AUGUST 2024

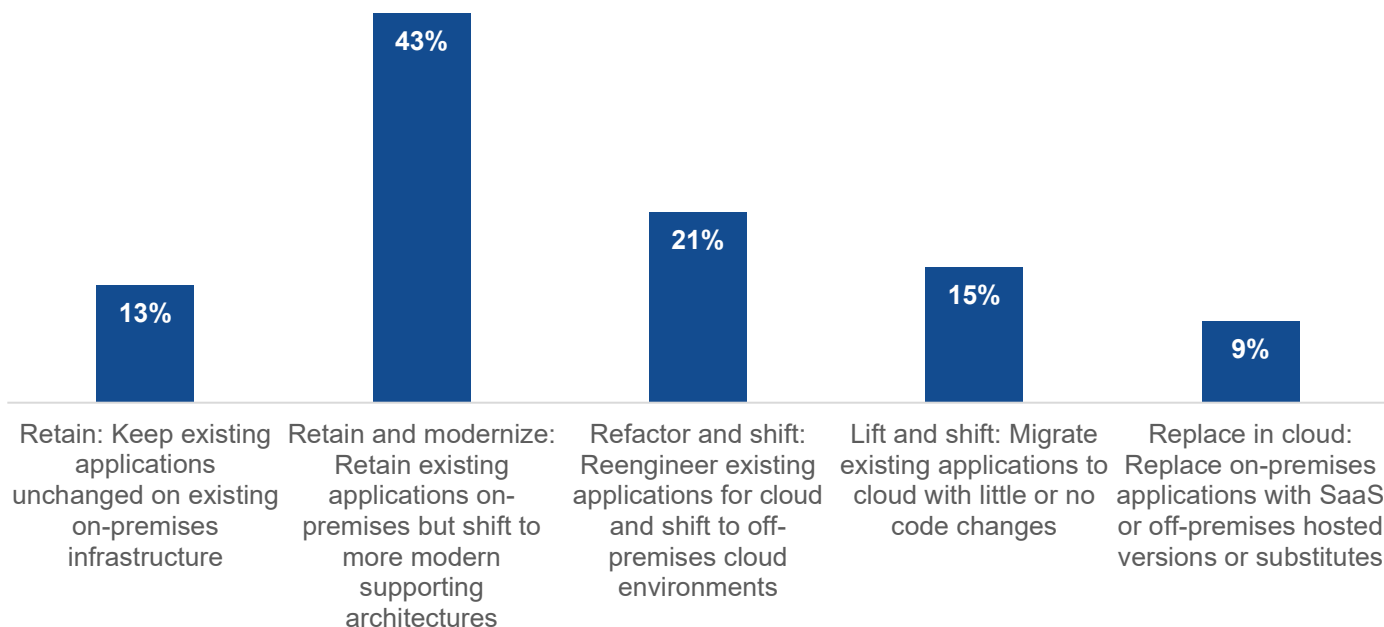# Red Hat OpenShift Virtualization Expands Enterprise Capabilities

Torsten Volk, Principal Analyst

Overview

The ability to deploy, run, and manage containers and virtual machines (VMs) side-by-side can enable developers to quickly add cloud-native capabilities to traditional virtualized apps. Adding containerized microservices and cloud-native components to existing web apps facilitates gradual modernization of legacy apps, instead of forcing organizations into risky mass-migration projects. This unified approach toward managing hybrid applications via a single platform and set of tools is the foundation for consistent, efficient, secure, and compliant deployment and management, as it avoids silos between traditional and cloud-native applications.[1]

Figure 1. Organizations want to be able to retain and modernize their application workloads without the pressure and risk of shifting to a microservices architecture immediately.

Which of the following best describes your organization's most common strategy for its existing on-premises applications when migrating to public cloud? (Percent of respondents, N=350)

Retain: Keep existing applications unchanged on existing on-premises infrastructure	Retain and modernize: Retain existing applications on-premises but shift to more modern supporting architectures	Refactor and shift: Reengineer existing applications for cloud and shift to off-premises cloud environments	Lift and shift: Migrate existing applications to cloud with little or no code changes	Replace in cloud: Replace on-premises applications with SaaS or off-premises hosted versions or substitutes
13%	43%	21%	15%	9%

---

[1] Source: All references and charts in this Brief, unless otherwise noted, are from Enterprise Strategy Group Complete Survey Results, *Multi-cloud Application Deployment and Decision Making*, May 2023.

