Private Cloud Stack Comparison — 2026 Edition

A neutral, technical-decision-maker comparison of nine enterprise private-cloud stacks, scored on a twelve-axis rubric aimed at VMware-exit candidate selection • Published 2026-05 • All prices are public list, pre-discount, unless noted • Cited sources listed in §8

> **How to read this document.** Every stack below is presented with the same twelve-axis rubric, the same weighting, and the same tone — including the one behind this document (Aether). VMware Cloud Foundation 9 is included in the deep dives, feature matrices, and cost model as the reference baseline for what an enterprise customer is leaving, but it is intentionally excluded from the §7 weighted scoring matrix on the basis that the document's purpose is to compare *replacements*, not to score VCF against itself. The scoring methodology is transparent and the weights are exposed in §2 so you can re-score with your own priorities.

---

1. Executive Summary

Nine stacks are compared in this document: **Aether** (Incus + IncusOS + Operations Center + Migration Manager), **VMware Cloud Foundation 9** (the reference baseline), **OpenStack**, **Nutanix Cloud Infrastructure**, **Proxmox VE**, **HPE VM Essentials**, **Red Hat OpenShift Virtualization**, **SUSE Harvester / SUSE Virtualization**, and **OpenNebula 7.2**.

No single alternative is a feature-for-feature drop-in for VCF 9. VCF 9 retains the widest feature footprint of any private-cloud product in existence — vSphere + vSAN ESA + NSX + Avi + VCF Automation + VCF Operations + HCX + VKS in one SKU — and every alternative trades something off. The question is therefore not "which stack equals VCF 9" but "which gaps are acceptable in exchange for which benefits (usually cost, openness, operational simplicity, or exit from Broadcom lock-in)."

The five things to internalise before picking a path:

1. **Broadcom's pricing reset is the trigger — vendor lock-in fear is the actual driver.** Customers consistently cite lock-in as their #1 reason to move off VMware, even when quoted pricing is eye-watering. Pricing brings them to the table; openness decides where they go. VCF list is **$350/core/year** subscription with a 16-core-per-CPU minimum and a 72-core-per-order minimum[1][2]. Publicly-reported customer-level increases range from 150% to more than 1,000%[3][4]. Computershare (24,000 VMs), Beeks Group (~20,000 VMs) and AT&T have all challenged or exited[5][6][7].
2. **SDN is where alternatives hurt the most.** NSX's in-kernel distributed firewall is the industry reference. Only three alternatives come close: OpenStack with OVN, Nutanix Flow (OVS-based), and Aether (OVN ACLs *plus* CSP-level global rules that are equivalent to NSX Federation's cross-site policy model — and additionally Aether's NDR layer adds a Layer-7 IDS / IPS with reactive OVN-ACL containment that no other stack in the comparison offers natively). Proxmox, Harvester, HPE VME (without CX 10000 switches), OpenShift Virtualization, and OpenNebula all have less mature east–west security stories in practice — OpenNebula's security groups are iptables-on-tap and "not supported for Open vSwitch" per the 7.2 docs.
3. **The CISO's veto: non-immutable hypervisor OS is a real enterprise deal-breaker, and this rubric weights it as the single heaviest axis (Lifecycle / OS at 3.0×, above even Vendor-independence at 2.5×).** ESXi is a stateless hardened appliance with no SSH, no package manager, and a minimal attack surface. Enterprise CISO departments — especially in regulated industries (finance, healthcare, government, critical infrastructure, defence) — actively reject VMware replacements that ship a mutable general-purpose Linux as the hypervisor base, because it *adds* attack surface rather than preserving the stripped-down appliance model. **Nutanix AHV is CentOS-based and mutable**; **Proxmox is Debian**; **HPE VME is hardened Ubuntu**; **OpenStack hosts run RHEL/Ubuntu**; **OpenNebula 7.2 nodes run RHEL/AlmaLinux/Ubuntu/Debian/SLE/openSUSE**. The three stacks that genuinely preserve the "appliance OS" mental model are **Aether (IncusOS — Debian-13 mkosi-built immutable image, A/B atomic updates via systemd-sysupdate, signed-UKI measured boot into TPM PCRs, LUKS, ZFS, no SSH)**, **OpenShift Virtualization (RHCOS rpm-ostree A/B)**, and **SUSE Harvester (SLE Micro transactional)**. For many CISO-led procurements this is the single largest architectural filter applied *before* cost or feature comparisons begin — which is why the rubric weights it at 3.0×.
4. **"DRS" is mostly gone — but not entirely.** Aether, Nutanix (ADS), HPE VM Essentials (Morpheus intelligent placement), and OpenNebula (**OneDRS** — ILP-solver-based, predictive, the most algorithmically sophisticated of the bunch) all ship continuous load-based rebalancing in 2026. The KubeVirt-based stacks (OCP-V, Harvester) now use the Kubernetes descheduler with PSI metrics — coarser and reactive. Proxmox and OpenStack have HA and placement hints but no continuous rebalancer.
5. **Migration from VMware is a solved problem — but rule migration is not.** Red Hat MTV/Forklift, Nutanix Move, Aether's Migration Manager, HPE VME's migration tool, Proxmox's ESXi import wizard, and OpenNebula's OneSwap (with `delta`-mode for low-downtime powered-on VMs since 7.0+) all move VMs effectively. **Aether is the only stack in this document that also migrates NSX-V and NSX-T firewall rules — both east-west DFW and north-south edge/gateway rules** — every other alternative, OpenNebula included, leaves the customer to re-author micro-segmentation and perimeter policy by hand.
6. **Cost differences are enormous and decisive at scale.** For a **100-node / 12,800-core** illustrative cluster, three-year licensing ranges from **~$300k** (OpenNebula Standard) to **~$46M** (VCF 9 with Live Recovery). OpenNebula Standard + HA ~$302k; Canonical OpenStack Advanced ~$450k; OpenNebula Premium + HA ~$474k; OCP-V Engine ~$570k; Proxmox Premium + PBS ~$850k; OpenNebula Premium + 1–2 add-ons ~$430k–$890k; HPE Morpheus Enterprise ~$1.5M; Aether at $80/core ~$3.07M; SUSE Harvester Suite ~$3.45M; Red Hat OpenShift Platform Plus ~$4.13M; VCF 9 base list ~$13.4M; **Nutanix NCI-Ultimate at $575/core**

list reaches ~$22.1M — more than VCF 9 base; VCF 9 with DR / Live Recovery reaches **~$46M**. Note: Broadcom is withdrawing the lower-tier VVF SKU from parts of EMEA as of December 2025, removing the "cheap VMware" option for those customers. Detail in §6.

WHO THIS DOCUMENT IS FOR

Architects, platform engineering leads, and CTOs evaluating what to replace VMware Cloud Foundation with, or what to build a greenfield private cloud on in 2026. An executive summary is included but the body assumes a technical reader.

WHAT THIS DOCUMENT IS *NOT*

A reseller battlecard. A Gartner Magic Quadrant. A performance benchmark. Vendor advocacy. Where a stack is genuinely strong it gets credit; where it is weak — including Aether — it is said so.

Final ranked Fitness scores (full methodology in §2, full matrix in §7). VCF 9 is the reference baseline that the customer is leaving and is therefore not included in the ranked scoring matrix — it appears in the Master Feature Matrix (§4), the Parity Matrix (§5), and the Cost Model (§6) as the comparator for what is being replaced.

Rank	Stack	Fitness Score	One-line verdict
1	Aether + Incus + IncusOS + OC + MM	84	Dell-resold, fully open Apache-2 substrate, immutable IncusOS, single-pane UI across 14 domains incl. **NDR (L7 IDS/IPS)** and **Shepherd (K8s-style)**, OVN DFW + CSP-level federation, NSX-V/NSX-T firewall rule migrator, continuous DRS, OVN + active-active HAProxy LB, integrated per-tenant billing + self-service portal.
2	Red Hat OpenShift Virtualization	77	Kubernetes-native VMs on RHCOS (immutable, rpm-ostree A/B). Strong roadmap. Heavy for pure-VM shops. OKD upstream exists but full OCP is Red Hat.
3	OpenStack	73	Most open, most proven at scale — perfect vendor-independence score pushes it up. Mutable RHEL/Ubuntu hosts hurt under the heavy Lifecycle weight.
3	OpenNebula 7.2	73	~18-year mature open IaaS, very strong OneDRS (ILP-solver), Zones + VDCs + Federation multi-tenancy, sovereign-cloud / AI Factory positioning, Veeam + Ceph backup, OneSwap migration. Mutable, general-purpose Linux hypervisor OS (CISO veto axis), no in-kernel DFW for KVM, no native LB, no NDR.
5	Nutanix Cloud Infrastructure	72	Closest operational experience to VMware. Best day-2 (LCM). Non-immutable CentOS-based AHV penalises it on the 3.0×-weighted Lifecycle axis; HCL lock-in + $575/core Ultimate list hurt on cost and vendor-indep.
6	HPE Morpheus Enterprise	71	Real HPE support, strong multi-tenancy via Morpheus. Closed source + HPE-hardware-adjacent lock-in drags vendor-indep score down; mutable Ubuntu HVM hurts under heavy Lifecycle weight.
7	SUSE Harvester	67	Open HCI on SLE Micro transactional (immutable). SDN weakest in the set; 2025 price hike hurt.
7	Proxmox VE	67	Cheapest, fully GPL — openness boosts vendor-indep to 10. Mutable Debian, no DRS, no DFW, shallow multi-tenancy remain the feature gaps and now hurt harder under the heavier Lifecycle / SDN weights.

Scores are computed from the weighted rubric in §2. They should be interpreted as "how well does this stack replace what the customer is leaving" — they are not absolute quality ratings. Re-weighting for edge, VDI, pure-K8s, sovereign-cloud, or cost-driven scenarios will move the numbers; the §7 matrix exposes every per-category cell so readers can re-score under their own priorities.

2. Methodology & Scoring

Each stack is scored 1–10 on twelve evaluation categories. **Lifecycle / OS is weighted at 3.0× — the single heaviest axis in this rubric, above even Vendor-independence (2.5×)** — because enterprise CISO procurements treat replacing ESXi (immutable, signed-boot, TPM-measured, no-SSH appliance) with a general-purpose mutable Linux hypervisor (apt/dnf, SSH, writable root) as a binary veto, not a tradable feature. Vendor independence (2.5×) and Cost (2.0×) sit just below, reflecting that kill-switch protection and total cost of ownership are the next-most-decisive criteria when picking an open-substrate cloud platform. SDN & DFW (2.0×), Operations / fleet UI (1.5×), and Ecosystem / support (2.0×) round out the heavier-weighted axes. The final fitness score is the sum of weighted category scores divided by the maximum possible (205), expressed as 0–100.

Category	Weight	What the score measures
Hypervisor, HA, DRS, live migration	1.5×	Core compute: host-failure recovery, continuous load balancing, online mobility, maturity of the underlying hypervisor.
Storage (distributed / snapshots / replication)	1.5×	Scale-out block fabric, snapshot performance, sync/async replication, external-array integration breadth.
SDN & distributed firewall	**2.0×**	Overlay transport, distributed east-west firewall, micro-segmentation model, multi-tenant VRF/VPC. *NDR / L7 IDS is scored separately as an enabler in the deep dive but does not bump this axis — it is detection / response, not SDN.* **Weighted at 2.0×** because network-policy depth is one of the two largest functional decisions for the buyer (the other being storage).
Vendor independence / exit	**2.5×**	Open-source substrate, hardware lock-in, contractual and data exit paths, ability to keep operating without the commercial vendor.
Cost	**2.0×**	License list price, basis of charging (core / socket / node / instance), minimums, support tiers. 10 = cheapest.
Migration from VMware	1.3×	First-party tooling, CBT-accelerated incremental sync, scale of mass migrations, network-policy carry-over.
Multi-tenancy depth	1.2×	Projects/orgs, domains, per-tenant quotas, network isolation, RBAC granularity, self-service portal.
Backup & DR	1.0×	Native backup, dedup, air-gap/tape, synchronous/asynchronous replication, DR runbooks.
Lifecycle / OS / patching	**3.0×** — single heaviest axis	Upgrade ergonomics, immutable/A-B images, rollback safety, cluster-wide rolling updates, live patching. **Weighted at 3.0× — the single heaviest axis in this rubric, above even Vendor-independence (2.5×)** — because enterprise CISO procurements explicitly replacing ESXi treat the immutable-appliance-OS criterion as a *binary veto*, not a tradable feature. "We are exiting VMware and replacing an immutable, no-SSH, signed-boot, TPM-measured, A/B-image hypervisor appliance with… a general-purpose Linux distribution with apt and SSH" is the sentence the CISO refuses to sign off on; everything downstream of that filter is secondary.
Operations / fleet UI	**1.5×**	Day-2 pane-of-glass, multi-cluster fleet management, observability, breadth of functional domains in one UI. **Weighted at 1.5×** because real day-2 operating cost is driven by how many separate consoles your team must master to run the stack.
Kubernetes integration	1.0×	First-party K8s distribution or K8s-style orchestration, CAPI providers, pod/VM co-tenancy quality.
Ecosystem / enterprise support	**2.0×**	Size of installed base, 24×7 support availability, partner ecosystem, ISV integrations. **Weighted at 2.0×** because for any enterprise production stack the question "who do I call at 03:00 in production" is non-negotiable.

Maximum weighted score: 10 × (1.5+1.5+2.0+1.3+1.2+2.0+1.0+**3.0**+1.5+1.0+2.5+2.0) = **205**. Final fitness = (weighted total / 205) × 100, rounded.

Scoring inevitably involves judgement on individual cells; reasonable people can disagree on any one of them. The §7 matrix exposes every per-category score so readers can recompute under their own weights.

3. Stack Deep Dives

3.1 Aether + Incus + IncusOS + Operations Center + Migration Manager

COMMERCIAL / OPEN-CORE KVM + CONTAINERS

A five-component stack: Aether (Go-based multi-tenant control-plane web app) on top of Incus (the LXC-lineage KVM + OCI + system-container hypervisor), IncusOS (immutable A/B appliance OS on cluster nodes), Operations Center (multi-cluster provisioning & lifecycle), and Migration Manager (VMware-to-Incus migrator with NSX-V/NSX-T firewall rule import). Sold exclusively through Dell Technologies.

