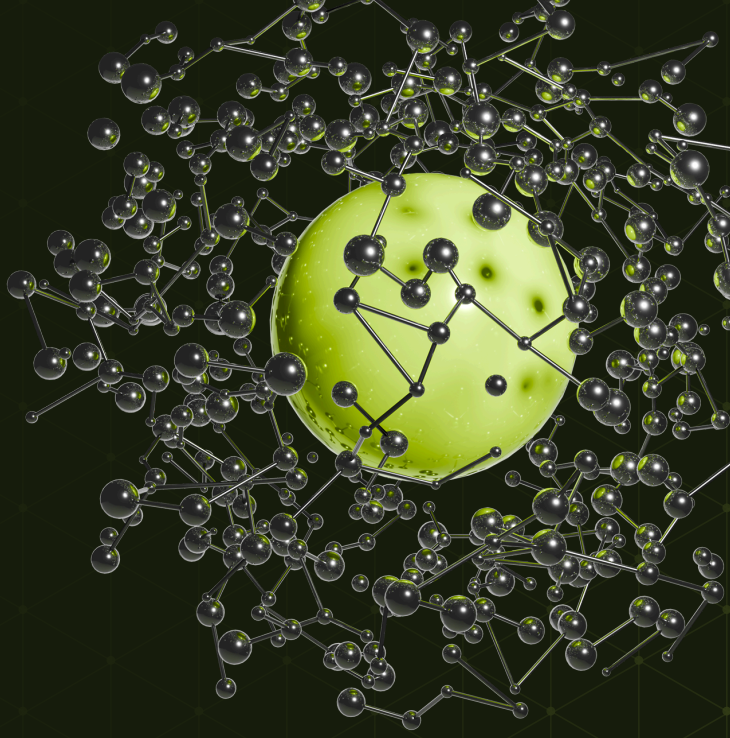
codilime



Solution offering

Migrating enterprise VMware at scale

End-to-end migration off VMware ESXi — virtualisation platform, storage and network foundation — delivered as one programme. We are hypervisor-agnostic: Proxmox VE, Nutanix AHV, OpenStack, OpenNebula, OpenShift Virtualization and beyond. We pick the target that fits your constraints, and extend it with code-level engineering where off-the-shelf falls short.

FOR

Enterprise IT and infrastructure teams leaving VMware

SCOPE

Hypervisor, storage, network, automation, integration

TARGETS

Proxmox, Nutanix, OpenStack, OpenNebula, OpenShift Virtualization, others

DELIVERY

Phased — assessment, PoC, platform build, migration, network rework, support

TEAM

Infrastructure, network and software engineers, and DevOps

PROOF AT SCALE

Numbers from a live programme

The figures below come from a **live open-source migration programme** we deliver for a Fortune-500 enterprise — ESXi to Proxmox VE with an EVPN/VXLAN overlay. The methodology, the team and the platform patterns carry over to Nutanix, OpenStack, OpenNebula or OpenShift Virtualization; the numbers are this programme's.

ESTATE IN SCOPE

10,000+

VMs in a single global programme

NETWORK SEGMENTS

~16M

VXLAN headroom against the ~4,094 VLAN ceiling

NODE RESTART

6x

~90 s → ~15 s per node; full cluster reload under 8 min

HYPERVISOR LICENCE

$0

Open-source stack; support, operations and storage separate

CLUSTER SCALE

32-node

Target scale, reached within the programme

PLATFORM CODE

Custom

Proxmox SDN extensions, built in-house

THE PROBLEM

Why enterprises are leaving VMware now

After the Broadcom acquisition, VMware pricing and packaging changed in ways that are hard to justify at scale — above all for lab, dev, test and large-tenant environments that use a fraction of the bundle. Many of these environments have also outgrown the classic network design underneath: they need more than the ~4,094 isolated networks the VLAN standard allows, and NSX, the add-on that lifts that limit, comes at an extra licence cost on a bill already in question.

The exit is technically achievable — mature alternatives exist on open-source and proprietary lines — and it is a programme: an evaluation that holds up on real workloads, a platform build that matches your operational model, a migration your users do not notice, and in most cases a network foundation rebuilt at the same time.

