The path must span all major hypervisors and cloud platforms, not just VMware, and it must move workloads without disruption. The destination must unify compute, storage, and networking into a cohesive platform that extends hardware life, supports new demands like private AI, and simplifies operations across core, edge, and cloud. Without both elements working together, consolidation is impossible, and enterprises remain trapped in a cycle of sprawl, rising costs, and instability.

These categories—rising prices, inflexibility, performance gaps, roadmap uncertainty, future readiness, security and compliance gaps, and lack of interplatform movement—are present across VMware, Nutanix, Hyper-V, Scale Computing, Proxmox, OpenStack, and cloud IaaS providers such as AWS and Azure.

The evidence is consistent: instability is a direct result of vendor design philosophies, not isolated missteps. They focus on speed to market using isolated development teams and claim integration behind a standard GUI. Enterprises need to recognize these patterns for what they are, before they can choose a better path forward.

THE CASE FOR A DEFINED PATH

The weaknesses in legacy platforms reveal one consistent truth: without a reliable way to migrate workloads, consolidation is impossible. Enterprises cannot accept extended downtime, yet most migration efforts depend on planned outages measured in hours or days. This makes leaders hesitate, and projects often stall before they ever begin.

A Path That Spans All Platforms

The path must span all major hypervisors and cloud platforms, not just VMware. Enterprises rarely run a single platform anymore. Mergers, acquisitions, and tactical projects have left IT with estates that include Hyper-V, Nutanix, Scale Computing, Proxmox, OpenStack, and workloads running in AWS or Azure. A path that assumes VMware as the only starting point does not address this reality.

To be effective, the migration layer must accept any of these sources and move them to a common target without disruption. It must move not only data but full workloads, including the operating system, applications, and configurations that production depends on. This capability allows IT to retire redundant platforms, eliminate licensing overhead, and simplify support models. Without it, consolidation stalls, and instability persists.

A Path That Minimizes Cutover Time

A defined path must operate while production runs. It must synchronize data in the background, protect application I/O, and provide administrators with control over the migration pace. Bandwidth must be managed intelligently so that workloads remain responsive. Approvals, logging, and detailed reporting are required to meet audit and compliance obligations. Above all, the path must shrink the cutover window to seconds or minutes so that migrations can align with existing change windows instead of requiring special outages.

Limitations of Current Tools

Current migration tools are narrowly scoped. Most assume VMware is the source and provide only limited utilities for importing into another stack. These tools do not manage performance, do not support compliance reporting, and cannot span the diversity of platforms in modern enterprises. As a result, they increase risk rather than reduce it.

Cirrus Data as a Defined Path

Cirrus Data provides a detailed example of what a true migration path should deliver. The company's software-only solutions, Cirrus Migrate Cloud and MigrateOps™, together enable the seamless migration of any Windows or Linux block storage solution, whether physical or virtual, across nearly any hypervisor, private cloud, or public cloud while maintaining production uptime.
This means enterprises can move workloads from VMware, Hyper-V, Nutanix, Scale Computing, Proxmox, OpenStack, or even public cloud IaaS into a common target.

The platform's Intelligent QoS (iQoS) allows for the leveraging 24 X 7 migration schedule with impact to production. It enables you to set the yeilding level so your applications maintain priority.. iQoS monitors disk activity in real time and will adjust its pace depending on the level of production activity. When systems are below the targeted level, the migration will consume the available bandwidth to accelerate progress. This dynamic control allows data to stay synchronized until the exact moment of cutover. Combined with approval workflows and audit-ready reporting, enterprises gain the assurance that migrations can be planned, tracked, and executed without impacting performance, violating compliance or governance requirements.

Cirrus Data's MigrateOps automates the migration process. Administrators define operations in simple YAML files that capture system details, integration points, network settings, and scheduling parameters. The automation reduces manual errors, makes migrations repeatable, and allows IT to scale from small pilots to estate-wide programs. Additionally, MigrateOps allows you to seamlessly link together your migration automations with other automations that are often managed externally, such as:

- OS Automation
- Application Automation
- Public/Private Cloud Platform Automation
- Virtualization Platform Automation
- Enterprise Storage Management Automation
- SAN Management Automation

Each MigrateOps operation consists of a configuration and a recipe.
A recipe is a data mobility operation template provided by Cirrus Data Cloud, providing instructions on WHAT tasks need to be executed to achieve a particular data mobility goal.

A configuration is a User-Defined configuration that you will create to provide instructions on HOW and WHERE these tasks should be secured.

