



Replatforming Virtual Infrastructure: Navigating Cost and Compliance Challenges

A practical guide for regulated mid-market IT leaders balancing cost shocks, compliance, and operational risk

Submitted by:

Anthony Mini - President

Rob Judd - Director Data Center Operations

March 25, 2026

Introduction

Mid-market organizations in regulated sectors (healthcare, finance, and others) are facing a pivotal decision point in their IT infrastructure strategy. A wave of industry changes – notably the overhaul of licensing models by a dominant virtualization vendor after its acquisition by a semiconductor giant – is forcing CIOs to rethink long-standing platform choices. Many IT leaders who have run on-premises or co-located virtualization stacks (instead of hyperscale public clouds) are seeing steep increases in licensing costs and new bundling requirements that compel them to pay for features they don't need. In some reported cases, companies are being quoted 3× to 10× higher prices upon renewal, a shock that few budgets can absorb.

This financial jolt, combined with strategic and operational pressures, is driving an industry-wide re-evaluation of virtualization platforms. Organizations are increasingly exploring alternatives – from open-source hypervisors to specialized private cloud solutions – in search of cost relief and better alignment with their needs. At the same time, CIOs in regulated industries must weigh these choices against strict security, compliance, and performance requirements that often limit wholesale moves to public cloud providers.

This white paper provides a comprehensive analysis of the key drivers and challenges for replatforming virtual infrastructure in mid-market, regulated environments. It is organized into ten sections, each addressing critical factors for decision-making and operations.

1. Strategic Drivers for Replatforming Virtual Infrastructure

Several strategic motivations are pushing CIOs to consider replatforming their virtual infrastructure. Foremost among them is the dramatic change in the vendor landscape and cost structure for incumbent virtualization solutions. After the acquisition and policy changes by a new owner, many customers have been hit with “take-it-or-leave-it” licensing terms and sharp price hikes, leaving mid-sized enterprises feeling “handcuffed” – either accept the increases or find an alternative. This has created an urgent strategic imperative: reduce dependency on a single vendor to escape an unpredictable cost trajectory and potential lock-in. The goal is not only immediate cost savings (addressed in detail in the next section) but also long-term control over the infrastructure roadmap. CIOs want to avoid being at the mercy of unilateral vendor changes in the future.

Another driver is the desire for modernization and alignment with cloud-native trends. The IT landscape is evolving toward containerization, DevOps, and hybrid cloud models. Some organizations see the current disruption as an opportunity to pivot – potentially bypassing traditional hypervisors entirely in favor of container platforms or cloud services. While a full move to public cloud may be infeasible for regulated firms (as discussed later), there is strategic value in adopting platforms that integrate better with cloud-native tooling (e.g., Kubernetes) and offer a cloud-like user experience on-premises. The status quo platform was often built for a client-server era and may lack the agility and API-driven automation that modern IT teams seek.

Competitive and operational considerations also play a role. Mid-market companies must remain agile and cost-effective to compete. If the current virtualization environment is leading to bloated costs or

slowing down deployment times (due to complex processes or limited automation), it directly impacts the business. Newer platforms promise streamlined operations – for example, simpler provisioning, built-in services (like object storage or load balancing), and easier scaling – which can translate into faster time to market for internal projects. CIOs are thus strategically motivated to evaluate solutions that can increase operational efficiency and innovation.

In summary, the strategic drivers for replatforming include cost control, avoidance of vendor lock-in, modernization of technology stack, and improvements in agility and efficiency. Together these create a compelling business case to explore alternatives even for organizations that have been comfortable with their virtualization status quo for years.

2. Cost Pressures and Licensing Model Disruptions

The most immediate catalyst for change has been the surge in costs and licensing model disruptions introduced by the incumbent platform's new ownership. Many organizations have found that what used to be a predictable expense has turned into a budget-busting line item. The vendor's shift from perpetual licenses to subscription-based, per-core licensing (with large minimums) has drastically increased costs for many mid-market IT shops. For example, under the new model a customer who previously needed one CPU license for a single server might now be forced to purchase a 72-core subscription bundle – even if their server has far fewer cores – representing a nine-fold increase in licensed capacity (and cost) for the same workload. This kind of change has been described as essentially a “Broadcom tax” on smaller deployments. Even larger environments are seeing similar effects; the reduction of the core allowance per CPU (from 32 to 16 cores) instantly doubled the license count needed for modern high-core-count processors.