Licensing pressure

Broadcom-era pricing forces a TCO reset. Open-source and HCI-based alternatives move from “interesting” to “required”.

VLAN ceiling

802.1Q caps a network at ~4,094 segments. Large multi-tenant and customer-case environments have used them up.

L2 scale and the NSX bill

Scaled L2 isolation on stock VMware means NSX, adding cost and complexity on top of licensing.

WHAT WE DELIVER

One estate out, a platform in

Every layer under pressure today has a replacement chosen to fit your constraints — open-source or licensed — and extended in source where it needs to be.

TODAY — UNDER PRESSURE

VMware estate

- vCenter — admin experience
- NSX — L2/L3 overlay, licensed
- ESXi — hypervisor, licensed
- vSAN or third-party storage
- VLAN namespace — ~4,094 ceiling
- Physical fabric — vendor-bound

Licence reset, VLAN exhaustion, the NSX bill

WHAT WE DELIVER

Platform-agnostic replacement

- Self-service portal — CMP or integrated
- IaC and GitOps — Ansible, Terraform, Pulumi
- Your hypervisor — Proxmox, Nutanix, OpenStack, OpenNebula, OpenShift Virtualization
- Storage — Ceph, HCI or SAN, matched to the platform
- EVPN/VXLAN fabric — ~16M segments, vendor-neutral
- Optional SmartNIC / DPU offload for AI workloads

Open-source or licensed, chosen to your constraints, extended in source when needed

OUR DELIVERY METHODOLOGY

Six stages, every migration

The same six-stage path runs on every migration. Scope and duration scale with the estate; the stages stay the same.

1

Assessment

The alternatives evaluated against your networking, scale, licensing and operational requirements.

2

Proof of concept

The shortlisted platform validated at representative scale, with any custom extensions it needs confirmed.

3

Platform build

Cluster, storage and automation, fully as code, integrated into your existing portals and workflows.

4

Migration

Workloads moved from ESXi to the target platform, the topology preserved to keep the risk low.

5

Network rework

EVPN/VXLAN overlay, physical switching refreshed, GitOps-driven network automation.

6

Support

Managed services for the extended platform, so your team always has the people who built it.

TARGET PLATFORMS

We match the target to your constraints

The right replacement for VMware depends on your licensing stance, operational model, network requirements and existing tooling. The list is illustrative — we work with other platforms as needed, including vendor-supported distributions of the open-source options.

Proxmox VE

Open source

KVM and LXC with native Ceph integration, source code open to deep customisation. Strong fit when TCO and extensibility lead the decision and your team is ready to own the stack.

Nutanix AHV

Proprietary

HCI-native, KVM-based hypervisor with Prism management and Flow networking. Strong fit when operational simplicity, bundled support and an all-in-one HCI experience come first, in exchange for a proprietary platform.

OpenStack

Open source

Modular IaaS, carrier-grade multi-tenancy, Neutron networking with rich SDN options. Strong fit for telco, service-provider and large multi-tenant environments that already think in IaaS primitives.

OpenNebula

Open source

Lightweight and opinionated, with low operational overhead. Strong fit for private cloud and edge, where a smaller ops team carries the stack and OpenStack's breadth exceeds the need.

OpenShift Virtualization

Proprietary, on OSS

KubeVirt on OpenShift: VMs as first-class Kubernetes workloads. Strong fit when the platform team is already Kubernetes-native and wants one control plane for containers and VMs.

... and others

Case by case

Harvester (SUSE), OLVM (Oracle), HPE Morpheus, Red Hat OpenStack Services on OpenShift, vendor-supported distributions of the above, or cloud-repatriation paths. Each goes through the five criteria below.

