



The VMware Exit Playbook

*A PHASED FRAMEWORK FOR MIGRATING SAFELY,
WITHOUT DISRUPTION*

STEVE MCDOWELL, CHIEF ANALYST

AUGUST 2026

EXECUTIVE SUMMARY

It's a challenging time for VMware customers. Since completing its acquisition of VMware in November 2023, Broadcom has eliminated perpetual licensing, discontinued standalone products, and moved to bring every customer to subscription-only bundles. Renewal quotes now commonly come in at two to twelve times the prior price. Most recently, Broadcom announced that general support for vSphere v8 will end in October 2027, regardless of the customer's contract renewal date.

Gartner's [2026 Strategic Roadmap for VMware Modernization](#) quantifies where this is headed: by 2029, 55% of enterprises will have migrated 100 percent of their VMware workloads to an alternative platform, and 67% are already actively or passively pursuing one. Furthermore, the Nutanix 2026 Enterprise Cloud Index reveals that 82% of organizations state their on-premises infrastructure is not fully ready to support AI workloads, making this migration a critical window to future-proof the stack.

None of this makes the exit any easier. Most unsuccessful VMware migrations stem from IT teams treating the exit as a simple infrastructure swap, overlooking the application-level dependencies that determine whether workloads survive the move. Teams sequence workloads for technical convenience, and critical dependencies surface only during the cutover window, when it's too late to plan around them.

The result can be downtime, performance regression, and a lengthy post-cutover period of instability, leading to business-level disruption and increased IT complexity.

Done correctly, however, a migration can help IT organizations increase efficiency by lowering costs and simplifying operations. Many enterprises view the migration away from VMware as the ideal opportunity to address broader application modernization efforts.

This playbook outlines a five-phase framework for executing a VMware migration without those outcomes. It covers assessment and risk mapping, defining the target state, workload segmentation and prioritization, organizational readiness, application dependency mapping, downtime avoidance and rollback planning, and program governance. It closes with a practical checklist IT leaders can use to gauge whether their exit is ready to move from planning to execution.

WHY VMWARE EXITS FAIL

Many failed VMware exits share the same root causes, and most are avoidable. This section examines the market pressures driving these decisions, the planning mistakes that derail otherwise well-resourced migration programs, and the five-phase framework that underpins the rest of this playbook.

The failure points for a migration are not complicated, but can be immensely impactful:

- **Focusing on infrastructure, not applications:** Standing up the new cluster and moving virtual machines onto it in whatever order is technically convenient often optimizes for the wrong unit of risk.
- **The lift-and-shift trap:** A VMware exit is more than a hardware refresh or hypervisor swap. Treating it as such forfeits most of the value of the move and transplants years of accumulated technical debt into a clean environment.

These traps don't only apply to migrations to a non-VMware ecosystem; they can also affect VMware to VMware migrations. For example, migrating from vSphere 8 to vSphere 9 requires deploying a full management domain just to manage license compliance, and it also requires VMware Cloud Foundation (VCF) Operations as a manageability plane.

Let's look at these in more detail.

FOCUSING ON INFRASTRUCTURE, NOT APPLICATIONS

The most common mistake in a VMware migration is starting with infrastructure rather than applications. A virtual machine is not what the business depends on; an application is. An application often consists of several VMs, a shared database, and a handful of integrations.

Infrastructure-first migrations often discover this during a cutover window, when an unmapped dependency can take down a service no one expected to be affected. Application-centric planning, sequencing, and grouping workloads by the business services they support can prevent that.

THE LIFT-AND-SHIFT TRAP

A VMware exit should be more than a change of hardware or hypervisor. Moving every VM as-is, without reconsidering its size, architecture, ownership, relevance, or future requirements, preserves many of the problems the migration could have solved.

Organizations that realize greater value from the transition use it to support broader modernization and efficiency goals. They right-size overprovisioned workloads, retire systems that are obsolete or no longer owned, consolidate where appropriate, and rearchitect the smaller set of applications whose legacy constraints no longer apply.

The migration also creates a natural decision point for application modernization. For each workload, organizations should determine whether a traditional VM remains the best operating model or whether

alternatives such as containers, cloud-native services, software-as-a-service platforms, or managed infrastructure would better meet future requirements.

Not every application needs to be transformed, but every application should be evaluated before its existing design is reproduced in the target environment.


