

CHOOSING INFRASTRUCTURE FOR THE NEXT DECADE

VIRTUALIZATION, CLOUD, AND THE COMING AI BUILDOUT

SUMMARY

The deadline for deciding about VMware is now a fixed point on the calendar. On October 11, 2027, general support for vSphere 8 ends. It is the last VMware release that can be licensed perpetually. Technical guidance continues into 2029, but the practical decision arrives sooner. Every organization running VMware must answer a direct question: Do we commit fully to VMware Cloud Foundation (VCF), deploy the less feature-rich VMware vSphere Foundation (VVF) bundle, or move on?

The VMware deadline is the entry point to this discussion, not the whole story. The platform an organization settles on now is the one most likely to carry its cloud-native modernization and AI workloads for the next decade. So, the VMware-specific question really reveals a convergence decision, and it's as much a choice about an operating model as about technology.

There is no universal right answer. The goal is to have a platform and an operating model matched to an organization's workloads, its readiness, and its AI ambition. And every organization is different, with different priorities and technical decision drivers.

This research brief examines the looming decision that many organizations will face and the convergence of factors that will require IT organizations to approach infrastructure modernization differently. It also provides a checklist for CIOs and IT leaders to help guide modernization strategies. Finally, it looks at the competitive landscape and discusses where solution providers such as VMware, Nutanix, and Red Hat stand, considering their strengths and how their offerings align with specific customer needs.

THE END-OF-SUPPORT DATE FORCES THE DECISION

For many organizations, the renewal date is what appears to set the timeline. That's when the bill arrives and the conversation can no longer be deferred, so until then it's easy to put off. But the date that actually governs is October 11, 2027, when general support for vSphere 8 ends. After that, environments run without security patches, and few enterprises will accept that exposure for what could be a year or more. For those in

regulated industries, the problem is more immediate, since unpatched infrastructure can put them out of compliance as soon as support lapses.

In the context of VMware, many organizations see real value in VCF 9.1, but paying for the full VCF stack can be difficult to justify if they're leveraging only part of its capabilities. While there is no market data that measures platform consumption, customer frustration over paying for unused capabilities has become a recurring theme in industry reporting.

Broadcom's strategy over the past couple of years has been to move customers toward the complete platform. By its own reporting, roughly 87% of its top 10,000 customers have adopted VCF, while VMware revenue has continued to grow through license conversions. That focus made sense, because these large enterprises were the customers most likely to use enough of the platform to justify the broader subscription.

Now, however, the company has shifted its focus to the commercial enterprise and mid-market. This is the vast majority of VMware's (reported) approximately 300,000 customers, many of whom are serviced by the channel and have leveraged vSphere to deliver operational agility.

As mentioned, for larger environments already using much of the VMware portfolio and willing to lock into the platform, VCF can make sense. For organizations with simpler needs, Broadcom's answer is the VVF bundle, though that carries its own capacity limits and licensing overhead. Organizations must make a choice regardless, yet the most common posture so far has been to wait. Teams are evaluating, delaying, and negotiating with Broadcom rather than committing to a course of action. While this is understandable, waiting also comes at a price. Delay compresses timelines, narrows the available options, and tends to erode negotiating leverage as the deadline closes in.

Deciding deliberately beats delay.

The first step is to acknowledge the difficulties that are sure to follow from this decision. Server virtualization is far more than a hypervisor. It encompasses the management plane, the automation, the networking and storage integration, and the operational habits built up over years. This is why switching is a considerable undertaking. And that's true regardless of the path chosen, whether that means moving to an alternative vendor, adopting a lighter VMware bundle like VVF, or committing to the full VCF stack.

THE BIGGER DECISION: INFRASTRUCTURE CONVERGENCE AND OPERATING MODELS

Moor Insights & Strategy (MI&S) believes this discussion requires a broader framing. If a company chooses to migrate from VMware, the target platform will, in all likelihood, also serve as its cloud-native platform and its AI platform. Virtualization, Kubernetes, AI infrastructure, and AI operations are converging onto shared foundations. That's why treating them as separate initiatives misses the point; increasingly, they're different expressions of one infrastructure strategy.

Choosing for virtual machines alone means making a ten-year decision with one eye closed.

This convergence is happening for practical reasons. New applications increasingly ship as containers. AI workloads run on Kubernetes across CPUs and GPUs. Data gravity pulls compute toward where the data already lives. And operations teams would rather run one control plane than stitch together three.

