IDC

SPOTLIGHT

Sponsored by: Red Hat and Portworx

This Spotlight examines the evolution of enterprise virtualization, from modern control planes to unified VM and container platforms.

Virtualization at a Crossroads: Building a Modern Virtualization Platform

November 2025

Written by: Gary Chen, Research Director, Software-Defined Compute

AT A GLANCE

KEY TAKEAWAYS

- Enterprise interest is growing for unified platforms that combine virtualization and containers, simplifying infrastructure management and supporting modern workloads like AI.
- Modern virtualization control planes like KubeVirt can not only offer new features and operating models but also operate very differently from legacy systems.

Introduction

Today's virtualization market is undergoing rapid changes amid shifting vendor dynamics. Enterprises are reevaluating their virtualization environments, often something they haven't done since they first deployed hypervisors 10–20 years ago. However, today's environments are much more complex, with vastly different factors for consideration:

- Infrastructure is moving to a cloud operating model where it is pooled, agile, software defined, and accessed programmatically by DevOps and developers.
- The use of containers and the Kubernetes platform for modern workloads is growing. While containers do not necessarily replace virtualization, they can place new and different demands on virtual infrastructure.
- A growing number of modern workloads are highly distributed and scalable and have a microservices architecture. Among modern tasks, AI workloads will play a significant role in new deployments and will come with their own set of unique characteristics, such as the use of accelerators.

As enterprises evaluate virtualization options, they must grapple with both the migration of existing workloads that have various levels of technical debt and the future fit for new and modern workloads. One option gaining interest is KubeVirt, an open source project that extends Kubernetes to the management of virtual machines (VMs). While Kubernetes is known as a platform for managing containers, it was designed to be extensible to manage any arbitrary resource. Several other projects extend Kubernetes to manage many different kinds of resources besides containers, with KubeVirt focusing on VMs. While KubeVirt enables Kubernetes to manage VMs, the VMs themselves run on KVM, an open source hypervisor that is part of Linux. KVM was integrated into the mainline Linux kernel in 2007 and has since been used in a wide variety of deployments, underpinning various virtualization offerings, OpenStack cloud software, and many public clouds.

Benefits

Using Kubernetes to manage VMs has drawn interest from enterprises for several reasons:

- A key requirement for users looking for a virtualization platform is long-term viability. Hypervisors are the foundation of modern datacenters, so users not only desire a technically capable virtualization platform but also one that will be around for the long haul. Containers have established themselves as the industry standard for modern compute operations, and Kubernetes has become the control plane for them. Kubernetes is widely adopted and has a large and active community, making it a safe bet that it will have a strong future.
- The entire KubeVirt stack, comprising KVM, Kubernetes, and KubeVirt, is open source, which can provide insurance against vendor lock-in. KVM and Kubernetes are available from a wide variety of vendors and service providers, with multiple vendors also offering KubeVirt.
- Kubernetes offers modernization benefits. Most virtualization infrastructures are many years or decades old and have naturally accumulated technical debt in some areas. Technical debt especially comes into play as the industry undergoes major technology shifts, such as cloud, containers, and AI. While KubeVirt vendors are improving legacy virtualization compatibility to make migrations of existing VMs as easy as possible, KubeVirt can also offer modernization capabilities for VMs, such as integrating with GitOps workflows and CI/CD pipelines, taking advantage of this and other inherent platform features, and a modern API that Kubernetes offers.
- Kubernetes offers efficiency and simplification with a unified platform for VMs and containers, simplifying management and attached subsystems, such as storage and networking.

