Isovalent Enterprise Platform & the Future of Network Virtualization

Enabling migrations, connectivity, security, and visibility for Virtual Machines

VMware Economics Shifting Fast

Enterprises are under pressure to control costs and regain strategic flexibility

Unpredictable licensing costs

Recent Broadcom changes have led to 200–500% cost increases in some enterprise environments.

Lock-in and limited agility

Long-standing VMware contracts limit cloud flexibility and delay modernization.

Business risk from infrastructure dependency

A single-vendor stack increases exposure to pricing or roadmap shifts.

- VMware licensing cost increased by 150% CIO
- Potential price increase of 1,050% for AT&T The Register
- Some customers reporting 5x cost increase Reddit Discussion
- Price hikes causing customers to seek alternatives Volico
- Significant disruption for enterprise customers CIO Dive

© 2026 Cisco and/or its affiliates. All rights reserved.

Migration Drivers Enterprise IT

Unlock faster modernization, lower TCO, and secure hybrid networks

Reduce VMware Dependency

- Avoid expensive license renewals and regain control over infrastructure spend.
- Migrate on your own timeline, no need for big-bang transitions.
- Increase flexibility in cloud and platform strategy.

Accelerate App Modernization

- Shift VM workloads to Kubernetes without refactoring.
- Run traditional and modern workloads side-by-side during transition.
- Shorten time-to-value for infrastructure modernization projects.

Simplify Security and Operations

- Apply one consistent policy model across VMs and containers.
- Gain full visibility into network traffic and access patterns.
- Reduce operational overhead with a unified control plane.

© 2026 Cisco and/or its affiliates. All rights reserved.

Kubernetes Universal Platform Containers & VMs

Enterprises trust Kubernetes. Now they want to do more with it.

Operational expertise already in place

Platform and infrastructure teams are already running Kubernetes successfully for modern workloads.

Common abstraction layer across on-prem and cloud

Kubernetes enables a consistent operational model and tooling experience across environments.

Future-proof Platform With Automation And Portability

Kubernetes helps future-proof your infrastructure with automation and flexibility.

Much like how vSphere once simplified VM operations, Kubernetes now provides a powerful platform for automating the orchestration, scaling, and management of containerized applications, ensuring consistent deployment and high availability across diverse environments.

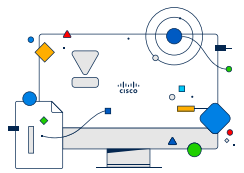

Traditional VM Workloads

Seamless Migration and Integration.

Modern Applications

Rapid Deployment and Scaling.

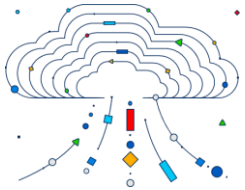

AI / ML Workloads

High Performance and Adaptability.

Kubernetes

Bare Metal | vmware by Broadcom | aws

Azure | Google Cloud | Red Hat OpenShift

© 2026 Cisco and/or its affiliates. All rights reserved.

Isovalent Networking for Virtualization Private Networks

Single model for advanced virtual machine networking

Purpose-built for KubeVirt in Kubernetes

Run traditional VMs in Kubernetes with native networking and security tailored for virtual workloads.

Enterprise-ready for virtualization

Extend network visibility, segmentation, and policy enforcement to KubeVirt and OpenShift Virtualization.

Connecting all environments and Simplifying Migration

Enable secure communication between different hypervisors making it easier to migrate workloads without downtime or rework.

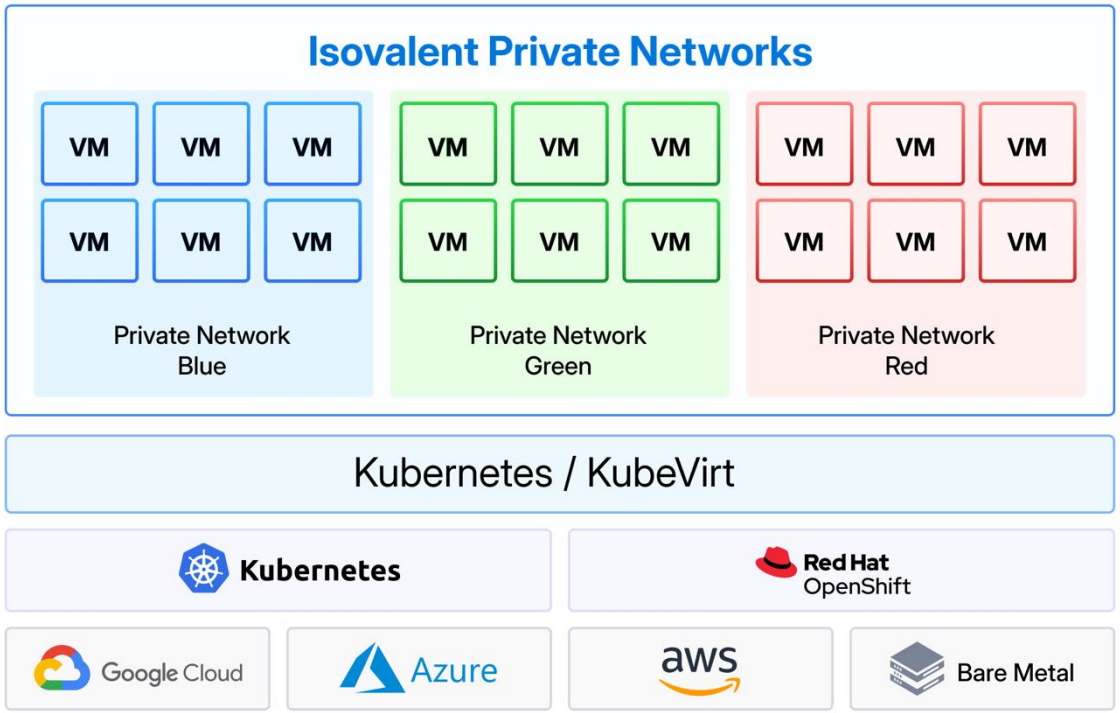

* **OpenShift Virtualization** and **KubeVirt** are confirmed integrations. Other partners are under evaluation. **No ETA.**

© 2026 Cisco and/or its affiliates. All rights reserved.