ARCHITECTURE

Aether is the customer-facing control plane — a Go application with a PostgreSQL backend and an HTML/Bootstrap web UI — that owns tenants, billing, RBAC, SDN firewall rules, backups, load balancers, the **Shepherd** Kubernetes-style workload orchestrator, and the **NDR** (Layer 7 IDS / IPS) layer. Under it, each customer cluster runs **IncusOS** on every node: a Debian-13-based `mkosi`-built immutable appliance image, Secure Boot with a signed Unified Kernel Image measured into TPM PCRs, LUKS-encrypted root, ZFS, A/B updates driven by `systemd-sysupdate`, no SSH on cluster nodes, and management exclusively through the Incus REST API (or the `incus admin os` proxy to IncusOS's unix-socket API).[58] **Incus** is the hypervisor — a Canonical-lineage fork of LXD, Apache-2.0 licensed, now maintained by LinuxContainers under Aleksa Sarai and Stéphane Graber — running KVM/QEMU VMs, OCI containers, and LXC system containers with a dqlite-replicated cluster state. **Operations Center** (FuturFusion) federates many Incus clusters: it registers IncusOS servers, drives A/B updates, provisions clusters via Terraform, and proxies an inventory mirror. **Migration Manager** (FuturFusion) handles the VMware-to-Incus VM transfer via VDDK + CBT warm-incremental disk copy, batch / queue management, per-batch worker VMs, and post-import OS fixups; the Aether control plane wraps Migration Manager and additionally imports NSX-V and NSX-T firewall rules into OVN ACLs as part of the cutover, carrying security policy across with the VMs.

HYPERVISOR / HA / DRS / LIVE MIGRATION

KVM via Incus with **online live migration**, quorum-based clustering through dqlite, and automatic VM restart on node failure. Aether ships **continuous, load-based DRS**: a background rebalancer monitors CPU and memory utilisation across the cluster and live-migrates workloads off hot nodes onto cooler nodes, together with affinity/anti-affinity placement and initial scheduler decisions. Incus supports vGPU passthrough, PCI passthrough, and both virtio and SR-IOV networking.

STORAGE

Three serious options: **Ceph RBD** (most common in production), **lvmcluster** (shared block via LVM with a cluster lock manager — lighter-weight HCI), and local ZFS/Btrfs for non-HA workloads. Snapshots and volume copies are first-class Incus primitives; cross-cluster replication uses Ceph RBD mirroring or Incus instance snapshots to a remote. **Instance backup to S3** uses Incus's native `backup_s3_upload` extension — data streams directly from the cluster to the object store, bypassing the Aether app server entirely.

SDN

OVN with Geneve overlays, distributed routing, distributed NAT, and — critically — **security groups enforced as OVN ACLs in OpenFlow on every node**, which is architecturally equivalent to NSX's distributed firewall model. Aether exposes security policy at three scopes: **CSP-level global firewall rules** that apply across every cluster and every tenant (the equivalent of NSX Federation's cross-site distributed security policy), **cluster-level policies**, and **tenant-level policies**. Multi-tenant VRFs are modelled as OVN logical routers; tenant virtual clusters use Incus projects (`?project=tenant-{slug}`).

Load balancing is enterprise-grade and built-in. Aether deploys an active-active pair of HAProxy instances behind an OVN load balancer VIP as an integrated service: the OVN LB handles L4 fan-out to both HAProxy nodes, HAProxy handles full L4 and L7 (HTTP/HTTPS/TCP modes, round-robin / least-connections / source-hash algorithms, HTTP and TCP health checks, sticky sessions, SSL termination and passthrough, rate limiting, access control, and advanced security/performance options). A Keepalived/VRRP mode is also supported for non-OVN networks. This is comparable in function to VMware Avi Load Balancer for the typical North-South + East-West LB use case, and is managed from the Aether UI rather than being a separately-licensed add-on.

MULTI-TENANCY

Aether is multi-tenant by design and the model is unusually complete for an enterprise-grade VMware replacement: each tenant gets an Incus project (with `?project=tenant-{slug}` URL scoping), an explicit CPU/memory/storage quota envelope, a granular RBAC role set (`full_admin`, `is_ffsdn_backup_admin`, NDR admin, shepherd admin, etc.), isolated OVN networks per tenant, **integrated per-tenant billing with showback / chargeback**, and a **polished customer-facing self-service portal** baked into the same Aether application. Firewall policy is hierarchical across **three scopes** — CSP-global, cluster-level, and tenant-level — so a service provider can layer mandatory CSP rules on top of customer rules without the customer being able to disable them. Multi-cluster operation is centralised in the same Aether control plane (one

logical pane managing many Incus clusters) rather than relying on federation of independent instances. Depth is comparable to Nutanix Projects, materially better than Proxmox Pools or bare Harvester, and on par with OpenStack's Keystone domains/projects on the tenant-isolation axis while exceeding it on the integrated-billing-and-portal axis.

BACKUP & DR

Built on Incus's `backup_s3_upload` (Incus 6.15+). S3-only target model: Dell ECS / ObjectScale / PowerScale S3, MinIO, AWS S3, Wasabi all work. Backup data streams directly from the Incus cluster to the object store. Restore flows through the Aether app (Incus has no native S3 import). DR via Ceph RBD mirroring or instance replication at the Incus layer. **Commvault is certified and integrated** as an enterprise-grade third-party backup target for Aether / Incus today — meaning customers standardising on one of the industry's top-tier backup platforms are covered out of the box. Veeam and Dell backup-product integrations are on the near-term roadmap.

MIGRATION FROM VMWARE

FuturFusion Migration Manager is the first-party VM migrator: vCenter source registration, Incus target, batch-based migration queue, worker-VM model for VDDK/CBT disk copy, post-import OS fixups for Linux and Windows. **Uniquely in this comparison, Aether also migrates firewall rules from NSX-V and NSX-T**: the security-policy import carries **both east-west DFW rules and north-south edge firewall rules** (gateway, T0/T1, NAT) into OVN ACLs alongside the VM migration, so the landing environment arrives with its full micro-segmentation *and* perimeter policy already in place. No other stack in this document ships a first-party NSX rule migrator.

OPERATIONS / FLEET

Aether is itself the fleet UI — it manages many Incus clusters from one pane of glass. Operations Center is the lifecycle / provisioning layer for IncusOS. A **single Aether UI covers 14 functional domains**: tenants, billing, RBAC, compute (VMs / OCI / system containers), storage, OVN SDN, distributed firewall (at CSP / cluster / tenant scopes), enterprise load balancing (OVN+HAProxy active-active), Commvault-integrated backup, Shepherd K8s-style orchestration, Migration Manager (with NSX-V/NSX-T rule import), Operations Center (IncusOS fleet lifecycle), **NDR (Layer 7 IDS / IPS with reactive containment)**, and a customer-facing self-service portal. **The functional breadth in one UI is matched only by Nutanix Prism Central and HPE Morpheus Enterprise** among the stacks in this document — and Aether is arguably broader than both, because it natively includes *enterprise load balancing, Kubernetes-style orchestration, VMware migration, and L7 IDS / IPS* in the same pane, where Nutanix and Morpheus both require separate products or consoles (Move, NKP, MKS, Veeam, third-party NDR appliances, etc.) for those functions. VCF 9 by contrast still splits across four UIs (vSphere Client + VCF Operations + VCF Automation + SDDC Manager); OpenStack, OpenNebula, and Proxmox are fragmented or single-zone-scoped by design.

NDR — LAYER 7 IDS WITH IPS-STYLE AUTOMATED CONTAINMENT

Aether ships an out-of-band **Layer 7 NDR layer** on top of the OVN substrate. Default engine is OSS — **Zeek + Suricata + RITA**, with Arkime on-demand for full-PCAP forensics — and the vendor brain is pluggable (Vectra, Darktrace, ExtraHop, Corelight, Cisco SNA). Detection is out-of-band over OVN port mirroring (ERSPAN to a sensor container per IncusOS host in `project=internal`, no inter-host sensor traffic); coverage is every Layer 7 protocol Suricata and Zeek parse — HTTP, TLS (with JA3 and JA4 fingerprints), DNS, SMB, RDP, Kerberos, MQTT, SSH, FTP, SMTP, and 20+ more. Five abuse.ch threat-intel feeds ship enabled by default (Feodo Tracker, URLhaus, SSLBL, ThreatFox, JA3 fingerprints) plus custom MISP / OTX / AbuseIPDB / generic-URL sources; the Aether server fetches feeds from the public internet and pushes the rendered Suricata files to every sensor — sensors themselves never touch the internet. An Aether-curated daily-updated subset of upstream Suricata SIDs gates auto-quarantine. Enforcement is reactive **OVN-ACL quarantine** at the instance scope (whole-instance block; per-5-tuple is on the Tier-0 roadmap), with phased rollout (log_only → soft → hard → power). On hard / power auto-quarantines an Arkime forensic session can auto-start. No other stack in this comparison ships a native L7 IDS / IPS with automated containment at this scope.

KUBERNETES / SHEPHERD

Rather than bundling a Kubernetes distribution, Aether ships **Shepherd** — a declarative Kubernetes-style orchestration engine for Incus. Shepherd exposes familiar primitives: **Flocks** (namespaces), **InstanceGroups** (Deployments — declarative N-replica specs with image, CPU, memory, and rolling-update strategy), **ServiceEndpoints** (Services — OVN load balancers with stable VIPs routing to healthy instances), **HerdRules** (HorizontalPodAutoscalers — CPU/memory-driven auto-scaling with separate scale-up/down cooldowns), **ConfigMaps**, **Secrets** (encrypted at rest, injected via Incus metadata), and a parameterised **App Catalog**. The reconciler loop enforces desired state with self-healing, rolling updates, automatic rollback on failed health checks, and anti-affinity placement. For customers who want standard Kubernetes, K8s guest clusters can still run as Incus VMs — but Shepherd is explicitly positioned as a Rancher+Kubernetes alternative: the K8s primitives without the K8s operational tax.

LIFECYCLE / OS — THE IMMUTABLE-APPLIANCE STORY AND THE CISO WIN

IncusOS is the modern part of the stack and the architectural feature most directly addressed at the CISO who is signing off the VMware exit. It preserves the ESXi appliance model rather than regressing from it: a Debian-13-based `mkosi`-built immutable appliance image with **A/B atomic updates driven by** `systemd-sysupdate` (the inactive A/B slot receives the new `usr-verity` + `usr-verity-sig` + UKI; reboot into the new slot promotes it, automatic rollback on a failed boot), **Secure Boot measured into TPM PCRs** (signed UKI with embedded PCR signature, attested to the management plane), **LUKS-encrypted root**, **ZFS**,

minimal service surface, **no SSH on cluster nodes** — all management goes through the authenticated Incus REST API or the `incus admin os` unix-socket proxy. Cluster-wide upgrades are orchestrated from Operations Center: live-migrate workloads off one node at a time, stage the A/B image, reboot into the new slot, repeat. Application sysexts (e.g. inside the appliance for Operations Center or Migration Manager hosting) update via sysext-swap with live restart; core-OS updates stage to the inactive A/B slot and require a reboot, but the rolling-drain-and-reboot pattern with live migration keeps cluster workloads online throughout the upgrade window. This is the dominant reason FuturFusion's Lifecycle/OS score is 9/10 on the 3.0×-weighted axis in §7 and it is the largest single contributor to its lead over the mutable-Linux alternatives in this comparison.[58]

LICENSING & PRICING

Aether is licensed at $80/core/year including support, and is resold exclusively through Dell Technologies. Customers purchase Aether as a Dell product, on Dell paper, under Dell support contracts, with Dell as the single point of contact for escalation, SLAs, and lifecycle. Incus, IncusOS, Operations Center, and Migration Manager are open source (Apache 2) and free to run standalone; the Aether subscription layers the commercial control plane and support on top. At list, for a 1,280-core cluster, three years of Aether is $307,200.

STRENGTHS

- **Genuinely single pane of glass across 14 functional domains** — compute, storage, SDN, distributed firewall, L4/L7 load balancer, backup, Kubernetes-style orchestration (Shepherd), VMware migration (incl. NSX rule import), tenants, billing, RBAC, IncusOS fleet lifecycle, **NDR (Layer-7 IDS / IPS)**, and self-service portal — all from one Aether UI. Matched in this comparison only by Nutanix Prism Central and HPE Morpheus Enterprise; broader than both on LB, K8s, migration integration, and L7 IDS / IPS; materially broader than VCF 9 (4 UIs), OpenStack (fragmented), OpenNebula (FireEdge covers compute/storage/networking/OneKE/OneFlow but lacks LB, DFW, NDR, integrated billing portal, and NSX-rule importer), Proxmox (federated), and the K8s-based stacks (console + ACM/Rancher + GitOps + backup operators).
- **Sold and supported exclusively through Dell Technologies.** Customers get a Dell SKU, Dell paper, Dell support, and Dell's global enterprise delivery organisation — one of the largest IT support footprints in the industry — so the "ecosystem and support" concern that applies to boutique alternatives does not apply here.
- Priced well below Nutanix, full Red Hat OpenShift Platform Plus, SUSE Harvester Suite, RHOSO, and VCF 9 — at $3.07M / 3 years it is ~14% of Nutanix NCI-Ultimate ($22.08M), ~23% of VCF 9 base list ($13.44M), and a small fraction of VCF 9 with Live Recovery ($46M); but materially above HPE Morpheus Enterprise ($1.5M), Proxmox Premium ($850k), OpenShift Virtualization Engine ($570k), and OpenNebula ($300k–$890k).
- **Modern, immutable OS model — critical for CISO sign-off.** IncusOS's Debian-13-based `mkosi`-built immutable image with A/B atomic updates via `systemd-sysupdate`, signed UKI measured into TPM PCRs, LUKS-encrypted root, ZFS, and no SSH on cluster nodes is ahead of Proxmox, OpenStack, OpenNebula, Nutanix AHV, and HPE VME on the heaviest-weighted axis in the rubric. Core-OS updates stage to the inactive A/B slot and require a reboot, but the rolling-drain-and-reboot pattern with live migration keeps cluster workloads online throughout. For enterprise CISO departments evaluating VMware replacements, a mutable general-purpose Linux hypervisor (AHV's CentOS-based base, Proxmox's Debian, HPE VME's Ubuntu, OpenNebula's RHEL/AlmaLinux/Ubuntu/Debian/SLE) is often a veto criterion before feature or cost comparisons even begin — IncusOS passes that filter.
- OVN distributed firewall *plus* CSP-level global rules deliver both the NSX DFW model (east-west micro-segmentation per host) and the NSX Federation model (cross-site/cross-tenant security policy) in one stack.
- Continuous load-based DRS, live migration, and HA are all built-in.
- **Shepherd** — a Kubernetes-style orchestrator for Incus — delivers declarative Deployments, Services (OVN LB), HPAs, ConfigMaps, Secrets, and an app catalog without requiring a separate Kubernetes distribution or Rancher stack to operate.
- **Enterprise load balancing is built in, not bolted on**: OVN LB frontend + active-active HAProxy backend covers L4 and L7, HTTP/TCP health checks, sticky sessions, SSL, rate limiting, and access control — managed from the Aether UI, no Avi license required.
- **Only stack in this comparison that migrates NSX-V and NSX-T firewall rules** alongside the VMs — **both east-west DFW rules and north-south edge/gateway rules**. Customers carry their full security policy (micro-segmentation + perimeter) with them in a single cutover.
- Open-source substrate (Incus, OVN, Ceph, IncusOS) means zero kill-switch risk — you can keep operating without Aether if you had to.
- Multi-tenancy is genuinely built in rather than grafted on.
- Certified, integrated enterprise backup via **Commvault** on top of Incus's native S3 streaming backups (Veeam and Dell backup integrations on the roadmap).