> *For example, "Migrating from Hypervisor A to VergeOS" is a recipe and a YAML file containing the host name, cloud regions and compute preferences, etc. is a configuration.*

MigrateOps recipes are actively being developed and added to the Cirrus Data Cloud platform.
In general, there are two categories of recipes:

- **automation-centric** recipes that facilitate storage integrations and migration at-scale. These recipes focus on resource allocation and preparation, migration, cutover and OS/ application remediation.
- **compute-mobility-centric** recipes that facilitate compute platform (physical/cloud/ hypervisor) integrations and OS/application integrations. These recipes handle the provisioning, migration, and remediation required to move workloads from one compute platform (hypervisor/cloud) to another.

Cirrus Data already has recipes available for migrating to VergeOS making it easier than ever to customize, and automate your migration.

Every step is logged, and every decision point can be pre-approved or gated by change management controls.

The most critical outcome is cutover time. Cirrus Data solutions reduce cutovers to seconds or minutes, not hours, by keeping workloads synchronized until the moment of switchover. Enterprises align migrations with existing maintenance windows rather than negotiating notable outages. Projects that were once unmanageable become achievable, and consolidation efforts can proceed without compromising production.

Always Ready Migration is a defining advantage. Migration is not a one-time event; it is an ongoing capability. If an enterprise acquires a new company or business unit, that environment—regardless of its platform—can be folded into the unifying destination through Cirrus Data. This prevents sprawl from re-emerging and ensures the long-term sustainability of consolidation.

In practice, Cirrus Data demonstrates how the defined path should work: spanning all platforms, protecting performance, automating operations, and minimizing downtime. By providing these capabilities, it proves that consolidation across diverse estates is not only possible but repeatable at scale.

THE CASE FOR A UNIFYING DESTINATION

A defined migration path solves part of the problem. Once workloads move, they must land on an infrastructure platform that eliminates the same weaknesses that caused instability in the first place. Without a unifying destination, migrations achieve little more than shifting workloads from one set of issues to another.

Requirements for a Unifying Destination

A unifying destination must unify virtualization, storage, networking, and AI under one operating model. The platform should operate from a single codebase rather than multiple products hidden behind a management GUI. This design reduces complexity, removes integration gaps, and delivers consistent performance. It must extend the usable life of hardware, allowing enterprises to run standard x86 servers and storage without being constrained by strict compatibility lists or forced refresh cycles.

The destination must also deliver strong resilience. Fault tolerance, automated recovery, and distributed data protection are non-negotiable. Security and compliance must be built in, not bolted on, with native multi-tenancy, tenant isolation, encryption, and full audit logging integrated into daily operations. Performance must scale up for core data centers and scale down for edge or remote offices, providing IT with a single platform that fits all locations. AI capabilities must be native to the platform, allowing organizations to run training and inference against their own data without setting up separate environments or relying on cloud providers.

Limitations of Current Destinations

Current platforms claim to provide a unified infrastructure, but in practice, they rely on multiple software modules developed by different teams. Integration is achieved at the GUI level, not at the code level. This approach introduces complexity, inefficiency, and inconsistent behavior across features. Hardware compatibility lists continue to shorten refresh cycles, driving up cost. AI readiness is left to add-ons or separate clusters, further fragmenting the infrastructure.

Changing the Infrastructure Model

A unifying destination set the legacy model on its head. It provides a single, cohesive infrastructure layer that reduces operating costs by eliminating stacked products and tool sprawl. It provides IT with a single system to operate across core, edge, and cloud-connected locations. It supports both current workloads and future requirements, such as private AI, without requiring separate deployments.

Enterprises that pair a universal migration path with a unifying destination break free from instability. They gain control of cost, extend hardware life, simplify operations, and position their infrastructure to meet both today's challenges and tomorrow's demands.

VergeOS as the Unifying Destination

VergeOS represents the design model that legacy platforms have failed to deliver. It unifies virtualization, storage, networking, and AI under a single operating system built from one codebase. Unlike other platforms that integrate loosely coupled modules behind a GUI, VergeOS provides true cohesion. It delivers consistent performance and predictable operations across core data centers, remote offices, and edge sites. Its architecture directly addresses the shortcomings outlined earlier, and its partnership with Cirrus Data extends its capabilities to complete the picture of enterprise consolidation.

