

Compare on the two filters that decide it.

A deeper technical evaluation of every credible VMware alternative against the two filters a regulated migration turns on — an immutable, attested host and first-party NSX rule migration — and the FuturFusion mechanism that passes both.

01 Executive summary

Every enterprise leaving Broadcom VMware faces the same shortlist — Nutanix, Red Hat, SUSE, Proxmox, HPE, XCP-ng, OpenStack, FuturFusion — and the temptation to compare on feature checklists. For a regulated estate two technical filters decide the outcome before any grid: whether the destination host is **immutable and attested**, and whether the platform **migrates the NSX firewall-rule corpus** first-party. This document is the technical basis for both filters and how each alternative measures against them, then the FuturFusion mechanism in full.

02 The two filters, technically

Filter 1 — immutable, attested host. The destination host OS must be a read-only signed image with atomic A/B updates, measured boot into TPM PCRs, and no interactive shell — so a compromised, un-patchable legacy guest cannot reach or persist on the host. A mutable, SSH-reachable Linux or Xen dom0 fails this outright; a merely “hardened” host without image-immutability and attestation is a partial pass at best.

Filter 2 — first-party NSX rule migration. NSX-V and NSX-T east-west DFW and north-south gateway rules must carry across with the VMs, into an enforced policy model, with tag/group membership resolved so posture follows identity. Moving VMs without the rule corpus leaves a security gap at cutover and a multi-month manual rebuild. HCX-style IP/MAC preservation is the necessary third property, since re-addressing breaks licence keys and address-based rules.

03 Where each alternative falls

Immutable but no rule migration — Red Hat, SUSE. RHCOS and SLE Micro clear Filter 1, but MTV and Harvester move VMs and leave the NSX corpus to be rebuilt by hand; both carry KubeVirt overhead for pure-VM workloads. **Neither filter — Proxmox, HPE, XCP-ng, OpenStack.** Mutable, SSH-capable control hosts fail immutability; none provides first-party NSX rule migration. **Closest — Nutanix.** AHV is hardened but not image-immutable or attested; Move imports security groups only, not the full DFW and edge corpus; multi-tenancy and DR sit behind additional tiers.

04 Component architecture

The five components and their responsibility boundaries:

Component	Role	Licence
FF Aether	Multi-tenant control plane — UI, RBAC, billing, DFW management, migration orchestration, NDR control, Shepherd, tenant portal.	Commercial
FF Incus	The hypervisor — QEMU/KVM VMs + OCI containers + LXC system containers; manages OVN networks and storage pools; dqlite-replicated cluster state.	Apache-2.0
FF HypervisorOS	The appliance OS — mkosi-built immutable Debian-13, A/B atomic updates, no SSH.	Apache-2.0
FF Operations Center	Multi-cluster fleet lifecycle — registers HypervisorOS nodes, drives A/B updates, provisions clusters via Terraform.	Apache-2.0
FF Migration Manager	VMware-to-Incus VM transfer — VDDK + CBT warm-incremental disk copy, worker-VM model.	Apache-2.0

Boundary discipline. FF Aether speaks to FF Incus, Operations Center, and Migration Manager exclusively through their published REST APIs — no private side-channels, no patched-upstream forks, no Aether-specific Incus build. A

customer who decides to keep operating the four open-source components without FF Aether can do so; nothing deployed becomes unreachable.

05 Twenty-one domains across six planes

FF Aether covers twenty-one functional domains, organised across six architecture planes.

Plane	Domains
Compute	Compute · Cluster Operations · Instance Metrics · Health · Capacity Planning
Storage	Storage · Backup
Network	SDN · IPAM · Load Balancer
Security	Distributed Firewall · NDR · Certificates · RBAC
Tenancy	Billing · Self-service portal · Tenants
Automation	Shepherd · Migration Manager · Operations Center · Plugins

Every domain is exposed both in the FF Aether UI and as an authenticated REST API (§16). The deep dives that follow cover the domains an evaluator interrogates first.

06 Platform architecture and the management/data-plane split

FF Aether is the management plane: it owns authoritative state for policy (ACLs, firewall rules), load-balancer configuration, tenant definitions, RBAC, billing, migration orchestration, and NDR control, and it continuously drives the data plane into conformance. It is a single Go binary with a PostgreSQL backend, reached at `https://<aether>:8443`, and it holds no packet path — an FF Aether outage is a management-plane event only; running VMs keep running, OVN keeps forwarding, firewall rules keep enforcing, and NDR keeps detecting.