WHAT DISRUPTION LOOKS LIKE

When a VMware exit goes wrong, it does so in a small number of predictable ways:

- **Downtime and service interruptions:** typically caused by a dependency that was not documented before the cutover.
- **Performance regressions:** occur when workloads that ran well on the old cluster degrade on the new platform due to sizing mismatches, storage latency differences, or network topology changes.
- **Broken application dependencies:** items such as hardcoded IP addresses, license servers, or authentication chains that fail silently after the move are common.
- **Operational instability after cutover:** the migration appears to succeed on cutover night but then destabilizes over the following weeks as edge cases surface.
- **Lifting and shifting into complexity:** Migrating between technologies, such as moving from VMware virtual machines to a Kubernetes-based virtualization platform, can require entirely new skill sets.

THE FIVE PHASES OF A SAFE VMWARE EXIT

This playbook organizes the exit into five phases. Each phase builds on the deliverables from the previous one, and each has a defined exit criterion before the program moves forward. Planning well and following the playbook can speed the transition while also minimizing “day 2” operational friction.



Each phase has a defined deliverable and exit criterion before the program moves to the next.

Figure 1: Five Phases of a Safe VMware Migration

ASSESSMENT & RISK MAPPING

The first 30 days set the trajectory for everything that follows. This section covers what to inventory, which business drivers should shape the plan, how to classify risk, and the tools and deliverables that transform raw discovery data into a workable roadmap.

INVENTORY BEFORE INTENT

Before making any decision about a target platform, the team managing the migration needs a complete and up-to-date picture of what it is currently running. That inventory should cover:

- **Virtual machines**, including power state, resource allocation, and actual utilization.
- **Containers** and any workloads already running on Kubernetes or other container platforms, including VMware Tanzu, alongside traditional VMware virtual machines.
- **Hosts and clusters**, including hardware age, warranty status, and licensing terms.
- **Storage infrastructure**, including arrays, protocols, and replication relationships.
- **Network services**, including load balancers, firewalls, and segmentation policies.
- **Third-party integrations**, including backup, monitoring, and identity systems, that have hard dependencies on vCenter APIs.

UNDERSTAND THE BUSINESS DRIVERS

The technical plan should follow the business driver, not the other way around. Business needs should shape the timeline, acceptable risk tolerance, and which workloads move first:

- **Contract renewal timelines** typically drive infrastructure migrations, often tied to the expiration of an Enterprise License Agreement. vSphere 8 presents a different kind of deadline. Broadcom has announced that general support for vSphere 8 will end in October 2027, requiring organizations to migrate to VMware vSphere Foundation 9 (VVF), VMware Cloud Foundation 9 (VCF), or an alternative platform before then, regardless of when their existing ELA expires. Organizations that remain on vSphere 8 beyond that point risk losing access to ongoing support and security patches, increasing the operational and security risks of delaying migration.
- **Budget pressures**, including whether the organization needs to demonstrate cost savings within a specific fiscal year.
- **Strategic modernization goals**, including growing pressure to free up capacity and budget for AI infrastructure investment. IT leaders must use this migration as a catalyst to adopt a unified platform that natively supports virtual machines, Kubernetes, and Agentic AI. This ensures a fast, efficient path forward that accelerates AI outcomes while minimizing future risk.
- **Risk reduction objectives**, including exposure to a single vendor with a volatile licensing model.

RISK CLASSIFICATION FRAMEWORK

Every workload in the inventory should be classified into a risk tier before segmentation begins. A simple three-tier model works for most organizations:

Tier	Definition	Migration Posture
High Cost/Complexity	Workloads with high operational and resource costs, or high operational complexity	Prioritizing workloads with the highest cost and complexity can lead to the fastest ROI.
Mission-critical	Direct revenue impact or regulatory exposure if disrupted	Move last, with extensive parallel validation and the longest rollback window
Business-important	Meaningful operational impact but tolerant of a short, planned outage	Move in the middle waves, once the process has been proven on lower-risk workloads
Development and test	Limited or no production impact if disrupted	Move first, to build team confidence and surface process gaps cheaply

A risk score based on criticality, technical complexity, and dependency density ranks each workload within its tier, rather than treating all mission-critical applications as equally risky.