Publicly available analyses and anecdotal reports paint a stark picture of the financial impact. Gartner noted that the licensing changes have “increased run costs by a factor of two or more” for many customers. Forrester researchers have cited anecdotal instances of 400% to 700% increases in costs – and in some extreme cases up to 800%. One large enterprise (a telecommunications company) even calculated that the new policies would raise its virtualization spending by 1,050%, an untenable jump that led it to threaten dropping support for tens of thousands of VMs. Such dramatic hikes are not universal, but double or triple previous costs have become alarmingly common.

The root causes are both the new per-core metric (often requiring more licenses for the same hardware) and the bundling of products into “suites” that eliminate cheaper editions. Under the revised scheme, customers who only needed a hypervisor and basic management find they must subscribe to a comprehensive cloud stack bundle. This “all-in-one” bundling forces organizations to pay for a broader selection of features than necessary, effectively charging for shelfware that provides no real value to them. As one industry expert put it, the new per-core + bundle approach is “masking that you are also responsible for paying for a lot of shelfware” that was previously optional.

Beyond pure licensing fees, support and maintenance costs have spiked as well. Annual support renewals that historically cost ~22% of license value are being increased into the 25–30% range by the vendor. This means over a typical 5-year period, a customer might pay 125–150% of the original license cost just in support fees – a substantially higher total cost of ownership than before. Moreover, under the new regime,

standard support may have been trimmed down, pressuring customers to pay for premium support tiers to get the level of service they used to receive. Mid-market organizations, which often rely heavily on vendor support due to lean internal teams, are seeing these support contracts become another source of budget pressure.

All these cost disruptions are forcing CIOs to rethink their investments. Many organizations simply cannot justify the skyrocketing expenses for virtualization, especially when it consumes budget that could be spent on digital transformation or customer-facing innovation. As a result, there is intense pressure to find cost-effective alternatives. Some are negotiating fiercely for concessions, but the new vendor has taken a notoriously rigid “take it or leave it” stance, even implying that refusing the new terms could result in license termination or support cutoff. In effect, the vendor appears “comfortable ‘firing’ smaller accounts” that won’t pay the new premiums. This hardline approach leaves mid-market IT leaders with a stark choice: pay exorbitantly more, or explore migrating to a different platform. Unsurprisingly, we are seeing an exodus begin – with companies actively evaluating other hypervisors, hyper-converged stacks, and cloud solutions purely because of cost-saving imperatives. In sections to come, we will examine how these alternatives stack up, but it is clear that the economic pressure alone is a primary driver for change.

3. Capital vs. Operational Expenditure Trade-offs (CapEx vs. OpEx)

When considering a move away from a traditional virtualization platform, CIOs must also grapple with the financial model implications – specifically, capital expenditure versus operational expenditure. Many mid-market firms have long followed a CapEx model for infrastructure: they purchase server hardware, hypervisor licenses, and related equipment upfront, then depreciate those assets over time. Shifting to a new solution, however, may open the door to different consumption models. For instance, adopting a cloud-like platform or an Infrastructure as a Service (IaaS) from a provider could turn what used to be a large periodic capital outlay into a steady monthly operational expense. Each approach has its benefits and challenges.

CapEx investments provide a degree of cost predictability and control that can be especially valuable in regulated industries. Upfront investments are expected and consistent, which is beneficial for budgeting, and the company owns the asset outright, maintaining full control over its use, security, and update cycle. This ownership aligns well with compliance needs – as discussed later, many firms feel more secure knowing their data and systems reside on hardware they control. CapEx also allows leveraging an asset for its full lifespan with no ongoing subscription fees, potentially yielding lower TCO if utilization is high and the equipment’s life is maximized.

On the other hand, OpEx models (like cloud subscriptions or managed service agreements) offer attractive flexibility. There are smaller upfront costs and faster approval cycles, since one isn’t committing a huge sum at once. The pay-as-you-go nature means you can scale resources up or down as needed without having stranded capacity (no need to overprovision for peak and leave hardware idle). This agility can be financially efficient and better align with operational demand – you pay for what you actually use. Additionally, OpEx shifts much of the maintenance responsibility to the provider: when using a cloud or managed service, the provider handles hardware refreshes, patching, and infrastructure upgrades as part of the service. For mid-market IT teams under strain, offloading these tasks can free up staff to focus on higher-value projects.