Target/Planned Compute Platforms Vision

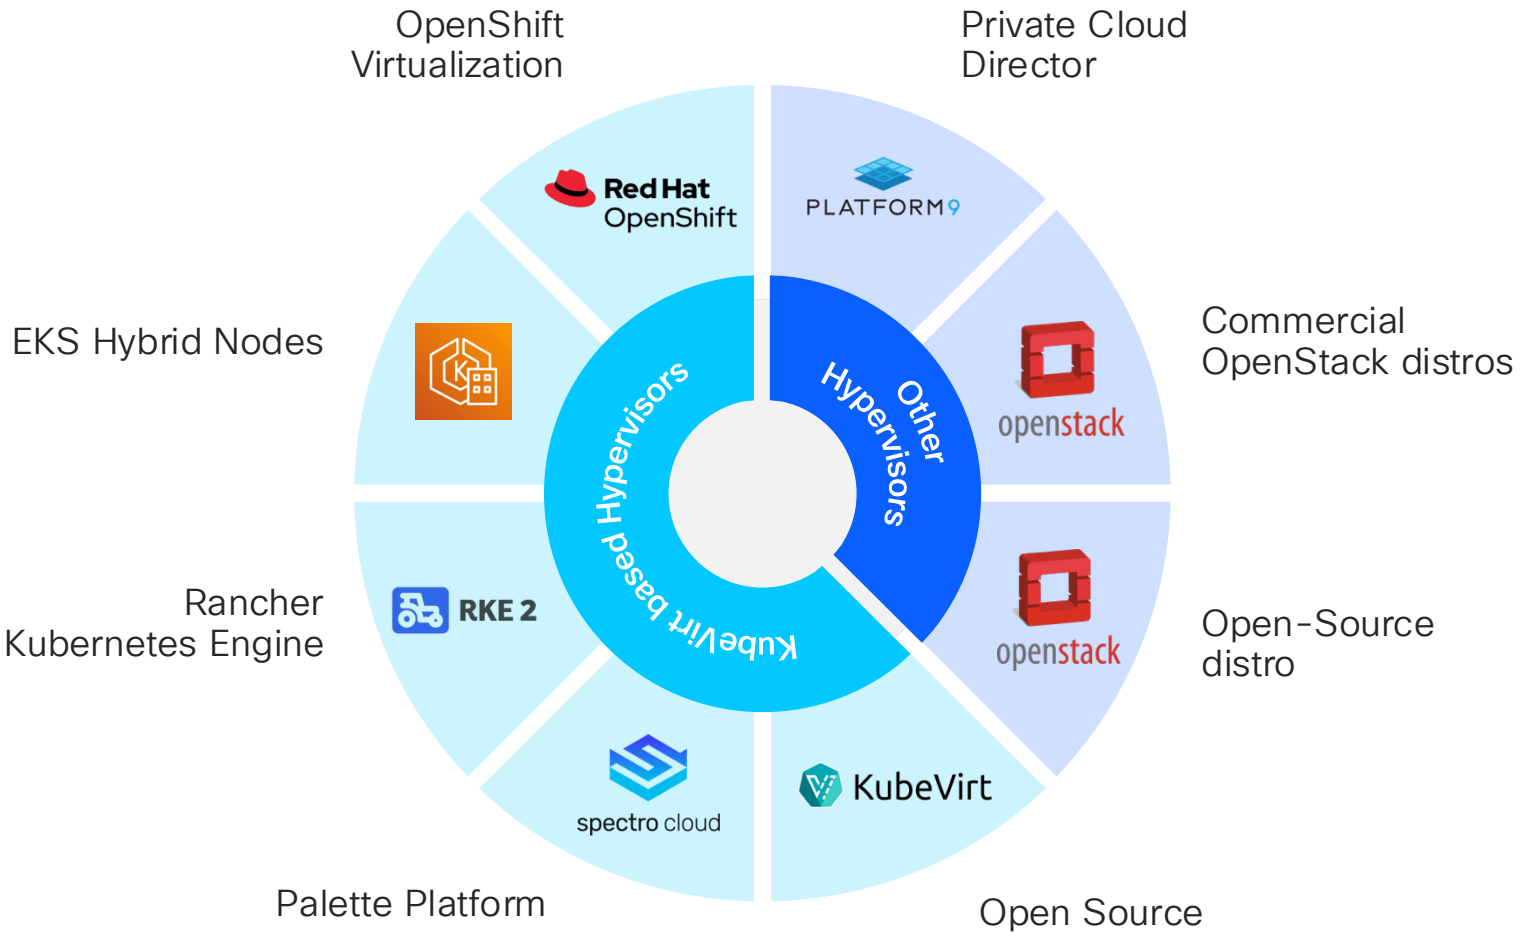

Other Targets:

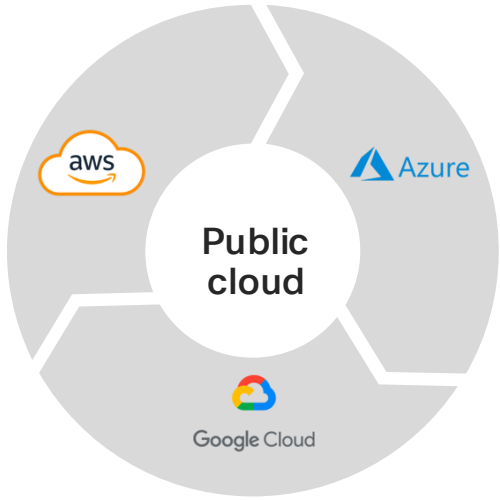

* OpenShift Virtualization and KubeVirt are confirmed integrations. Other partners are under evaluation. No ETA. **Shared under NDA.**

© 2026 Cisco and/or its affiliates. All rights reserved.

Two Paths for Your VM Migration

Choose the model that aligns with your modernization goals, team structure, and operational readiness.

