

# Migrating VMs from VMware or Anywhere to KubeVirt and oVirt with Cirrus Data Solutions MigrateOps

## Executive Summary

Organizations moving away from VMware should not replace one platform dependency with another. They need a migration operating model that can move workloads from VMware, other hypervisors, public clouds, private clouds, and managed infrastructure to the platform that best fits the business. OpenShift/KubeVirt and OLVM/oVirt are strong destination choices for teams standardizing on open virtualization, Kubernetes-based operations, or Linux-centered private cloud. The migration layer must provide a practical, repeatable path to the selected platform.

Cirrus Data Solutions MigrateOps is a preferred solution because it treats migration as live workload mobility rather than a one-time export-and-import project. It brings continuous replication, destination-neutral compute transformation, performance assessment, intelligent quality of service (QoS), recipe-driven automation, approval gates, validation, rollback or reverse migration, and post-cutover data protection and disaster recovery (DR) into a single operating model. Its automation can also be used by AI agents to help plan and execute migration workflows.

## Best-Fit Use Cases

- VMware estate reduction or exit to OpenShift/KubeVirt or OLVM/oVirt or 12 other supported hypervisor destinations
- Migration from another hypervisor or cloud into an open virtualization private cloud
- Repatriation from public cloud to customer-owned virtualization
- Platform modernization where workloads move in waves, not all at once
- Enterprise applications with clustered nodes, shared disks, pass-through storage, or strict identity requirements

## Migration Process to OpenShift/KubeVirt and OLVM/oVirt

- 1 Discover source VMs, disks, networks, dependencies, performance, and owners.
- 2 Map each workload to the target platform: KubeVirt or oVirt VM spec, cluster, storage domain, network, boot mode, disk controller, and guest driver requirements.
- 3 Define the migration wave in a MigrateOps recipe, including the source inventory, target mappings, quality of service policy, validation checks, approval gates, and cutover rules. Add post-migration disaster-recovery configurations as required.
- 4 Synchronize workload data while production stays online.
- 5 Run preflight and transformation steps for the destination hypervisor.
- 6 Approve cutover, finalize synchronization, boot at the destination, validate the application, and quarantine or preserve the source according to policy.

## Demo Videos

For short videos, visit the Cirrus Data Solutions YouTube channel: [youtube.com/@CirrusData](https://www.youtube.com/@CirrusData)

For technical “See it in Action” demos, play these:

[RedHat OpenShift Virtualization \(KubeVirt\)](https://www.youtube.com/watch?v=W8AwXLlc4Zl)

( <https://www.youtube.com/watch?v=W8AwXLlc4Zl> )

[Oracle Linux Virtualization \(oVirt\)](https://www.youtube.com/watch?v=18RglBBTAw0)

( <https://www.youtube.com/watch?v=18RglBBTAw0> )

## Why CDS MigrateOps: The Cirrus Data Difference

**Destination-neutral mobility:** With MigrateOps, any destination can be chosen based on cost, compliance, performance, and operational requirements—not migration-tool limitations.

**Live migration model:** Block-level synchronization runs while the source workload remains online, so the outage window is limited to final consistency, destination activation, and validation.

**Intelligent Quality of Service (iQoS):** Migration can automatically use idle bandwidth and back off when production demand rises, optimizing the balance between migration speed and application impact.

**Assessment-based placement:** Observed throughput, IOPS, latency, CPU, memory, burst behavior, and change rate automate the target sizing and storage-class selection by AI agents.

**Automated compute transformation:** Boot, driver, controller, network, guest customization, IP preservation or reconfiguration, OS modernization, and application modifications are handled as planned workflow steps instead of post-cutover scramble.

**Minimal post-migration application remediation:** The workflow accounts for IP addresses, network interface names, disk paths, hostnames, cluster identifiers, device ordering, storage labels, and mount points. Preserving interface names is particularly important for Oracle RAC environments.

**Cluster-aware migration:** Cluster roles can move, fail over, or fail back during migration, and the target model supports migrating all nodes instead of pinning

a clustered application into a single node throughout the migration.

**Complex storage support:** Shared disks, guest-visible block devices, raw or pass-through mappings, application-owned paths, and clustered storage dependencies are all supported.

**Block, file, and object coverage:** VM disks, NAS shares, CIFS/NFS exports, object repositories, and S3-compatible data can be migrated with a single tool.

**Two-way mobility:** Reverse migration and repatriation reduce the risk of switching platforms by providing a path back if needed.

**Protection after cutover:** The same platform usage continues into data protection and DR, avoiding a gap where workloads land on OpenShift/KubeVirt or OLVm/oVirt or 12 other supported hypervisor destinations before backup and DR are ready.

**Flexible deployment:** Cloud-service and the Cirrus Data Private Edition models support connected environments, dark sites, sovereign environments, air-gapped sites, and policy-restricted datacenters.

**Automation for scale:** Declarative recipes, YAML configurations, CLI, REST API, MCP, agent skills, structured output, and stable IDs make migration waves reviewable, repeatable, auditable, and AI-operable.

**Operational control:** Preflight, assessment, placement approval, synchronization, QoS selection, cutover approval, final sync, source shutdown, destination boot, and validation are explicit gates.

### The Bottom Line for Business

MigrateOps reduces migration risk by giving teams a consistent, repeatable migration practice across multiple destinations. It can mitigate outage windows, reduce manual remediation, preserve production performance, support enterprise storage and clustered applications, and avoid double-run cost by migrating on a predictable schedule. MigrateOps can also incorporate data protection and disaster recovery planning into the target operational design. For VMware alternatives such as OpenShift/KubeVirt and OLVm/oVirt, MigrateOps provides a practical path to open virtualization without requiring a different migration toolchain for each source or destination.

#### ABOUT CIRRUS DATA

Cirrus Data Solutions Inc. (CDS) is a global leader in automated data mobility technology and services. CDS distributes its solutions through systems integrators, managed service providers, channel resellers, and partners, including Dell Technologies, Microsoft, AWS, and many others. CDS works closely with Global Systems Integrators (GSIs), including Accenture, HCL, and others. The company's flagship data mobility-as-a-service offering, Cirrus Migrate Cloud, is also available on the Microsoft Azure Marketplace and Amazon Web Services (AWS) Marketplace. CDS is headquartered in Syosset, New York. It also has sales and support offices in Atlanta, Chicago, New York, Dallas, Denver, London, Melbourne, Munich, and Tampa.

Cirrus Data, the Cirrus Data logo, Cirrus Migrate Cloud, Cirrus Migrate On-Premises, Cirrus Data Cloud, MigrateOps, cMotion, iQoS, and TDI are trademarks of Cirrus Data Solutions, Inc. All other trademarks are the property of their respective owners.



**Cirrus Data Headquarters**  
6800 Jericho Turnpike Suite 213W  
Syosset, NY 11791  
United States

[info@cirrusdata.com](mailto:info@cirrusdata.com) | [www.cirrusdata.com](http://www.cirrusdata.com)

1.516.888.3397