The data plane is a fleet of FF Incus clusters on FF HypervisorOS nodes. FF Aether reaches each cluster over the Incus REST API under mutual TLS; Operations Center and Migration Manager are reached the same way. Because every interface is a published REST API, a customer can continue operating the four open-source components without FF Aether and nothing becomes unreachable.

Boundary discipline. There is no Aether-specific Incus build and no patched-upstream fork. The contract between FF Aether and the substrate is REST + JSON + mutual TLS. This is the property that makes the platform auditable: the control plane can be inspected, replaced, or removed without stranding workloads.

07 FF HypervisorOS — the immutable, attested appliance

This is the feature that most changes the security posture of a heritage estate, and the reason a compromised legacy guest cannot become a platform compromise.

Construction and immutability

- **Debian-13, mkosi-built into a usr-verity image.** The image is a sealed base with a separate `usr-verity-sig`. The result is immutable at runtime: there is no `apt`, no writable `/usr`, and no shell account that can install a package onto a running node. Immutability is a property of the image, not a configuration that can be edited back in.
- **A/B atomic updates via `systemd-sysupdate`.** The inactive slot receives the new `usr-verity`, its signature, and a signed Unified Kernel Image (UKI); a controlled reboot promotes it; a failed boot rolls back automatically. There is no in-place upgrade and no half-upgraded node state.

- **Secure Boot with measured boot into TPM PCRs.** The UKI is signed, PCR predictions are signed and embedded, the boot chain extends measurements into the TPM, and the resulting PCR values are attested to the management plane — a tampered boot chain is detectable rather than silent.
- **LUKS root sealed to TPM PCRs**, so disks at rest are useless off-platform; **ZFS** for the writable state partition.
- **No SSH on cluster nodes.** Management is the authenticated Incus REST API or the `incus admin os` proxy to the node's unix-socket API. There is no interactive shell to compromise, phish a key for, or leave a web-shell on.

Cluster-wide upgrade flow (workloads stay online)

Operations Center stages the new image to the inactive slot; FF Aether drains the node by live-migrating every instance to cluster peers respecting affinity rules; the node reboots into the new slot; on a successful boot the slot is promoted and the previous slot retained as the rollback target; the node rejoins and accepts workloads back, and the cycle moves on. There is no maintenance window in which the estate is down.

Mapped to the threat model. The documented hypervisor-ransomware chain (UNC3944) works by reaching vCenter, enabling SSH on the ESXi host, and encrypting datastores from below. On FF HypervisorOS the host-persistence step is structurally unavailable: there is no SSH to enable and no writable system partition to plant a payload in, regardless of credential exposure. The host-to-guest VSOCK path abused by CVE-2025-22225 is on a current, measured host rather than an end-of-life one with no fix coming.

08 Preserving the legacy guest – FF Migration Manager and QEMU device mapping

Why a guest that VCF 9.x deprecates boots and runs unmodified, and what exactly changes underneath it.

FF Migration Manager remaps each VM onto a fresh FF Incus/QEMU hardware profile chosen to preserve everything the guest can observe. Memory is preserved 1:1 in bytes and the **source UUID is preserved** so guest-side licensing keyed on UUID keeps validating. The machine is `q35` + `UEFI` by default; **legacy BIOS is used when the source's LegacyBoot flag is set**, Secure Boot is propagated as instance config, and a TPM device is attached when the source had one. CPU is presented as a flat vCPU count.

Source (VMware)	Target (FF Incus / QEMU)	Why the guest doesn't notice
Machine type (various)	q35 (legacy BIOS if LegacyBoot)	Firmware path matches the source; old boot loaders still find their disk
Disk controller	virtio-scsi (virtio-blk fallback for older guests)	Post-import driver injection; fallback covers guests without virtio-scsi
vmxnet3 NIC	virtio-net, MAC preserved exactly	<code>hwaddr=<source MAC></code> per NIC; DHCP leases and per-MAC rules still bind
CPU sockets/cores	Flat vCPU count	Guest sees the same core count; no re-licensing on socket change
Memory / UUID	1:1 bytes / UUID preserved	UUID-keyed licences continue to validate

FF Migration Manager runs post-import scripts in the migration window: GRUB fixup, virtio driver injection, and the `incus-agent`. For a guest family the point of decision is not “does it boot” but “is it supported”: FuturFusion commits in writing to Windows Server 2003/2008, older Windows client (XP/7), RHEL 4–6 and derivatives, and other KVM-class legacy guests confirmed case-by-case — the same list VCF 9.x is deprecating.