Depending on the size of the VMware estate, some workloads may need to be migrated to other VMware versions while the broader migration is underway. For example, vSphere 8 workloads will lose support in October 2027, which may necessitate migrating some workloads or applications to vSphere 9 as an interim solution.

FINDING THE UNKNOWN UNKNOWNNS

Nearly every assessment reveals untracked workloads. This has become a larger problem with the rise of shadow AI, as business units stand up GPU-enabled VMs or AI tooling outside the infrastructure team's visibility, in addition to the more familiar categories:

- **Shadow IT and shadow AI deployments** that never went through a standard provisioning process; Nutanix's [2026 Enterprise Cloud Index](#) shows that 79% of responding organizations encounter AI applications or agents implemented by non-IT employees.
- **Orphaned virtual machines** with no identifiable owner or current business function.

- **Legacy integrations** built by staff no longer with the organization.
- **Undocumented dependencies** between applications that surface only under traffic analysis.

ASSESSMENT TOOLING

Manual inventory does not scale beyond a few hundred VMs and can miss the dependencies that matter most. Most assessment programs combine several categories of tools:

- **Discovery and inventory platforms** that scan the environment and build an up-to-date asset list without relying on outdated CMDB records.
- **Application dependency mapping tools**, typically agentless, that observe traffic and build a topology showing what talks to what.
- **Infrastructure dependency mapping tools**, including those that automatically catalog and retain critical networking properties (such as IP and MAC addresses), as well as security tags and firewall properties.
- **CMDB reconciliation** to compare the discovered environment against the existing inventory system of record and close the gap.
- **Automated assessment platforms**, including target-platform-native tooling that translates the discovered environment into sizing and readiness recommendations.

Vendors in this space report first-pass discovery times as fast as under an hour for a full environment scan. Be careful to treat vendor-reported deployment speed as a starting point for evaluation, not a guarantee across every environment.

Using tools such as these reduces cutover risk while avoiding the high cost of third-party professional services typically required for complex migrations.

DELIVERABLE: MIGRATION READINESS SCORECARD

By Day 30, the assessment should be complete. The output of this phase should be a single artifact that the entire program uses to guide the migration. That scorecard should include:

- **A risk-ranked workload inventory** covering every VM, container, and dependency discovered.
- **An executive dashboard** summarizing scope, timeline, and top risks in terms a non-technical sponsor can act on.
- **A draft migration roadmap** sequencing workloads into waves.
- **Named owners** for every workload and every phase of the program, so accountability does not default to the infrastructure team alone.

DEFINING THE TARGET STATE

Before evaluating any platform, the migration team needs a clear statement of what it requires from that platform. This section outlines the business and operational requirements to document up front, how to define success, and how to decide which workloads should move.

BUSINESS REQUIREMENTS

The decision of where to migrate VMware workloads should start with requirements. Document the non-negotiables before evaluating any platform:

- **Availability objectives**, expressed as uptime targets and recovery time objectives by workload tier.
- **Performance requirements**, including any workloads with specific latency or IOPS floors.
- **Security and compliance needs**, including frameworks the organization must attest to.
- **Sovereignty requirements**, including any data residency or jurisdictional constraints that limit viable platforms.
- **Future workload support**, including recognizing upcoming modernization and AI-led digital transformation efforts; while many AI transformation projects are in early planning stages, enough is known about the operational characteristics to plan for future needs.

OPERATIONAL REQUIREMENTS

Alongside business requirements, define what the operations team needs the new platform to support on day one.

- **Unified Manageability**, with a platform that unifies the management of containers and virtual machines.
- **Monitoring and observability** at parity with what the team relies on today.
- **Backup and recovery** processes that meet current retention and restore-time commitments.
- **Disaster recovery** capabilities, including whether DR needs to span the transition period.
- **Automation and orchestration**, including whether existing infrastructure-as-code investments carry over.

SUCCESS CRITERIA

Define what a successful migration means before it starts, not after. Vague success criteria are among the most common reasons a technically complete migration is still perceived internally as a failure:

- **A specific, written definition** of what constitutes a successful migration for the program as a whole.

- **Operational KPIs**, such as incident rate and mean time to resolution, are benchmarked against the current environment.
- **User experience metrics** for the application owners and end users who will judge the migration by whether anything felt different.
- **Cost and efficiency targets**, since cost is usually the initiating driver and must be tracked, not assumed.