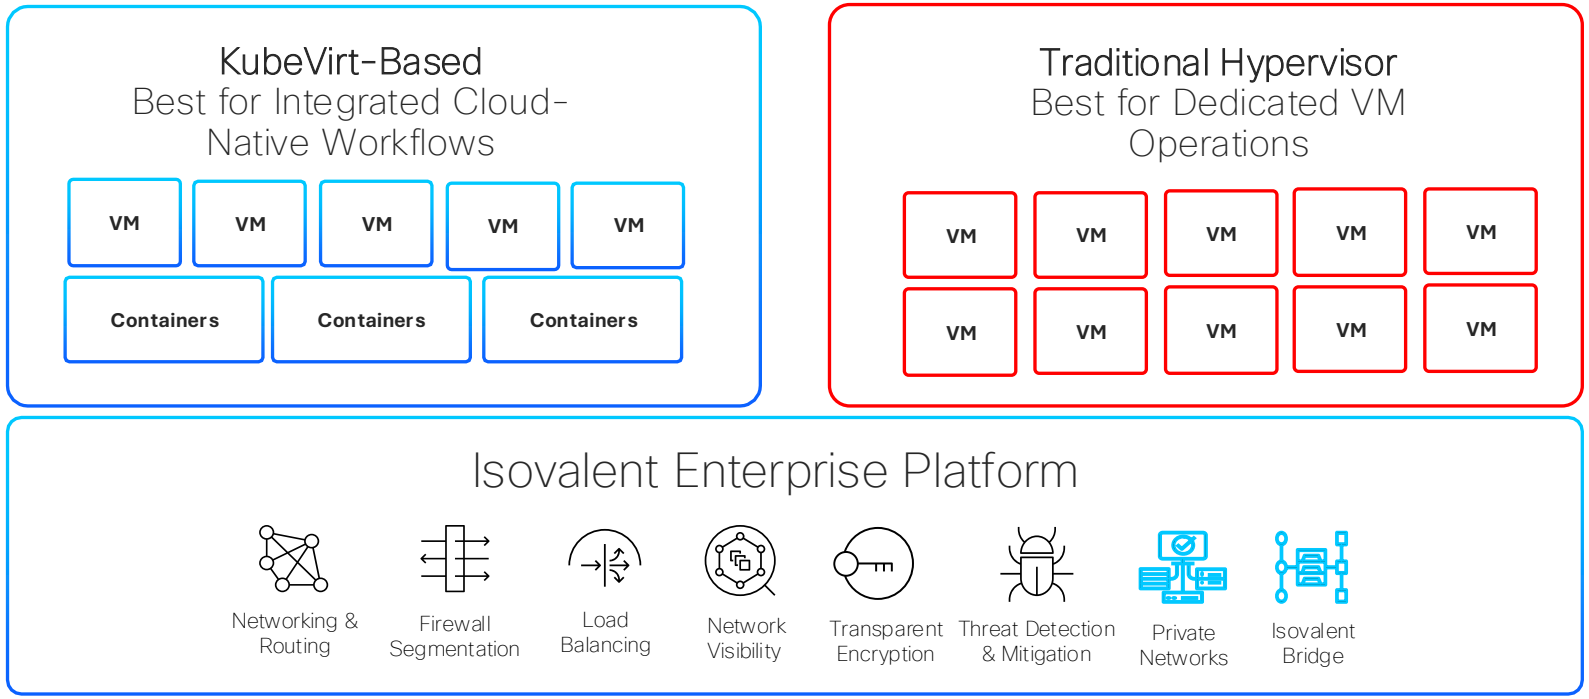

The Isovalent Enterprise Platform supports both paths, enabling secure, incremental migration with consistent policies and visibility, no matter which way you choose.

© 2026 Cisco and/or its affiliates. All rights reserved.

Network Migration The Hidden Barrier

Migration is more than storage and compute. It’s about preserving connectivity.

Applications rely on static IPs and traditional topology

Most traditional applications are tightly coupled to specific network configurations, creating network challenges when they run on modern, dynamic environments.

Re-IP causes risk and rework

Changing IP addresses often requires updating firewall rules, DNS entries, application configs, and documentation, introducing significant risk and disruption.

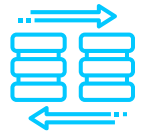

Full subnet migrations often mean downtime

Moving entire subnets typically involves service interruptions and asynchronous cutovers that impact availability.

© 2026 Cisco and/or its affiliates. All rights reserved.

Isovalent Network Bridge Capabilities

Seamless workload communication between VM and Kubernetes environments

Enables communication between VM and K8s workloads

- Enables distributed application components during migration
- Minimizes downtime
- Provides businesses with flexible migration strategies
- Ensures service continuity without interruption

Maintains IP identity and avoids re-architecting apps

- Apps can migrate seamlessly without modifying firewall rules, DNS records, or code
- Accelerates project timelines
- Minimizes migration risks

Supports virtual environments with bare metal on the roadmap

- Flexibility to connect workloads across diverse infrastructure
- Enables organizations to modernize selectively
- Maximizes existing infrastructure investments

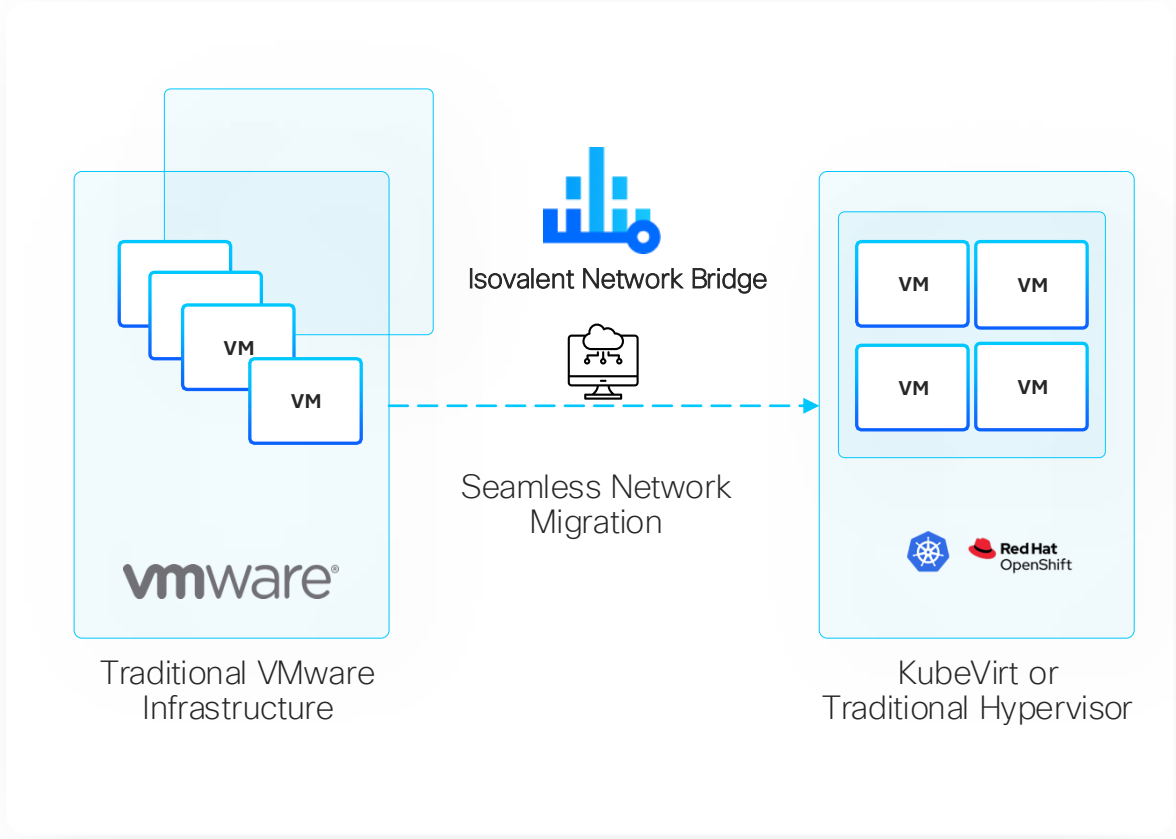

* OpenShift Virtualization and KubeVirt are confirmed integrations. Other partners are under evaluation. No ETA. **Shared under NDA.**

© 2026 Cisco and/or its affiliates. All rights reserved.