Research by TechTarget's Enterprise Strategy Group shows that 43% of organizations prefer to retain existing applications on-premises but at the same time shift them to more modern architectures. This demonstrates how important it is to offer organizations a gradual, controlled, and hybrid approach to application modernization.

Red Hat's OpenShift 4.16 adds several strong enterprise-grade features that aim to enable OpenShift Virtualization to allow organizations to modernize at their own pace by running legacy apps and modern microservices apps side-by-side on the same platform.

New Enterprise Grade Virtualization Capabilities Added in OpenShift 4.16

Unified Disaster Recovery as Part of GitOps

OpenShift now supports synchronous data replication and automated failover between geographically close locations for containers and VMs. Red Hat promises that this new capability covers the vast majority of common VM setups, but admins should still verify the proper replication and failover for their specific workloads.

OpenShift 4.16 also brings a tech preview of the policy-driven definition and management of disaster recovery via declarative code statements. This enables DevOps teams to define the desired state of the entire cluster of VMs and containers to ensure automated recovery and consistency across regions.

Performance and Scalability

OpenShift Virtualization now allows the addition of virtual CPUs to VMs during runtime (without the need for a reboot). In combination with the tech preview of memory overcommit without negative performance impact, this enables great VM density and more efficient hardware utilization.

VM live migration now adheres to affinity rules to enable DevOps teams to control which hosts VMs can move while maintaining the desired performance and isolation. Certain workloads may perform better when running on the same hosts, while others may need physical separation for performance or security reasons, or they may require the presence of GPUs on their host.

The new ability to run latency-sensitive applications that require consistent performance opens use cases in the areas of telecommunications networking, financial trading, and industry control systems. This is achieved by OpenShift's ability to talk directly to the Linux kernel to prioritize CPU cycles based on workload requirements.

Unified Multi-Cluster VM-Monitoring

Collecting VM metrics and health data from different clusters in different locations and then aggregating them is a key capability required for consistent, and therefore scalable, VM management. This enables the IT team to monitor all VMs across multiple OpenShift clusters from a single dashboard and to centrally resolve issues across all VMs, from one central interface. For example, an admin may be responsible for 50 OpenShift clusters across different regions but will receive a unified dashboard to look at all VMs by health status, resource utilization, or performance, and define alerts depending on these metrics and the type of VMs in question.

Server, Storage, and Network Integration

OpenShift 4.16 adds several key integrations with hardware platforms by vendors like Dell, HPE, Lenovo, and IBM. These integrations enable VMs running on OpenShift to take advantage of these platforms' specific capabilities like snapshots, cloning, network policies, quality of service controls, and GPU-based AI training and inference. Therefore, these integrations enable enterprises to leverage past investments in specialized hardware and infrastructure, while modernizing applications on a modern Kubernetes-based platform.

One Fish, Two Fish, Redfish

The new Redfish integration is exciting to see as it demonstrates the significant role this standard can finally play when it comes to enabling scalability of cloud-native apps. OpenShift can now evacuate workloads based on Redfish alerting of problems of the underlying server host. It can also update the BIOS, GPUs, network interface cards, and DPUs of hosts to specific versions, attach "data images" to allow OpenShift to update device drivers and manage the BIOS. This simplifies the deployment and management of OpenShift deployments on bare metal (at the edge).

Simplified Migration

Ansible Migration Factory is a reference implementation of how to leverage Ansible Automation Platform for determining the requirements of virtualized workloads and for the subsequent automated VM migration to OpenShift Virtualization. This provides DevOps teams with the ability to consistently onboard VM workloads, based on predefined criteria and with minimal manual effort.

Conclusion

This release shows that Red Hat is clearly serious about capturing any and all VM-based workloads that fall by the wayside during the transition from VMware to Broadcom. But is OpenShift ready to take on all enterprise workloads? Not yet, but the platform provides enterprises with an option to gradually consolidate their application portfolio, based on the requirements of individual applications, while Red Hat is hard at work on getting OpenShift Virtualization completely ready for prime time.