MODERNIZE, RE-PLATFORM, OR MIGRATE?

Not every workload should move as-is. Every workload in the inventory should receive one of four dispositions before entering a migration wave:

Disposition	When It Applies
Retire	The workload has no remaining business value, or its function has already been absorbed elsewhere
Replace	A SaaS or managed service now covers the function better than a self-hosted VM ever will
Re-platform	The application changes shape as part of the move, most often by containerizing a workload that no longer needs a full VM
Migrate	The application moves largely as-is, because it works and rewriting it adds no value at this stage

Running this exercise before segmentation prevents wasted migration effort on workloads that should never have been moved in the first place.

WORKLOAD SEGMENTATION & MIGRATION PRIORITIZATION

Not every workload should move in the same wave or for the same reason. This section explains how to group workloads by risk and complexity, which candidates should be included early or late in the program, and the sequencing mistakes that undermine an otherwise sound plan.

SEGMENTATION CRITERIA

With a disposition assigned to every workload, the remaining candidates for migration and re-platforming need to be grouped into waves. Five criteria guide that grouping:

- **Business criticality**, derived directly from the risk tier established during the assessment.
- **Return on Investment**, including recouped licensing and operational costs.
- **Technical complexity**, including how many components the application spans and how tightly they are coupled together.
- **Dependency density**, or the number of other workloads and systems that connect to this one.
- **Compliance sensitivity**, including any regulatory constraints on how or when the workload can be moved.

MOVE FIRST

The *move-first* cohort should be selected for two reasons: these workloads are safe to move without extensive risk mitigation and moving them first delivers the fastest return.

- **Prioritize the highest-cost workloads first.** Retiring the largest share of licensing spend earliest delivers the greatest time-to-value for the program and strengthens the executive case for continued investment.
- **Development and test environments** that tolerate disruption and provide the team with a low-stakes environment to prove the process.
- **Low-dependency services**, where the blast radius of an unexpected issue remains contained.
- **Internal business applications with a limited user base**, where any disruption is easier to absorb and recover from.

MOVE LATER

Regulated workloads, legacy applications with poorly understood internals, systems with high dependency density, and anything tied to specialized infrastructure should be deferred to later waves.

The migration team can address these workloads once the team has proven the process on lower-risk workloads and has a validated rollback plan for anything that matters.

COMMON PRIORITIZATION MISTAKES

Segmentation can fail in a few recognizable ways:

- **Sequencing by technical ease rather than by business risk reduction**, which front-loads easy wins that do not reduce program risk.
- **Ignoring business impact** when a workload appears simple from the infrastructure side but supports a critical process.
- **Underestimating operational dependencies** that become visible only after dependency mapping is complete.

BUILDING MIGRATION WAVES

Group workloads into waves of manageable size, typically defined by risk tolerance (rather than a fixed VM count), and include explicit success criteria and a go/no-go gate at the start of each wave.

Executive sign-off at each gate keeps the program accountable to the criteria defined during the assessment, keeping schedule pressure from quietly eroding them.




Figure 2: Building Migration Waves

ORGANIZATIONAL READINESS

A migration can be technically sound yet still fail if the organization is not ready to operate the new environment. This section covers the skills, operating model, tooling, and change management work that must occur alongside the technical planning, not after it.

SKILLS READINESS ASSESSMENT

Before moving the first workload, assess the gap between the team's current VMware-centric skill set and the target platform's requirements. Inventory current skills, define the new platform's operational demands, and prioritize training to close the gap.

Don't assume that familiarity with virtualization concepts translates automatically across environments. While some alternative platforms offer a user experience familiar to VMware users, others, especially container-centric ones, will require entirely new skills and training.

OPERATING MODEL CHANGES

Some responsibilities transfer directly to the new platform, while others do not. Document explicitly what remains the same, what changes, and what new responsibilities and workflows the operations team inherits, so the transition does not surface as a series of unpleasant surprises during on-call rotations.

TOOLING READINESS

Operational tooling needs to reach parity with the current environment before cutover, not after. Confirm that monitoring, backup, disaster recovery, and automation and orchestration tooling, including infrastructure-as-code providers, function correctly on the target platform under real-world conditions.