Virtualization costs are complex and extend well beyond raw license costs. They can be highly dependent on individual scenarios, so this document cannot provide any definitive conclusions on this matter. Some key areas to consider when modeling costs are:

- The VM density and server utilization achievable using various CPU and memory virtualization techniques, especially under load, will affect the overall efficiency. The robustness of cluster scheduling and rebalancing features is also a key factor to consider. These can give users more confidence in running servers at higher utilization, knowing that appropriate remediation capabilities are in place to deal with load changes. These capabilities allow more VMs to be packed onto the hardware, which is important to consider for projected hardware and license costs.
- The sophistication of management tools in enabling more efficient infrastructure management is especially important as the scale grows. The style of management is also a vital consideration because enterprises are moving to cloud operating models that incorporate approaches such as DevOps and platform engineering. And last but not least is the cost of training — there is always a learning curve associated with any new platform, and this is especially true when transitioning to new operating models.
- Organizations must consider the costs and risks of downtime and security breaches. Many factors, beyond just the technical features of the product, can influence these aspects, including the maturity of the third-party tool ecosystem and experience with operating the platform.

- Ripple effects on the rest of the stack, such as storage, networking, backup, and management software, could impact visualization costs. In some scenarios, these effects could lead to cost savings, such as when the platform includes built-in functionality that saves on having to purchase a third-party tool and integrate it. Other scenarios might add to the costs; for instance, having to swap tools or reintegrate an existing tool into a new platform. Migrating to platforms with a rich partner ecosystem can offer more opportunities to leverage existing investments, as some of these assets will be supported under the new and old platforms. Keep in mind that many cross-platform tools have different levels of support and functionality for different platforms.

Migrating from an existing platform obviously incurs its own costs; however, it can also lead to attached hardware and software expenses if those components also require changing. Aligning those refresh cycles and license terms to the migration event can have a significant impact on overall costs.

Trends

Migration Strategies

Most enterprise users are still evaluating and forming plans for their long-term VM strategies. Those considering virtualization diversification are generally looking at some form of dual platform operation, at least for the foreseeable future. The strategies they are considering fall into three primary approaches:

- **A conservative, migrate-nothing approach:** Highly risk-averse users can be reluctant to tinker with well-established systems that are operating effectively, so they prefer to leave most VM workloads running where they are. However, they would establish new platforms, such as a new hypervisor, containers, a unified VM and container platform, or public cloud for any new workloads, preventing growth on their existing VM platform. Eventually, applications on the legacy platform will become deprecated, which will shrink the platform slowly and naturally over time.

- **An aggressive migration:** This scenario occurs when organizations under budget pressures or want to transition as quickly as possible are willing to take more risks in terms of downtime, security, migration/refactoring costs, operational disruption, or reduced functionality to migrate faster. In some cases, they may face more up-front costs in exchange for future savings. For very large deployments, these migrations can still take some time, and it is possible that not everything can be realistically migrated, but companies aim to migrate as much as they can.

- **A prepare-and-see approach:** In this case, organizations want to implement an alternate virtualization platform today to get it operationalized in their organizations and gain experience with it. While this may be starting small, it gives the platform time to prove itself. The expectation is that these platforms will improve over time in terms of technical capabilities, robust migration tools, and overall market adoption. As a result, organizations will migrate more workloads when they reach certain milestones and comfort levels, rather than adhering to a set timetable.

VM and Container Convergence

Since server virtualization started, a new form of compute has emerged that is adding a novel wrinkle to future considerations: the container. Containers and Kubernetes have become the industry-standard platform for modern applications, and this has many implications for the future of server virtualization since these two technologies are intertwined. First, most modern containers run in VMs for security and manageability reasons. Second, there are many mixed-mode apps in the enterprise, where one part of the app is in a VM (e.g., the database), while other parts (such as

the front end) are in containers. Having VMs and containers in different silos makes deployment across these boundaries difficult, and operations that cross this boundary, such as monitoring/observability, are significantly more complicated. Thus the industry has been trending toward a convergence of VM and container platforms into a single control plane that can provide much better coordinated operations between VMs and containers. This can also simplify the infrastructure architecture, with unified and shared subsystems for storage, networking, security, and other areas. While there are many different approaches to converging VMs and containers, the fact is that both technologies will be around for the foreseeable future, and eliminating silos between them will be beneficial.