HOW WE CHOOSE

Five criteria, weighed in this order

1 Licensing and TCO
Subscription model, per-core or per-socket, open-source or proprietary, exit cost

2 SDN and L2 scale
VLAN ceiling, overlay support (VXLAN, EVPN), per-tenant isolation at scale

3 Storage model
HCI, external SAN or NAS, Ceph-native or proprietary distributed — against existing ops skills

4 Proven scale
Documented production scale — nodes per cluster, VM density, large customer references

5 Ops maturity and APIs
Day-2 tooling, REST, Terraform and Ansible providers, integration with existing portals

WHAT WE BRING

Capability for whichever platform you pick

The target platform is one decision. The programme around it — network rework, automation, custom extensions, integration with existing tooling, day-2 operations — is where a VMware migration succeeds or stalls. These capabilities come with every hypervisor.

OUR MOAT

Code-level platform engineering

We extend the hypervisor in source

When the chosen hypervisor has gaps — **SDN behaviour, integration hooks, scale thresholds, operational semantics** — we extend it directly in source. On one recent programme we added per-tenant VLAN pools, selective cluster restart logic and a netlink-based orchestration daemon on top of the stock hypervisor. We close gaps where they are.

- Patches upstream
- Custom SDN extensions
- Daemons and APIs
- We maintain what we build

EVPN/VXLAN network foundation

BGP EVPN control plane for the VXLAN overlay, per-tenant VLAN pools, L3-fabric underlay design. Network infrastructure is the core of what we do.

Physical network refresh

Switching vendor evaluation, end-of-life replacement (Arista to Juniper or SONiC, for example), an open-standards fabric that keeps your options open.

Storage design and build

Ceph (hyperconverged or dedicated), vendor HCI distributed storage, external SAN or NAS — chosen to fit the target hypervisor.

IaC and GitOps automation

Ansible for provisioning, GitOps for network state, Terraform or Pulumi where supported. Repeatable and version-controlled from day one.

Portal and API integration

The new platform wired into existing automation portals, ticketing, CMDB and customer-facing tooling, so the end-user workflow stays the same.

Managed services and long-term support

Once it ships, we stay on — essential when the platform carries custom extensions, so the code you depend on stays maintained.

VIRTUALIZATION
Your pick
Open-source or proprietary

STORAGE
Ceph, HCI, SAN
Matched to the platform

NETWORK
EVPN/VXLAN
Vendor-neutral fabric

AUTOMATION
Ansible and GitOps
IaC end to end

HOW WE DE-RISK

Migration discipline on every wave

The six stages on page 3 describe what we do. The six controls below describe how we keep the programme safe. An enterprise VMware exit succeeds or fails on execution — dependency discovery, backup integrity, wave gating and post-cutover support. Each discipline runs on every engagement and produces a concrete artefact you own.

OUR POSTURE

Programme-safe by design

Bounded waves, documented gates, rehearsed fallback

We run **six explicit risk disciplines** on every engagement, fitted to your estate. Waves are scoped so the impact of any one is contained. Gates are documented so every step forward is defensible. Fallback paths are rehearsed before they are needed.

- Bounded waves
- Stage-gate progression
- Verified backup
- Hypercare window

1. Discovery and dependency mapping

An authoritative inventory extracted into a CMDB or IaC source of truth. Topology and configuration documentation generated automatically. NetFlow and sFlow runtime observation fused with configuration data, so application-level dependencies are captured alongside VM-to-VM ones.

2. Wave-model selection

Five candidate patterns — by criticality, tenant, network zone, workload class, or combined — chosen per engagement from the dependency graph and your risk tolerance. The live programme is led by network zone, aligned to the EVPN/VXLAN fabric rework.

3. Gated execution and rehearsal

A stage gate between waves: soak period, automated regression, incident budget, your sign-off. A pre-production rehearsal environment mirrors the target configuration and change set. IaC-driven change management covers cutover and every reconfiguration after it.

4. Contained blast radius

Waves scoped tight — a single tenant or location. IaC re-provisioning backed by per-VM snapshots where appropriate. The source hypervisor stays live during coexistence, with documented fallback and wave-level RTO and RPO. Each wave hardens the automation for the next.

5. Data and backup integrity

A verified, restorable backup before every wave. Checksum and consistency validation on the target. Backup and DR tooling re-platformed and validated on the new stack, with VADP-era dependencies identified early and replaced deliberately.