BUILDING INTERNAL CHAMPIONS

A center of excellence, staffed by a small pilot team that goes deep into the new platform before the broader organization, provides later waves with institutional knowledge and troubleshooting experience that vendor documentation cannot replace.

That team also becomes the natural owner of the knowledge transfer strategy across the rest of the organization.

APPLICATION DEPENDENCY MAPPING

This is the phase most programs shortchange, and the one most responsible for cutover surprises. This section explains why dependency mapping matters, the traps that even experienced teams fall into, and how to build a map the rest of the program can rely on.

WHY DEPENDENCY MAPPING MATTERS

Hidden dependencies determine migration sequencing, and getting the sequencing wrong is what turns a planned maintenance window into an unplanned outage with real business continuity implications.

Everything in this section exists to surface those dependencies before they surface on their own during a cutover.

COMMON DEPENDENCY TRAPS

A handful of dependency types account for most cutover failures, including:

- **Hardcoded IP addresses** and hostnames embedded in application configuration files.
- **Legacy authentication services** that applications assume will always be reachable at a fixed address.
- **License servers** tied to MAC addresses or other hardware fingerprints that fail the moment a VM moves.
- **Shared databases** that multiple applications depend on but that no single team fully owns.
- **Backup and monitoring integrations** with hard dependencies on vCenter APIs that silently stop working after migration.

DISCOVERY TECHNIQUES

No single technique catches everything, so effective dependency mapping combines several:

- **Traffic analysis** to observe real east-west and north-south communication patterns over time.
- **Application-level tracing** to follow a request through its full dependency chain.
- **Log analysis** to catch intermittent or scheduled dependencies that traffic analysis alone does not reveal.
- **Vendor documentation reviews** for third-party systems when internal knowledge may have been lost.

DOCUMENTATION GAP

By the time a migration starts, existing architecture diagrams are usually outdated. Systems drift from their documentation continuously, and no one updates the diagram when a developer adds a new integration under deadline pressure.

Every assumption in an existing dependency diagram must be validated against live discovery data before it is trusted for migration sequencing.

DELIVERABLE: APPLICATION DEPENDENCY MAP

The output of this phase should feed directly into wave planning:

- A validated map of service relationships across the estate.
- Direct inputs into migration sequencing, so applications that depend on each other move together.
- Risk mitigation plans for every dependency identified as fragile or poorly understood.

DOWNTIME AVOIDANCE & ROLLBACK PLANNING

How a migration is executed determines whether it is disruptive. This section covers migration patterns, runbook design, and the technical and operational rollback planning that must be in place before a cutover window, not improvised during one.

MIGRATION PATTERNS

Three broad patterns cover most migration execution approaches, with the right choice varying by workload:

- **Maintenance-window migrations:** the workload goes offline for a defined period and then comes back online on the new platform.
- **Live migration approaches:** the workload moves with minimal to no perceived downtime.

- **Replication-based cutovers:** data replicates continuously to the target ahead of a brief final cutover window.

CUTOVER RUNBOOK DESIGN

Every wave needs a written runbook. A complete runbook outlines pre-checks to confirm readiness, the exact execution steps in order, validation checkpoints after each major step, and escalation paths when things do not go as planned.


ROLLBACK PLANNING

Rollback capability must be designed in advance. Define in writing how the team rolls back the VM itself, any data that changed during the migration window, and any network configuration changes before the first production workload moves.

OPERATIONAL ROLLBACK PLANNING

Technical rollback capability is useful only when paired with clear operational authority. Define who has the authority to call a rollback, the escalation process for reaching that person during the migration window, and how long the team keeps the source infrastructure live as a safety net.

Shortening that window too early to reclaim licensing costs sooner is one of the most common self-inflicted causes of an unrecoverable failed migration.



Rollback criteria and decision authority are defined before cutover, not improvised during one.

Figure 3: Operational Rollback Planning

COMMUNICATION PLANNING

A cutover plan needs a corresponding communication plan that covers:

- **Executive stakeholders:** need summary-level status.
- **Application owners:** need enough detail to validate their systems.
- **End users:** need to know what to expect.
- **Support teams:** need to know what changed before the first ticket comes in.

GOVERNANCE & CUTOVER

Moving workloads is only one measure of progress. A migration succeeds when applications remain stable, users can work without disruption, and operations teams can effectively support the new environment. This section explains how to govern the program, evaluate each migration wave, and determine when to increase or reduce the pace.