09 The network substrate – OVN overlay and distributed firewall

The network maps cleanly onto concepts an NSX-T engineer already holds:

- **Management plane – FF Aether** (analogue of NSX Manager): authoritative policy, LB config, tenants, RBAC; JWT-authenticated REST + UI.

- **Control plane — OVN** (analogue of NSX CCP+LCP): the Northbound DB holds logical intent (switches, routers, ACLs, LBs), `ovn-northd` compiles to Logical Flows, the Southbound DB holds compiled flows and per-chassis port-binding.
- **Data plane — Open vSwitch on every host**: a local `ovn-controller` subscribes to the Southbound DB, translates the flows that concern its host into OpenFlow, and programs the kernel OVS data path; stateful rules use kernel `conntrack`.

Tenant transport is a **Geneve overlay** over the customer underlay; north-south follows a T0/T1 mapping familiar from NSX-T; multi-site is **OVN Interconnect**. Forwarding and enforcement are upstream Linux kernel code paths.

What OVN does, and what it does not. OVN provides the L2/L3 overlay and the distributed-firewall ACL data path — and nothing above it. Layer 4–7 load balancing, Layer-7 network detection and response, and the migration of an NSX-V/NSX-T rule corpus are not OVN functions, and the OVN community has no roadmap to add them. Those are FuturFusion’s own layer, built on top of the OVN substrate: the FF Aether SDN engine (§7), the L4–7 load balancer and Layer-7 NDR (§9), and the first-party rule import in the migration data path (§8). This division is the point — the open substrate handles transport and stateful ACLs; FuturFusion supplies the enterprise network services and the NSX migration that no OVN-based platform provides out of the box.

10 The SDN engine and per-vNIC ACL model

How FF Aether replaces a mature NSX-T deployment without asking the customer to rebuild their microsegmentation or abandon their CI/CD.

The rule-pusher pattern — keep your pipeline

Mature NSX-T shops run a custom SDN controller that owns the source of truth for microsegmentation and pushes it through the NSX Policy API from a pipeline. FF Aether is designed for exactly this: the customer’s controller authenticates once, receives a short-lived JWT, and creates/updates/deletes rules against a small stable REST surface scoped by cluster (with a global-rules scope for fleet-wide rules). FF Aether acknowledges synchronously, persists to its authoritative store, and asynchronously reconciles every affected Incus cluster. The customer keeps their change control, pipeline, IP/subnet scheme, group-naming, and review process — they are re-terminating, not rebuilding.

Four-stage reconciler and the per-vNIC effective ACL

FF Aether holds the authoritative rule store and drives a fleet of Incus clusters into conformance through a four-stage idempotent reconciler. Crucially, it materialises a **per-vNIC effective ACL** — only the rules a given vNIC actually participates in are pushed to that port — rather than a flat global ruleset, which keeps the OVS flow tables small and the enforcement path fast on every host.

Four-tier precedence — baseline the tenant cannot weaken

Scope (highest → lowest)	Authored by	Purpose
Provider Global	CSP / platform operator	Applies across every cluster and tenant; tenants cannot disable or override. (Analogue of NSX Federation cross-site policy.)
Tenant Global	Tenant administrator	Across every network the tenant owns; layers on top of Provider Global.
Per-Cluster	Cluster admin	Every tenant on that specific cluster.
Workload-Group	Tenant	Finest grain — application microsegmentation on a tagged instance set.

The scopes compose but cannot punch upward: Provider Global is always in force, and a Workload-Group rule cannot override it. **Metadata-driven dynamic groups** resolve tag membership into address sets so posture follows identity, not IP — the mechanism that lets migrated NSX security-group semantics survive intact.

11 Migrating the estate — the data path in detail

FF Migration Manager is the Apache-2 VMware-to-Incus transfer engine; FF Aether wraps orchestration and adds the first-party NSX-V/NSX-T rule import that OVN does not provide.

