CLOUDIFICATION'S

C12N.CLOUD

OPEN SOURCE
PRIVATE & EDGE
CLOUD PLATFORM

openstack. POWERED

cloudification.io/c12n-private-cloud
sales@cloudification.io

Nov 11, 2026

VMWARE TO OPENSTACK MIGRATION GUIDE

A Practical Roadmap for Moving from VMware to c12n.cloud

Introduction: Your Migration Roadmap

VMware's shift to subscription-based licensing has prompted many infrastructure teams to reassess their long-term virtualization strategy. OpenStack provides an open-source foundation to build private cloud infrastructure that allows organizations to retain compatible hardware and support existing workloads.

A successful migration involves more than transferring virtual machines. The dependencies of supported applications, storage, networking, availability requirements, and operational processes all influence the migration approach.

This guide outlines the main stages of a VMware-to-OpenStack migration, from assessing the existing environment and selecting migration tools to preparing c12n.cloud, moving workloads, and validating the target platform. It is intended for infrastructure teams, technical leads, and IT decision-makers evaluating a transition from VMware to OpenStack-based cloud such as c12n.

Migration at a glance: Infrastructure blueprint analysis -> Dependency mapping -> Target architecture -> Environment preparation -> Parallel deployment -> Workload migration -> Performance and failover testing -> Post-migration support.

The execution activities for workload migration and acceptance are covered in detail in Chapter 05.

CONTENTS

Introduction 1

Contents 2

VMware to OpenStack: What Changes? 3

Assessing Your VMware Environment 4

Choosing a Migration Strategy and Tools 5

Preparing Your c12n Private Cloud 7

Migration Execution Playbook 8

VMware Migration Economics 10

Migrating with Cloudification 11

Sources and Further Reading 12

Appendix A. VMware Migration Readiness Checklist. 13

Appendix B. Cutover, Acceptance & Rollback Checklist 14

Why Cloudification? 15

01. VMware to OpenStack: What Changes?

VMware vSphere is built around ESXi and vCenter, with products such as NSX and vSAN providing additional functions. OpenStack takes a cloud-native approach instead, integrating compute, networking, storage, and identity management through modular services. In c12n.cloud, KVM runs the virtual machines, while OpenStack provides the management layer.

As a result, the target environment works with cloud-native resources such as instances, images, flavors, volumes, networks, and projects. Teams can administer these resources through dashboard UI (Horizon), Standardized APIs, and Infrastructure as Code (IaC) tools such as Ansible or Terraform (OpenTofu). Mapping VMware features to their OpenStack equivalents is only part of the migration work: existing operational workflows need to be reviewed as well.

The table below summarizes how the main VMware capabilities map to OpenStack and c12n.

VMware capability	OpenStack / c12n option	Migration consideration
ESXi and vCenter	KVM; Nova, Horizon, Keystone	Prepare guest OS drivers and adapt administration
vSAN	Ceph or a supported Cinder storage backend (30+ vendors)	Reassess capacity, latency, storage policies, data movement and QoS
NSX networking	Neutron and additional network/security components such as FWaaS or VPNaaS	Map routes, segmentation, firewalls, and load balancing
vMotion	OpenStack live migration	Moves VMs within OpenStack; VMware-to-OpenStack transfer uses non-OpenStack migration tools
VMware HA	Masakari	VM recovery service available as an optional deployment
VMware DRS	Nova Scheduler / Placement; Watcher where deployed	Validate required placement and optimization behavior
VMware HCX	Tools such as Coriolis	Select a cross-platform transfer and cutover method

A feature-by-feature mapping is a starting point. The key question is how applications actually use these capabilities, including automated recovery, storage policies, network segmentation, identity integration, and observability. For each capability, record the required outcome and test it in the target environment.

Reuse opportunities. Existing x86 servers, network switches, external storage, guest operating systems, and selected IP ranges may be retained. Once its workloads are migrated, an ESXi host can be rebuilt for KVM. VMware datastores and virtual networking definitions, however, need to be converted or reconfigured for the new platform based on Ceph SDS and Neutron with OVN respectively.