The trade-off, of course, is that OpEx approaches can sacrifice some control. By renting infrastructure or platform services, a business gives up direct ownership and must trust the provider for reliability and security. Costs, while smooth, can vary month to month based on usage, and it can sometimes be harder to predict or explain those fluctuations. Moreover, companies in highly regulated sectors often discover that CapEx ownership feels safer for compliance – they “know where their data is at all times” when it’s on their own servers. In contrast, OpEx models require diligence to ensure the provider’s compliance and security measures meet all requirements.

Many mid-market CIOs are now performing a CapEx vs. OpEx analysis in the context of virtualization replatforming. For example, if they choose an open-source hypervisor on self-owned hardware, they remain in a CapEx posture (buying servers and perhaps purchasing support contracts, but largely investing upfront). If they decide to use a hosted private cloud service or a managed hypervisor-as-a-service, that might shift costs to OpEx via monthly fees. There is no one-size-fits-all answer – some organizations with tight capital budgets may prefer an OpEx model to turn large purchases into manageable operating expenses, whereas others with available capital might invest in CapEx to avoid recurring charges and have more control. The key is aligning the financial model with corporate financial strategy and risk tolerance. CIOs must also communicate these trade-offs in financial terms to their CFO and leadership team, translating IT choices into business impacts (e.g., depreciation vs. expense on the balance sheet, and the implications for cash flow and flexibility).

In sum, replatforming is not just a technical decision but a financial one: mid-market IT leaders need to weigh whether buying and owning infrastructure or “renting” it via services will best meet their cost and compliance objectives.

4. Migration Friction and Data Gravity

Switching virtualization platforms is not as simple as flipping a switch. There is significant migration friction involved, owing to differences in hypervisor technologies, file formats, and integrated tooling. CIOs must plan for the technical challenges of moving dozens or hundreds of VMs off their existing platform and onto a new one.

One major hurdle is the conversion of VM disk formats and configurations. Each hypervisor has its own way of encapsulating virtual machine disks and settings (for example, VMware’s VMDK disks vs. KVM’s QCOW2 format). Converting these reliably requires tools or complex processes. While there are migration utilities (e.g., VMware’s exporter or third-party tools to convert VMDK to other formats), it’s rarely seamless. In many cases, VMs may need to be exported, copied, and imported into the new system, potentially incurring downtime. One example from an alternative platform approach: administrators used an open-source virt-v2v tool to automate offline conversion of VMware images to KVM format. This works, but it’s an extra step that must be scripted and verified for each VM, and certain aspects like guest OS drivers might need adjustment post-conversion.

Beyond VM images, there are also networking and security constructs that often don’t translate automatically. Virtual switches, VLAN configurations, or software-defined networking rules from the old environment will likely need to be recreated in the new one. For instance, a company using NSX micro-segmentation or distributed virtual switches in VMware cannot automatically port those to a different

platform – administrators must manually rebuild equivalent network policies and security groups in the new one. Similarly, automation scripts (PowerCLI, vCenter Orchestrator workflows, etc.) and backup systems tied into the old environment’s APIs will not work out-of-the-box on a different platform. A research report noted that backup systems based on VMware’s APIs could not be reused and had to be replaced or re-architected when moving to an alternative virtualization stack. This means new backup solutions or methods might be required in the target platform (covered later in more depth, but crucial to flag during migration planning).

Data gravity is another critical concept: wherever large amounts of data reside, they tend to “anchor” workloads due to the difficulty and cost of moving that data. Mid-market firms often have terabytes or petabytes of VM images, databases, and archives sitting in their current virtualization environment (perhaps on a SAN or local storage). Migrating this data to a new platform means copying all those bytes, which can be time-consuming and expensive. In a co-location or on-prem scenario, it may involve setting up temporary storage or networks to transfer data between old and new systems. If considering a cloud migration for some workloads, egress fees come into play – cloud providers charge significant fees to pull data out, so any exit from a cloud or even between clouds can rack up costs. CIOs need to account for these in their ROI analysis. Sometimes, the sheer volume of data suggests a phased migration or a hybrid approach (keeping some data in place to avoid the gravity problem).