AI Workloads

As we enter the AI era, the number of AI workloads will grow extremely fast, driving new requirements across the entire stack, including VM and container platforms. While AI workloads are very modern and thus lean toward containers, AI will be grafted onto nearly everything in existence, including bare metal and VMs. AI will have many infrastructure implications, including increased demands on east-west networking, data management, efficient access to very large data sets, and the virtualization of accelerators for efficient use. The path of an early-stage technology is often unpredictable, so customers need to consider the adaptability and flexibility of vendors and platforms. With this rapid pace of innovation, AI platforms are evolving into highly flexible and pluggable frameworks that can accommodate a wide range of model serving, inference infrastructure, AI development environments, and templated, automated setups of common patterns.

Considering Red Hat and Portworx

Red Hat OpenShift Virtualization

Red Hat OpenShift Virtualization offers a modern way to deploy, manage, and scale VMs using the proven infrastructure of Red Hat OpenShift. Every OpenShift edition includes the KubeVirt-based virtualization capability, enabled through the OpenShift Virtualization operator. A special edition of OpenShift, OpenShift Virtualization Engine, offers virtualization-only capabilities for organizations targeting VM migration use cases. This edition can eliminate unnecessary complexity and cost for those who only want to manage VMs for now while ensuring a clear future path toward containers when needed. All OpenShift product editions share the same code base, with the restrictions only in the license; so upgrading to higher editions does not require any redeployment, only a license change. OpenShift Virtualization Engine also does not include any RHEL entitlements, reducing costs for companies that are migrating existing VMs that already have an OS license. Red Hat Advanced Cluster Management for Virtualization is also made available at a reduced cost for VM-only multicluster management.

Key features of OpenShift Virtualization include:

- Simplified support subscription model, supporting up to 128 cores per CPU socket pair
- A tailored web-based management console, specifically for virtualization administrators, that reduces the level of Kubernetes expertise required to manage VMs on OpenShift
- Support for key enterprise virtualization features such as memory overcommit, live migration, and storage live migration

- Support for deploying OpenShift Virtualization via Red Hat's OpenShift cloud partners, such as the Red Hat OpenShift Service on AWS, Microsoft Azure Red Hat OpenShift (public preview), OpenShift Dedicated on Google Cloud (technical preview), Red Hat OpenShift on IBM Cloud, and Oracle Cloud Infrastructure, enabling hybrid cloud VM deployments
- Integration with other key Red Hat solutions, including Red Hat Ansible Automation Platform for large-scale automation and Advanced Cluster Management for extended orchestration across clusters

In addition, Red Hat OpenShift Virtualization can leverage Red Hat's robust partner ecosystem for OpenShift to enhance capabilities ranging from storage and backup to disaster recovery (DR) and networking. The Migration Toolkit for Virtualization (MTV) is an included Red Hat tool that enables the migration of VMs into OpenShift Virtualization clusters.

As organizations look to the future, this foundation also makes it easier to adopt advanced capabilities that go beyond virtualization. By upgrading to the broader OpenShift application platform, organizations gain a future-ready foundation for gradual modernization and adoption of containers, AI/ML, and serverless technologies. Red Hat OpenShift AI expands the modernization journey by enabling teams to build, train, deploy, and monitor AI applications directly within the OpenShift platform, providing a consistent future-ready foundation.

How Portworx Storage Supports OpenShift Virtualization

Storage and data management is one of the most important subsystems in modern server virtualization, enabling key functions such as clustering, live migration, high availability, and DR. However, existing virtualization workloads may assume a certain storage infrastructure that may differ from many container storage assumptions. While Kubernetes has added many storage capabilities over the years, especially for stateful workloads, the storage system and the CSI driver (a Kubernetes standard interface for storage) determine much of the storage functionality.