- **Cost efficiency** is delivered through VergeOS's software-defined model, per-server licensing, and flexible deployment options. Customers can purchase multi-socket servers with petabytes of capacity without incurring software licensing penalties. VergeOS replaces multiple stacked products—hypervisor, SAN storage, network overlays, and separate AI modules—with one operating system, consolidating costs into a single, predictable platform.
- **Hardware freedom** comes from VergeOS's support for standard x86 servers and storage. There are no restrictive compatibility lists that shorten useful life. Organizations extend the life of their existing infrastructure, add new servers on their own schedule, and avoid wasting capital on premature refreshes.
- **Scalability and flexibility** are inherent in VergeOS's architecture. It supports large multi-node clusters for core data centers and runs just as effectively in two- or three-node configurations for edge sites, remote offices, and venues. IT can consolidate on a single operating system that adapts to both large and small environments.
- **Architectural consistency** is built into VergeOS. From the start, it was designed as one codebase for virtualization, storage, networking, and now private on-premises AI. Customers avoid the fragmented module strategy of other vendors and gain a platform with a clear, consistent roadmap independent of shifting bundles or acquisitions.

AI readiness is a native capability of VergeOS. It supports GPUs for training and inference directly on the platform, allowing enterprises to run AI workloads against sensitive data without relying on AWS or Azure. VergeOS can also run GPU-less for initial testing and validation, giving organizations a low-cost entry point to private AI strategies.

Integrated security and compliance are standard in VergeOS. Tenant isolation, distributed data protection, global inline deduplication, encryption, and audit logging are part of the core operating system, not bolt-on modules. Security is applied consistently across the environment, simplifying operations and reducing cost.

Seamless workload migration is achieved through VergeIO's partnership with Cirrus Data. Together, they provide a universal migration path that spans VMware, Hyper-V, Nutanix, Scale Computing, Proxmox, OpenStack, and public cloud IaaS. Cirrus Data handles synchronization, bandwidth management, and cutover, enabling enterprises to move workloads into VergeOS with minimal downtime and no data loss.

This combination creates an infrastructure foundation that is cost-effective, resilient, secure, and prepared for both current workloads and the emerging demands of private AI.

Before vs. After Cost Model (illustrative)

Category	Current State (Legacy Platforms)	Target State (VergeOS + Cirrus Data)
Licensing	Per-core fees; escalating support contracts; forced bundles	Per-server licensing; predictable cost; no penalties for multi-socket or high-capacity servers
Hardware Lifecycle	Strict HCLs shorten life to 3–4 years	Standard x86 runs 5–7 years; refresh on IT's schedule
Cloud Spend	Monthly IaaS bills with no asset value	In-house infrastructure with capitalized assets
Tools / Modules	Multiple products (hypervisor, SAN, DR, AI add-ons)	Single OS with virtualization, storage, networking, AI
Migration	Limited tools, VMware-centric, downtime measured in hours	Universal path across platforms, downtime in seconds/minutes

ROI AND TCO ANALYSIS

The financial impact of consolidating diverse platforms into VergeOS, enabled by Cirrus Data, is immediate and measurable. Enterprises no longer carry the weight of multiple hypervisors, hardware compatibility restrictions, and recurring cloud rent. Instead, they gain a single operating model that extends hardware life, eliminates redundant licensing, and reduces operational overhead.

Licensing and Support

Legacy platforms stack costs across hypervisors, storage, networking overlays, and add-ons for backup or AI. Per-core and per-socket licensing models drive costs higher as hardware scales.

VergeOS simplifies this through a per-server licensing model with no penalties for multi-socket servers or petabyte-scale deployments. Enterprises consolidate stacked products into a single license, reducing software costs by 40–60% in the first renewal cycle.

Hardware Lifecycle and Portability

Strict compatibility lists from legacy vendors force premature refresh cycles, typically every 3–4 years. VergeOS avoids this by running on standard x86 servers and storage, extending usable life to 5–7 years or longer. Its deep abstraction capabilities decouple workloads from the underlying hardware, making the environment portable across generations of servers and storage devices. This abstraction means enterprises can repurpose existing assets, mix hardware generations, and migrate workloads seamlessly without disruption.

Cirrus Data complements this portability by ensuring workloads transition smoothly during phases, allowing enterprises to consolidate onto VergeOS without discarding functioning hardware.
The combined effect is a two- to three-year extension of server life across the estate, delivering substantial TCO reductions while giving IT control over refresh timing.

Cloud Economics

Cloud IaaS creates recurring spend with no residual value. AWS and Azure costs rise with egress fees, storage growth, and hourly compute rates. By consolidating into VergeOS, enterprises bring steady-state workloads back in-house, converting unpredictable operating expenses into predictable, depreciable assets.

VergeOS not only matches but exceeds the cloud's claim of simplicity due to the cohesiveness of its single codebase and its deep abstraction capabilities. Cirrus Data accelerates repatriation, cutting egress timelines from months to weeks. This shift alone can save millions over a three-year horizon for enterprises with large cloud footprints.