WEAKNESSES

- No CPU-lockstep Fault Tolerance — consistent with every other alternative in this document, but a gap vs. VCF 9's vSphere FT.

- Aether-branded installed base is smaller than VMware or Nutanix — substantially offset by Dell's global partner reach and by certified tier-1 ISV integrations (Commvault today; Veeam and Dell backup solutions on the roadmap). The long-tail niche ISV ecosystem is still younger than Broadcom's.
- No NVMe memory tiering or confidential-compute (SEV-SNP / TDX) surfaced at the Aether control-plane layer yet, although the Incus substrate supports both.
- Shepherd is a Kubernetes-*equivalent*, not Kubernetes itself. Shops with heavy investment in Helm charts, existing operators, or CNCF-conformant tooling may still want to run upstream K8s (possible, but as Incus guest clusters rather than a first-party distribution).
- Operations Center and Migration Manager are young products; mileage on edge-case VMware environments varies.
- No bundled Private-AI-Foundation or Data-Services-Manager equivalent (DBaaS) out of the box.

3.2 VMware Cloud Foundation 9 (reference) COMMERCIAL TYPE-1 ESXI

Broadcom's integrated private cloud: vSphere 9 + vSAN ESA + NSX + VCF Operations + VCF Automation + HCX + vSphere Kubernetes Service + (optional) Avi Load Balancer, Private AI, Data Services Manager, and Live Recovery. GA 17 June 2025.[8]

ARCHITECTURE

VCF 9 collapses the historic VMware portfolio into two SKUs: **VCF** (the full stack) and **VVF** (vSphere Foundation, lighter). Standalone vSphere Standard / Enterprise Plus are legacy, capped at vSphere 8 Update 3, EOL October 2027.[9] New in 9: a **unified VCF Installer** replacing Cloud Builder, **VCF Automation** (ex-Aria Automation, absorbing vCloud Director concepts), and **Fleet Management** inside VCF Operations, which centralises lifecycle, certificates, passwords, tags, and drift across the estate. SDDC Manager is **on the deprecation path** — its functions are migrating into VCF Operations.[8][10]

HYPERVISOR / HA / DRS / LIVE MIGRATION

The reference. vSphere HA, DRS (continuous, cluster-wide vMotion-driven load balancing), Storage DRS, Fault Tolerance (lockstep), EVC, non-disruptive vMotion and Storage vMotion. VCF 9 adds **AMD SEV-SNP + Intel TDX confidential computing** and **NVMe Memory Tiering** extending DRAM with NVMe.[8]

STORAGE

vSAN 9 ESA is the strategic path. Global (cluster-wide) deduplication, native B+ tree snapshots with no delta chains (up to ~200 snapshots/VM), stretched clusters across two AZs with a witness, disaggregated **vSAN Storage Clusters**, and Enhanced vSphere Replication. **vVols are deprecated in 9.0** — a major signal to customers on Pure, NetApp, Dell external arrays that Broadcom wants them on vSAN.[11]

SDN

NSX 9 with Geneve overlay, in-kernel Distributed Firewall (DFW) with expanded L7, T0/T1 gateways, and — new in VCF 9 — a **native Virtual Private Cloud (VPC) construct** with **Transit Gateway** (both centralized and distributed). NSX Federation spans sites with distributed security policy. Avi Advanced Load Balancer is an optional add-on; NSX Edge provides basic LB/NAT.[12][13]

MULTI-TENANCY

VCFA Organization → Project → vSphere Namespace → VKS Namespace. Far better than classic vSphere, but complex and heavily VCFA-dependent. Customers coming from vCloud Director are being forced onto VCFA.[14]

BACKUP & DR

No native backup product. VDP has been gone for years and was never replaced. Real-world backup is 100% third-party (Veeam, Cohesity, Rubrik, Commvault, Dell PowerProtect). DR is **VMware Live Site Recovery 9.0.4** (formerly SRM) + Enhanced vSphere Replication, integrated into VCF Operations, with RPOs as low as 1 minute. VLR is **sold separately**.[15]

MIGRATION FROM VMWARE

N/A — it is VMware. HCX (now VCF Operations HCX) is included and handles bulk, live, and OS-assisted moves inside the VMware universe.

OPERATIONS / FLEET

Four primary UIs with active consolidation: vSphere Client, VCF Operations (mandatory in 9, central), VCF Automation, and (still present, but being absorbed) SDDC Manager. Fleet Management unifies config-drift detection, tag management, cert renewal, and password rotation across the VCF estate.[10]

KUBERNETES

vSphere Kubernetes Service (VKS), formerly TKG Service, running on vSphere Supervisor. VCFA 9.0.1 added native VKS cluster lifecycle + backup/restore. Tanzu Application Platform and higher Tanzu layers were descoped/sold by Broadcom.[16]

LIFECYCLE / OS

vLCM image-based management is mandatory for new clusters in VCF 9. Declare desired state (ESXi version, firmware, drivers); vLCM enforces it. **ESX Live Patch** extends live patching to vmkernel, user-space, and NSX components — updates without reboot or VM migration — a genuine 2025 feature.[17]

LICENSING & PRICING

Perpetual licensing ended January 2024. Subscription-only. **VCF list ~$350/core/year**; **VVF list ~$135/core/year**. 16-core-per-CPU minimum; 72-core-per-order minimum enforced April 2025. Smallest legal VCF new-order is 72 × $350 = $25,200/year at list.[1][2][18] Trade press has reported 150%–1,200% real-customer increases.[3][4] AT&T filed court evidence showing a 1,050% offer; Computershare migrated 24,000 VMs to Nutanix; Beeks Group moved ~20,000 VMs to OpenNebula.[5][6][7]

Live Recovery is the brutal add-on. If the customer wants VMware Live Recovery (SRM + Live Site Recovery) for DR — and almost every enterprise VCF deployment does — the effective per-core cost rises to approximately **$1,200/core**: the customer licenses VCF at both the *active* site and the *standby* site (2× the base VCF core count), and pays an additional per-VM Live Recovery subscription on top. A DR-enabled VCF 9 deployment therefore commonly lands at 3–3.5× the headline $350/core figure.

VVF is being withdrawn from parts of EMEA. In December 2025, Broadcom began pulling VMware vSphere Foundation (VVF) — the lower-tier $135/core SKU — from multiple EMEA countries, effectively forcing those customers onto VCF at $350/core. Trade press reported one customer's bill jumping from ~$130,000/year on VVF to ~$1.3M/year on VCF — a 10× increase driven purely by SKU withdrawal. Smaller customers who cannot justify the full VCF feature set are the most affected. The cost model in §6 therefore excludes VVF as a viable comparison in 2026.

STRENGTHS

- Feature breadth nothing else matches: compute, HCI storage, SDN, LB, K8s, automation, ops, DR, DBaaS, GPU, confidential compute, AI foundation — one SKU.
- The reference implementations for HA/DRS/vMotion, vSAN ESA, and NSX DFW.
- ~80% of enterprise VM workloads still run on ESXi as of 2024–2025. Largest ecosystem on the planet.

WEAKNESSES

- Post-Broadcom pricing shock is the dominant reason this document exists.
- Forced bundling — small shops cannot buy just vSphere.
- Channel and partner turmoil since early 2024.
- Documented customer exodus (Computershare, Beeks, AT&T dispute). Analysts project 100,000+ customers may exit over the medium term.[5][6][7]
- Deep lock-in (VCFA, NSX Federation, vSAN ESA, VKS). HCX helps you move *inside* VMware, not out.
- vVols deprecation signals disinvestment in third-party array interop.

3.3 OpenStack OPEN SOURCE KVM + IRONIC

Modular, API-driven open-source IaaS. Family of independent services (Nova, Neutron, Cinder, Keystone, Glance, Swift, Heat, Horizon, Ironic, Octavia, Magnum, Manila, Designate, Masakari) sharing Keystone identity and a message bus. Current releases: 2024.1 Caracal, 2024.2 Dalmatian, 2025.1 Epoxy, 2025.2 Flamingo.

ARCHITECTURE & COMPONENTS

OpenStack's core is 30+ microservices. The canonical set is Nova (compute), Neutron (SDN, now OVN by default), Cinder (block, ~80 driver backends), Glance (images), Keystone (identity, domains, projects, RBAC), Swift (object), Heat (orchestration), Horizon (reference dashboard, being replaced by Skyline), Ironic (bare-metal), Octavia (LBaaS), Magnum (K8s-as-a-service), Manila (file), Designate (DNS), Barbican (secrets), Masakari (VM HA), Placement (scheduler). Control-plane services run HA behind HAProxy/keepalived or, in modern distros, inside Kubernetes. Every service exposes its own REST API and scales independently — the reason OpenStack scales to 100k+ cores and also the reason it has a heavy-ops reputation.

HYPERVISOR / HA / DRS / LIVE MIGRATION

Hypervisor-agnostic in theory, **KVM/libvirt in practice**. Ironic for bare metal. Live migration (pre-copy, post-copy, auto-converge) with shared or block storage; vGPU live migration landed in Caracal. VM HA is not in Nova — it is delivered by **Masakari**, which uses pacemaker host monitors that notify Masakari API to evacuate VMs to healthy nodes in a "failover segment." **No DRS equivalent**; Watcher provides rule-based audits but is nowhere near vSphere DRS.

STORAGE

Cinder abstracts ~80 drivers (Pure, NetApp, Dell, HPE, IBM, Huawei, Infinidat, etc.). **Ceph RBD is the de facto default for greenfield OpenStack**, integrated across Nova / Cinder / Glance / Manila, delivering distributed block, S3/Swift object via RGW, and CephFS. Snapshots, clones, replication (v2.1), multi-attach. Manila handles shared NFS/CIFS/CephFS. Cross-site via Ceph RBD mirroring or driver-level replication.

SDN

Neutron ML2/OVN is now the recommended default. OVN replaced the legacy OVS + DHCP/L3/metadata agents with a db-driven control plane and per-host `ovn-controller`. Geneve overlays, distributed virtual routing, distributed NAT/DHCP, and — critically — **security groups implemented as OVN ACLs in OpenFlow on every compute node**. This is the architectural equivalent of NSX DFW and is one of the three serious distributed-firewall stories in this document. Overlapping tenant IPs and per-tenant VRFs are standard. Octavia (Amphora or OVN provider) for LBaaS.

MULTI-TENANCY

Best in this document. Keystone's model is Domains → Projects → Users/Groups → Roles. Every resource is project-scoped. Since Yoga, **Secure RBAC (SRBAC)** adds system/domain/project scopes and a reader/member/manager/admin hierarchy. Domains allow hosting providers to federate separate customer identity silos behind one cloud. Quotas exist per project for every resource type.

BACKUP & DR

Cinder has native volume backups to Swift/Ceph/NFS/S3. Nothing in upstream OpenStack is a full application-consistent backup suite. Production shops run **Trilio TrilioVault**, Commvault, or Veeam Kasten for K8s. Cross-site DR is usually Cinder replication + Heat templates.

MIGRATION FROM VMWARE

No single first-party tool. Mature third-party options: **CloudBase Coriolis** (commercial, CBT-based incremental), **virt-v2v**, Platform9 Arcus, Mirantis Migration Tool. The OpenInfra Foundation formed a VMware Migration working group in 2024; over 80% of member vendors had already done VMware migrations by mid-2024.[19]

OPERATIONS / FLEET

Horizon is dated; **Skyline** is the new React UI gaining adoption. Most operators use the `openstack` CLI or Terraform/OpenTofu. Fleet-level multi-cloud management is not built in — Red Hat ships ACM+ArgoCD, Canonical ships Juju+Landscape, Mirantis ships Mirantis Container Cloud.

KUBERNETES

Magnum provisions K8s clusters (modern driver uses Cluster API). Zun (container-as-a-service) is effectively dead. Real shops run Cluster API with the OpenStack provider, Red Hat OpenShift on OpenStack (RHOSO is literally OCP on OpenStack on OCP), or Rancher/RKE2. Cinder-CSI and Manila-CSI are first-class.

LIFECYCLE

Two named releases per year. Since Caracal, **SLURP (Skip-Level Upgrade Release Process)** lets operators upgrade annually instead of every 6 months — a genuine improvement. Host OS is typically RHEL/CentOS Stream (RHOSO runs on RHCOS), Ubuntu 24.04 (Canonical), or Ubuntu under Mirantis.

LICENSING & PRICING

OpenStack itself is Apache 2 and free. You pay vendors for support, lifecycle, and opinionated installers:

- **Red Hat OpenStack Services on OpenShift (RHOSO 18.0)** — per populated 2-socket server; quote-only. Historical RHOSP SKUs ran ~$3,000–$5,000/socket-pair/year pre-discount; RHOSO additionally requires OpenShift entitlements for control-plane nodes.[20]
- **Canonical Charmed OpenStack** — ~$1,500/year/node Advanced, ~$4,275/year/node fully managed, plus $75k–$150k onboarding services.[21]
- **Mirantis OpenStack for Kubernetes** — quote-only, 30–50% below RHOSP.
- **StarlingX** — free, Wind River Studio commercial support.
- **Rackspace OpenStack Flex** (April 2025) — fully-managed, consumption-priced.[22]

STRENGTHS

- Genuinely open, no single-vendor kill switch. OpenInfra Foundation membership grew materially in 2024 as VMware refugees arrived.
- Massive scale proven — CERN, Walmart, Yahoo!JAPAN, GEICO (dropped VMware for OpenStack in 2024).[23]
- OVN + Ceph + Ironic is a hyperscaler-class stack for a zero-license-cost substrate.
- SLURP finally made annual upgrades realistic.
- Deepest multi-tenancy model in the document.

WEAKNESSES

- "OpenStack is hard to operate" is still largely true without a commercial distro.
- No first-class DRS, no first-class backup suite, no first-class fleet manager.
- Upstream breakage between minor releases still occurs; policy overhaul broke third-party scripts.
- Vendor support fragmentation — each distro is meaningfully different.

3.4 Nutanix Cloud Infrastructure (NCI) + Cloud Manager (NCM) COMMERCIAL AHV / KVM HCI

Hyperconverged platform: same x86 nodes run compute and a distributed storage fabric (ADSF). Default hypervisor is AHV (Nutanix's KVM fork). Prism Central + NCM deliver multi-cluster management, Flow (micro-seg + SDN), Files/Objects/Volumes for unified storage, NDB for DBaaS, Move for VMware migration, Nutanix Disaster Recovery for DR.

ARCHITECTURE

Each node runs a Controller VM (CVM) that owns local NVMe/SSD and joins the **Acropolis Distributed Storage Fabric** — a scale-out, shared-nothing, log-structured block store. The hypervisor is **AHV** (CentOS-based KVM fork) by default; ESXi and Hyper-V are still supported. Prism Element manages one cluster; **Prism Central + NCM** manage many and add self-service, cost governance, intelligent ops, and security central. Unified storage add-ons: Files (SMB/NFS scale-out), Objects (S3 with WORM), Volumes (external iSCSI). Data services: NDB (fleet DB-as-a-service for Postgres/Oracle/SQL/MySQL/MariaDB). **NC2** is the separate SKU that projects the same stack onto AWS/Azure bare metal.