What's easy to miss is that this isn't just a technology decision. It's an operating model decision. A traditional virtualization model puts the infrastructure team in charge. A cloud or self-service model shifts more responsibility to developers. Platform engineering introduces a dedicated team to build and operate shared services. AI adds another layer, with data scientists and AI engineers consuming infrastructure in very different ways. Each model changes who owns the infrastructure, who consumes it, and how it's operated.

Organizations rarely struggle because they picked the wrong technology. More often, they struggle because the technology doesn't align with how they're organized to work. The best platform isn't necessarily the one with the longest feature list. It's the one that fits the organization's operating model today *and* supports where it wants to go tomorrow.

THE AI BUILDOUT: FROM GPU CLUSTERS TO AGENTIC SYSTEMS

The decision about what to do with VMware is playing out in the context of the massive AI buildout. [McKinsey projects](#) roughly \$6.7 trillion in datacenter capital spending by 2030, with about \$5.2 trillion of that going to AI-capable facilities and AI driving close to 70% of new datacenter demand. [Goldman Sachs models](#) around \$765 billion in annual

AI capex in 2026, more than doubling to \$1.6 trillion by 2031. Whatever the precise figure, the direction is unmistakable.

Meanwhile, AI adoption is mainstream but still early. [McKinsey's 2025 State of AI report](#) found that 88% of organizations use AI in at least one business function, 62% are experimenting with agents, and 23% are running agentic AI systems somewhere in the business. Yet nearly two-thirds haven't scaled enterprise-wide. So the appetite is real, but the maturity is uneven.

Kubernetes is quietly becoming the substrate for this work. [CNCF's annual survey](#) reports that 82% of container users run Kubernetes in production, and 66% of organizations hosting generative AI use it for inference. The container platform and the AI platform are turning out to be the same. More than that, Kubernetes is emerging as the common abstraction layer beneath modern infrastructure, whether a workload begins life as a container or a virtual machine. This common abstraction layer is a major driver of the convergence this paper describes.

The agentic shift raises the bar for infrastructure. AI is no longer just a GPU cluster running a model. Agentic systems need multiple models, orchestration, vector databases, governance, access to private data, and policy controls working together. That asks far more of the underlying platform. It must integrate compute, data, Kubernetes, networking, and governance, not simply schedule GPUs.

It's also worth discussing where GPUs are actually required. While training and large-scale model-serving use GPUs, a good deal of agentic work, including orchestration, retrieval, vector search, and inference on smaller models, runs well on modern CPUs. This means that an AI-ready platform must address not only GPU scheduling but also general-purpose infrastructure needs. The CPU-powered infrastructure an organization already runs tends to carry more of the agentic load than GPU-first framing would suggest.

There's a parallel pull toward private AI. In the vast majority of MI&S's conversations with enterprise leaders, repatriation is a topic of discussion, driven by data gravity and sovereignty, inference economics, cost, and governance. This is not about a mass exodus from the cloud, but rather a normalization. Taken together, all of this means the right platform has to span virtualization, cloud-native, and AI or agentic workloads on a single foundation.

FOUR PATHS FORWARD, AND WHY COEXISTENCE IS STRUCTURAL

Before looking at the infrastructure options, it's important to recognize that there is no longer a "do nothing" path. Even organizations that decide to stay with VMware are facing meaningful change, because VCF 9 isn't an in-place upgrade. It introduces a new management model, new operational requirements, and a migration to the VCF 9 architecture. The decision is no longer whether to change. It's what kind of change an organization is prepared to undertake. That said, Broadcom's smaller VVF bundle does give some organizations a lighter way to stay.

With this in mind, there are four paths for the VMware decision, with AI as a layer that crosses all of them rather than a fifth option on its own.

The first is to stay with VMware. The appeal is twofold: a mature stack, and avoiding a move to another vendor. However, staying does not mean standing still, because VCF 9 introduces 180-day mandatory compliance reporting, a required VCF Operations console, and a management domain built on new Fleets and Instances constructs. It also installs mandatory VCF 9 components whether a team needs them or not. The practical effect is that organizations can no longer wrap a VCF entitlement around an existing deployment. They have to deploy VCF 9 and migrate their environment to it. That represents meaningful operational change, even without changing vendors. For large environments already using much of the VMware portfolio, that tradeoff can make sense; for others, it probably makes less sense.