VM Migrations Not an Event, a Methodology

Understand the phased approach that enables successful execution

Assessment

- Inventory your current VM workloads and gather OS and licensing info.
- Estimate migration effort and infrastructure requirements.
- Evaluate migration cost and resource needs.

Dependency Mapping

- Identify technical and business dependencies between apps.
- Group workloads based on network constraints and application logic.
- Identify business critical and latency-sensitive applications.

Timeline

- Plan staggered rollouts: first movers, pilot groups, and core systems.
- Define clear timelines, rollback strategies, and test cycles.
- Avoid big-bang cutovers that risk downtime.

Migration Execution

- Use tools for replication and network extension.
- Monitor SLAs, performance, and stability.
- Complete cutover while maintaining security posture and availability.

Personnel & Resources

- Self-managed, partner-supported, or delivered via PSO.
- Evaluate AI-driven tooling to automate planning and execution.
- Align project teams and stakeholders for execution support.

© 2026 Cisco and/or its affiliates. All rights reserved.

Use Cases

Use Cases Migration

Full VMWare Migration

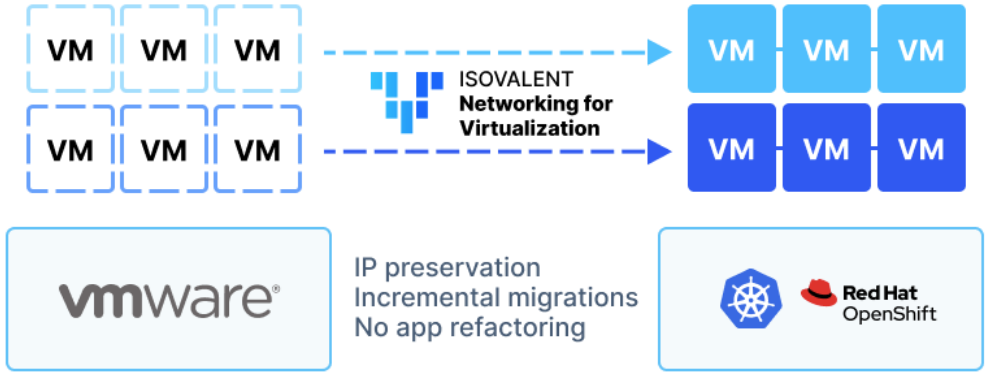

Partial VMware Migration

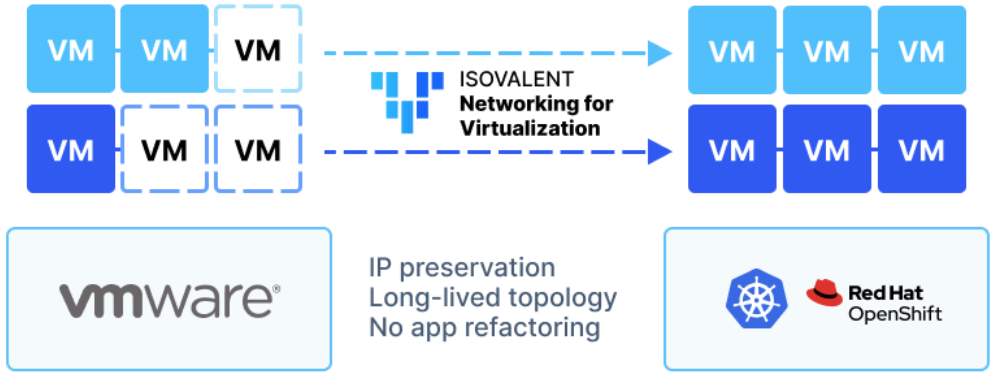

Security Policy Migration

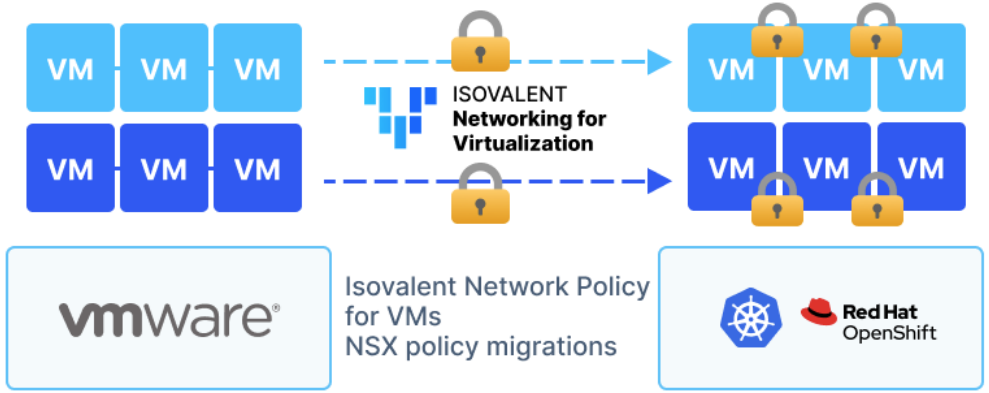

Network Visibility

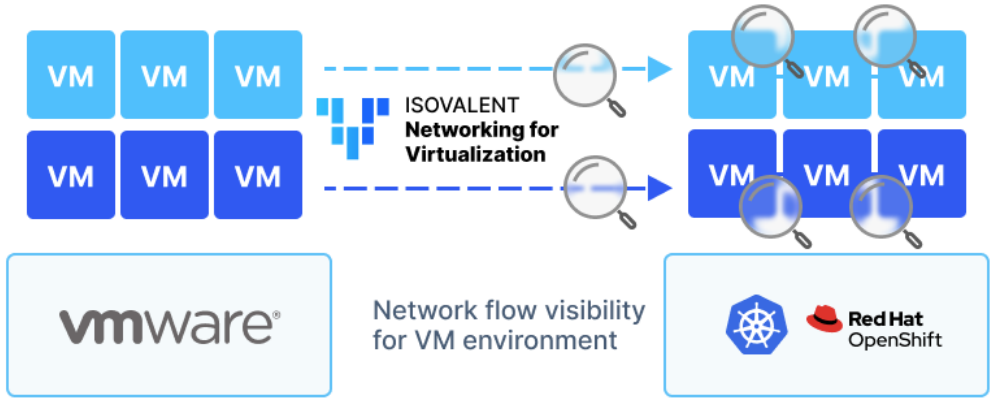

© 2026 Cisco and/or its affiliates. All rights reserved.