HYPERVISOR / HA / DRS / LIVE MIGRATION

AHV = KVM/QEMU with a Nutanix-authored Acropolis management plane. Parity vs. ESXi: live migration, VM HA, snapshots, templates, affinity/anti-affinity, vGPU (NVIDIA), UEFI/Secure Boot, vTPM. Gaps: no Fault Tolerance (lockstep), weaker NIC teaming, smaller ISV plugin ecosystem. VM HA is automatic; "Guarantee HA" reserves enough cluster capacity to survive N failures. **DRS-equivalent = Acropolis Dynamic Scheduling (ADS)**: the Lazan service runs every 15 min, detects >85%-for-15-min hotspots, and live-migrates two VMs at a time to rebalance. Reactive and throttled, not continuous like vSphere DRS, but present.

STORAGE

ADSF is Nutanix's core differentiator. Data is written locally first (data locality) then async-replicated to 1 or 2 peers (RF2/RF3). Features: inline/offline compression, post-process dedupe, **EC-X erasure coding** (~1.5× capacity gain), ILM tiering NVMe→SSD→HDD, auto-healing, capacity-aware rebalancing. Snapshots are COW at the vDisk level, instant, effectively unlimited. Replication: Async (≥1 hour RPO), NearSync (1–15 min), Metro (0 RPO, active-passive <5ms RTT). Unified storage: Files (SMB/NFS), Objects (S3 + WORM), Volumes (iSCSI LUNs).

SDN

Nutanix Flow has two products: **Flow Network Security** (identity-aware category-based micro-segmentation enforced by OVS flow rules on each AHV host — the NSX DFW equivalent) and **Flow Virtual Networking** (overlay SDN on OVS/OVN + Geneve with VPCs, per-VPC virtual routers, overlapping subnets, integrated NAT, policy-based routing, IPsec VPN, L2 VXLAN stretch, BGP peering to fabric). Load balancing is "bring your own" (F5, Citrix, integrated ALB).

MULTI-TENANCY

Prism Central implements **Projects** (tenant boundary with quota, users, default network) and **Categories** (metadata taxonomy driving Flow and protection policy). ~20 built-in RBAC roles plus custom. No domain concept; cross-cluster RBAC inheritance is limited. For service-provider-grade multi-tenancy, Nutanix pushes NCM Self-Service (ex-Calm) as the portal.

BACKUP & DR

DR is one of Nutanix's strongest stories: **Nutanix Disaster Recovery** (formerly Leap/Xi) with Async/NearSync/Sync and orchestrated Recovery Plans (boot order, network remap, script hooks, tested failover/failback). Backup: HYCU (literally born inside Nutanix), Veeam, Commvault, Rubrik, Cohesity all have AHV agents. **Mine** is Nutanix's turnkey backup appliance integrating with Veeam/HYCU/Commvault.

MIGRATION FROM VMWARE

Nutanix Move — free Linux VM appliance, agentless for most guests, agent-required for Windows. CBT-based incremental sync, cut-over downtime is the last delta + reboot. Move 6.0 (2024) added VCF/NSX security policy import into Flow. Nutanix reported **68,000 migrations via Move in May–July 2025, +70% YoY**.[24] Gaps: no vSAN stretch, no NSX-T logical router conversion end-to-end, no RDM, Windows occasionally needs hand-holding.

OPERATIONS / FLEET

Prism Central is genuinely best-in-class: VM lifecycle, storage, networking, DR, Flow, NDB, reports, capacity, cost. X-Play for low-code automation ("when X happens, do Y"), NCM Intelligent Ops for right-sizing, NCM Cost Governance (ex-Beam) for showback. Up to 300+ clusters and 25,000+ VMs per Prism Central.

KUBERNETES

Nutanix Kubernetes Platform (NKP) built on the acquired D2iQ Kommander assets, replacing Karbon. Multi-cluster management, GitOps via Flux, observability, runs on AHV / ESXi / bare metal / public cloud. Nutanix CSI exposes Volumes/Files/Objects to K8s. Partners with Red Hat so OpenShift is fully supported as a guest.

LIFECYCLE

LCM (Life Cycle Manager) is Nutanix's strongest day-2 feature: one-button upgrade for AOS, AHV, firmware (BIOS/BMC/HBA/NIC), and every Prism service. Rolling, transparent. LTS channels for conservative customers. AOS on a quarterly feature-release cadence plus monthly security.

LICENSING & PRICING

Sold **per physical CPU core** with tiered SKUs. NCI-Starter, NCI-Pro, NCI-Ultimate, NCI-Edge, NCI-Compute, NCI-D. NCM layers Starter/Pro/Ultimate for self-service and cost governance. Files, Objects, NDB separately licensed. Subscription 1/3/5 yr, minimum 3 nodes, license must cover **every populated core in the cluster** — the biggest gotcha on large-core modern CPUs. Nutanix does not publish list prices on its website, but distributor price books and CDW/Insight reseller SKU listings expose them. For SKU `NT-NCP-NCI-ULT-PR` (NCI-Ultimate with Production Support): **$575/core for 1 year, $1,724/core for 3 years, $2,873/core for 5 years** — i.e. **$575/core/year** list. NCI-Pro is lower, Starter lower still. Mission Critical support tier is $1,987/core (3-yr), Federal Mission Critical is $2,860/core (3-yr). NCM layers are additional.[25][26][46] For the 12,800-core reference cluster, Ultimate Production list works out to **$22,080,000 over three years** — more than VCF 9's base list for the same footprint, a fact that surprises customers assuming Nutanix is the "cheap alternative" to VMware. Nutanix raised list prices in 2024 and again in 2025 via a "PnP 2.0" low-list/low-discount regime.

STRENGTHS

- Best-in-class day-2 ops — LCM, Prism Central, one-click upgrades.
- Genuinely turnkey: storage + hypervisor + DR + micro-seg in one stack.
- AHV removes a hypervisor license line-item entirely.
- Strongest built-in DR story in this document.
- Real, measured VMware migration momentum (68k VMs in Q2 2025 alone).[24]
- Consistent hybrid-cloud extension via NC2 on AWS/Azure.

WEAKNESSES

- **AHV is NOT an immutable hypervisor OS.** AHV is based on CentOS 7 (earlier CentOS 6) — a traditional mutable Linux distribution with a writable root filesystem, SSH access, package manager, and no A/B image model. Secure Boot is a guest-VM feature in AHV, not a hypervisor-level feature. This is a material regression versus ESXi (stateless hardened appliance) and often a deal-breaker for enterprise CISO departments replacing VMware — swapping an immutable appliance for a general-purpose Linux adds attack surface rather than preserving it. Nutanix LCM is excellent for upgrade ergonomics, but the underlying OS model is architecturally older than IncusOS, RHCOS, and SLE Micro.
- **Hardware HCL lock-in.** Nodes must be on the Nutanix HCL (NX, Dell XC, HPE DX, Lenovo HX, Cisco UCS-HX, Fujitsu). Swapping parts outside HCL voids support.
- **Per-core pricing** bites hard as CPU core counts explode — licence cost scales with silicon, not workload.
- Opaque pricing and the 2024–2025 list hikes frustrate buyers.
- Multi-tenancy lags OpenStack for true service-provider use cases (no domains).
- AHV ecosystem still smaller than ESXi — niche ISVs are often vSphere-only.
- ADS is reactive and less sophisticated than vSphere DRS.
- "Escape velocity" is real: once you have 20 PB on ADSF, leaving means a full data migration — the same lock-in concern customers are fleeing at VMware.

3.5 Proxmox VE 8.x + PBS + PDM + SDN

OPEN SOURCE

DEBIAN + KVM/LXC

Debian 12 + KVM + LXC + corosync + pmxcfs + a fast web UI. Proxmox Backup Server for dedup backups. Proxmox Datacenter Manager hit v1.0 on 4 Dec 2025 for multi-cluster fleet. SDN plugin with Simple/VLAN/QinQ/VXLAN/EVPN zones.

ARCHITECTURE

PVE is Debian 12 wrapping KVM (VMs), LXC (system containers), a custom `pve-manager` web UI on port 8006, and a REST API. Cluster state is held in **pmxcfs**, a FUSE-mounted database-backed filesystem replicated in real time across nodes via **corosync**. Kernel 6.8/6.11, QEMU 9.x, LXC 6, ZFS 2.2. **PBS** is a separate Debian appliance providing deduplicated block-based backups. **Proxmox Datacenter Manager** reached v1.0 on 4 Dec 2025 after a long alpha/beta and federates multiple PVE clusters into one pane — but it is NOT a single control plane; each PVE cluster remains authoritative.

HYPERVISOR / HA / DRS / LIVE MIGRATION

KVM/QEMU for VMs, LXC for containers. HA Manager uses corosync quorum, watchdog fencing, and resource groups to restart VMs on surviving nodes. Online live migration for KVM (with local disks via `--with-local-disks`); LXC needs stop/start. Quorum needs 3+ nodes or a QDevice witness. **No proper load-based DRS**: the Cluster Resource Scheduler (CRS) only makes static placement decisions at HA recovery or manual trigger. This is a real gap vs. vSphere.

STORAGE

Unusually broad matrix: ZFS (local, snapshots, replication), LVM/LVM-thin, directory, NFS, CIFS, iSCSI (raw or LVM-over-iSCSI), **Ceph RBD** and **CephFS** (hyperconverged — Proxmox ships Ceph Reef/Squid packages and installs OSDs/MONs from the GUI), GlusterFS, ZFS-over-iSCSI. ZFS storage supports async replication (snapshot-based, ~1-min RPO) between nodes. The killer feature is Ceph integration: one cluster runs compute and storage, and the GUI manages OSDs/pools/CRUSH natively.

SDN

The Proxmox SDN plugin (production since 8.1) has **Simple / VLAN / QinQ / VXLAN / EVPN zones**. EVPN uses FRRouting with BGP-EVPN for L2/L3 overlays and symmetric IRB. IPAM (built-in, NetBox, or phpIPAM) + dnsmasq DHCP + PowerDNS plugin. The new nftables-based firewall has datacenter / node / VM / vNIC / VNet rule scopes with auto-generated IPsets. **This is not an NSX DFW.** Rules are per-host nftables against tap interfaces, no centralised DFW control plane, no tag/attribute-based policy, no east-west service insertion, no native LB or NAT gateway service. Micro-segmentation is achievable at vNIC level but must be modelled manually.

MULTI-TENANCY

Users from realms (PAM, PVE-local, LDAP, AD, OIDC). Groups, roles (including custom), and pools (grouping VMs/storage for shared ACLs) on a hierarchical path. **No real "project" abstraction with quotas.** Cannot natively cap CPU/RAM/disk per tenant. Shallower than OpenStack, Nutanix, vCD, or Aether.

BACKUP & DR

Proxmox Backup Server is excellent: source-side dedup with variable chunks, client-side AES-256-GCM, incremental-forever, verification, namespaces, prune/GC, sync jobs (offsite), and **LTO-5+ tape**. Tight GUI integration with PVE — schedule, browse, file-level restore, instant restore (live-mount a backup and migrate). DR: ZFS replication, PBS sync/tape, or third-party.

MIGRATION FROM VMWARE

PVE 8.2 (Apr 2024) added an **ESXi Import Wizard** built on libnbd/FUSE mount of the ESXi datastore. Imports one-by-one or in small batches (≤4 disks concurrent). Limitations: no vSAN, no encrypted VMDKs, snapshots slow imports dramatically, guests must be powered off or use risky live-import, no vCenter orchestration — you add ESXi hosts directly. Fleet moves still typically use third-party (Veeam, Vinchin, `virt-v2v`).

OPERATIONS / FLEET

Per-cluster web UI is mature, fast, and consistent. Multi-cluster was historically a gap — now filled by **PDM v1.0 (Dec 2025)** with remotes, consolidated dashboards, cross-cluster live migration, backup/update status, tested at >5000 remotes / >10,000 guests. Still, PDM is federated, not a true single control plane.

KUBERNETES

No official K8s integration. Community-maintained Cluster API providers — the IONOS Cloud `cluster-api-provider-proxmox` (CAPMOX) tracks CAPI 1.9 and K8s 1.31+. Talos, Rancher, k3s-on-VMs, Kamaji all work but you assemble them yourself.

LIFECYCLE

Debian rolling apt updates from three repos (`pve-no-subscription` , `pve-enterprise` , `pvetest`). Major version upgrades are classic Debian-style in-place. No A/B immutable image. Hosts are "regular Linux" — powerful and familiar, but you own patch hygiene.

LICENSING & PRICING

Software is GPL, free to run with **zero feature gate**. Commercial subscriptions are **support + access to the "enterprise" stable repo**, priced **per physical CPU socket per year** (customer-friendly — no per-core scaling). Current list (proxmox.com, 2025–2026):[27]

- PVE Community — **€120/socket/yr** (enterprise repo, forum only)
- PVE Basic — **€370/socket/yr** (3 tickets, 1 business day)
- PVE Standard — **€550/socket/yr** (10 tickets, 4-hour response)
- PVE Premium — **€1,100/socket/yr** (unlimited, 2-hour, remote SSH)

PBS is priced per server, not per socket: Community €560, Basic €1,120, Standard €2,240, Premium €4,480 per year.[28] PDM is currently free.

STRENGTHS

- Radically low cost. Zero software lock-in — it's just Debian + KVM + Ceph.
- Very mature product (first release 2008). Large active community.
- Deeply integrated hyperconverged Ceph managed from the GUI.
- Per-socket pricing on a public page — the most transparent in the document.
- Excellent PBS dedup + encryption + tape story.
- Genuinely useful SDN with EVPN.

WEAKNESSES

- No load-based DRS.
- No true distributed firewall comparable to NSX-T DFW, Nutanix Flow, or OVN ACLs.
- **Debian-based mutable hypervisor OS.** Proxmox runs on a standard Debian 12 installation with apt package management, SSH, and a writable root — same architectural model as Nutanix AHV and HPE VME, and a regression vs ESXi. Enterprise CISO departments that rejected Nutanix for this reason will reject Proxmox on the same grounds. For regulated shops that want a minimal-attack-surface appliance model, only Aether (IncusOS), OpenShift Virt (RHCOS), and Harvester (SLE Micro) apply.
- Multi-tenancy is shallow — no quotas, no project abstraction.
- PDM is brand-new (Dec 2025) and federated rather than a true single control plane.
- K8s integration is community-driven, not first-party.
- Enterprise RBAC integrations (SAML, fine-grained IdP claims) less polished than VMware/HPE/Red Hat.
- Support is EU-timezone ticket-based; 2-hour SLA only at €1,100/socket Premium.

3.6 HPE VM Essentials (ex-Morpheus KVM) COMMERCIAL KVM + MORPHEUS

HPE acquired Morpheus Data in Sep 2024. HPE VM Essentials became a standalone SKU on 24 Dec 2024, a stripped-down KVM-only entry point under the broader HPE Morpheus Enterprise cloud management platform, explicitly priced as a VMware alternative.