Core c12n base services include Nova, Neutron, Cinder, Glance, Keystone, Placement, Horizon, and S3/Swift. Additional services such as Designate, Masakari, Manila, Octavia, Cyborg, Barbican, Ironic and Watcher can be integrated when workload and operational requirements call for them.

02. Assessing Your VMware Environment

Begin with an inventory of the existing VMware environment. Identify workloads and dependencies, then compare allocated resources with actual utilization during normal and peak periods. This assessment establishes the migration scope, technical requirements, and acceptance criteria.

Inventory and application dependencies

The inventory should cover five areas:

Area	What to capture
Virtual machines	OS and version, vCPU/RAM usage, BIOS/UEFI, disks, controllers, snapshots, special devices
Applications	Business owner, upstream/downstream services, databases, authentication, scheduled jobs
Network	VLANs, IPs, DNS, gateways, MTU, firewalls, load balancers, external connections
Storage	Datastore types, used capacity, IOPS, latency, change rates, shared disks, backup dependencies
Continuity	Criticality, planned downtime window, RTO, RPO, restore method, acceptance tests

Dependencies should be mapped at the application level. A web tier, application server, and database, for example, may require a coordinated migration window. If these components move separately, document the routes and service connections that must remain available between VMware and OpenStack.

Network, storage, and compatibility checks

Next, confirm the network identity each application needs after cutover. Some services can keep their IP or possibly even MAC address; others need new addressing and a scheduled DNS, routing, or load-balancer change. Review NSX rules and security boundaries before recreating network segments in Neutron.

For storage, compare usable capacity and performance rather than raw disk totals, and include each workload's write rate when estimating transfer time. Guest compatibility is worth checking early, particularly for legacy OS releases, shared-disk clusters, passthrough devices, and licensing tied to virtual hardware or MAC/IP addresses.

Recovery targets and ownership

For critical workloads, define three values: the permitted planned outage, the recovery time objective (RTO), and the recovery point objective (RPO). RTO is the time required to restore service after a failure; RPO defines the acceptable amount of data loss, measured in time.

Then assign clear ownership: The application owner should approve functional tests, while the infrastructure owner verifies resource and connectivity requirements.

Assessment output: a prioritized VM inventory, application dependency map, capacity baseline, compatibility exceptions, and proposed migration waves with named owners. These deliverables then inform tool selection and target architecture design.

03. Choosing a Migration Strategy and Tools

With the assessment complete, you can select a migration method for each workload based on its downtime window, data change rate, and technical constraints. For larger VMware environments, a combination of methods is often the most practical approach.

Cold or warm migration?

Method	How it works	Best suited to
Cold migration	Stop the source VM, transfer/convert its disks, then start it in OpenStack	Smaller workloads and applications with a generous maintenance window
Warm / near-live migration	Transfer data while the source runs, synchronize changes, then cut over for minimum downtime	Larger workloads or tighter outage windows, where supported by the tool

Warm migration still requires a cutover window for final synchronization, application startup, network changes, and acceptance testing. Databases with high change rates may require application-aware coordination or separate database replication. Legacy guests or specialized clusters may be better rebuilt on new instances, with their application data migrated separately.

Organize migration waves

Start with a representative pilot. Then migrate lower-risk production environments before dependent and business-critical workloads. Group applications by dependency and scheduled business windows. If ESXi hardware will be repurposed for OpenStack, record which hosts can be released after each wave.

Compare available tools

Several tools can handle the transfer itself:

Tool	Main strength	Planning consideration
virt-v2v	Open-source guest and disk conversion to KVM	Suited to planned offline transfers; check supported guest and VMware access modes
MigrateKit	Open-source VMware-to-OpenStack migration with pre-copy capabilities	Validate source/target versions and operational requirements
Coriolis	Commercial migration and replication orchestration	Evaluate connectors, network/storage mappings, support, and project cost

Test the selected tool against your VMware versions, guest systems, and OpenStack storage configuration before using it for production workloads.
The pilot then provides reliable figures for transfer throughput, final synchronization time, guest adaptation effort, and application validation. Use these measurements to refine the remaining migration windows for subsequent waves.