6. Hypercare and handover

An elevated support window of 30–90 days right after cutover. In steady state your IT owns L1 and we keep L2 — breakage, improvements, platform evolution, further migration. Automation and operational knowledge are handed over at cutover.

EXTEND THE PLATFORM — OPTIONAL

Day 2, self-service and the adjacent layers

The migration lands a new hypervisor, and a production platform needs more: a self-service layer for end users, observability the network team can trust, regression pipelines for the fabric, and accelerated data paths where workloads demand them. These are productised offerings we bring alongside the core migration, or add afterwards.

OPTIONAL SELF-SERVICE LAYER

Cloud Management Portal

The vCenter experience, kept after vCenter

A **multi-tenant self-service portal** over Proxmox, OpenStack and other hypervisors, under one UI. It brings back the admin experience teams lose when they leave vCenter — tenant hierarchy, RBAC, team quotas, scheduled operations, VM lifecycle, SSO through Keycloak — and a pluggable backend, so the same portal follows whichever hypervisor you pick.

- Multi-tenant and RBAC
- Pluggable hypervisor backends
- Quotas and scheduling
- SSO / Keycloak
- Ansible integration

VM detail and metrics

Console, snapshots, quotas, scheduling, automation, usage and audit per VM. Live CPU, memory, disk and network metrics with anomaly detection and right-sizing hints.

Templates on real clusters

VM templates registered against real Proxmox clusters: pluggable hypervisor backends, Cloud-Init provisioning, enforced per-template resource floors.

Usage and chargeback

Monthly running and provisioned hours, vCPU-h, RAM GB-h, disk GB-h and cost per project. Showback and chargeback built in, exportable as a report.

OPTIONAL WORKSTREAMS

Commonly added to a VMware-exit programme

Around the platform rebuild, a programme of this scope can pull in modular engineering streams, each added when it helps.

CODILIME OFFERING

Observability, AI-enhanced

End-to-end visibility across network, fabric, hypervisor, OS and application layers, on Prometheus, Grafana or Splunk. Per-tenant dashboards, SLO tracking, and an optional AI layer for anomaly detection and correlated alerting across the stack.

CODILIME OFFERING

QA automation for networks

Regression-test pipelines for the fabric and the hypervisor platform. VXLAN/EVPN state, tenant isolation and cluster behaviour verified on every change, so drift is caught before production.

CODILIME OFFERING

Data-plane acceleration

DPDK, VPP, SmartNIC and DPU, FPGA and GPU data planes for tenants with high-throughput, low-latency or secure-networking workloads that need more than stock vSwitch performance.

CODILIME OFFERING

AIOps and GenAI for networks

Intelligent incident response, runbook automation and GenAI-assisted troubleshooting on top of the observability stack, turning the new platform's operational data into faster MTTR.

REFERENCE ENGAGEMENTS

Delivered for Tier-1 vendors and enterprises

A selection of production engagements that exercise what a VMware exit needs — open-source hypervisor migration at scale, carrier-grade virtualised infrastructure, multi-cloud self-service, SmartNIC and DPU offload, and AI-driven operations. Further case studies are available under NDA.

FLAGSHIP — LIVE PROGRAMME

Fortune-500 enterprise

ESXi → Proxmox VE at production scale

A multi-year open-source migration off VMware onto **Proxmox VE with an EVPN/VXLAN overlay**. CodiLime owns the platform extensions, SDN design, CI/CD for the fabric and the migration tooling. The figures on page 2 come from this programme — live, in production and expanding.

Proxmox VE · EVPN / VXLAN · Ansible and IaC · SDN extensions

TIER-1 TELECOM

Carrier-scale OpenStack NFVI

IMS platform · 14M subscribers

Automated deployment and monitoring of virtualised network functions on OpenStack for a Tier-1 mobile operator's IMS core — carrier-grade open-source virtualisation at subscriber scale.