ESTABLISHING A MIGRATION PMO

A program of this scope often requires a program management office (PMO) with clear executive sponsorship, dedicated management capacity, and technical leaders empowered to make timely decisions on sequencing, cutover, and rollback.

The PMO should align technical execution with business priorities, maintain accountability across teams, and ensure that schedule pressure never compromises operational readiness.

WAVE GOVERNANCE

Every migration wave should follow a consistent governance process. Before cutover, a formal go/no-go review should confirm application readiness, reassess dependencies, validate the rollback plan, and verify that support teams are prepared. Change control should maintain the wave's scope once execution begins.

After cutover, the wave should remain under heightened observation until predefined stability criteria are met. Only then should it be considered complete.

METRICS THAT MATTER

The number of VMs moved and the migration pace are useful measures of activity, but they do not demonstrate success. A wave is successful when the affected business services operate reliably after cutover and the new environment does not create an unsustainable support burden.

Measurements should therefore focus on operational outcomes, including:

- Application availability and performance
- Incident volume and severity
- User experience and support requests
- Backup, recovery, monitoring, and security-control health
- Time required to resolve post-cutover issues
- Rollback frequency
- Stability over an agreed observation period

Migration velocity remains important for forecasting and capacity planning, but it should be considered alongside these measures. A program that moves workloads quickly while increasing incidents or degrading service is accumulating risk, not making meaningful progress.

WHEN TO ACCELERATE

Wave size or frequency can increase when applications consistently meet post-cutover stability thresholds, incident rates remain within agreed tolerances, dependencies are accurately identified, and operations teams can absorb the added workload. Repeatedly validated rollback plans and several uneventful waves provide further evidence that the migration process is mature enough to scale.

WHEN TO SLOW DOWN

The program should slow its pace when incident rates rise, application performance degrades, users report recurring problems, or support teams accumulate a growing backlog. Frequent rollbacks and previously unidentified dependencies also indicate that assessment, sequencing, or testing needs improvement.

These conditions call for corrective action before the next wave. It's important not to consider slowing down as a failure of the program, but rather a governance decision that protects service continuity and prevents unresolved problems from being repeated at greater scale.

CHOOSING A VMWARE ALTERNATIVE

With the forced march toward subscription bundles, staying with VMware can be highly restrictive and costly. Broadcom's consolidation leaves most customers with only two primary choices if staying within the VMware ecosystem: VMware Cloud Foundation 9 (VCF) or VMware vSphere Foundation 9 (VVF).

At the same time, options such as Nutanix Cloud Platform become attractive.

THE VVF TRAP AND THE "CAPACITY TAX"

Many organizations view VVF 9 as the logical landing spot to avoid the massive footprint of a full VCF 9 deployment. However, VVF 9 has significant drawbacks:

- **Price hikes & capacity taxes:** List prices have effectively doubled. VMware has moved away from simple socket-based pricing toward a complex core-plus-TiB capacity tax.
- **Restricted functionality:** Although VMware has eased many availability restrictions on VVF 9, VVF 9's functionality remains more limited than VCF 9's. Organizations that need to support private AI, advanced container environments, NSX, Azure VMware Solution (AVS), or other modern workloads may require VCF 9. The economics can also favor VCF 9 when deployments exceed VVF 9's included vSAN capacity. As a result, organizations adopting VVF 9 should evaluate whether their anticipated requirements could necessitate a move to VCF 9 shortly afterward, potentially altering the cost and complexity of their VMware strategy.
- **Loss of centralized management:** VVF 9 deprecates Enhanced Linked Mode (ELM). Customers are left managing one vCenter per site with no central linking unless they upgrade to VCF 9.