Use Cases Cloud and AI

Multi Tenant Segmentation

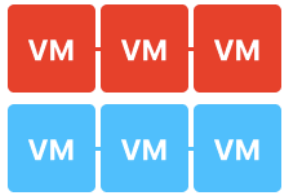

Overlapping IPs
IPv4 and IPv6 support
VRF segmentation

Red Hat OpenShift

Multi Site

Red Hat OpenShift

Availability and resiliency
Disaster recovery
Hybrid cloud

Red Hat OpenShift

Secure AI Factory (2H2026)

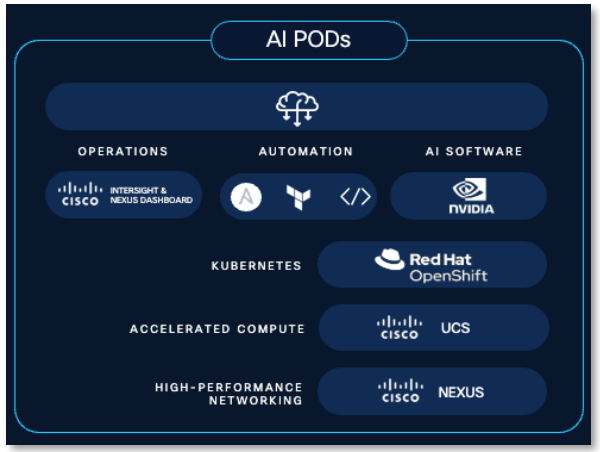

Edge Compute (Committed)

Cisco
Validated Design

* OpenShift Virtualization and KubeVirt are confirmed integrations. Other partners are under evaluation. **No ETA.**

© 2026 Cisco and/or its affiliates. All rights reserved.

Summary of Capabilities

Isovalent Networking for Virtualization At-a-Glance

Features and capabilities

L2/L3 Networking	
VPC support	Yes (IPv4 and IPv6)
VLAN access to Fabric	Yes (L3)
EVPN access to Fabric	Yes (L3 VNI)
SRv6 access to Fabric	Planned*
Network extension to VMware NSX	Yes (Isovalent Network Bridge)
Dynamic routing	Yes – BGP (EVPN use case)
Multi-NIC support	Yes, without dependencies on the pod network
VRF support	Yes, including control plane and dataplane separation
Multi-site	Multi-AZ in 1.19 Multi-site / Multi-region planned (ETA 1.20)*
Validated scale limits	5000 Nodes, 200K Endpoints

Security	
L3/L4 Distributed Firewall	Yes (Isovalent Network Policies)
L7 Distributed Firewall	Planned*
ConnTrack preservation for Live Migration	Planned (ETA 1.20)*
NSX Policy Migration	Yes (via 3P solution)

Observability	
Timescape	Yes
Splunk	Yes

© 2026 Cisco and/or its affiliates. All rights reserved.

Roadmap items and timelines subject to change.

Cisco Confidential. Shared under NDA.

Isovalent Networking for Virtualization

Additional components

Miscellaneous / Add-ons	
Load Balancing	Isovalent Load Balancer
NAT (Egress GW)	Planned*
WAF and IDS	Planned (ETA 1.20)*
Runtime Security	Planned*

© 2026 Cisco and/or its affiliates. All rights reserved.

Roadmap items and timelines subject to change.

Cisco Confidential. Shared under NDA.

Migration Ecosystem 1P/3P Integrations

FABRIX By ReSTNSX

Capabilities:

- **Firewall-style policy:** Author Isovalent policy as source / destination / port / action rules — no manifests, no kubectl, no YAML.
- **Reporting & compliance:** Audit-ready evidence, control posture, and exportable PDF / CSV reports out of the box.
- **Hubble flow observability:** See real east-west flows and verdicts, then convert any dropped flow into an allow rule in a click.
- **Inventory & objects:** Universal endpoint & identity inventory with reusable security objects across policy.

Supported

Red Hat Migration Toolkit for Virtualization

Capabilities:

- **Isovalent Networking for Virtualization** automatically detects the source VM IP and injects it in the right configuration stanza post-migration.
- **Multi-Provider Ingestion:** Directly migrates VMs from VMware vSphere, an others.
- **Warm Migration & Zero-Downtime Transfer:** Pre-copies VM disk data in the background to minimize business disruption during cutover.
- **Automated Batching & Ansible Integration:** Orchestrates mass VM transfers with automated network, storage, and resource mapping at scale.

Supported

Cisco Helix

Workload Mobility Solution

Capabilities:

- Transforms workload mobility from a reactive, labor-intensive event into a **durable**, strategic capability.
- Reduces hidden friction and manual labor through composable tools and explicit resource mapping.
- Ensures business continuity during platform transitions with network identity preservation and topology consistency.
- Provides continuous mobility assessment via the Mobility Score to proactively identify and mitigate migration risks.

Planned

© 2026 Cisco and/or its affiliates. All rights reserved.

cisco

Hypervalent Sales

VMware Exit Sales Play

Isovalent Networking for Virtualization Internal Site

Hypervalent Sales Portal

NEW!!

Improved navigation and access to all INV assets

[Isovalent Networking for Virtualization – SharePoint]

© 2026 Cisco and/or its affiliates. All rights reserved.

VMware Exit Sales Play FY2027

- Customer segment
  Enterprise, Service Providers (IT), All Verticals
- Priority industries
  All verticals
- Priority geographies
  Global with focus on regions with high VMware adoption and cloud-native initiatives
- Environments supported:
  VMware ESXi (vSphere with VSS, VDS or NSX) – only applicable to the migration use case
- Competitive Products – Network Extensions / Migrations
  VMware HCX, Google Cloud Hybrid Subnets, Azure Extended Network

  Competitive Products – Network Virtualization Products
  OpenShift OVN-K, Nutanix Flow Networking, OpenStack Neutron, VMware NSX (as a legacy solution being migrated from)

Key Use Cases

Migration from VMware

Primary Buyer: Network Infrastructure Team
Customer utilizes VMware vSphere with VSS, VDS or NSX and has a strategic initiative or urgent need to migrate workloads off VMware due to cost pressures resulting from licensing changes or upcoming renewals, OR customer has applications with network dependencies tied to fixed IP addresses which make a traditional "lift and shift" approach challenging or disruptive.