ARCHITECTURE

Three layers: (a) the **HVM hypervisor** — KVM on a hardened Ubuntu base with clustering, HA, live migration baked in; (b) the **VME Manager** (formerly Morpheus VM Manager) — appliance providing web UI, REST API, multi-tenant RBAC, automation, and migration/orchestration managing one or more HVM clusters *and* existing vSphere clusters side-by-side; (c) optional upgrade path to **HPE Morpheus Enterprise** adding multi-cloud broker, Terraform/IaC, self-service catalog, Morpheus Kubernetes Service (MKS), hybrid DBaaS. Manager sizings: Small (2 vCPU/12 GB → 1 cluster), Medium (4/16 → 3 clusters), Large (4/32 → 10 clusters).

HYPERVISOR / HA / DRS / LIVE MIGRATION

KVM-based HVM with QEMU. Native cluster HA, live migration, and **intelligent workload placement** — Morpheus's placement engine acts as a DRS-equivalent, weighted by CPU/RAM/affinity. Minimum HCI cluster is 3 hosts. GPU, PCI-E, NVMe, USB passthrough added in v8.0.5 (Apr 2025), closing a gap vs. vSphere. Host fencing via quorum/watchdog.

STORAGE

Storage-agnostic and leans on HPE's portfolio: **Alletra MP B10000** (FC or NVMe-oF), Alletra 9000, 3PAR, MSA, Nimble/Alletra 6000 via iSCSI/FC; **HPE SimpliVity** as HCI (dedup, compression, integrated backup); generic NFS/iSCSI/FC SAN. Ceph supported. **Snapshots, thin provisioning, and replication come from the array** (Peer Persistence, 3PAR Remote Copy) rather than from the hypervisor — a different model from Proxmox/ZFS/Ceph.

SDN

Historically the weakest leg of the Morpheus KVM story (Linux bridge + OVS + VLANs). HPE is now filling the gap with two partnerships announced at HPE Discover 2025:

- **Juniper-sourced zero-trust SDN** for the HVM hypervisor — east-west micro-segmentation at the hypervisor level with centralized policy.
- **HPE Aruba Networking CX 10000** — AMD Pensando DPU per port, orchestrated from the Morpheus UI via Aruba Fabric Composer, delivering a **distributed stateful firewall in switch silicon** across ESXi, HVM, and bare metal. Marketed as NSX-equivalent at 10× scale/performance and a fraction of the cost — with the catch that it requires buying CX 10000 switches.

LB/NAT comes from external appliances (F5, Avi) or cloud-broker plugins rather than a native VME service.

MULTI-TENANCY

Genuinely strong, inherited from Morpheus. First-class Tenants, Groups, Clouds, Users, fine-grained Roles with dozens of policies, SSO via SAML/OIDC, plus resource-visibility controls (public/private at tenant level) for sharing networks/images/clouds. Quotas, chargeback/showback, and self-service catalogs per tenant. Far deeper than Proxmox pools; comparable to vCD / OpenStack projects.

BACKUP & DR

VME ships with **integration hooks** rather than a native engine. The flagship partner is **Veeam**: on 24 Jun 2025 HPE and Veeam deepened the partnership; in March 2026 **Veeam GA'd agentless host-level image-based backup for VME** (CBT-accelerated incrementals, VSS app-consistent quiescing, cross-hypervisor restore ESXi↔HVM, DR to AWS/Azure, Kasten for K8s). **HPE Zerto** adds continuous data protection / near-zero RPO DR. SimpliVity bundles bring their own integrated backup. No PBS-equivalent in-box.

MIGRATION FROM VMWARE

A key selling point. The VME Manager directly registers vCenter, inventories existing VMware VMs, presents them alongside HVM VMs. **VM Essentials 8.0.8 (mid-2025)** shipped a native migration tool that defines migration plans mapping VMware VMs to HVM target clusters/networks/datastores and executes bulk cutovers (VirtIO/QEMU-guest-agent prep required, VMware Tools removed). Veeam Software Appliance is a second path. More mature than Proxmox's import wizard for fleet moves.

OPERATIONS / FLEET

Single pane for **both HVM and vSphere** from VME Manager. HPE Morpheus Enterprise extends this to AWS/Azure/GCP/Nutanix/OpenShift. Rich automation (tasks, workflows, library), cloud-broker semantics, self-service catalog inherited from Morpheus. Large Manager handles up to 10 HVM clusters; bigger fleets move to Morpheus Enterprise.

KUBERNETES

Morpheus Kubernetes Service (MKS) — CNCF-certified — on top of HVM VMs, with cluster autoscaling, lifecycle management, HPE CSI tying storage back to Alletra MP. Veeam Kasten for backup. A real K8s story, not a community bolt-on.

LIFECYCLE

HVM hosts run a hardened Ubuntu-based OS managed through Morpheus; VME Manager ships as an all-in-one ISO (since 8.0.8). Upgrades orchestrated from the Manager. Not immutable/A-B, but more appliance-feeling than bare Debian, and HPE holds patch cadence/QA under paid support.

LICENSING & PRICING

Commercial, closed-source. **Per CPU socket per year** — a deliberate swipe at Broadcom. Two meaningful SKUs: the stripped-down **VM Essentials tier at $600/socket/yr** (KVM hypervisor only, basic management), and the full **HPE Morpheus Enterprise tier at ~$2,500/socket/yr** (cloud management platform, multi-tenancy, self-service catalog, IaC, MKS Kubernetes, multi-cloud broker, automation engine).[29][30] **For the enterprise private-cloud buyer replacing VCF 9 — the target of this document — Morpheus Enterprise is the realistic SKU**, not VM Essentials. The $600 entry-level tier is for small VMware-refugee shops that only need basic KVM management; the feature depth an enterprise needs (true multi-tenancy, RBAC, self-service, full automation) lives only in Enterprise. The comparisons in §6 therefore use the Enterprise tier exclusively. 5-year E-LTUs exist in the HPE Store, and bundling into HPE Private Cloud Business Edition / GreenLake changes the math further.

STRENGTHS

- Aggressive per-socket pricing (~$600/socket/yr list).
- HPE global enterprise support — real advantage in large regulated customers.
- Morpheus maturity in multi-tenancy, RBAC, self-service, and IaC is best-in-class for a VMware alternative.
- **Dual-hypervisor management** — HVM + vSphere from the same console — a genuinely useful transition story.
- Strong K8s (MKS + HPE CSI) and DR (Veeam + Zerto) partnerships.
- Built-in DRS-equivalent intelligent placement.

WEAKNESSES

- **Hardened-Ubuntu HVM is NOT immutable.** HPE's HVM runs on a hardened Ubuntu base — "hardened" in the sense of minimised packages and locked-down defaults, but still a traditional mutable Linux with apt, SSH, and a writable root. For CISO-led procurements in regulated industries this is typically an immediate veto against the "ESXi replacement" pitch, because it *adds* attack surface. Contrast with IncusOS (Debian-13 mkosi A/B + systemd-sysupdate + signed UKI + TPM + LUKS + no SSH), RHCOS (rpm-ostree A/B), or SLE Micro (transactional).
- Very new as a standalone SKU (GA Dec 2024) — real-world footprint is small vs. vSphere or Proxmox.
- Closed source — no community path, no audit.
- **Hardware-adjacent lock-in**: strongest SDN story depends on HPE Aruba CX 10000 switches; richest storage story depends on HPE Alletra/SimpliVity/3PAR.
- Native SDN/DFW inside HVM still maturing.
- Upgrade path to "full" Morpheus features requires Morpheus Enterprise (~$2,500/socket).
- Backup is partner-driven — Veeam plug-in only GA'd March 2026.
- Smaller ecosystem than Proxmox's Debian/KVM community.

3.7 Red Hat OpenShift Virtualization

COMMERCIAL KUBEVIRT ON OCP

OCP + KubeVirt operator + CDI + HyperConverged operator + OVN-Kubernetes + ODF (Rook/Ceph). VMs run as QEMU/KVM pods on RHCOS worker nodes. Red Hat's aggressive 2024–2026 VMware-replacement pitch.

ARCHITECTURE

OpenShift Container Platform (OCP) + RHCOS (rpm-ostree immutable) + HyperConverged operator pulling in KubeVirt + CDI (disk image population) + SSP (templates / tunings) + Cluster Network Addons Operator + Hostpath Provisioner for local disks. Storage typically ODF (Rook/Ceph). CNI is OVN-Kubernetes. Each running VM is a QEMU/KVM process wrapped in a `virt-launcher` pod; `virt-handler` runs as DaemonSet. VM definitions are CRDs (`VirtualMachine` , `VirtualMachineInstance` , `VirtualMachineInstanceMigration`). Red Hat also ships **OpenShift Virtualization Engine (OVE)**, a leaner VM-only SKU that strips out developer console, Service Mesh, Pipelines, and container workload support.

HYPERVISOR / HA / DRS / LIVE MIGRATION

KVM/QEMU via libvirt. Live migration is a KubeVirt primitive; HA is Kubernetes rescheduling — if a node goes NotReady, VM pods are rescheduled (requires RWX storage and eviction strategy `LiveMigrate`). **No true VMware DRS equivalent.** Red Hat's answer is the Kube Descheduler Operator with the LongLifecycle profile plus, since OCP 4.17/4.18, a tech-preview `DevKubeVirtRelieveAndMigrate` profile that uses **PSI (Pressure Stall Information)** CPU metrics from Prometheus to evict VMs from over-utilised nodes (defaults evict at 10% above cluster mean, switches to PSI once cluster CPU passes 80%). Reactive, eviction-based, re-evaluated on interval — not continuous flow-level DRS.[31][32]

STORAGE

Strongly prefers **RWX block-mode** storage for online live migration (two virt-launcher pods must share the PVC during handover). In practice that means ODF (Ceph RBD in RWX block mode) or a certified third-party CSI (Pure, NetApp Trident, Dell PowerStore/PowerMax, Lightbits NVMe-over-TCP, IBM Spectrum Fusion). CDI handles image import/clone/upload; CSI VolumeSnapshot gives PITs; ODF provides replication (RBD mirroring, Metro/Regional DR via ACM's DR add-on). Trade-off: Ceph-backed VM storage is functional but needs more spindles/NVMe to match vSAN or PowerStore on small-block random I/O, so heavier VM shops increasingly pair OCP-V with NVMe/TCP arrays rather than ODF.[33]

SDN

Default CNI is OVN-Kubernetes. OCP 4.17 introduced **User Defined Networks (UDN)** as Tech Preview; **OCP 4.18 (early 2025) promoted UDN to GA**. UDN delivers custom L2/L3/localnet segments, primary or secondary, with stable lifetime VM IPs, persistent MACs across live migration, and BGP advertisement so a UDN prefix is reachable from the physical fabric. Micro-segmentation via NetworkPolicy, AdminNetworkPolicy / BaselineAdminNetworkPolicy, MultiNetworkPolicy for UDN. EgressFirewall, egress IP, egress service exist. MetalLB for external LBs. Caveat: powerful but k8s-shaped, not NSX-shaped — a traditional vSphere admin will have to relearn.[34][35]

MULTI-TENANCY

OpenShift Projects (namespaces with metadata), full k8s RBAC, ResourceQuota, LimitRange, NetworkPolicy. ACM adds cluster-set and placement-rule multi-cluster tenancy. OCP 4.18 added dedicated Virtualization roles in ACM Fleet Management for cross-cluster VM admin. Strong within the k8s model, weaker than vCD / vRA for hosters used to per-tenant portals.

BACKUP & DR

OADP (OpenShift API for Data Protection) — Red Hat's productised Velero — is the first-party answer. OADP includes a KubeVirt plugin that freezes guest filesystems via QEMU guest agent, orchestrates CSI VolumeSnapshots, and captures VM CRs. Restore handles original/alternate namespace/storageclass. Policy-based DR via ACM's DR add-on + ODF RBD mirroring. Third-party: Veeam Kasten K10, Trilio TVK, Commvault, Cohesity, IBM Storage Defender.

MIGRATION FROM VMWARE

Migration Toolkit for Virtualization (MTV), upstream Forklift — free Red Hat operator. MTV 2.7 (2025) added an integrated OpenShift console UI, warm migration with CBT, cold migration, plan/test/cutover workflow, VDDK transfer, vCenter inventory. MTV 2.8/2.9 improved concurrency, network mapping UX, and added Red Hat Virtualization / oVirt, OpenStack, and KubeVirt-to-KubeVirt sources. Bundled with OVE at no extra charge.[36]

OPERATIONS / FLEET

Single cluster: OpenShift console with a dedicated Virtualization perspective. Multi-cluster: **Red Hat Advanced Cluster Management (ACM) 2.14+** with a Virtualization Fleet dashboard (aggregate capacity, per-VM health, cross-cluster live migration, VM search and policy). GitOps via OpenShift GitOps (Argo CD) and Pipelines (Tekton) — a genuine differentiator versus vSphere automation.[37]

KUBERNETES

Trivial: OCP is Kubernetes. Pods and VMs run side-by-side on the same workers, share the same CNI/CSI/RBAC, and can talk on the same Services. A VM can be a Service backend or sit behind a Route/Ingress. `virtctl` provides console/VNC/SSH access.

LIFECYCLE

Hosts run **RHCOS, an rpm-ostree immutable image**. Upgrades are cluster-scoped via the Machine Config Operator: stage new image, cordon/drain, VMs live-migrate off, host reboots into new image (A/B ostree). One `oc adm upgrade` moves everything — K8s, OCP operators, KubeVirt, ODF. ~1 OCP minor every 4 months, 18 months full support per release, Extended Update Support (EUS) on even minors.

LICENSING & PRICING

- **OpenShift Platform Plus** (OCP + ACM + ACS + ODF Essentials + Quay) — per 2-core "core pair" or per bare-metal socket. Red Hat does not publish list prices; third-party-cited figures show Platform Plus / Container Platform at roughly **$13,766/year per 1–2 sockets (Premium)** — the figure Red Hat itself used in its post-Broadcom comparison. A core-pair Premium subscription is typically €2,000–€5,000/year in EMEA depending on tier and discount.[38]
- **OpenShift Virtualization Engine (OVE)** — VM-only SKU launched late 2024 explicitly for VMware migrations. Per socket-pair up to 128 cores. **MW03816 (Standard, 1y, 1–2 sockets / 128 cores) ≈ $1,903.99**; MW03815 (Premium) sits higher.[39][40][41]
- OCP-V is included at no extra cost in Platform Plus — many enterprises already own the right to run it.