THE NUTANIX CLOUD PLATFORM ADVANTAGE

Nutanix provides a flexible, proven alternative that gives organizations greater control over infrastructure strategy without requiring a wholesale re-architecture of the workloads they manage:

- **Modular licensing without a "Capacity Tax":** Organizations can purchase what they need today and scale compute and storage as demand changes. Rather than locking customers into rigid, one-size-fits-all bundles, Nutanix licensing aligns more closely with actual consumption and growth.
- **Enterprise-grade hypervisor:** Nutanix AHV provides a secure, resilient virtualization foundation for enterprise workloads and integrates with Nutanix Cloud Platform. Its familiar VM operating model allows infrastructure teams to continue using established virtualization practices without adopting container orchestration solely to run conventional applications.
- **Integrated, policy-aware migration:** Nutanix Move treats migration as a coordinated process that accounts for the workload and its supporting configuration. Networking and security policies (including applicable NSX-T rules) can be mapped during the transition, reducing risk and the manual rework often associated with infrastructure migrations. This flexibility supports a range of destination strategies, whether organizations retain existing third-party storage investments from providers such as Dell, Everpure, and NetApp or migrate workloads to AWS, Microsoft Azure, or Google Cloud.
- **Operational simplicity:** Prism and Nutanix Central provide unified management for compute, storage, and networking across on-premises and supported public-cloud environments. This consolidated operating model reduces administrative complexity and improves visibility across the infrastructure estate.
- **Modern applications:** Beyond VM consolidation, Nutanix provides a path to modern applications with Nutanix Kubernetes Platform. Integrated data and application services

enable organizations to support containerized applications and emerging AI use cases without creating separate infrastructure and operational silos.

- **Enterprise AI and the Nutanix Agent Gateway:** Nutanix Enterprise AI enables infrastructure teams to unify infrastructure, Kubernetes, and data through a centralized control plane, while the built-in agent gateway provides cost control and governance capabilities to help manage users of autonomous agents.

Together, these capabilities position Nutanix as more than an alternative to VMware's licensing model. The platform provides a unified, enterprise-grade foundation for virtual machines, containers, hybrid-cloud operations, and emerging AI workloads.

CONCLUSION

Licensing costs are what start most VMware exit conversations, but they should not be the only metric the organization uses at the end of one.

Done well, a migration can also be a welcome catalyst for modernization opportunities deferred for years, operational simplification by retiring workloads nobody needed, and platform standardization that eliminates the crippling operational overhead of running fragmented virtualization, container, and AI stacks.

VMWARE EXIT CHECKLIST

Before moving from planning to execution, confirm that the program can answer yes to each of the following:

- **Assessment complete:** every VM, container, and dependency across the estate has been inventoried and risk-scored.
- **Risks identified:** mission-critical workloads and their dependencies are documented and understood, not assumed.
- **Teams are prepared:** skill gaps are addressed, tooling reaches operational parity, and a center of excellence is in place.
- **Rollback plans validated:** technical and operational rollback procedures have been tested.
- **Governance established:** a PMO, wave gates, and clear metrics are in place before the first production wave begins.

NEXT STEPS

Organizations live at different stages, and may start from different places:

- **A migration assessment engagement** for organizations that need Day 0-30 inventory and risk mapping with outside support.

- **A platform readiness workshop** for organizations that have selected a target platform and need help defining requirements and success criteria.
- **Pilot migration planning** for organizations ready to move a first, low-risk wave and prove the process before scaling it.
- **A professional services engagement** for organizations that need hands-on execution support across the full five-phase program.

The teams that navigate a VMware exit cleanly are rarely the ones that move fastest. Successful teams treat the migration as an application-centric, risk-managed program.

The five-phase approach outlined here requires more time up front. It also makes the difference between a migration the organization remembers as disruptive and one it barely notices.

Stop guessing about your renewal costs and migration path. Take control of your infrastructure strategy today. Learn more about how Nutanix can help you streamline migrations and modernize your infrastructure by visiting [Nutanix.com/vmware](https://nutanix.com/vmware).

© Copyright 2026 NAND Research. NAND Research is a registered trademark of NAND Research LLC, All Rights Reserved.

This document may not be reproduced, distributed, or modified, in physical or electronic form, without the express written consent of NAND Research. Questions about licensing or use of this document should be directed to info@nand-research.com.

The information contained within this document was believed by NAND Research to be reliable and is provided for informational purposes only. The content may contain technical inaccuracies, omissions, or typographical errors. This document reflects the opinions of NAND Research, which is subject to change. NAND Research does not warranty or otherwise guarantee the accuracy of the information contained within.

NAND Research is a technology-focused industry analyst firm providing research, customer content, market and competitive intelligence, and custom deliverables to technology vendors, investors, and end-customer IT organizations.

Contact NAND Research via email at info@nand-research.com or visit our website at nand-research.com.