As mentioned, Broadcom now positions VVF 9 as a smaller bundle for teams that mainly want core virtualization without the full Cloud Foundation stack. It delivers that through vSphere and vSAN, while adding basic Kubernetes support via the vSphere Kubernetes Service. The tradeoffs are real, though. VVF caps the included vSAN capacity and leaves out NSX, VCF Automation, and VMware Private AI Foundation. And licensing still routes through a VCF Operations appliance, with the same subscription reporting that VCF requires. So, while VVF can lower the entry cost for a smaller vSphere-centric environment, it also narrows the cloud-native and AI headroom that this decision is ultimately about.

The second path is to virtualize on another platform in a like-for-like move. This preserves the virtual machine operating model and minimizes near-term disruption, which is its primary advantage. The tradeoff is that it remains VM-centric and defers the modernization question. This choice fits organizations that are prioritizing cost, continuity, and a familiar operating model.

The third path begins the move from virtualization toward cloud-native on another platform. This is the modernization path, opening the door to containers and AI, but it also requires new skills, new operational practices, and some application rearchitecture. This approach fits organizations that are genuinely ready to modernize rather than simply talking about it.

The fourth path goes all-in on public cloud. It offloads infrastructure management and offers real elasticity, balanced against cost at scale, data gravity, and the repatriation dynamic described earlier. It fits cloud-native workloads and organizations that have already committed to a cloud operating model.

Across all four, the AI lens remains the same. Each path should be evaluated based on its AI and agentic readiness, access to GPU infrastructure, proximity to data, and the governance it can enforce.

Regardless of which path an organization chooses, the work starts well before the first workload moves. Teams have to inventory the environment, map the application dependencies built up over the years, and separate the workloads that move cleanly from those anchored by integrations, operational requirements, or licensing constraints. Storage and networking rarely translate one-for-one, so data stores, addressing, and security policies often need redesign rather than a simple lift. The harder part is doing all of this while keeping the business running, which is why any credible plan moves in phases, preserves a rollback path, provides for retraining the operations team, and runs the old and new environments side by side for a stretch. That overlap is what turns coexistence from a nice-to-have into a structural part of the plan.

In practice, most organizations won't cleanly pick one path. They'll spend years in a hybrid-modernization middle, where existing VMs remain VMs, ERP stays virtualized, new applications arrive as containers, and AI workloads run on Kubernetes and GPUs. Thus, the useful question isn't "VM or cloud-native?" It's which platform gives the organization the most flexibility to run all of these at once and evolve at its own pace. The platforms a CIO is most likely to evaluate fall into a short list.

Platform	Virtualization	Cloud-Native (K8s)	AI / Agentic	Best-Fit Buyer
VMware VCF	vSphere + vSAN	VKS (formerly Tanzu)	VMware Private AI	Fully utilized VMware private cloud-centric estates
VMware VVF	vSphere + vSAN	VKS (basic)	vSphere GPU; no Private AI Foundation	Smaller vSphere-centric estates
Nutanix	AHV (license included)	Nutanix Kubernetes Platform	Enterprise AI plus agentic solution	Organizations wanting one platform for VMs, K8s, and AI anywhere
Red Hat	OpenShift Virtualization (KubeVirt)	OpenShift (native)	OpenShift AI	Container-first, developer-heavy shops
Microsoft	Hyper-V / Azure Local	AKS / Arc-enabled	Azure AI (hybrid via Arc)	Microsoft-centric, Azure-aligned shops
Public cloud	Cloud-native VMs	Managed K8s (EKS / AKS / GKE)	Cloud AI services	Cloud-committed, low data gravity shops

WHAT THE RIGHT INFRASTRUCTURE PLATFORMS HAVE IN COMMON

Step back from the vendors for a moment. The platforms capable of carrying an organization through the end of this decade share a recognizable set of attributes, and those attributes inform a practical, vendor-neutral evaluation framework. A durable platform runs and migrates existing VM estates. It treats Kubernetes and cloud-native workloads as first-class rather than bolted on. It's ready for AI and agentic work, with GPUs, model serving, orchestration, vector data, and governance. It flexes across on-premises, cloud, and edge. It presents consistent operations through a single control plane. It enforces data sovereignty, security, and governance. And it fits the operating model and automation an organization can realistically run.

By applying these criteria, an organization can assess any option on its merits. The framework is deliberately vendor-neutral, and it sets up a fair way to look at all of the potential solutions.