STRENGTHS

- IBM/Red Hat financial muscle and 24/7 global support.
- First-class, free VMware migration tool (MTV/Forklift).
- OCP-V bundled into existing OCP contracts — many shops already own it.
- Integrated Ceph-backed storage via ODF.
- Full GitOps / CI/CD / service mesh story.
- ACM multi-cluster fleet view.
- Very active 2024–2026 roadmap (UDN GA, BGP, load-aware descheduler).
- Real momentum: Red Hat reported 178% YoY growth, with flagship wins such as Emirates NBD migrating ~9,000 VMs.[42]

WEAKNESSES

- OCP is a heavy platform for shops that "just want to run VMs." OVE helps but still carries the OCP substrate.
- Steep learning curve for vSphere admins.
- No native DRS — descheduler is coarse, eviction-based.
- ODF/Ceph VM performance lags purpose-built block SANs.
- No NSX-class DFW — NetworkPolicy + UDN + ANP cover most but not all cases.
- Full Platform Plus pricing is high at scale and quote-only.
- UDN is GA but still maturing.
- VMware's own benchmark comparison claims 1.5× VM density on vSphere 8 — implicitly confirmed by Red Hat's roadmap focus on CPU management.

3.8 SUSE Harvester / SUSE Virtualization

OPEN SOURCE KUBEVIRT ON RKE2

Open-source HCI distribution: SLE Micro 6.1 + RKE2 + KubeVirt + Longhorn + Multus + (optional) Kube-OVN. Managed from SUSE Rancher. Renamed "SUSE Virtualization" in 2025 along with Longhorn → "SUSE Storage".

ARCHITECTURE

Bare-metal appliance install. **SLE Micro 6.1** (immutable, transactional Linux managed by Elemental/rancherd) boots on each node and runs **RKE2**. On top: KubeVirt, Longhorn (distributed replicated block), Multus + Linux bridge and (as of v1.7/v1.8) **Kube-OVN** for VM networking, CDI, kube-vip for control-plane HA, Prometheus/Grafana, and the embedded Harvester UI. As of v1.7–v1.8 (late 2025), RKE2 ~v1.34, KubeVirt 1.6.x, Longhorn 1.10.x, embedded Rancher 2.13 for single-cluster management. Multi-cluster managed by pointing an external **SUSE Rancher** at Harvester as "infrastructure clusters"; Rancher then provisions RKE2/K3s **guest** clusters on top — giving a classic HCI + on-demand K8s story.

HYPERVISOR / HA / DRS / LIVE MIGRATION

KVM/QEMU via libvirt. Live migration native; Harvester exposes it as a UI button. HA is Kubernetes rescheduling with Longhorn replicas ≥ 2. **DRS: historically absent.** Harvester v1.7.0 introduced the **Virtual Machine Auto Balance add-on**, which deploys the upstream Kubernetes Descheduler with policies to evict/re-migrate VM pods off over-utilised nodes — the first built-in DRS-ish capability. Eviction-based and coarse, not real-time flow-aware. NVIDIA vGPU + MIG supported in v1.7+.[43]

STORAGE

Primary is **Longhorn** (now SUSE Storage): K8s-native distributed block from per-volume Go engine processes replicating synchronously across N nodes. Thin provisioning, snapshots, recurring snapshots, backups to S3/NFS targets, cross-cluster DR via backup-restore, DR volumes, and **engine V2 (SPDK/NVMe-oF)** with much better performance (GA-ish in Longhorn 1.7+). External storage via a Third-Party Storage add-on (CSI) — iSCSI/NFS/vendor — with feature caveats. The V1 data path was a known bottleneck under heavy write workloads; V2 closes most of the gap.

SDN

Until v1.5/1.6 networking was deliberately simple: Multus + a harvester-network-controller using Linux bridges over bonded uplinks for VLAN networks. v1.6/1.7 added **Kube-OVN** as an optional add-on giving VPCs, subnets, distributed gateways, security groups, EIPs, and a VPC NAT Gateway (GA in v1.8 for external connectivity). VM micro-segmentation via subnet ACLs and K8s NetworkPolicies. **Notable weakness: there is no NSX-class distributed firewall product.** The VPC/ACL story is newer, less mature, and lacks the tooling (flow tracing, IDS/IPS, groups-based policy at scale) that NSX admins expect. For many enterprise shops this is the biggest gap vs. vSphere+NSX.[44]

MULTI-TENANCY

Rancher Projects + namespaces + RBAC. A Project groups namespaces with user/group role bindings; Harvester exposes this as the VM-pool tenancy primitive. ResourceQuota, NetworkPolicy, and Kube-OVN VPCs add network isolation. Strong for teams-within-one-org; thinner than vCD/OpenStack for true service-provider billing portals.

BACKUP & DR

Built-in: Longhorn VM backups to an NFS or S3 backup target (scheduled or on-demand), VM snapshots, restore to same or new cluster, and the Rancher Backup operator for Rancher/Harvester cluster state. Velero 1.16+ with the Longhorn snapshot class works for Velero workflows (announced in Harvester KB, May 2025). Third-party: Kasten K10, Trilio, Storware vProtect.

MIGRATION FROM VMWARE

vm-import-controller (initial v1.2+, maturing through v1.6/1.7) is a Harvester add-on that registers a `VmwareSource` and a `VirtualMachineImport` CR per VM. It logs into vCenter, validates datacenter/storage/network mapping, powers off the VMware VM, exports via vCenter APIs, converts disks into Longhorn volumes, reconstructs the KubeVirt VM. Less polished than MTV/Forklift — no warm CBT, no concurrent-plan UI — but covers common cases and is actively maintained.

OPERATIONS / FLEET

Single cluster: embedded Harvester UI (also a tab in Rancher when registered). Fleet: **SUSE Rancher + Rancher Fleet** for GitOps-at-scale + Rancher Continuous Delivery. Monitoring via embedded Prometheus/Grafana. The operational win is one Rancher console across Harvester HCI + all downstream guest k8s clusters (on-prem RKE2 and cloud EKS/AKS/GKE).

KUBERNETES

Native. Harvester is literally a K8s cluster and is explicitly designed as a Rancher infrastructure cluster onto which you push guest clusters via `cluster-api-provider-harvester`. VMs + pods coexist on the same nodes. Common pattern: legacy VMs + RKE2 guests on the same hardware/network/storage.

LIFECYCLE

Each node runs SLE Micro 6.1 (transactional, btrfs snapshots, snapper-based rollback) under Elemental. Cluster-wide upgrades orchestrated by the Harvester upgrade controller: download new ISO image, apply to each node sequentially with live-migration drain, rollback on failure. Strictly one minor at a time. ~2 minor releases/year.

LICENSING & PRICING

Harvester itself is Apache 2 — you can run it free, community-supported. Commercial support is sold under **SUSE Rancher Prime**, unified in 2023/24 as a single pricing model across Rancher Prime, Longhorn, Harvester, NeuVector. In **2025 SUSE changed that model from per-node to per physical-core-pair / per 4 vCPU** — widely reported as a 4–9× price increase. Portainer analysis (Oct 2025) quotes:[45]

- Rancher Prime bare-metal 1–2 socket / up-to-64-core subscription: from ~$2,000/yr → **$7,300–$9,800/yr**
- **Suite Edition** (the tier that actually includes Harvester/SUSE Virtualization and Longhorn/SUSE Storage): **$9,800–$13,100/yr** for the same bare-metal footprint
- 16-core/32-vCPU VM: ~$2,000/yr → ~$19,200/yr Standard or ~$25,600/yr Priority

SUSE does not publish list prices on its own site — request-for-quote only. To get supported Harvester you need the Suite tier.

STRENGTHS

- Fully open-source; you can run it without a SUSE contract.
- Single integrated bare-metal appliance — no separate OS install, no K8s bring-up, no storage bolt-on.
- Rancher pane-of-glass across HCI and guest K8s is genuinely unique.
- Longhorn is pleasant to operate and very observable.
- Strong edge story (small footprints, single-node clusters supported).

WEAKNESSES

- **SDN/security story is the weakest of the mainstream VMware alternatives.** No NSX-class DFW; Kube-OVN VPC newer, rough edges.
- Smaller customer base and support footprint than OCP / Nutanix / VMware.
- Longhorn V1 performance limits for heavy enterprise writes (V2 better but newer).
- Young product (v1.0 mid-2022). Each release still brings architectural changes.
- Max ~200 pods/node gives a practical VM density ceiling lower than ESXi.
- **2025 SUSE per-vCPU licensing change** erased Harvester's "cheap vs OpenShift/VMware" positioning at Suite tier.
- VM import controller less mature than Red Hat MTV.
- Scale ceiling lower than OCP+ACM in practice.

3.9 OpenNebula 7.2 COMMERCIAL / OPEN-SOURCE KVM APACHE-2

Open-source IaaS platform founded in 2005 by CTO Ignacio M. Llorente and CSO Rubén S. Montero, now stewarded by OpenNebula Systems S.L. (Spain). Current GA is **7.2** (April 2026[48]). For production VM workloads it is KVM-only since the vCenter driver was discontinued in 7.0 (LXC system containers remain available as a non-VM workload type). Strong sovereign-cloud / AI Factory positioning with NVIDIA Blackwell, NVLink/NVSwitch, and Spectrum-X validation in 7.2.

ARCHITECTURE

Two server roles: **Front-end** (the OpenNebula daemon, scheduler, OneDRS, OneFlow, FireEdge UI, Sunstone UI, gRPC and REST APIs, MySQL/MariaDB metadata) and **Nodes** (KVM/QEMU hosts). The front-end manages nodes over SSH for control-plane and monitoring operations. Federation across multiple OpenNebula Zones gives shared identity, RBAC, and resource visibility across data centres. The current 7.2 platform-certified host operating systems for both front-end and nodes are **Red Hat Enterprise Linux 9 / 10, AlmaLinux 9 / 10, Ubuntu Server 22.04 / 24.04 LTS, Debian 12 / 13, SUSE Linux Enterprise 15, openSUSE 16**[47] — all mutable, apt/dnf-managed, SSH-accessed traditional Linux distributions. No A/B image, no Secure-Boot-measured-into-TPM-PCR boot chain.

HYPERVISOR / HA / DRS / LIVE MIGRATION

KVM/QEMU is the only supported production hypervisor — the legacy vCenter driver was **discontinued in 7.0** (per docs: "Support for VMware vCenter has been discontinued")[49]. LXC remains as an alternative for system containers. Live migration is supported; **7.2 added Storage Live Migration** between LVM and file-based datastores with zero downtime. HA is VM restart on host failure, driven by the OpenNebula scheduler. No CPU-lockstep Fault Tolerance.

OneDRS is OpenNebula's strongest feature on this axis and one of the most algorithmically sophisticated DRS implementations in the open-source private-cloud space. It uses an **integer linear programming (ILP) solver** over real-time and forecasted host/VM resource consumption, with two policies — *Packing* or *Load Balancing* — and three automation levels: *Manual*, *Partial*, *Full*.[50] Initial placement is also driven by OneDRS, with the same metrics.

STORAGE

Production choices: **Ceph RBD** (recommended; OpenNebula has a long Ceph track record and certified integration support is a paid add-on), iSCSI / iSER, NFS, LVM (block) on shared SAN, local file. Snapshots and clones are first-class; the 2025 line added **incremental Ceph backups**.[51] NetApp partnership announced with 7.0; certified array support extends through NetApp ONTAP, Pure Storage (Everpure), and others. Storage Live Migration LVM↔file landed in 7.2. DR is built around Ceph RBD mirroring across two sites with a documented failover/failback runbook.

SDN

OpenNebula's Virtual Networks support **802.1Q VLAN, VXLAN, OVSwitch, and bridged networking**, with optional BlueField-DPU validation in 7.2 for hardware-accelerated dataplane offload. **Security groups are implemented as host-side iptables / ebtables rules on the VM tap interface** — not an in-kernel distributed firewall in the NSX-DFW or OVN-ACL sense. Documented verbatim in 7.2: *"Security groups are not supported for Open vSwitch"*[52]; the NSX-V / NSX-T DFW integration that once existed only ran via the vCenter driver, which has been discontinued. No native VPC / Transit-Gateway analogue and **no native L4 / L7 load balancer** — load balancing is community `addon-eywa`, HAProxy in a service-template appliance, or a third-party appliance.

MULTI-TENANCY

Mature: **Zones → VDCs (Virtual Data Centers) → Groups → Users**, with **federation across multiple Zones** sharing user accounts, groups, and permissions. A VDC is a fully-isolated virtual infrastructure environment scoped to one Zone, with its own quotas, networks, and access control. Federation lets a service provider span multiple data centres or geographies under one identity plane while keeping each Zone an independent OpenNebula deployment. RBAC is fine-grained with custom roles and per-resource ACLs.

BACKUP & DR

Native VM backup engine with a Restic backend by default, full and incremental backups, and (since 2025) **incremental Ceph backups**.[51] The 7.0 release added a deeper Veeam integration as a first-party partnership, giving customers a tier-1 enterprise backup option. DR via Ceph RBD mirroring two-site replication with documented failover/failback runbooks.

MIGRATION FROM VMWARE

OneSwap is OpenNebula's first-party VMware-to-KVM migration CLI[53] — connects to vCenter, lists data centres / clusters / VMs, and converts selected VMs into OpenNebula instances along with their network and storage resources. Transfer pipeline uses `virt-v2v` and `qemu-img` by default; **VDDK is an optional transfer accelerator**. OVA imports supported. 7.x introduced a **delta-mode lower-downtime workflow for powered-on VMs**; 7.0 added hybrid mode for multi-disk transfers. OpenNebula claims ~96% automatic-conversion success. Not a CBT-managed warm-sync pipeline in the HCX / Migration Manager sense, and no NSX rule import.

OPERATIONS / FLEET UI

UI lineage is **Sunstone (legacy)** + **FireEdge (modern)**, with 6.10 finalising the migration of Sunstone functionality onto the FireEdge server[54]. The 7.x line continues to polish FireEdge as the primary day-2 UI. Operations breadth covers compute, storage, virtual networks, security groups, marketplaces, OneFlow service templates, OneKE clusters, OneDRS, federation, RBAC, multi-tenancy, and backup. **OneForm** is the lifecycle-automation primitive. A **gRPC API** in addition to REST gives the platform a higher-concurrency control plane for scale operators.

KUBERNETES

OneKE is OpenNebula's production Kubernetes solution — opinionated cluster templates running on OpenNebula VMs, lifecycle managed from the OpenNebula UI. A **KaaS** offering is on the roadmap. **SUSE Rancher Prime Embedded** is available as an optional integration add-on.

LIFECYCLE / OS

Hypervisor and front-end nodes run **traditional, mutable, general-purpose Linux distributions** — RHEL / AlmaLinux 9/10, Ubuntu 22.04/24.04, Debian 12/13, SLE 15, openSUSE 16[47] — managed with `apt` or `dnf`, with SSH on every node, writable root filesystems. OpenNebula Systems' enterprise offering includes meaningful operational mitigations: **OneForm** automated cluster provisioning, **Maintenance Packs** and LTS releases for a predictable enterprise cadence, and — for customers on Ubuntu Pro, RHEL with extended security, or SLE — kernel livepatching as a distro feature. The 7.2 line adds vTPM, AMD SEV-SNP, and Intel TDX confidential-compute primitives surfaced at the control plane, plus mandatory 2FA on Sunstone. **The structural gap vs immutable-appliance hypervisors remains, however**: no A/B atomic image-based updates, no signed-UKI measured boot into TPM PCRs, no LUKS-encrypted root by default, no no-SSH appliance posture. For enterprise CISO-led replacements of ESXi-as-appliance, the Ubuntu-Pro-hardening counter-pitch is typically rejected in favour of stacks that genuinely preserve the appliance model.