Implement Virtualization on top of Kubernetes for Greenfield Deployments

Primary Buyer: Platform Team
Customer has an existing or planned Kubernetes platform such as OpenShift Virtualization, EKS Hybrid Node or other, and plans to deploy new workloads in this environment.

Deal Accelerators:

- Broadcom licensing changes or impending renewal creating significant cost pressure.
- Applications with complex network dependencies which hinder traditional migration methods.
- Desire to reduce vendor lock-in, increasing IT flexibility and agility.
- Existing Kubernetes adoption or strong cloud-native strategy.
- Need for consistent networking and security policies across environments with both VMware and Kubernetes.

Disqualification Criteria:

- Customer has no plans to migrate off VMware or adopt Kubernetes as a target environment, has workloads that are out of scope for migration such as certain databases, VDI solutions or COTS applications which are not certified for the target environment, or if Kubernetes, OpenShift or another PaaS is already running on VMware.

© 2026 Cisco and/or its affiliates. All rights reserved.

VMware Sales Play Basic Migration Use Case

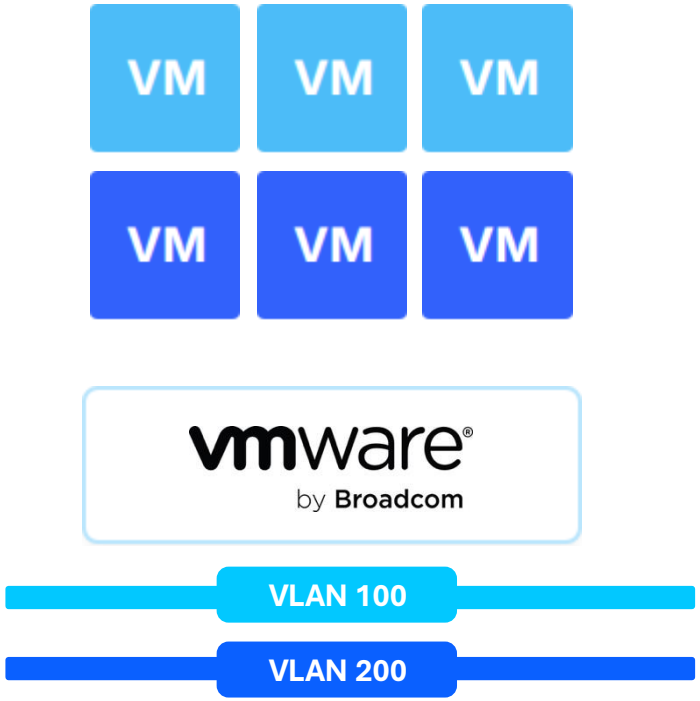

Architectural Reality

Most existing VMware deployments run on ”bare vSphere” relying entirely on traditional VLANs without NSX, vRA, VCF, or other VMware components.

This means these environments are primed for a direct migration to OpenShift Virtualization and KubeVirt using VLANs at the destination.

© 2026 Cisco and/or its affiliates. All rights reserved.