What is preserved vs what changes underneath

Preserved — the workload sees no difference


- ▶ IPv4/IPv6 address; MAC on every NIC
- ▶ Default gateway (until cutover); DHCP lease still valid
- ▶ DNS / hostname / asset records; reverse DNS / monitoring
- ▶ MTU on the wire; firewall rules (per-IP/per-MAC)
- ▶ Licence bindings (per-MAC / per-UUID); application config (no re-IP)

Changed — invisible to the workload

- ▶ Hypervisor: ESXi → FF Incus / QEMU+KVM
- ▶ vNIC driver: vmxnet3 → virtio-net
- ▶ Storage: VMFS → enterprise array / local pool
- ▶ L2 transport: VXLAN/Geneve → VLAN/OVN
- ▶ Management: vCenter → FF Aether API

Two-phase disk import

Background sync uses VMware CBT (Changed Block Tracking) to drive an incremental copy while the source VM keeps running, repeating every ~10 minutes (`config.background_sync_interval`) with no downtime. **Final import** runs in a migration window: the source is powered off, the last delta is copied, post-import scripts run (GRUB fixup, virtio drivers, `incus-agent`), and the VM boots on FuturFusion. Typical downtime is a few minutes per VM. Because CBT copies only allocated blocks, a 2 TB thick-provisioned VMDK with 200 GB used lands as a ~200 GB thin volume — roughly 1.8 TB saved at rest.



FUTURFUSION SUITE — THE PLATFORM THAT PASSES BOTH FILTERS

In-guest network preservation — no re-IP

The guest’s existing configuration survives through **MAC preservation plus interface-name aliasing**. On Linux a worker reads the existing config (netplan, NetworkManager, `/etc/network/interfaces`, systemd-networkd, or `ifcfg-*`) and writes a udev rule pinning the original interface name to the preserved MAC. On Windows the registry’s MAC↔InstanceGUID mapping is captured offline and a first-boot PowerShell script re-applies the original GUID to the new virtio NIC, so existing TCP/IP bindings latch on. DHCP guests keep their lease (identical MAC); static-IP guests keep their IP (interface name preserved).

Subnet cutover — three source topologies, one target primitive

Per-subnet migration preserves IP and MAC by extending the source broadcast domain into the FuturFusion rack for the cutover. Three source-side topologies converge on one target:

Source topology	Bridge mechanism	FuturFusion target
Physical VLAN (dvPortGroup)	None — fabric already L2-stretched	OVN tenant network, <code>localnet</code> -extended during migration; VMs land with original IP + MAC and Aether OVN ACLs from day one
NSX-V VXLAN (Logical Switch)	DLR L2 Bridge on a designated ESXi	
NSX-T Geneve (Segment)	Edge Bridge profile on an Edge node	

The same Migration Manager batch model and disk-copy mechanism apply to all three, and the distributed firewall is in force on the target from the first booted VM — there is no window of degraded protection.

12 NDR — out-of-band Layer-7 detection

Detection is out-of-band over OVN port mirroring (ERSPAN-style): one sensor container per host in the `project=internal` namespace, fed selected logical-port traffic. The sensor sits beside the data path, so sensor failure never affects forwarding. Coverage is opt-in by instance metadata (e.g. `user.nsxt` contains `pci` or `user.tier == prod`); combined with the NSX-T tag importer — which rolls every NSX tag into a single `user.nsxt` key — the customer’s existing tagging drives NDR scope without re-tagging. The default detection stack is **Suricata** (signature IDS, ET Open), **Zeek** (protocol parsing), **RITA** (beacon/exfil analysis), and **Arkime** (session capture), with reactive OVN-ACL quarantine driven from FF Aether.

13 High availability, healing and operations

The cluster is the unit of availability; FF Aether is the unit of management. Supported maximum is **50 nodes per cluster**; larger estates deploy multiple clusters under one FF Aether and one Operations Center pane. dqlite consensus runs on a small fixed quorum (`cluster.max_voters` default 3, any odd number ≥ 3), with up to five non-voting standbys holding a replica. Per-node minimum is 4 GiB RAM, 50 GiB disk, one NIC.

Auto-evacuation is built into Incus (`cluster.healing_threshold`): after `cluster.offline_threshold` seconds of missed heartbeats the leader double-confirms loss over ICMP, then migrates the failed host’s instances to survivors (instances on shared storage restart on peers). Partial degradation (a lost PSU, predictive disk errors) drives a graceful operator-initiated evacuate — the same shape as a vCenter maintenance-mode drain. **Anti-affinity** is enforced by the `user.ha-group` tag, scoped per project so two tenants’ identically named groups do not collide. FF Aether itself runs active-passive for management-plane resilience.