Portworx Enterprise is a software-defined storage and data management platform specifically designed for Kubernetes environments. Portworx abstracts any storage system or cloud storage with a consistent set of enterprise-grade storage capabilities that integrate natively into Kubernetes via CSI. Portworx has been designed to support any application running on Kubernetes, from containers to AI/ML workloads, and now extending to VMs running on Kubernetes. Portworx has taken common legacy virtualization workflows and concepts like live migration, metro DR, and datastores and created similar workflows for VMs running on OpenShift Virtualization, including:

- Consistent storage capabilities across multiple storage systems and multicloud or hybrid cloud environments
- Support for very large stateful disks that are common for VM workloads
- RWX (i.e., ReadWriteMany) block volumes, key for enabling enterprise VM features such as live migration and high availability
- Support for key storage features for enterprise virtualization, such as storage live migration, storage resource balancing, and storage I/O control
- Thin provisioning for storage efficiency
- Volume encryption integrated with key management servers

- Kubernetes and OpenShift Virtualization VM-aware backups, including file-level backup for Linux guests with the Portworx Backup add-on
- Synchronous and asynchronous DR with the Portworx Disaster Recovery add-on

Portworx specifically validates its storage solutions for OpenShift Virtualization through joint lab testing with Red Hat. This assures customers that Portworx storage will work seamlessly with OpenShift Virtualization to solve specific storage challenges when running VMs with Kubernetes.

By leveraging the capabilities from Red Hat and Portworx, companies can accelerate a successful transition from traditional virtualization to a more modern approach that unifies containers, virtual machines, and AI workloads under a Kubernetes control plane.

Challenges

While Kubernetes is a very modern control plane and KubeVirt allows Kubernetes to treat VMs as native objects, virtualization is a mature and established platform in datacenters, with much inertia. This is a historical challenge with any platform transition, where some organizations take the path of least resistance, reluctant to touch systems that are not broken and simply keep revising them. Conservative users may not migrate even if there are cost or feature benefits and instead stick with the comfort and predictability that comes from the longtime operational experience of their existing platform.

Enterprises with large installed bases of VMs often assess KubeVirt and Red Hat OpenShift Virtualization by comparing the capabilities of these platforms with their current platform. Capabilities continue to evolve over time as they are adapted to the Kubernetes model and validated in enterprise environments. Because Kubernetes is designed differently from traditional virtualization platforms, certain functions may be implemented or experienced in different ways. These differences can reflect Kubernetes' strengths as a modern platform but also pose challenges when migrating from an existing VM environment.

Conclusion

With many decades of VMs in the install base, enterprises will need support for gradual migration and modernization to help preserve existing VM investments while transitioning to cloud-native architectures and agile workflows. Red Hat OpenShift Virtualization, in conjunction with Portworx storage, provides a modern, unified platform that bridges traditional VMs and containerized workloads, enabling organizations to simplify and modernize their infrastructure. To the extent that Red Hat and Portworx can address the challenges described in this paper, the companies have a significant opportunity for success in the virtualization market.

About the Analyst

Gary Chen, *Research Director, Software-Defined Compute*

Gary Chen is IDC's research director for Software-Defined Compute. His research focuses on server virtualization, container infrastructure and management, and cloud system software (system software used to build IaaS clouds, such as OpenStack).

MESSAGE FROM THE SPONSOR

Red Hat OpenShift Virtualization and Portworx offer a powerful and flexible solution for organizations seeking to modernize their virtualization strategies. By providing a unified platform for managing both VMs and containers, along with robust data management and migration tools, the solution helps organizations to streamline operations, protect critical data, and accelerate application deployment. With low-cost entry points like OpenShift Virtualization Engine and Portworx for OpenShift Virtualization, customers can take advantage of a cost-effective starting point for their VM migrations. This strategic combination empowers organizations to embrace new technologies while maintaining and optimizing their existing VM investments.

Ready to learn more?

red.ht/catlog-redhat-portworx

https://portworx.com/customers/siriusxm/

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com