OpenStack · KVM · Ansible

AI SILICON LEADER

DC-scale network orchestration

55K+ devices · GPU and AI-cluster farms

Vendor-agnostic orchestration across the client's global footprint: 3 data centres and ~200 offices refreshed per year, powering GPU clusters, AI factories and super-labs. Nautobot as the single source of truth, replacing legacy SaltStack.

Nautobot · Ansible · Python · Arista

GLOBAL HARDWARE VENDOR

Multi-cloud self-service orchestration

Provisioning: weeks → hours

A unified self-service portal spanning private DC and public cloud, with tenant quotas, policy enforcement and approval workflows — the admin experience CMP brings to VMware-exit customers.

Python · Terraform · RBAC / SSO

NETWORKING AND AI SILICON

Service-mesh offload to DPU

AI-factory security · SmartNIC

Cilium and eBPF policy, and the L7 service-mesh dataplane, offloaded from the host kernel onto SmartNICs. **100 Gbps line speed**, 400 Gbps target. Frees host CPU for AI workloads; integrates with Hypershield and Splunk.

C++ · DPDK · P4 · eBPF · BlueField

TIER-1 CYBERSECURITY

Whitelabel migration factory

Hundreds of enterprise migrations

The tooling a Tier-1 security vendor's professional-services team uses to move enterprise customers off Cisco, Zscaler, Fortinet and legacy versions: vendor-agnostic config converters, a portal and CI pipelines, packaged as whitelabel PS.

Python · Pydantic · GitLab CI · Grafana

AIOPS AND OBSERVABILITY

AI-based anomaly detection

Hypervisor, network and hardware telemetry

An ML-driven root-cause model correlating telemetry across the hypervisor, network devices and physical hardware. Fewer false positives, causal chains surfaced — it pairs directly with the observability offering on page 7.

Python · Prometheus · Grafana Loki

ABOUT CODILIME

The partner you are engaging

250+ engineers

15+ years in networks

Tier-1 vendor and enterprise clients

SELECTED CLIENTS — NETWORKING AND INFRASTRUCTURE LEADERS WE BUILD FOR

Cisco · Juniper Networks · Palo Alto Networks · Nutanix · NVIDIA · Equinix · AT&T · Altera · Swift Navigation · Selector · Linux Foundation

VOICES FROM OUR WORK

"

CodiLime's software engineers have been a great addition to our team by bringing a unique combination of experience in Go and network domain knowledge.

Kristine Antons
Senior Software Engineering Manager, Equinix

"

CodiLime's DevOps teams helped us reduce our development cycle time by 50%, and its software engineering helped us improve the quality of our product.

Harjot Gill
General Manager, Nutanix Xi Epoch

"

Over the last two years we have developed a strong partnership with CodiLime that has allowed our Professional Services team to expand its reach and ability to deliver more projects.

Alexandru Smeureanu
Service Delivery Manager, Global Scale Team, Palo Alto Networks

WHY CODILIME — OUR DIFFERENCE

Network infrastructure is our core

SDN, SD-WAN, EVPN, VXLAN, BGP, MPLS, L2/L3/L7 — every layer is our own work, and the fabric ships with the hypervisor.

Full stack within networking, UX to firmware

UI/UX, backend, system level, drivers, hardware and firmware — one team across every layer, a rare combination inside an infrastructure specialist.

We build platforms from zero

Mesh7 (microservices security, acquired by VMware), an XDR product built from scratch, SD-WAN products for vendors. When off-the-shelf falls short, we go into the source or start from the first line.

AI-assisted delivery, inside and for customers

GenAI and classic ML shipped in products, and AI-assisted engineering in-house — the CMP platform in this paper went from zero to demo in weeks, then was hardened with tests and architecture controls for production.

Platform-agnostic by design

We recommend the target that fits your constraints. Several hypervisors are in live delivery with us today.

We stay once it ships

Managed services for the environments we build — above all when the platform carries custom extensions, so the people who wrote them keep them running.

codilime

Let's talk about your VMware exit

Whether you are at evaluation, at PoC, or mid-delivery and stuck, we can help: assessment workshops, PoC delivery, full programmes, or advisory on a running migration.

Contact us →