Lastly, provider or vendor cooperation (or lack thereof) can add friction. Incumbent vendors rarely have incentives to assist customers in leaving. For example, while you can always get your data, you might find that your license to use certain tools (like live migration or replication features) expires if you don’t renew, making it harder to do a graceful transition. It’s important to carefully read contracts and perhaps negotiate migration support as part of any renewal or exit plan. Forward-thinking organizations maintain an “exit strategy” from day one – ensuring data is in standard formats and that leaving a service is not prohibitively difficult or costly. This includes verifying that any new provider or platform has documented procedures for extracting your data and VMs if needed, and that there are no hidden penalties for doing so.

In summary, migration friction and data gravity mean that CIOs must approach replatforming with careful planning and realistic timelines. Strategies like pilot migrations, dual-running environments, and swing hardware can mitigate some pain. Some companies choose to enlist expert partners or consultants who specialize in virtualization migrations to handle format conversions and cut-over scheduling. The key is not to underestimate the effort: moving critical workloads between virtualization platforms is a project that touches storage, network, compute, and application layers. With thorough preparation, testing, and an eye on where data gravity could snag progress, CIOs can overcome these frictions – but they must be baked into the project plan from the start.

5. Security and Compliance Constraints on Public Cloud Adoption

For CIOs in healthcare, finance, government, and other regulated sectors, any discussion of changing infrastructure inevitably raises the question: why not move to the public cloud? Hyperscale cloud providers (AWS, Azure, GCP) could eliminate the need to run your own virtualization stack, offering on-demand infrastructure with operational convenience. However, the realities of security and compliance

requirements make public cloud a non-starter or a limited option for many mid-market regulated organizations.

These firms operate under strict mandates such as HIPAA (health data privacy), PCI DSS (payment card security), GLBA (financial data protection), and a host of other laws and frameworks. Compliance requires demonstrable control over data, auditability of systems, and often an on-premises or single-tenant approach to guarantee isolation. Many companies thus find that keeping systems on-premises (or in private hosted environments) is the safest route to ensure compliance. In fact, businesses in highly regulated industries often feel the decision is already made for them: they “must house their applications on-premise,” because that’s the only way to be absolutely sure of data location and security. When your organization is subject to laws like HIPAA, you need to know where your data is at all times and maintain strict control over who accesses it. Public cloud, by its nature, involves a third-party controlling the infrastructure and often dispersing data across regions (unless specifically configured and verified otherwise). Even if a cloud provider offers compliance assurances and signs Business Associate Agreements (in the case of HIPAA), many mid-market IT leaders remain cautious.

Security concerns compound this hesitation. While large cloud providers invest heavily in security, the perceived risk of multi-tenancy and being a high-profile target remains a barrier. In contrast, an on-premises environment, though not invulnerable, offers a sense of control: one can deploy network segmentation, internal firewalls, and physical security measures tailored to their specific threat model without relying on someone else’s shared platform. Additionally, regulated companies often undergo rigorous audits. Having infrastructure in-house can make audit scope clearer and under direct oversight. In a cloud scenario, proving compliance might require relying on the provider’s attestations and ensuring they cover all concerns, plus ensuring your configuration meets requirements. Data residency adds another layer: while cloud regions exist, demonstrating and enforcing residency can be operationally complex.

That said, regulated businesses do use public cloud in constrained ways. Major cloud providers have obtained certifications that indicate they can support HIPAA, PCI, FedRAMP, and related requirements. However, using the cloud in a compliant manner still places responsibility on the enterprise to configure and manage it correctly. Mid-market organizations may lack the specialized skills to confidently secure and monitor public cloud to the standard required. And if a breach occurs in the cloud, it is still the organization’s breach – liability and impact remain with the company.

Given these factors, many CIOs conclude that a private or hybrid cloud model is more appropriate. They seek cloud-like operations (self-service, scalability) implemented on single-tenant or fully isolated infrastructure within on-prem or co-location environments. This allows them to meet compliance obligations and maintain oversight. For many mid-market firms, the practical reality is that core systems and sensitive data will continue to reside on infrastructure they control, and public cloud will remain adjunct for select use cases. This constraint shapes the solution set: the chosen platform must support strong security controls, robust compliance features (encryption, auditing, data locality), and typically private/hybrid deployment patterns rather than assuming a full hyperscale migration.