14 Storage architecture

FF Incus supports local pools (ZFS, Btrfs, LVM), shared block (clustered LVM with `lvmlockd` / `sanlock` on FC, iSCSI, or NVMe-oF; LINSTOR/DRBD; TrueNAS zvol over iSCSI), and external enterprise arrays over standard protocols — **Dell PowerStore**, **PowerScale**, **ECS**, **ObjectScale**, NetApp, Pure. Thin/sparse provisioning is the default on local pools; per-disk pool placement lets one VM mix pools (root on SSD, data on capacity) via the Migration Manager

Starlark `set_pool` scriptlet, and the same is available for VMs created directly in FF Aether. SAN connectivity layers REST-managed `nvme → multipath → lvm` services on HypervisorOS so multipathed devices can underlie clustered LVM. Standard-protocol integration with ECS/ObjectScale (S3) and PowerStore/PowerScale (iSCSI/NVMe-oF/NFS) is production-ready today; PowerEdge with BOSS is a validated HypervisorOS install target in PoC labs; Commvault has officially certified FF Aether/Incus workloads in its LTS release.

15 Security architecture, end-to-end

Layer	Control
At rest	HypervisorOS root LUKS-encrypted, key TPM-sealed; FF Aether PostgreSQL on encrypted storage with sensitive columns (cluster client certs/keys, Shepherd secrets, plugin creds) AES-GCM-encrypted at the application layer under the FF Aether master key; DB backups encrypted before S3 push.
In transit	HTTPS browser→FF Aether; mutual TLS FF Aether↔Incus/OC/MM with client certs decrypted in memory at use, never persisted in cleartext; OC→HypervisorOS via the <code>/os/1.0/*</code> proxy on the same mTLS basis; no SSH; inter-cluster OVN tunnels over Geneve.
Identity	Local + AD/LDAP + OIDC + Cloudflare Access; TOTP + WebAuthn/passkeys; SCIM 2.0 provisioning; multi-axis RBAC (scope × domain × role); JWT for the external REST API, session cookies for the UI.
Boot integrity	UEFI Secure Boot on every node; signed UKI with embedded signed PCR predictions; measured boot attested to the management plane; A/B rollback on failed boot.
Workload policy	OVN-ACL distributed firewall on every host at three scopes; Layer-7 NDR out-of-band with reactive OVN-ACL quarantine; vulnerability scanner surfacing CVE findings on managed workloads; audit log for every administrative action, forwardable to Splunk HEC or syslog.

16 Load balancing

FuturFusion ships first-party enterprise load balancing — not a third-party add-on — in two complementary tiers, both managed from the same FF Aether UI and REST API.

- **Lightweight L4 — OVN native.** A VIP on an OVN network with a simple backend pool, living in the OVN data plane itself with zero operational overhead; used by Shepherd service endpoints and any pipeline that just needs to distribute TCP/UDP to a pool. Backends can be any tenant instance, not only FuturFusion-managed ones; health probes use OVN active TCP probes.
- **Production L4 + L7 — FF Aether HAProxy appliance.** A managed HA pair per cluster covering the full production feature set — the Layer-7 capability OVN does not provide. This is FuturFusion’s own layer on the substrate, consistent with the network-substrate division in §7.

17 Identity and access

The identity stack is unusually complete for an open-substrate platform — most of what an enterprise needs from an on-prem IdP integration is in the box.

Area	Shipping today
Authentication	Local users (break-glass); Active Directory / LDAP bind-and-search with group→role mapping; OIDC SSO for any RFC-compliant on-prem IdP (Keycloak, Authentik, ADFS, on-prem Okta/Ping); Cloudflare Access (edge-enforced MFA via trusted access JWT).
Multi-factor	TOTP (RFC 6238) and WebAuthn / passkeys.
Provisioning	SCIM 2.0 for automated user/group lifecycle.
Authorisation	Multi-axis RBAC — scope × domain × role — so a role can be granted on a specific tenant, for a specific domain, at a specific permission level.
Session	JWT bearer for the external REST API; session cookies for the UI.

18 Integration and APIs

Every operational surface in the FF Aether UI — every action across all twenty-one domains — is also exposed as an authenticated REST API. Third-party tools and customer-built management planes can drive FF Aether end-to-end from their own pane of glass without ever touching the UI.

- ▶ The API lives under `/api/...` with JWT bearer authentication, documented in Swagger / OpenAPI, with the interactive Swagger UI shipped inside the application at `/swagger-ui/` — no separate documentation portal.
- ▶ The rule is enforced by the codebase: there is no UI-only path. Every operation a human can perform, a machine can perform — the property that lets an existing SDN controller or CI/CD pipeline retain its source of truth and re-terminate against FF Aether (§8).