04. Preparing Your c12n Private Cloud

Before production workloads move, the target environment needs to be built and tested. Plan capacity for both the first migration wave and any parallel VMware operation required during the transition.

Architecture and hardware

c12n.cloud combines OpenStack and KVM with Ceph or supported third-party storage, networking, automation, and observability. Cloudification's recommended production baseline uses six physical servers: three control-plane nodes and three hypervisor nodes. A separate bootstrap host is also required and may run as a VM. Storage topology and usable capacity are sized separately for the selected Ceph or external storage design.

For networking, Cloudification recommends 25 Gbit/s or faster for the data plane and at least 1 Gbit/s for management. Actual compute, storage, and network sizing follows from the assessment. If you plan to reuse VMware hosts, prepare enough initial c12n capacity to migrate a first group of VMs before rebuilding the freed ESXi servers for KVM. For larger clusters (30+ servers) Cloudification strongly recommends using dedicated physical storage networks and optionally own storage nodes.

Storage and networking

Ceph provides distributed block and object storage, while OpenStack can also integrate supported external arrays through Cinder and Manila drivers. For either option, confirm the storage driver's operations, expected latency, recovery behavior, and the conversion path from VMware disks. For Ceph, account for replication or erasure-coding overhead when determining usable capacity.

On the network side, configure the target networks, subnets, routes, security policies, and required load-balancing connections. Plan how applications will communicate across VMware and OpenStack during the phased migration, and document the IP or DNS changes required for each wave.

Operational readiness

Before the pilot, make sure teams can provision instances, attach volumes, access logs, and observe resource health. Credentials, monitoring integration, validate backup, and recovery procedures should also be configured by this point. c12n uses tools including Prometheus, Grafana, OpenSearch, and GitOps automation, while backup and disaster recovery may involve third-party solutions such as Storware or Trilio.

Finally, identify the optional OpenStack services required by the target architecture, confirm they are installed, and test the relevant functions, including HA or load balancing where applicable.

Target-ready criterion: A representative VM can be deployed, connected, monitored, backed up, restored, and tested on the intended compute and storage hardware.

05. Migration Execution Playbook

The six points below cover execution within Cloudification's broader eight-phase migration journey. Each activity has a defined objective and completion criterion, so the team knows when it is ready to move forward.

Prepare the source and target

Objective: Create an executable migration wave with known dependencies and a working recovery route.

- Freeze the wave inventory and map each source VM to target compute, storage, and network resources.
- Check VMware snapshots, disk layout, BIOS/UEFI, boot settings, and guest operating system support. Plan VMDK conversion as required by the target storage backend.
- Prepare guest drivers and operating system configurations. Windows commonly needs VirtIO storage/network drivers, while Linux may require initramfs and bootloader checks. Also review VMware Tools and virtual hardware-specific configuration.
- Confirm a usable backup and agree on application-level acceptance tests.
- The guest OS preparation can be done fully automatically when using migration tools.

Ready when: Every workload in the wave has an owner, migration method, recovery point, and target configuration.

Configure the migration path

Objective: Establish reliable connectivity between VMware and OpenStack.

- Grant the required vCenter/ESXi and OpenStack API access with appropriate scope.
- Configure source and target endpoints, storage mappings, target networks, and any migration workers.
- Check transfer bandwidth, expected data volume, and the impact on production storage traffic.

Ready when: A test transfer succeeds, and source/target permissions and mappings are documented.

Run the pilot

Objective: Verify that the selected workflow can migrate representative applications.

- Transfer Windows and Linux guests where both exist in the estate.
- Test boot, attached disks, network adapters, application transactions, and dependencies.
- Record transfer speed, final synchronization duration, issues, and actual downtime.

Ready when: Application owners approve the tests, and the wave schedule reflects the measured results.

Transfer workloads

Objective: Populate target instances while preserving application data consistency.