6. Risks of White-Labeled VPS Offerings and Unvetted Cloud Providers

In the scramble to find alternatives to the incumbent virtualization suite, some organizations might be tempted by “quick fix” solutions – small service providers offering to host your VMs, or resellers pitching a white-labeled virtualization environment promising similar functionality at a lower cost. While there are credible regional cloud providers and managed services, CIOs should approach lesser-known or unvetted providers with caution.

One concern is security and compliance due diligence. Major cloud providers publish extensive documentation of controls and compliance certifications (SOC 2, ISO 27001, PCI, HIPAA, etc.). Smaller providers should be able to demonstrate equivalent evidence. If a vendor cannot produce independent audit reports or compliance attestations, that’s a red flag. A white-labeled VPS offering may be a repackaged open-source stack running on unclear infrastructure. Without transparency, you cannot validate data isolation, encryption, patching discipline, or incident response maturity. For regulated sectors, entrusting data to an uncredentialed provider can create compliance exposure and real security risk. Verification is non-negotiable: request certifications, audit reports, architectural documentation, and incident procedures.

Reliability and business continuity are also risks. Hyperscalers and mature providers invest heavily in redundancy; smaller providers may operate in a single data center with limited failover. If their service goes down, you may be stuck. If they go out of business, you may have limited time to extract data and rebuild elsewhere. CIOs should assess provider financial stability, support capability (24/7 with meaningful SLAs), and contractual exit provisions, including data export formats and timeframes.

White-labeled solutions from MSPs can be viable, especially if built on known platforms, but they introduce “support indirection” (you call the MSP, the MSP calls the platform vendor) and sometimes lag in upgrades. Also watch for vendor lock-in at the provider level: some environments make export or migration painful. Exit strategy matters: prefer standard formats, documented export procedures, and clear responsibilities in the contract (who patches what, who owns incident response, who owns backups).

This doesn’t mean small providers are bad; many are excellent. It means CIOs should apply strict due diligence: ask for customer references in similar regulated industries, run pilots with non-critical workloads, validate security controls, and ensure contractual clarity. Otherwise, you risk trading one problem (high cost) for a worse one (downtime, security failure, or compliance violation).

7. Technical Complexity and Staff Readiness for Open-Source Tools (e.g., Proxmox)

Open-source virtualization platforms (KVM-based stacks such as Proxmox VE, oVirt, XCP-ng, etc.) are attractive because they reduce or eliminate hypervisor licensing fees. The catch: they can introduce operational complexity and require staff readiness that many mid-market teams underestimate.

Open-source tools are powerful but often less turnkey. Proxmox, for example, is a wrapper around mature components (Linux, KVM, Ceph), and many advanced capabilities require deeper Linux knowledge. Networking, storage, and clustering can involve command-line configuration, tuning, and careful change control. Compared to the polished wizard-driven workflows VMware teams may be used to, the shift can feel abrupt. High availability, quorum, fencing, and multi-node cluster behavior may require more hands-on engineering and troubleshooting. Upgrades and lifecycle management can also be less “one click,”

especially across a cluster, and staff must be comfortable with maintenance sequencing and rollback plans.

Support models differ. You can buy enterprise support for many open-source stacks, but the experience may not match the expectations set by longstanding commercial ecosystems. Without paid support, you're relying on community forums and documentation, which can be excellent but not SLA-backed. For teams without deep Linux virtualization expertise, the cost savings on licensing can be partially offset by added labor, training, or external consulting. Many organizations bridge the gap by pairing open-source platforms with MSP support or commercial subscriptions to secure enterprise-style responsiveness and accountability.

Feature and ecosystem differences matter too. VMware's ecosystem is huge; tooling integrations, third-party products, and established runbooks are common. With open-source stacks, you may need more custom scripting, tool integration (monitoring, reporting, backups), and internal engineering. That's not a deal-breaker — but it's a different operational posture.

The practical takeaway: open-source platforms can be an excellent strategic move, but CIOs should treat it as both a technology transition and a people/process transition. Run a lab, pilot real workloads, invest in training, and decide up front whether you need commercial support or a partner. "Free software" can still be expensive if you don't budget for the human side of operating it well.

