

VMware Migration Assessment Checklist

25 control points for evaluating a VMware exit, from inventory and dependencies to sovereignty, GPU workloads, and cost trajectory.

How to use this checklist

This 25-point checklist gives infrastructure leaders, platform engineers, and CIO offices a structured way to evaluate readiness for a VMware exit. It is the same starting framework Aenix uses during the **Platform Readiness Assessment** — a 14-day engagement that converts the checklist into a tailored migration plan.

Use it in three ways:

- 1. **Self-assessment** — walk through with your platform team in one workshop. Each unchecked box becomes a discovery item.
- 2. **RFP scoping** — share with bidders so they answer against the same surface.
- 3. **Board readiness** — the section summaries map directly to a board-grade migration brief.

Each checkpoint maps to a structural decision in your destination architecture. An "unknown" answer is more useful than a guess — it surfaces the workstream that needs assessment first.

Coverage map

Section	Checkpoints
1. Inventory	4
2. Dependencies	4
3. Network & storage redesign	4
4. Multi-tenancy	2
5. AI / GPU workloads	3
6. Sovereignty & compliance	3
7. Operational readiness	3
8. Cost trajectory	2

Section 1 Inventory

- ☐ **Workload count by tier** — Total VM count broken down by criticality tier (mission-critical / important / non-critical). Include uptime SLO per tier.
- ☐ **OS and kernel mix** — Distribution of guest OS versions, kernel versions, end-of-life status. Identify kernels older than 5 years.
- ☐ **Resource shape per workload** — vCPU, RAM, attached storage, network bandwidth peak/avg per VM. Flag oversized and undersized.
- ☐ **Application-layer ownership** — Named owner team per workload, last-modified date, deployment cadence (manual / scripted / GitOps).

Tip. Workloads without a named owner are migration-cost amplifiers. Resolve ownership before estimating effort.

Section 2 Dependencies

- ☐ **vSAN dependencies inventoried** — Workloads relying on vSAN-specific features (datastore-level snapshots, deduplication, RAID-on-storage). List a replacement strategy per workload.
- ☐ **NSX configurations exported** — Distributed firewall rules, micro-segmentation policies, edge gateway configs. Note: many cannot be ported 1:1.
- ☐ **vCloud Director (vCD) tenant model documented** — Organization VDCs, catalog structure, role assignments, vApp templates. Map to the destination tenant model (e.g., Cozystack Tenant CRD).
- ☐ **Custom integrations cataloged** — vRealize Automation/Operations, third-party backup (Veeam, Rubrik), monitoring agents, license servers. List vendor support for the destination platform.

Warning. NSX distributed firewall semantics rarely map 1:1 to Kubernetes NetworkPolicy or Cilium. Plan for **redesign**, not lift-and-shift, of network security policy.

Section 3 Network & storage redesign

- ☐ **Network topology documented** — VLANs, IP plan, MTU, BGP/OSPF if applicable, ingress/egress edges. Map to the destination CNI (Cilium / Calico) and gateway model.
- ☐ **Storage classes mapped to destination** — vSAN tiers, datastore policies, IOPS commitments. Map to LINSTOR / Ceph / Longhorn / equivalent storage classes.
- ☐ **Snapshot and replication SLAs** — Current snapshot schedules, replication targets, RPO/RTO commitments. Identify gaps in the destination storage backend.
- ☐ **L2 stretched / multi-site networking** — Cross-site L2, vMotion across DCs, stretched clusters. Most do not survive migration; identify workloads that genuinely require this.

"We use vSAN for everything" usually means 4–6 distinct storage profiles. Surface them all — destination platform sizing depends on accurate IOPS/throughput per class.

Section 4 Multi-tenancy

- ☐ **vCD organization model translated** — Each vCD Org maps to a destination tenant boundary (Tenant CRD, namespace bundle, project). Map quotas, network isolation, and IAM per Org.
- ☐ **Self-service surface defined** — Portal vs API vs YAML-driven self-service for tenants. Define what tenants can and cannot do (provisioning, scaling, network policy, storage class selection).

Section 5 AI / GPU workloads

- ☐ **GPU inventory** — vGPU vs passthrough, GPU model (A100 / H100 / L40S / Blackwell), per-card workload assignment. Note vendor-locked features (e.g., MIG profiles).
- ☐ **vGPU to KubeVirt vGPU mapping** — Confirm the destination supports the GPU class and slicing model needed. Validate driver versions.
- ☐ **AI workload architectural fit** — Inference vs fine-tuning vs training; multi-tenant isolation requirements; sovereignty constraints on model weights and training data.

Warning. VMware's vGPU licensing terms differ materially from NVIDIA's direct licensing under Kubernetes. Validate with an NVIDIA partner before committing to a migration timeline.

Section 6 Sovereignty & compliance

- ☐ **Regulatory scope identified** — DORA / NIS2 / GDPR / sectoral / national mandates that apply to workloads in scope. Map per-workload, not per-environment.
- ☐ **Encryption key custody** — Current key management (vendor-managed / customer-managed / HSM-backed). Define the destination key custody model and supplier transparency requirement.
- ☐ **Audit-trail completeness** — Who-did-what-when log retention, immutability, exportability. Map to the destination platform audit pipeline (e.g., VictoriaLogs + immutable backend).

Section 7 Operational readiness

- ☐ **Runbook inventory** — Documented runbooks per workload class (provisioning, backup, restore, failover, patch). Identify tribal-knowledge gaps.
- ☐ **On-call coverage** — Current rotation, escalation tiers, MTTA/MTTR targets. Validate that the destination platform fits on-call team capacity.
- ☐ **Knowledge transfer plan** — Skills uplift required for the destination stack (Kubernetes, KubeVirt, Cozystack operations). Hours per engineer plus training format.

Tip. Most VMware-native ops teams need 60–120 hours of structured Kubernetes training before independently operating a Kubernetes-based destination. Build this into the migration timeline, not after it.

Section 8 Cost trajectory

- ☐ **TCO inputs gathered** — Current VMware licensing (VCF subscription tier, host count, support tier, expiration date). Hardware refresh cycle. Operations cost (FTE, support, datacenter share).
- ☐ **Commitment expirations mapped** — VMware contract end date, hardware warranty end date. These set the migration window. Earlier expirations act as a forcing function.

Pair this checklist with the **Cloud Repatriation TCO Worksheet** to model destination economics quantitatively.

Next steps

Option 1 — Self-remediate using this checklist. Walk through with platform, ops, and compliance teams. Convert unchecked boxes into a backlog. Most organizations find 8–12 unchecked boxes that need discovery work.

Option 2 — Platform Readiness Assessment with Aenix (14 days). Aenix runs the checklist with you, models destination architecture, and produces a phased migration plan. Fixed-fee engagement.

Option 3 — Full migration engagement. Multi-month build with architecture, build, migration waves, and operational handover. Scope set by Phase 2 of the assessment.

Schedule a discovery call. aenix.io/services/platform-readiness-assessment/

Aenix is the team behind Cozystack (CNCF Project) and offers the Aenix Platform — our commercial productized offering based on Cozystack. aenix.io