CONFIDENTIAL COMPUTE / SOVEREIGN CLOUD / AI

7.2 added control-plane surfacing of vTPM, AMD SEV-SNP and Intel TDX. Mandatory 2FA can be enforced globally on Sunstone. The 2025–2026 product strategy leans hard into **sovereign cloud** and **AI Factory** positioning, with NVIDIA Spectrum-X Ethernet validation, Grace Blackwell GPU validation, and 7.2 orchestration of NVLink and NVSwitch topologies via NVIDIA Fabric Manager.[55]

LICENSING & PRICING

Community open-source release is Apache-2 free; Enterprise Edition is a commercial subscription. The 2025 published model is **component-based**: front-end subscription cost + per-managed-node cost, with progressive discounts above 50 nodes and multi-year commitments[56]:

- **Standard** — €7,000 / yr front-end (≈$8,750) + €700 / yr per managed node (≈$875). 10-node packs €5,600. 9×5 EST/CET support, unlimited tickets.
- **Premium** — €11,000 / yr front-end (≈$13,750) + €1,100 / yr per managed node (≈$1,375). 10-node packs €8,800. 24×7 EST/CET support, unlimited tickets.
- **HA front-end add-on** — €3,500 / yr Standard or €5,500 / yr Premium.
- Optional integration add-ons (Ceph, NFV/Edge, AI Factory, OneKE / Kubernetes, SUSE Rancher Prime Embedded) priced per node.

For the 12,800-core reference cluster: ~$302k Standard + HA / ~$474k Premium + HA / ~$430k–$890k with 1–2 add-ons over 3 years before progressive discounts. After typical >50-node + multi-year discounts (15–25% off), ~$340k–$720k.

STRENGTHS

- **~18 years of production maturity** — public reference deployments at significant scale (Beeks Group: 20,000+ VMs + 3,000+ bare-metal servers migrated off VMware in 2024).[57]
- **OneDRS** — the most algorithmically sophisticated DRS in the open-source private-cloud space (ILP-solver, predictive, pack/balance, manual/partial/full).
- **Multi-zone federation** — independent OpenNebula deployments under shared identity across data centres or geographies.
- **Strong sovereign-cloud and AI Factory positioning** — NVIDIA Spectrum-X / Blackwell / NVLink/NVSwitch validation, vTPM + SEV-SNP + TDX, 7.2 GPU orchestration.
- **OneSwap migration tool** with delta-mode powered-on workflow.

- **Veeam first-party backup partnership** from 7.0; native Ceph incremental backups.
- **Materially cheaper than most commercial alternatives at the same scale.**
- **Fully open Apache-2 substrate**; single commercial vendor (OpenNebula Systems) but no kill-switch.

WEAKNESSES

- **Mutable hypervisor OS** — RHEL / AlmaLinux / Ubuntu / Debian / SLE / openSUSE, all with apt/dnf, SSH, writable root. Often a binary veto in regulated CISO procurements.
- **No in-kernel distributed firewall for KVM** — security groups are iptables-on-tap, "not supported for Open vSwitch" per docs. No NSX-DFW analogue; prior NSX integration was via the now-discontinued vCenter driver.
- **No native L4 / L7 load balancer** — addon-eywa, HAProxy templates, or third-party appliances only.
- **No NDR / L7 IDS / IPS** — third-party SPAN-fed appliances only, no native integration.
- **No NSX rule migration** — security policy must be re-authored.
- **Smaller global enterprise support footprint** than Red Hat, Broadcom, Nutanix, or Dell — OpenNebula Systems is headquartered in Spain with regional partners.
- No native dedup / tape (Veeam-dependent); no CPU-lockstep Fault Tolerance.

---

4. Master Feature Matrix

Legend: **Yes** native / first-party / strong **Partial** available but limited / via add-on **No** missing or weak *3rd* third-party-dominant

Capability	Aether	OpenNebula	VCF 9 (ref)	OpenStack	Nutanix	Proxmox	HPE VME	OCP-V
Type-1 hypervisor maturity	KVM (Incus)	KVM (libvirt)	ESXi — ref	KVM	AHV (KVM)	KVM	KVM (HVM)	KubeVirt
Cluster HA (auto VM restart)	Yes	Yes	Yes	Masakari	Yes	Yes	Yes	k8s resched
Continuous DRS / load rebalance	Yes	OneDRS (ILP)	Yes — ref	Watcher only	ADS reactive	No	Intelligent placement	Descheduler/PSI
Live VM migration	Yes	Yes + Storage Live Mig	vMotion	Yes	Yes	Yes	Yes	Yes
Fault Tolerance (CPU lockstep)	No	No	Yes	No	No	No	No	No
Confidential compute (SEV-SNP/TDX)	Incus-level	7.2 vTPM+SEV+TDX	Yes	Yes (driver)	Partial	Kernel-level	Kernel-level	KubeVirt
Native distributed HCI storage	Ceph/lvmcl	Ceph RBD	vSAN ESA	Ceph	ADSF	Ceph	SimpliVity/SAN	ODF/Ceph
Instant snapshots (no delta chain)	Ceph/ZFS	Ceph	ESA B+tree	Ceph/driver	ADSF COW	ZFS/Ceph	Array	CSI
Sync (0-RPO) replication	Ceph mirror	Ceph mirror	vSAN stretch	Ceph mirror	Metro	Ceph stretch	Array	ODF Metro
External array CSI breadth	via Incus	Ceph/NetApp/Pure	vVols deprecated	~80 drivers	Volumes/ext	All	HPE-centric	Trident/Pure/Dell
Overlay SDN (Geneve/VXLAN)	OVN	VXLAN/OVS	NSX Geneve	OVN	FVN/OVN	SDN plugin	OVS + Juniper	OVN-K + UDN
Distributed firewall (in-kernel / per-host)	OVN ACLs	iptables-on-tap	NSX DFW — ref	OVN ACLs	Flow	nftables only	CX10k / Juniper	NetworkPolicy/UDN
Multi-tenant VPCs / VRFs	Per-tenant OVN	VLAN/VXLAN/VDC	NSX VPC	Per-project	FVN VPC	Manual zones	CX10k	UDN
L4/L7 LB (native)	OVN+HAProxy A/A	addon-eywa/BYO	Avi add-on	Octavia	BYO	BYO	BYO	MetalLB
NDR / L7 IDS-IPS (native)	Aether NDR (Zeek+Suricata+RITA)	None	NSX IDS / IPS	None	None	None	None	None
Project / tenant construct	Aether tenant	Zones/VDCs/Groups	VCFA Project	Keystone	PC Projects	Pools only	Morpheus tenants	OCP Projects
Per-tenant quotas & chargeback	Yes (native billing)	Showback/WHMCS ext.	VCFA	Yes	NCM Cost Gov	No	Morpheus	ResourceQuota
Native backup engine	S3 upload	Native + Ceph incr.	None	Cinder-backup	Mine/HYCU	PBS	Veeam GA'd	OADP/Velero
Dedup / tape / air-gap	S3-dependent	*Veeam*	*3rd (Veeam/Rubrik)*	*3rd (Trilio)*	HYCU/Mine	PBS + tape	*Veeam*	*Kasten/Veeam*
DR orchestration (runbooks)	Script	Ceph mirror blueprint	Live Recovery	Heat/3rd	Recovery Plans	Manual	Zerto	ACM DR

Capability	Aether	OpenNebula	VCF 9 (ref)	OpenStack	Nutanix	Proxmox	HPE VME	OCP-V
First-party VMware migrator	Migration Manager	OneSwap (delta+VDDK)	HCX (intra)	*Coriolis/3rd*	Nutanix Move	ESXi import	VME 8.0.8	MTV/Forklift
CBT-accelerated warm migration	VDDK + CBT	Delta mode (not CBT)	HCX	Coriolis	Yes	No	Yes	Yes
NSX rule migration (E-W + N-S)	Yes — unique	No	*N/A (it is VMware)*	No	No	No	No	No
Single-pane multi-cluster UI	Aether	Per-Zone + Federation	VCF Ops	3rd-party	Prism Central	PDM 1.0	Morpheus	ACM
Functional breadth in one UI (ops+SDN+LB+backup+K8s+migration+LCM+NDR)	14 domains	~10 (no LB/DFW/NDR/NSX-import)	4 UIs split	Fragmented	Prism+NCM	PBS separate	Morpheus CMP	Console+ACM
Immutable / A-B host OS	IncusOS	apt/dnf mutable	ESXi image	RHEL/Ubuntu	AHV (updates)	Debian apt	Ubuntu	RHCOS
One-click firmware & LCM	OC	OneForm provisioning	vLCM + Fleet	Distro-dep	Nutanix LCM	Manual	VME Manager	MCO
First-party K8s distribution	Shepherd (K8s-style)	OneKE + KaaS (roadmap)	VKS	Magnum	NKP	Community CAPI	MKS	OCP (native)
Vendor-independent substrate	Apache 2	Apache 2	Broadcom	Apache 2	HCL-locked	GPL	Closed	OKD upstream
Global 24×7 enterprise support	Dell (exclusive reseller)	OpenNebula Systems	Broadcom	Red Hat/Canonical	Nutanix	Proxmox GmbH	HPE	Red Hat/IBM
Published list pricing	$80/core	€7k+€700/node (Std) / €11k+€1.1k/node (Prem)	$350/core	Distro-dep	Quote only	€120–1100/socket	$600/socket	OVE SKUs

---

5. VCF 9 Feature Parity Matrix

This matrix asks the inverse question: *for each VCF 9 capability, how does each alternative deliver it (if at all)?* This is the list a migration architect actually uses.

VCF 9 feature	Aether	OpenNebula	OpenStack	Nutanix	Proxmox	HPE VME	OCP-V	Harvester
vSphere HA	Incus HA	Scheduler restart	Masakari	AHV HA	ha-manager	HVM HA	k8s resched	k8s resched
vSphere DRS (continuous)	Built-in	OneDRS (ILP)	Watcher	ADS (reactive)	None	Morpheus IP	Descheduler/PSI	Auto Balance
vMotion / Storage vMotion	Yes	Yes + Storage Live Mig	Yes	Yes	Yes	Yes	Yes	Yes
vSphere Fault Tolerance	No	No	No	No	No	No	No	No
vSAN ESA (HCI block)	Ceph/lvmcl	Ceph RBD	Ceph RBD	ADSF	Ceph	SimpliVity	ODF	Longhorn
vSAN stretched cluster	Ceph stretch	Ceph stretch	Ceph stretch	Metro	Ceph stretch	Peer Persist.	ODF Metro	Longhorn DR
ESA global dedup	Ceph compress	Ceph compress + incr.	Ceph compress	Dedup+EC-X	Ceph compress	SimpliVity dedup	ODF compress	No
NSX Distributed Firewall	OVN ACLs	iptables-on-tap	OVN ACLs	Flow	nftables	CX10k/Juniper	NetworkPolicy+UDN	Kube-OVN ACLs
NSX IDS / IPS (L7 detection)	Aether NDR	None	None	None	None	None	None	None
NSX VPC / Transit Gateway	OVN VPC	VLAN/VXLAN per VDC	Neutron	FVN VPC	No	CX10k	UDN+BGP	Kube-OVN VPC
NSX Federation	CSP global rules	Federation = identity	OVN-IC	Flow multi-site	No	Juniper	ACM	No
NSX-V / NSX-T rule migration (E-W + N-S)	Yes — E-W & N-S (unique)	No	No	No	No	No	No	No
Avi LB (enterprise ADC)	OVN+HAProxy A/A	addon-eywa / BYO	Octavia	BYO	BYO	BYO F5/Avi	MetalLB	kube-vip
VCF Automation (self-service IaaS)	Aether + Shepherd	FireEdge + OneFlow	Horizon + Heat	NCM Self-Service	No	Morpheus	OCP console	Rancher
VCF Operations (fleet mgmt)	Aether + OC	Per-Zone + Federation	3rd-party	Prism Central	PDM 1.0	Morpheus	ACM	Rancher
Fleet image-based LCM (vLCM)	OC + IncusOS	OneForm	Distro-dep	Nutanix LCM	apt	VME Manager	MCO	Upgrade ctl
App / kernel live-patch (no full reboot)	IncusOS sysext-swap for apps (OS = A/B reboot)	Kernel livepatch (distro-dep)	Kernel livepatch	No	Kernel livepatch	Kernel livepatch	MCO reboot	MCO reboot
HCX (bulk cross-cluster move)	Migration Mgr + NSX rules	OneSwap (delta+VDDK)	*Coriolis*	Move	Import wiz	VME 8.0.8	MTV	vm-import-ctl
vSphere Kubernetes Service	Shepherd (K8s-style)	OneKE	Magnum	NKP	CAPI (comm)	MKS	Native	Native
Data Services Mgr (DBaaS)	No	OneFlow templates	Trove	NDB	No	Morpheus DB	Operators	Operators
Private AI Foundation	No	Blackwell + Spectrum-X (7.2)	Kolla+GPU	GPT-in-a-Box	No	MKS+GPU	OpenShift AI	GPU + Rancher
VMware Live Recovery (SRM)	Scripted	Ceph mirror + Veeam	Heat/3rd	Nutanix DR	Scripted	Zerto	ACM DR	Backup/DR
vVols (ext-array policy)	via Incus	NetApp/Pure drivers	Cinder driver	Volumes	iSCSI/FC	Alletra	Trident	CSI

6. 3-Year Cost Model

Illustrative scenario. 100 physical nodes, 2 × 64-core AMD EPYC = 128 cores/node → **12,800 physical cores**, 200 sockets, 4 TB RAM/node, ~5 PB usable storage. 3-year commitment. **List prices, pre-discount, excluding hardware, power, floor space, integration services, and third-party backup.** Real enterprise discounts of 20–50% are typical for most vendors; Proxmox and Aether are closer to street prices because their lists are already low.