8. Platform Feature Gaps: Multi-Tenancy, Self-Service, Audit Logging, Monitoring, etc.

When evaluating alternatives, CIOs must identify which enterprise capabilities they currently rely on — and whether the new platform will match them out-of-the-box. Over time, many environments accumulate dependencies: role-based access control (RBAC), multi-tenant isolation, self-service provisioning portals, audit trails for compliance, monitoring/alerting integrations, and mature automation ecosystems. Some alternative platforms — particularly lighter-weight or open-source ones — may lack parts of this by default.

Multi-tenancy is a common gap. Platforms designed for single-organization use may not provide strong tenant isolation, quotas per project, or safe delegated administration without add-ons. That becomes a major issue for IT departments acting as internal service providers, or for MSPs hosting multiple customers. Without native tenancy controls, either admins remain the bottleneck or you risk over-permissioned access.

Audit logging and compliance evidence can also vary. Regulated environments often need strong "who did what, when" logs across UI and API actions. If an alternative's audit trail is limited, you may create a compliance blind spot and reduce your ability to investigate incidents.

Monitoring and alerting is another area where teams can get surprised. Some platforms provide basics but require external tooling for enterprise-grade visibility: capacity trends, predictive alerts, anomaly detection, SLA reporting. This isn't necessarily bad — mature monitoring stacks like Prometheus/Grafana or Zabbix can be excellent — but it adds integration work and operational overhead.

Networking and storage features can have similar gaps. Advanced SDN capabilities (micro-segmentation, policy-driven networking, multi-tenant overlays) may be absent or less mature. Storage integrations may be more manual, and enterprise SAN features may not map cleanly. Self-service portals and orchestration tools vary widely; some platforms focus strongly here, while others assume traditional admin-driven provisioning.

The right move is to build a requirements checklist before picking a platform: must-have features (RBAC, audit logs, MFA integration, monitoring hooks, backup compatibility, tenancy) and acceptable gaps (things you'll supplement). Some ecosystems (open-source cloud managers layered over hypervisors) exist specifically to fill these gaps, particularly around self-service and multi-tenancy. The worst-case scenario is saving on licenses but losing control and visibility — which can quickly turn into outages, security incidents, or audit pain. Feature gap analysis is where migrations are won or lost.

9. Key Infrastructure Considerations: Redundancy, Oversubscription, and Support

Beyond features and price, platform viability rests on infrastructure fundamentals: uptime, performance, and the ability to operate the system reliably under stress.

Redundancy and high availability come first. Any alternative must support automatic recovery from host failure, planned maintenance without major downtime, and ideally live migration or equivalent mechanisms. How that is implemented varies: shared storage clustering, hyperconverged replication, or distributed storage. CIOs should verify that HA behaviors are understood, tested, and operationally “boring” (boring is good). Also consider disaster recovery: can the platform replicate to another site or integrate with DR tooling?

Oversubscription — overcommitting CPU/memory assuming not all workloads peak simultaneously — drives virtualization economics. Different platforms handle scheduling and contention differently. CIOs need to understand how the new stack manages resource guarantees, priorities, and noisy-neighbor behavior, and whether it provides sufficient controls to protect critical workloads. Capacity planning practices may need to change, and monitoring becomes essential to avoid performance cliffs.

Support and ecosystem matter more than teams like to admit. With mature platforms, you get predictable escalation paths, a deep consultant ecosystem, and lots of learned tribal knowledge. With newer or open-source alternatives, you may need to secure paid support, partner services, or internal expertise to reach the same reliability. Hardware compatibility and drivers also matter: ensure your servers, NICs, and storage controllers are supported and well-tested in the target environment.

Operational readiness should be tested, not assumed. Run failure scenarios: “Host dies at 3 AM,” “storage latency spikes,” “cluster loses quorum,” “certificate expires,” “backup repository fills.” Confirm monitoring catches it, alerts route correctly, and your team can respond confidently. A well-designed migration isn't just getting workloads moved; it's ensuring the new platform can meet or exceed service levels under real-world chaos.