- For cold migration, shut down the source before transferring its disks.
- For warm migration, perform an initial copy while VMware remains active, then synchronize subsequent changes. This is normally done by migration tooling, but requires testing for specific guest OS versions.
- Reconcile the final dataset using the migration tool and, for transactional systems, any application-specific procedures.

Ready when: Target instances are staged for the agreed production window.

Cut over to OpenStack

Objective: Switch the production service within its approved maintenance window.

- Hold a Go/No-Go decision with the application and infrastructure owners, and verify the recovery point.
- Stop or restrict writes on the source, complete final synchronization, and start the target workload.
- Apply DNS, routing, IP, or load-balancer changes and run the agreed functional tests.
- Record acceptance or execute the approved recovery procedure before the decision deadline.

Ready when: The application owner accepts the migrated service, and the operations team can monitor its health.

Critical rollback rule. Once the target has accepted new production writes, the original VMware VM may no longer contain recent transactions. The recovery plan must therefore identify the authoritative dataset and specify how to preserve or restore those changes. Throughout cutover and recovery, prevent concurrent conflicting writes. Migration tooling simplifies the rollback by making it fully automatic.

Validate, stabilize, and hand over

Objective: Establish the migrated workload as an operable production service.

- Compare application response time, storage latency, and network behavior with the baseline.
- Verify backup restoration and relevant failure-recovery scenarios.
- Update inventory, monitoring, credentials, runbooks, and support ownership.
- Retain the source environment for the agreed stabilization period, and decommission it only after dependencies are cleared.

Ready when: The migration wave meets its acceptance criteria, and the next wave can proceed using the refined plan.

06. VMware Migration Economics

The following example illustrates the potential platform-cost difference for a VMware Cloud Foundation (VCF) environment with 15 servers. Each server contains two 24-core processors, resulting in 720 licensed cores. At an annual subscription rate of $350 per core, VMware software costs amount to $252,000 per year, including standard subscription support.

For c12n, the example assumes $40,000 for deployment, $16,000 for migration, and $7,000 for training, which brings one-time costs to $63,000. Technical support adds $54,000 per year. Both scenarios reuse existing compatible hardware.

Period	VMware VCF	c12n.cloud	Estimated savings
1 year	$252,000	$117,000	54%
3 years	$756,000	$225,000	70%
5 years	$1,260,000	$333,000	74%

Pricing note: This is an illustrative platform-cost scenario based on a publicly reported 2026 VCF pricing benchmark, not a customer quote. Actual rates and savings depend on the VMware product and contract, c12n project scope, staffing, and third-party services, so these items should be added to an organization-specific TCO model.
Source: AceMQ, VVF vs. VCF: Price per Core (updated October 2026).

To discuss a migration assessment and develop a cost model for your own infrastructure, visit the Cloudification VMware Alternative page.

07. Migrating with Cloudification

Case study: Building a private cloud for a FinTech company

An Asian FinTech company running VMware clusters needed an alternative infrastructure platform for its Linux and Windows workloads after changes in VMware licensing costs. Cloudification deployed c12n.cloud using OpenStack and Kubernetes, with Ceph storage, MAAS and Ansible for deployment, and Prometheus, Grafana, and OpenSearch for observability.

The deployment established a private cloud platform to replace the customer's existing VMware-based platform for these workloads.

From assessment to a migration plan

Cloudification provides OpenStack architecture consulting, private-cloud deployment, workload migration, and ongoing platform support. An initial engagement defines the migration scope: the VMware clusters and workloads involved, the connections that must remain available during the transition, and the services required from the target environment.

Responsibility boundaries. The infrastructure team owns target capacity, networking, storage, and platform access, while application owners validate business transactions, data consistency, and functional acceptance. Before production work begins, the migration plan assigns responsibility for the transfer, the cutover decision, and the recovery procedure.

Operations after handover. c12n configuration is maintained through GitOps in the customer's infrastructure. Organizations can operate the platform with their own team or engage Cloudification for ongoing support or managed private cloud operations. In either case, monitoring, incident response, upgrades, backup ownership, and escalation routes should be agreed as part of the project scope.