WHY NUTANIX MAPS TO THIS FRAMEWORK

Measured against that framework, MI&S believes Nutanix deserves a closer look. The reason isn't that it has a hypervisor, a Kubernetes platform, or AI software. Several vendors offer those capabilities. The difference is that Nutanix approaches virtualization, cloud-native infrastructure, and enterprise AI as parts of a single operational platform rather than separate technology stacks.

That alignment shows up across the portfolio. AHV provides a mature virtualization platform without a separate hypervisor license. Nutanix Kubernetes Platform brings virtual machines and containers under the same operational model, making the hybrid environment described earlier practical to manage rather than simply possible. Nutanix Enterprise AI and the company's broader agentic AI capabilities extend that same platform into enterprise AI by focusing on model orchestration, governance, private data access, token and inference cost control, and operational management rather than GPU infrastructure alone. The result is a common operating environment that spans on-premises infrastructure, cloud, and the edge. That same approach extends to storage, where integrations with enterprise storage providers including Dell, Everpure, and NetApp let organizations migrate to Nutanix while continuing to get value from the arrays they already own. For organizations expecting to run virtual machines, containers, and AI workloads side by side for years to come, reducing the number of operational silos may prove just as valuable as any individual technology.

That said, it's important to acknowledge the gravitational pull that comes with incumbency. Organizations deeply invested in VMware tooling, OpenShift and its developer ecosystem, or Azure-native operations may reasonably conclude that a different platform fits their existing skills and investments better. That's a valid outcome because the case here is about fit rather than superiority. A shop already standardized on OpenShift may land on Red Hat, a Microsoft-centric team may prefer Azure Local for its consistency with Azure, and a large enterprise already consuming much of the VMware portfolio may decide VCF remains the right answer, just as public cloud still fits cloud-native workloads with low data gravity. The framework in this paper is intended to help organizations make that decision deliberately, not prescribe a single answer.

A CIO DECISION CHECKLIST

The following questions turn the framework into a working checklist for the decision ahead.

Decision Area	Questions to Answer
Workload and application fit	What runs on VMware today, what is VM-bound versus container-ready, and are there hard dependencies to account for?
Economics and TCO	How does a full VMware renewal compare with the total cost of alternatives, including migration effort, third-party support as a bridge, and any hardware refresh forced by a tightened compatibility list?
Timeline and the cost of waiting	How much runway remains, and what does deferral cost in terms of lost options and leverage?
Operating-model readiness	Which of the four operating models can the organization actually run today, and which is it still building toward?
Cloud-native and AI roadmap	Will the platform support containers, agents, vector data, governance, and GPUs over the next three to five years?
Coexistence	Can VMs, containers, and AI run together and evolve at the organization's pace on one platform?
Data gravity, sovereignty, and compliance	Where must data live, and which regulatory constraints apply?
Ecosystem, support, and deployment model	How strong are the vendor roadmap and support, and is turnkey, OEM, or as-a-service delivery available, including through Dell or HPE GreenLake?
Lock-in and exit	How portable and open is the platform, how flexible are the contracts, and does the choice avoid simply trading one lock-in for the next?

CALL TO ACTION

The decision point is approaching, and October 11, 2027, is a fixed deadline. Don't allow the date to dictate decisions that should be driven by the organization's long-term needs. Waiting may feel like the least disruptive option, but as the deadline approaches, waiting reduces flexibility, compresses migration timelines, and weakens negotiating leverage.

Ultimately, this decision extends well beyond virtualization. It determines the operating model an organization will build around for years to come, influencing how it runs traditional applications, adopts cloud-native architectures, and scales AI. This is precisely why this paper argues that the evaluation should begin with the operating model rather than the technology itself.

No single platform will be the right answer for every organization. Existing investments, operational maturity, and long-term strategy all matter. The organizations that navigate this transition successfully won't necessarily all choose the same platform, but they will all make a deliberate decision based on where they are today and where they intend to be over the next decade. Measured against that framework, MI&S believes Nutanix is one of the strongest platforms that organizations should evaluate.

For more information, visit nutanix.com/vmware

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR

[Matt Kimball](#), Vice President and Principal Analyst, Datacenter Compute and Storage

PUBLISHER

[Patrick Moorhead](#), CEO, Founder and Chief Analyst at [Moor Insights & Strategy](#)

INQUIRIES

[Contact us](#) if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy." Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

Nutanix commissioned this paper. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

© 2026 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.