10. Snapshot, Backup, and Restore Realities Across Platforms

Data protection is non-negotiable. When changing virtualization platforms, backup and recovery is often where hidden complexity and costs show up. Many environments rely on backup tools integrated into the incumbent platform's APIs (for example, VMware's VADP ecosystem). When you move away, those integrations may not transfer cleanly. Some tools support multiple hypervisors, but feature parity may vary. In some cases, backup architectures must be redesigned.

Snapshots are also not interchangeable across platforms. Performance characteristics, snapshot chain limitations, and the ease of creating application-consistent snapshots can differ significantly. Teams accustomed to certain workflows (quick snapshots before patches, rolling back) need to re-learn caveats in the new environment. Snapshots are not backups, and this needs to be reinforced during transitions, especially if new tools make snapshots feel more accessible.

Restore workflows deserve special attention. It's easy to focus on "can we back up," and neglect "can we restore fast." CIOs should require test restores on candidate platforms to validate RTO/RPO expectations. During migration, ensure there is no coverage gap — every workload should be protected continuously, even if it lives temporarily in a dual-run phase.

Also watch for "new costs" that sneak into the design: backup server licensing, replication features, offsite storage, or DR orchestration that was previously bundled or already owned. Even when the hypervisor is "free," enterprise-grade backup often is not.

The best practice is to validate backup tooling support early, build a transitional backup plan (covering both old and new environments), perform restore drills, and keep legacy backups accessible for a period after migration (especially for less frequently used systems). If auditors or regulators are involved, demonstrate that backup integrity and recoverability were tested and documented. You can change hypervisors; you cannot gamble with recoverability.

Conclusion: Charting a Future-Ready Path – Key Considerations and Solutions

Faced with rising costs, changing vendor dynamics, and the unique constraints of regulated environments, mid-market CIOs are right to carefully chart their next steps. The challenges outlined – from strategic realignment and financial modeling to technical migration hurdles and feature gaps – may seem daunting. However, they also point toward a forward-looking solution space. The ideal path forward is one that delivers cloud-like agility and economics while preserving the security, control, and tailored support that mid-market regulated organizations require.

In evaluating options, CIOs should prioritize solutions that address the critical needs revealed by this analysis:

- **Cost efficiency with predictability:** Break the cycle of opaque pricing, forced bundles, and runaway renewals. The best alternatives offer straightforward pricing and reduce the proportion of IT budget consumed by virtualization overhead.
- **Multi-tenancy and self-service (where needed):** For IT departments acting as internal providers — and especially for MSPs — strong tenant isolation, RBAC, quotas, and self-service provisioning reduce bottlenecks and improve governance.

- **Audit logging and compliance readiness:** Detailed logging, strong identity integration, encryption, and access control maturity are essential for regulated environments.
- **Backup and DR as first-class capabilities:** The future-ready platform either integrates robust protection workflows or cleanly supports best-in-class third-party tooling, with proven restore performance.
- **Modern networking and SDN:** Native support for flexible virtual networking, segmentation, and hybrid connectivity reduces complexity and improves security.
- **Security automation and lifecycle management:** Rapid patching, security hardening guidance, and operational tooling reduce risk without increasing downtime.
- **Hardware flexibility:** Avoid shifting lock-in from software to appliances. Hardware-agnostic platforms preserve freedom and protect future sourcing decisions.
- **Operational ease and dependable support:** The platform should reduce cognitive overhead, not increase it. Strong support — vendor or partner — is often what turns a “good platform” into a safe enterprise decision.

While this paper has remained vendor-neutral, the market is clearly moving: platforms are emerging that combine open-source cores (often KVM-based) with enterprise-grade management, support, and cloud-like features such as tenancy, portals, logging, and integrated storage/networking. These solutions aim to deliver private/hybrid cloud “on your terms”: data remains in your control, compliance is achievable, but operationally it behaves more like modern cloud.

The bigger point: CIOs should treat the current disruption as an opportunity rather than purely a defensive reaction. With careful selection and disciplined migration planning, mid-market organizations can emerge with infrastructure that is cheaper, more strategically flexible, and better aligned to the next decade — hybrid models, containerization, edge workloads, and increased regulatory scrutiny.

Do the due diligence, model the costs honestly (including staff readiness and migration effort), test the operational realities (especially backup and HA), and then move decisively. Organizations that do this well won’t just avoid pain — they’ll buy themselves real leverage and optionality in an infrastructure market that has become far less forgiving.