VMware Sales Play Basic Migration Use Case

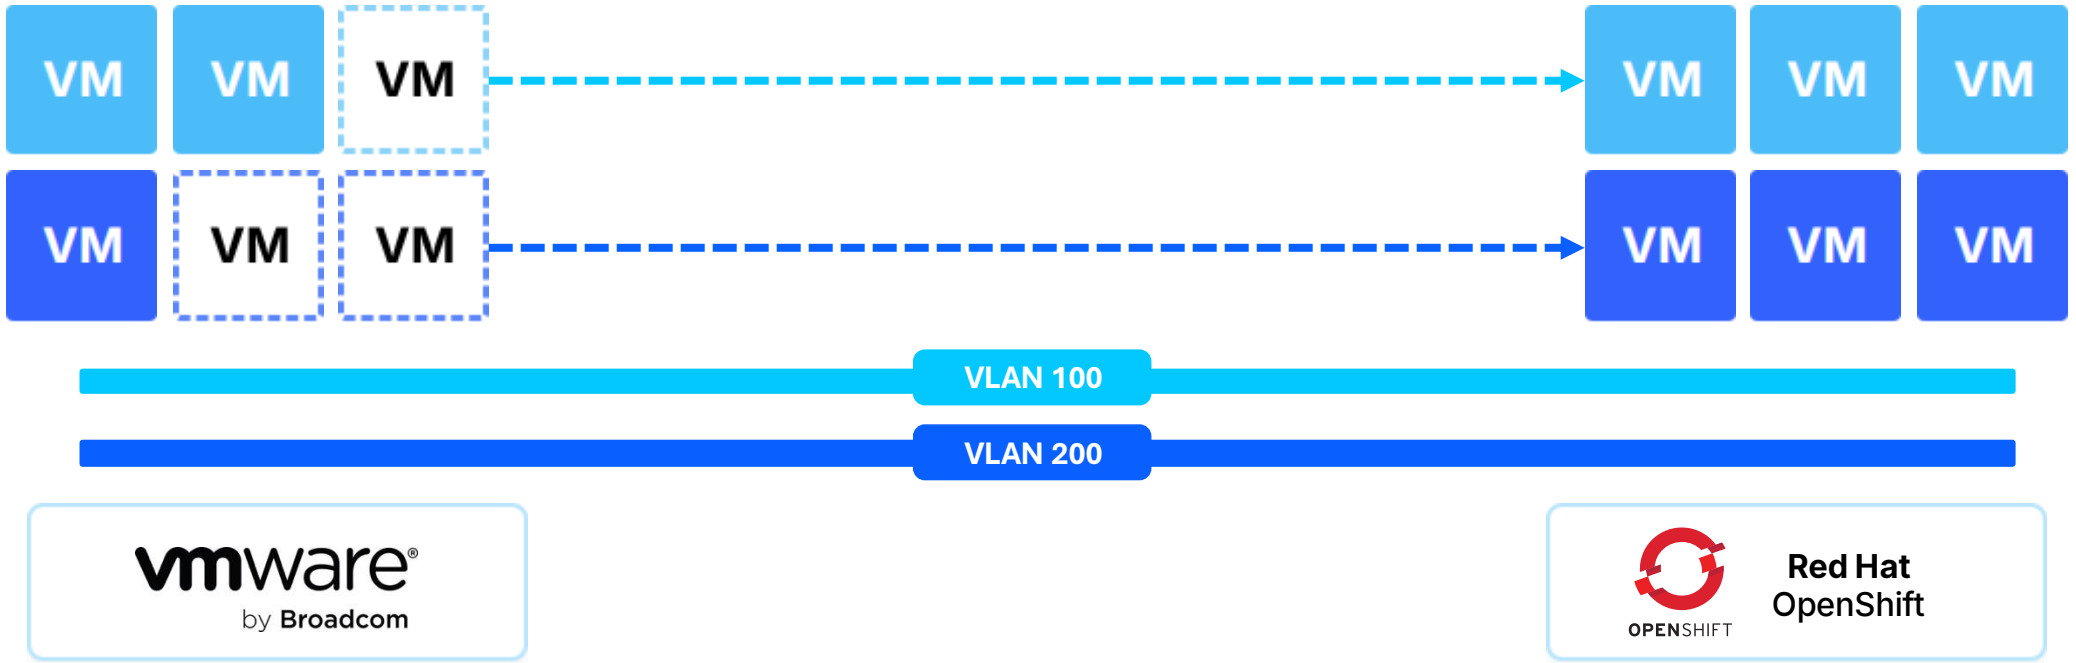

- VMs can be migrated to a VLAN-backed architecture in OpenShift Virtualization / KubeVirt that exposes the VMs to the exact same topology. Migrations proceed with IP preservation without having to deploy the Isovalent Network Bridge.
- This model allows the customer to implement Isovalent Network Policy for E-W micro-segmentation, something they do not get with just vSphere, plus deep visibility with Timescape, VPC segmentation with Private Networks, and all the benefits of Isovalent Networking for Virtualization.

© 2026 Cisco and/or its affiliates. All rights reserved.

Shift-on-Shift Current Support and Roadmap

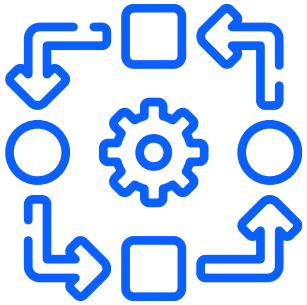

Supported

"Shift on Shift" K8s

Replace legacy VMware environments with OpenShift Virtualization or KubeVirt to establish a highly economical hypervisor layer. In this model, the VMs being provisioned are Kubernetes nodes themselves, running tenant K8s clusters on top of K8s-native virtualization.

Planned

Hosted Control Planes

In HyperShift the tenant K8s control plane runs purely as standard pods on the underlying bare-metal OpenShift layer, while only the tenant worker nodes run as VMs. This drastically reduces hypervisor tax, overhead, and cluster provisioning time.

© 2026 Cisco and/or its affiliates. All rights reserved.

License Mapping Red Hat OpenShift

OpenShift Platform Plus (OPP)
Full Platform

OPENSHIFT Platform Plus

- **Kubernetes** for containers
- **Virtualization** for virtual machines

OpenShift Container Platform (OCP)
Full Platform

OPENSHIFT Container Platform

- **Kubernetes** for containers
- **Virtualization** for virtual machines

OKE + Virtualization
Containers + Virtual Machines

OPENSHIFT Kubernetes Engine (OKE)

- **Kubernetes** for containers
- **Virtualization** for virtual machines

OVE (VMs Only)
Virtual Machines

OPENSHIFT Virtualization Engine (OVE)

- **For Virtual Machines** Only

Isovalent Networking for Kubernetes Advantage (INK)
License for containers and virtual machines

- **Consistent security** Apply and enforce policies across containers and virtual machines
- **Full-stack observability** Correlate and analyze telemetry across containers and virtual machines
- **Operational efficiency** Streamline operations and reduce complexity at scale

ISOVALENT **Networking for Kubernetes**

Isovalent Networking for Virtualization (INV)
License for virtual machines only

- **Advanced networking for virtual machines** High-performance connectivity, network security, and visibility built for virtual machines

ISOVALENT **Networking for Virtualization**

© 2026 Cisco and/or its affiliates. All rights reserved.

Isovalent Networking for Virtualization Licensing

Node Types	
Standard Node	A physical or virtual host running workloads on the Isovalent platform, capped at 16 cores / 64 GB RAM. Larger hosts are counted as multiple Standard Nodes based on the higher of the core or RAM ratio.
NEW! Virtualization Node	A bare-metal host running only VM workloads (no containers). Uncapped for CPU/RAM. If VMs and containers run on the same host, Standard Node sizing applies instead.

Isovalent Networking for Virtualization Offering

- Billed at 30 Units per Virtualization Node, per year.
- Available exclusively in the Advantage tier.
- Included automatically when Kubernetes Networking Advantage runs on Standard Nodes.

© 2026 Cisco and/or its affiliates. All rights reserved.

Isovalent Networking for Virtualization Sizing and Pricing

Environment — Obtain RVtools or similar CMDB export

Inventory — Count the number of physical ESXi hosts (from vHost tab or vInfo tab)

Sizing — ESXi host = Virtualization Node
(assumes same VM/host density and 100% of the VMs in-scope for migration)

Pricing — Allocate 30 units per Virtualization Node for ACV (list)

© 2026 Cisco and/or its affiliates. All rights reserved.

HSBC selects Isovalent Networking for Virtualisation to underpin HSBC's Internal Cloud (HIC) 2.0

Setting the Scene

HSBC were being strong armed into purchasing VMware subscription licenses to retain access to VMware's latest offerings, patches & upgrades, despite owning perpetual licenses and maintaining active support contracts for those licenses.

HSBC made the decision to not pay for VMware licenses twice, and in turn wholesale migrate entirely HSBC's internal cloud platform (presently 80/20 in terms of 'naked vSphere' versus NSX backed) to Red Hat Open Shift Virtualisation.

Through lengthy early access and design partnership, Isovalent established themselves as a key foundational capability offering HSBC; Layer 1 → Layer 7 Reference Architecture & Best Practice, Enhanced Enterprise Grade Connectivity, Security & Observability, above & beyond what is provided out of the box with native Red Hat tooling. Network Bridging & VM IP Preservation was descoped in late stages of the pursuit as HSBC needed to further accelerate their migrations, by +12 months, following Broadcom skullduggery (10 Commitment / $1.7B (+100% versus old contract)).