Operational Efficiency

It is unlikely that organizations run all of these platforms, but it is not uncommon for them to run at least two, especially now with the VMware upheaval. It is also not unusual for most enterprises to maintain at least one arrangement with a cloud IaaS provider. Running VMware, Hyper-V, Nutanix, Proxmox, and OpenStack in parallel, along with a public cloud provider, requires specialized staff and a fragmented toolchain.

VergeOS eliminates this duplication with a single system to learn, a unified set of tools to manage, and consistent operations across core, edge, and ROBO. Cirrus Data further reduces labor by automating migration tasks with YAML-defined operations and approval workflows. The result is leaner IT staffing requirements and faster time-to-value.

Risk and Downtime Costs

Downtime during migration is a hidden cost. Legacy import utilities require outages measured in hours or days, halting projects and raising risk.

Cirrus Data reduces cutovers to seconds or minutes, while VergeOS delivers integrated high availability and fault tolerance post-migration. Enterprises lower both migration risk and ongoing operational risk, resulting in fewer incidents, shorter recovery windows, and an improved compliance posture.

Unprecedented ROI:

When modeled across a three-year window, consolidation into VergeOS with Cirrus Data consistently demonstrates ROI within the first 12–18 months. Savings stem from:

- 40–60% software license reduction.
- 2–3 year extension of hardware refresh cycles.
- 30% or more reduction in server count and power consumption.
- Multi-million dollar cloud repatriation savings.
- Reduced staffing costs tied to platform specialization.
- Always Ready Migration: consolidation is not a one-time project. When the enterprise acquires a new company or absorbs a new business unit, its infrastructure platforms—regardless of type—can be seamlessly migrated into VergeOS through Cirrus Data. This ensures long-term consolidation value and prevents sprawl from returning.

Lower TCO:

The shift removes stacked fees, hardware waste, cloud rent, and duplicated operations. Enterprises gain a predictable cost model, lower risk, and an infrastructure foundation that is simpler to run and ready for AI.

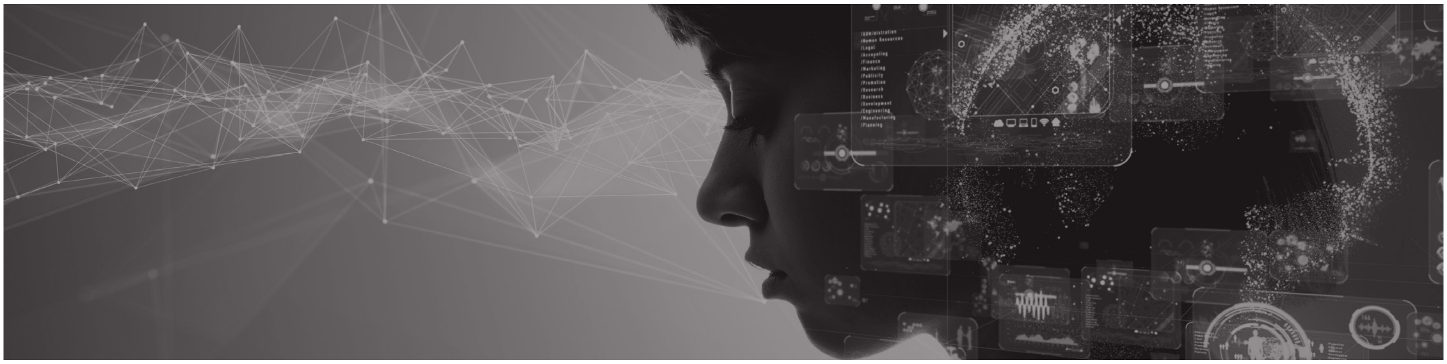

CONCLUSION

Legacy infrastructure is unstable by design. The vendors that dominate the market deliver higher prices, stricter hardware rules, limited scalability, and uncertain roadmaps. These weaknesses create sprawl across multiple hypervisors, cloud half-steps, and redundant platforms.

The way out requires both a universal migration path and a unifying destination. Cirrus Data provides the path, enabling seamless movement of workloads across all major platforms without disruption. VergeOS provides the destination, unifying virtualization, storage, networking, and AI in a single operating model that extends hardware life and lowers costs.

Together, they provide enterprises with a practical framework to exit instability. What makes this combination unique is its durability: consolidation is not a one-off initiative but an ongoing capability. As organizations grow, restructure, or acquire new entities, these environments can be seamlessly integrated into VergeOS with the help of Cirrus Data. This ensures that consolidation delivers an immediate ROI while continuing to protect against the return of sprawl. The result is infrastructure that is simpler to run, resilient against outages and cyber threats, and ready for the demands of private AI.