Stack	Pricing basis	3-yr license+support (list)	Notes
HPE Morpheus Enterprise	$2,500 / socket / yr × 200 × 3	**$1,500,000**	Per socket (not core). This is the enterprise-tier SKU (full Morpheus CMP, multi-tenancy, self-service, MKS K8s) — the realistic SKU for a VCF-replacement buyer, not the stripped-down $600/socket VM Essentials entry tier. Hardware lock-in implicit — most customers will also be buying HPE Alletra/CX 10000.[29]
OpenShift Virtualization Engine	$1,904 Std × ~100 units (128c/unit) × 3	**~$570,000**	OVE Standard tier, one unit per 128-core / 1-2-socket pair. Premium roughly 2×. OCP-V bundled in existing OCP Platform Plus is $0 marginal.[39][40]
Proxmox VE Premium + PBS	€1,100 / socket / yr × 200 × 3 + PBS	**~€790,000 ≈ $850,000**	Proxmox's highest enterprise tier — 2-hour response, unlimited tickets. Lower tiers exist (Standard €550, Basic €370, Community €120) but enterprise VCF-replacement buyers go straight to Premium for the SLA.[27][28]
Canonical Charmed OpenStack (Advanced)	~$1,500 / node / yr × 100 × 3	**~$450,000**	Excludes the $75k–$150k professional services onboarding that Canonical typically requires for greenfield.[21]
Canonical Charmed OpenStack (Fully Managed)	~$4,275 / node / yr × 100 × 3	**~$1,280,000**	Canonical operates the cloud for you.[21]
OpenNebula Standard + HA front-end	($8,750 + $875 × 100 + $4,375 HA) × 3	**~$302,000**	Front-end €7k/yr + €700/node/yr + €3,500/yr HA add-on. 9×5 EST/CET support, unlimited tickets. Pre-discount; >50-node + multi-year discounts apply.[56]
OpenNebula Premium + HA front-end	($13,750 + $1,375 × 100 + $6,875 HA) × 3	**~$474,000**	Front-end €11k/yr + €1,100/node/yr + €5,500/yr HA add-on. 24×7 EST/CET support. Pre-discount.[56]
OpenNebula Premium + 1–2 add-ons (realistic band)	Base + HA + (Ceph and/or K8s) add-on(s) × 3	**~$430,000 – $890,000**	Each integration support add-on (Ceph, OneKE/K8s, NFV/Edge, AI Factory, Rancher Prime Embedded) adds ~$131k Standard or ~$206k Premium over 3y at 100 nodes. After typical >50-node + multi-year discounts (15–25% off): ~$340k–$720k.[56]
Aether (Dell-resold)	**$80 / core / yr × 12,800 × 3**	**$3,072,000**	Per physical core, includes Dell support. Sold exclusively through Dell Technologies. Incus/IncusOS/OC/MM are free open source standalone.
SUSE Rancher Prime Suite	~$11,500 / node / yr × 100 × 3	**~$3,450,000**	Suite tier required for Harvester/SUSE Virtualization + Longhorn support. 2025 pricing change. Request-for-quote on SUSE site.[45]
Nutanix NCI-Ultimate (all features)	$575 / core / yr × 12,800 × 3	**$22,080,000**	Verified list for SKU `NT-NCP-NCI-ULT-PR` (NCI-Ultimate + Production Support): $575/core/yr, $1,724/core for 3 years, $2,873/core for 5 years. Per physical core, license must cover every populated core in the cluster. NCI-Pro and Starter tiers are lower but strip features; Mission Critical and Federal MC support tiers are higher. **More expensive than VCF 9 base list** for the same footprint.[25][26][46]
Red Hat OpenShift Platform Plus (Premium)	~$13,766 / socket-pair / yr × 100 × 3 (est.)	**~$4,130,000 (est.)**	Quote-only. Core-pair Premium €2,000–€5,000/yr reported in EMEA. OCP-V is bundled at no extra cost.[38]
Red Hat OpenStack Services on OpenShift	RHOSP $3k-5k/socket-pair + OCP entitlement	**~$4,500,000+ (est., quote-only)**	Full quote-only. Requires OpenShift subscriptions for the control plane in addition to OpenStack entitlements.[20]
VMware Cloud Foundation 9 (base)	**$350 / core / yr × 12,800 × 3**	**$13,440,000**	List, pre-discount. At a more realistic 30% enterprise discount: ~$9.4M. AT&T filed evidence of a 1,050% hike offer.[1][2][7]
VCF 9 with VMware Live Recovery (DR)	**~$1,200 / effective core / yr × 12,800 × 3**	**~$46,080,000**	When the customer wants VMware Live Recovery (SRM + Live Site Recovery) they pay VCF on both the *active* and *standby* sites (2× core count) plus a per-VM Live Recovery subscription. Effective per-core economics rise to approximately $1,200/core. Almost every enterprise VCF deployment needs DR — this is the realistic total, not the base list.

Observations from the cost model:

- At 12,800 cores / 3 years the spread is enormous: the cheapest tiers (**OpenNebula Standard at ~$302k**, Canonical OpenStack Advanced ~$450k, OpenNebula Premium ~$474k, OCP-V Engine ~$570k) sit at the bottom; the most expensive — **VCF 9 with Live Recovery at ~$46M** and **Nutanix NCI-Ultimate at ~$22M** — sit **40×–150× higher**. This is why the market is moving.
- **Aether at $3.07M** for 12,800 cores is materially more expensive than the per-socket and free-licence-plus-cheap-support stacks (HPE Morpheus Enterprise $1.5M, Proxmox $850k, OpenNebula $300k–$890k, OCP-V $570k) but well below SUSE Harvester Suite ($3.45M), Red Hat OpenShift Platform Plus (~$4.13M), full RHOSO (~$4.5M), Nutanix Ultimate ($22.08M), VCF 9 base ($13.44M), and VCF 9 with Live Recovery ($46M).
- **Per-core vs per-socket matters enormously** at modern CPU densities. Proxmox Premium and HPE Morpheus Enterprise come out ahead on absolute list because they price per socket; OpenNebula prices per node which has the same insulating effect against runaway core counts; everyone else prices per core, and 12,800 cores at $350/core or $575/core adds up very quickly.
- **Nutanix Ultimate at $575/core list is more expensive than VCF 9 base list** for the same footprint — this is the direct list-to-list comparison, and it surprises customers who assume Nutanix is the "cheap alternative" to VMware. Nutanix's real-world discount regime is aggressive, but so is Broadcom's when pressed.
- **VCF 9 with Live Recovery at ~$46M is 3.4× the VCF 9 base list** — and this is the realistic number for any enterprise that actually does DR, which is essentially all of them.
- **VVF (the lower $135/core tier) is being withdrawn from parts of EMEA as of December 2025** and is therefore not a viable comparison in this cost model. EMEA customers that formerly ran on VVF are being told by Broadcom to upgrade to full VCF — a ~10× list price jump.
- Real negotiated discounts compress this ranking. Assume 20–50% off for every commercial SKU above; Proxmox and Aether are already near-street.
- The list does not include third-party backup (Veeam, Commvault, PBS-equivalents, Trilio, Kasten), professional services, or hardware refresh — all of which can exceed software licensing at 3-year TCO.

---

7. Weighted Scoring Matrix

Per-category 1–10 scores, weighted per §2, producing the final Replacement Fitness number for each candidate stack. Every cell is a judgement call grounded in the deep dives and the cited sources; re-score with your own priorities by editing the weights. **VCF 9 is intentionally excluded from this matrix** — it is the reference baseline that the customer is leaving, not a candidate replacement, and including it as a ranked stack would be a category error. VCF 9 remains visible in the §4 Master Feature Matrix and §5 Parity Matrix as the comparator for what is being replaced.

Category (weight)	Aether	OCP-V	OpenStack	OpenNebula	Nutanix	HPE VME	Harvester	Proxmox
Hypervisor / HA / DRS (1.5×)	8	7	6	8	8	8	6	6
Storage (1.5×)	7	7	9	9	10	7	6	8
SDN & DFW (2.0×)	9	7	9	5	8	6	5	5
Migration from VMware (1.3×)	9	9	5	8	9	8	6	6
Multi-tenancy (1.2×)	9	8	10	9	7	9	6	4
Cost (2.0×, 10 = cheapest)	7	6	7	9	2	9	5	10
Backup & DR (1×)	8	7	5	8	9	6	7	9
Lifecycle / OS (3.0×) — single heaviest axis · CISO veto	9	9	5	4	8	7	8	5
Operations / fleet UI (1.5×) — incl. functional breadth	9	8	5	8	10	9	8	6
Kubernetes integration (1×)	7	10	7	7	8	7	10	4
Vendor independence (2.5×)	9	6	10	9	3	4	8	10
Ecosystem / support (2.0×)	9	9	9	7	9	7	6	6
Weighted total / 205	**173.0**	**157.3**	**150.5**	**150.2**	**148.6**	**145.2**	**138.0**	**137.6**
Replacement Fitness (0–100)	**84**	**77**	**73**	**73**	**72**	**71**	**67**	**67**

Under the rubric defined in §2 — Lifecycle/OS weighted at 3.0× (the single heaviest axis), Vendor independence 2.5×, Cost 2.0×, SDN/DFW 2.0×, Ecosystem 2.0×, Operations/UI 1.5× — **Aether leads at 84**, followed by **OCP-V at 77** (the only other stack with a true immutable A/B hypervisor OS, RHCOS), **OpenStack and OpenNebula tied at 73** (mutable Linux hurts them on the 3.0× Lifecycle axis but openness/cost partially offset), **Nutanix at 72** (best day-2 ops in the market but CentOS-based mutable AHV and per-core pricing hurt), **HPE Morpheus Enterprise at 71**, and **SUSE Harvester and Proxmox tied at 67**. The shape of the result is direct: under a rubric that rewards an immutable appliance-OS posture, the open stacks with immutable substrates (Aether's IncusOS, OCP-V's RHCOS, Harvester's SLE Micro) gain ground, and the mutable-Linux stacks (Proxmox, Nutanix AHV, HPE VME, OpenStack, OpenNebula) lose it. Aether's top position reflects the rare combination of strong vendor independence (9/10 — fully open Apache-2 substrate even though the commercial product is Dell-channelled), broad feature coverage across SDN, DRS, migration (VMs + NSX E-W + N-S rules), NDR (Layer-7 IDS / IPS — unique in this comparison), and immutable IncusOS lifecycle.

---

8. Sources

All URLs were verified during research in April 2026. Prices and SKUs change — treat every figure as a snapshot, not a guarantee. Where a figure is "quote-only" it is noted as such above.

1. Redress Compliance — VCF cost per core
2. Redress Compliance — VCF vs VVF licensing
3. fwdsystems — VMware Apr 2025 72-core minimum
4. Software Pricing Guide — VMware pricing after Broadcom 800–1,500% shock
5. The Register — Computershare migrated 24,000 VMs
6. The Register — Beeks Group to OpenNebula
7. The Register — AT&T 1,050% filing
8. Broadcom TechDocs — VCF 9.0 Release Notes
9. Vladan — vSphere 9 no Standard/Enterprise Plus
10. Broadcom TechDocs — VCF Fleet Management
11. StarWind — vSAN 9 ESA global dedup, snapshots, vVols deprecation
12. sdn-warrior — NSX VPC + Transit Gateway in VCF 9
13. VMware Blog — VMware Virtual Private Cloud in VCF 9
14. VMware Blog — VCFA Org → Project → Namespace hierarchy

15. VMware Blog — VCF 9 Enhanced vSphere Replication + Live Recovery
16. William Lam — VKS management in VCFA 9.0.1
17. William Lam — ESX Live Patch in VCF 9
18. VMware — vSphere Foundation product page
19. TechCrunch — OpenStack ready for VMware refugees
20. Red Hat — RHOSO 18 subscription planning
21. TrustRadius — Canonical Charmed OpenStack pricing
22. Rackspace — OpenStack Flex launch
23. The Register — GEICO drops VMware for OpenStack
24. Constellation Research — Nutanix vs VMware (Move migration numbers)
25. StoragePricing.Org — Nutanix pricing datapoints
26. Vendr — Nutanix marketplace pricing
27. Proxmox — Proxmox VE pricing
28. Proxmox — Proxmox Backup Server pricing
29. The Next Platform — HPE VM Essentials pricing and positioning
30. HPE — HPE Morpheus VM Essentials QuickSpecs
31. Red Hat developer — Load-aware rebalancing with OCP-V
32. Red Hat developer — Dynamic CPU rebalancing with descheduler
33. Lightbits — OCP-V storage FAQ and performance
34. Red Hat blog — User Defined Networks in OCP-V
35. Red Hat blog — What's new in OpenShift 4.18 (UDN GA)
36. Red Hat — MTV 2.9 documentation
37. Red Hat developer — Monitor OCP-V at scale with ACM
38. Medium / Mr.PlanB — True price of OpenShift in 2025
39. TechTarget — OCP-V roadmap and $1,903.99 OVE figure
40. CDW — OVE Standard SKU MW03816
41. CDW — OVE Premium SKU MW03815
42. TechTarget — Red Hat 178% YoY growth, Emirates NBD
43. SUSE — SUSE Virtualization / Harvester v1.7.0
44. Harvester docs — Kube-OVN VPC
45. Portainer — SUSE Rancher 2025 price hike
46. CDW / Insight reseller SKU listings — Nutanix NCI-Ultimate Production Support per-core list ($575/yr, $1,724/3y, $2,873/5y per SKU NT-NCP-NCI-ULT-PR); Mission Critical and Federal MC support tiers higher
47. OpenNebula docs — 7.2 Platform Notes (certified host OS)
48. OpenNebula — Release Schedule (7.0 Jul 2025, 7.0.1 Oct 2025, 7.0.2 Dec 2025, 7.2 Apr 2026)
49. OpenNebula docs — 7.0 Compatibility Guide (vCenter Drivers — Support has been discontinued)
50. OpenNebula docs — OneDRS Distributed Resource Scheduler (ILP solver)
51. OpenNebula — Incremental Ceph Backups
52. OpenNebula docs — 7.2 Security Groups (note: "not supported for Open vSwitch")
53. OpenNebula docs — OneSwap migration tool (delta mode + VDDK)
54. OpenNebula — 6.10 "Bubble" Release (FireEdge GA, vCenter driver demoted, OneSwap)
55. OpenNebula — 7.2 release notes (vTPM, SEV-SNP, TDX, Blackwell, NVLink/NVSwitch, Spectrum-X)
56. OpenNebula — Subscription Plans (current Standard / Premium tiers, per-node + per-front-end model, 10-node packs, HA add-on, >50-node and multi-year progressive discounts)
57. The Register — Beeks Group migrates ~20,000 VMs + 3,000 bare-metal off VMware to OpenNebula (Dec 2024)
58. IncusOS upstream — github.com/lxc/incus-os — Debian-13 base, `mkosi`-built immutable appliance image, A/B updates via `systemd-sysupdate`, signed UKI with TPM PCR-measured boot, LUKS, ZFS, unix-socket-only management API

Private Cloud Stack Comparison — 2026 Edition (9 stacks). Prepared for private-cloud architects evaluating post-Broadcom alternatives. Twelve-axis weighted rubric with **Lifecycle / OS at 3.0× — the single heaviest axis, encoding the enterprise-CISO immutable-hypervisor-OS binary veto**; SDN/DFW 2.0×, Vendor-independence 2.5×, Cost 2.0×, Ops/UI 1.5×, Ecosystem 2.0×, others lower. VCF 9 is the reference baseline being replaced; it appears in §4 Master Feature Matrix, §5 VCF 9 Parity Matrix, and §6 Cost Model as the comparator for what is being left behind, but is excluded from the §7 Weighted Scoring Matrix. Neutral rubric, transparent weights, cited sources.