Isovalent's Proposal

1. **Full Isovalent Enterprise Platform capabilities available in the platform today, and throughout contract duration**
2. **Roadmap alignment aligning Isovalent's pre-existing roadmap with HSBC's project deliverables & milestones**
3. **Software entitlement up to ~140,000 CPU Cores / 5 Years / $20,900,000**

Phased Adoption

To meet HSBC's rapidly accelerated migration timelines, HSBC will deploy an initial phase of HIC 2.0 with native Open Shift tooling to deliver a simplistic platform to repatriate VMs off of VMware onto Open Shift Virtualisation.

In parallel to this Isovalent are engaging with HSBC in an 'enablement' phase of work, to build labs, further design partnership, educate HSBC Teams on the Isovalent based approach.

As Isovalent is further embedded in HSBC, Isovalent Networking for Virtualisation will be deployed across HIC 2.0, anticipated rollout in CY26Q4 timeframe.

NOT A PUBLIC REFERENCE
NOT FOR EXTERNAL DISTRIBUTION
CISCO HIGHLY CONFIDENTIAL

© 2026 Cisco and/or its affiliates. All rights reserved.

Appendix

Roadmap Road to General Availability

2025 Aug-Nov — Shipped

CONTROLLED AVAILABILITY 1-4

- Support for VSS/VDS/NSX-backed networking
- Support for multiple VMware networks, including dot1q support
- Basic Cilium Network Policy and Timescape support
- Enhanced MTV integration and automation
- Enhanced Bridge and Timescape observability
- Optimized routing at the destination (East-West)
- Scale and performance validation (1st round)
- Product documentation – Preview

2025 December — Shipped

CONTROLLED AVAILABILITY 5

- High Availability (HA)
- Simpler Lifecycle Management (LCM) – Clife
- Enhanced status reporting for PrivNets
- Enhanced product documentation

2026 April-May — Shipped

CONTROLLED AVAILABILITY 6-7

- Local Access support (L3) - Preview
- EVPN integration (L3 VNI) – Preview
- Scale and performance validation (2nd round)
- Enhanced product documentation

General Availability (IEP 1.19)

Shipped

- Distributed Local Access (VLAN)
- Distributed Fabric Access (EVPN L3 VNI)
- Multi-homed VM support (multiple vNICs)
- Integration with Isovalent Load Balancer (ILB)
- IPAM integration (DHCP)
- Multi-AZ Isovalent Private Networks
- NSX DFW policy migration (3P solution)

© 2026 Cisco and/or its affiliates. All rights reserved.

Roadmap items and timelines subject to change.

Cisco Confidential. Shared under NDA.

High Level Roadmap

1.20 and Beyond

2026 - 2027

1.20 and beyond

- **Multi-cluster Private Networks**
- **ConnTrack state preservation across Live Migrations**
- K8S IPAM support in Private Networks
- BGP support in Local Access
- L7 for Isovalent Network Policy
- VPC access: SRv6, VPC peering.
- K8s-aware Policy-Based Routing
- Partnerships: SpectroCloud
- EVPN support in Cat9K
- Project Helix integration (workload mobility solution)

© 2026 Cisco and/or its affiliates. All rights reserved.

Roadmap items and timelines subject to change.

Cisco Confidential. Shared under NDA.

Isovalent Networking for Virtualization Vision

Mid-term (2027) and long-term plans (2027+)

Compute

- Isovalent Private Networks (VPC) become the default connectivity option for all workload form-factors in K8s (VMs, Containers, Serverless).
- Expand support for Compute at the Edge, Campus, and in native Public Cloud (including Tech Partners)

Migration

- Integration with Project Helix (workload mobility solution).
- Public Cloud use cases.
- 100,000 VM migration validation.

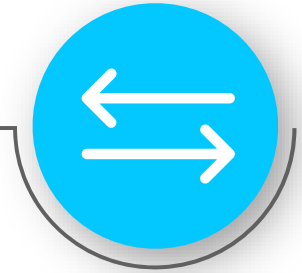

Security and Observability

- Support for the full range of IEP L3-L7 network security and observability capabilities.
- Runtime security support for Virtual Machines (universal agent).
- Deeper Splunk, Nexus Dashboard, and **Cisco Cloud Control** integration

Commercial

- Isovalent Private Networks (VPC) bundled in every DC, Edge, and AI CVD.
- Joint GTM campaigns with Red Hat.
- Increase market share 10X.

Cisco Better Together

- EVPN/SRv6 integration with in Cat9K, SONiC 8K, Hyperfabric.
- IPN + ISE/pxGrid integration for E2E identity (Campus, SDA, SD-WAN).
- DC, AI Pod, and Secure AI Factory compatibility and validation (CVD).

© 2026 Cisco and/or its affiliates. All rights reserved.

Roadmap items and timelines subject to change.

Cisco Confidential. Shared under NDA.

Common Objections - Business

Objection	Response
Is there really a TCO reduction play?	The destination stack carries real costs: OpenShift Virtualization, INV, and a migration partner engagement. The case depends on your VMware renewal reality. Broadcom's post-acquisition pricing has reset the break-even math for many customers. Where INV changes the equation is migration risk: automated, progressive network cutover reduces professional services exposure and eliminates the costly outage windows that inflate migration budgets.
Kubernetes is not ready for production VMs	KubeVirt has been running at scale in production since 2022 and ships as a fully supported component of OpenShift Virtualization. The networking layer, historically the most cited gap, is solved. Live migration, static IP preservation, multi-NIC support, and tenant isolation are production-ready in INV today, validated through POC engagements running since August 2025.
We have compliance requirements built on VMware micro-segmentation	INV's identity-based policy model provides L3/L4 (GA) through L7 (post-GA) enforcement that survives pod reschedules and live migrations without manual intervention. NSX-T objects are mapped directly: security groups become label selectors, DFW rules become Isovalent Network Policies. The migration of the policy framework is a defined project, not an open risk.
Network migrations always cost more than the vendor says.	That is why INV's architecture avoids a hard cutover entirely. VMs remain reachable on existing IPs throughout migration. Traffic shifts progressively to Kubernetes via EVPN host-route advertisement (for example) as workloads move, with the source environment remaining live until deliberate decommission. The model has been validated by a 3rd party and verifiable in a structured POC against your own environment.
We are already invested in another CNI	A CNI replacement requires planning and a maintenance window . That is not trivial. The justification depends on your workload profile. If you are adding VMs to Kubernetes, need hard multi-tenant isolation, or are executing a VMware migration, a general-purpose CNI will reach its limits. INV is purpose-built for that intersection. The migration path from OVN-Kubernetes is documented and field-tested on OpenShift 4.19 and 4.20.

© 2026 Cisco and/or its affiliates. All rights reserved.

NOT FOR EXTERNAL DISTRIBUTION

Thank you!