19 Lifecycle management

Lifecycle is unified across FF HypervisorOS, FF Incus, OVN, FF Aether, Operations Center, and Migration Manager — orchestrated rolling upgrades, channel-based version pinning, maintenance windows, and air-gapped operation, all drivable from the FF Aether or Operations Center UI.

Field (HypervisorOS update config)	Default	Effect
<code>channel</code>	stable	<code>stable</code> or <code>testing</code> — selects the upstream image stream.
<code>check_frequency</code>	6h	Any Go duration, or <code>never</code> to disable automated checks.
<code>auto_reboot</code>	false	When false, a staged update waits for a manual reboot; boot-time updates always reboot.
<code>maintenance_windows</code>	—	Allowed time windows for the periodic check loop.

Operations Center adds fleet-scale provisioning channels and version pinning, so an operator pins a known-good platform version across many clusters and promotes new versions on their own schedule — including fully air-gapped, with images staged internally.

20 Versus the other VMware alternatives

Two filters decide a regulated shortlist before any feature grid: an immutable, attested host, and first-party migration of the NSX rule corpus.

Platform	Immutable + attested host	First-party NSX rule migration	Legacy guest support	Result
FuturFusion	PASS — mkosi/verity, TPM-attested, no SSH	PASS — NSX-V/NSX-T → OVN ACLs	Committed, incl. 2003/2008, RHEL 4–6	Passes both
Nutanix AHV	Hardened, not image-immutable/attested	Move imports security groups only	2008 R2; 2003 best-effort	Fails
Red Hat OpenShift Virt.	RHCOS immutable	MTV: no rule migration	KubeVirt overhead	Fails filter 2
SUSE Harvester	SLE Micro immutable	No rule migration	Best-effort pre-2016	Fails filter 2
Proxmox / XCP-ng	Mutable Debian / Xen dom0, SSH	No rule migration	Uncertified for oldest	Fails both
OpenStack (self-built)	Mutable control hosts	Coriolis, 3rd-party	KVM guest-agnostic	Fails

21 The cost case

VMware component (bring-your-own licence)	List / core / yr	10,000 cores · 5-year*
VCF 9.x — compute platform	\$400	\$10M
NSX / vDefend — networking & security	\$200	\$5M
Avi — general-purpose load balancing (estimated)	\$70	\$1.75M
Storage & DR add-ons, core minimums	~\$100	~\$2.5M
Effective VMware wallet	~\$770	~\$19.25M
FuturFusion Suite — all-in	\$80	\$4M

*5-year total at 10,000 cores assuming a 50% enterprise discount on VMware list; FuturFusion at its \$80 all-in cost price, no discount. Avi is carried at an estimated \$70/core/yr — Broadcom publishes no list price for it. Even after that discount the platform-licensing saving is ~\$15.25M over five years (~79%), before migration, hardware, and rental/egress savings.

Hardware-neutral — you are not required to move to any particular vendor. FF HypervisorOS installs on any UEFI x86-64 server with Secure Boot and TPM 2.0. The OEM is prime and provides the migration-factory landing zone; the platform then runs on the customer's own hardware, whatever vendor that is.

Open source, and what you buy. The FuturFusion Suite is an enterprise-supported product built on Apache-2.0 projects — Incus, HypervisorOS, Operations Center, and Migration Manager — of which FuturFusion is the largest contributor and support provider. FF Aether is the only proprietary component; it speaks to the four supported builds solely through their published REST APIs, so the customer can operate the open substrate independently. Licensed and supported, not free software.

Sources. Platform documentation as cited — VMware NSX DFW, Red Hat MTV, Nutanix Move, SUSE Harvester, OpenStack Coriolis. FuturFusion (Aether) Technical Whitepaper, 2026. Threat context: Mandiant (UNC3944); Huntress (CVE-2025-22225). Cost: VCF list \$400/core, 50% discount applied; FuturFusion \$80/core all-in.

FuturFusion.io LLC · 551 S Interstate 35, Suite 300 · Round Rock, TX 78664-2820 · United States · Channel partner to a Tier-1 global infrastructure OEM — the OEM is prime; FuturFusion is named subcontractor. Technical statements reflect shipping code or are labelled as PoC; pricing figures are illustrative, not a quotation. British English throughout.