Discuss your migration scenario on the Cloudification VMware Alternative page.

Sources and Further Reading

- Cloudification, VMware Alternative
- Cloudification, c12n Private Cloud (architecture and FAQ)
- Cloudification, FinTech migration case
- Cloudification, Migration Tools
- OpenInfra Foundation, VMware to OpenStack Migration Guide

Appendix A. VMware Migration Readiness Checklist

Use this checklist during assessment and before approving the first production wave.
Project / environment: _____ Review owner: _____
Date: _____

Source estate and applications

□ Inventory VMware hosts and VMs, guest OS versions, virtual devices, disks, and snapshots.
□ Record measured resource utilization, storage change rates, and peak traffic.
□ Map application dependencies, business owners, and external integrations.
□ Define acceptable migration downtime, RTO, RPO, and application acceptance tests.
□ Identify unsupported guests, shared disks, passthrough hardware, and license constraints.

Target platform and strategy

□ Approve c12n capacity, storage, networks, identity, and required optional services.
□ Confirm parallel-operation capacity and a sequence for reusing freed ESXi hosts.
□ Test monitoring, backup restoration, and required recovery mechanisms.
□ Select cold/warm approaches and validate tool and guest compatibility.
□ Assign migration waves, decision owners, and escalation contacts.
□ Complete the representative pilot and approve the production cutover plan.

Outstanding blockers / decisions: _____

Appendix B. Cutover, Acceptance & Rollback Checklist

Use this checklist at the Go/No-Go meeting and during the migration window.
Migration wave: _____ Cutover window: _____
Decision owner: _____

Before cutover

□ Verify source health, target readiness, and required service connectivity.
□ Confirm the recovery point, approved outage window, and rollback decision owner.
□ Approve final-sync steps and the plan for stopping source writes.
□ Confirm who will change IP, DNS, routing, or load balancing and who will test the application.

Switch and accept

□ Stop or restrict source writes and complete the final consistent transfer.
□ Prevent conflicting production writes in source and target environments.
□ Start target instances, apply network changes, and verify critical user workflows.
□ Check transactions, integrations, performance, monitoring, and backups.
□ Record the application owner's acceptance or make the rollback decision by the deadline.

If a rollback is required

□ Identify the authoritative data source, including any new target-side writes.
□ Stop conflicting writes and follow the approved data recovery or reconciliation method.
□ Restore service, validate key application functions, and log the incident for the next wave.

Decision: □ Go □ No-Go □ Rollback Time: _____ Accepted by: _____
Follow-up actions / incident notes: _____

"OpenStack global footprint exceeds 45,000,000 compute cores" - OpenInfra Foundation

WHY CLOUDIFICATION?

We believe that the future of cloud infrastructure is **open**, **flexible**, and **cost-transparent**. The shift away from platforms like VMware is not just about avoiding rising costs - it reflects a broader demand for **autonomy**, **scalability**, and alignment with **modern development practices**.

Cloudification's c12n provides a compelling, enterprise-ready alternative that delivers on these expectations. Built upon a solid open-source foundation - OpenStack, Kubernetes, and Ceph - c12n stack gives organizations the ability to run a fully featured private cloud maintaining full control over the infrastructure. This means greater **cost efficiency, vendor independence,** and **seamless integration** with cloud-native tooling and modern DevOps workflows.

OpenInfra FOUNDATION

THE LINUX FOUNDATION

CLOUD NATIVE COMPUTING FOUNDATION

As active contributors to three major open-source foundations, we bring deep expertise in designing and operating OpenStack and Kubernetes environments for a wide range of use cases - including helping VMware customers migrate to vendor-neutral, license-free solutions.

c12n is already powering modern, scalable infrastructure in production at organizations across Europe, Asia and the USA. Whether you're planning a **migration from VMware**, building a **new sovereign cloud**, or simply seeking a **future-proof foundation** for your workloads, c12n offers the flexibility and transparency you need to move forward with confidence.

Request a Demo

Interested to see **c12n.cloud** in action? Drop us a message to sales@cloudification.io