Red Hat OpenShift Container Storage 4

Dynamic, shared, and highly available storage for OpenShift applications

Overview

Red Hat® OpenShift® Container Storage is persistent software-defined storage integrated with and optimized for Red Hat OpenShift Container Platform. It runs anywhere Red Hat OpenShift does: on-premise or in the public cloud. Built on Red Hat Ceph® Storage, the platform offers tightly integrated, persistent data services for Red Hat OpenShift and hybrid multicloud. Dynamic, stateful, and highly available container-native storage can be provisioned and deprovisioned on demand as an integral part of the Red Hat OpenShift administrator console.

Red Hat OpenShift Container Storage

Red Hat OpenShift Container Storage is engineered, tested, and qualified to provide data services for Red Hat OpenShift Container Platform on any infrastructure (Figure 1). OpenShift Container Storage runs as a Kubernetes service within Red Hat OpenShift and can also be decoupled and managed as a separate data store, delivering file, block, and object data for one or many OpenShift Container Platform clusters. Tightly bundled offerings including Red Hat OpenShift, Red Hat Middleware, and Red Hat OpenShift Container Storage can also help organizations automate data pipelines, enabling real-time data processing from edge to hybrid and multicloud platforms. This level of integration removes the guesswork from running Red Hat OpenShift across multiple platforms and provides the data storage functionality, data services, and data protection that enterprises require.

Key Benefits

- Agility, scalability, and consistency across hybrid cloud and multicloud environments, resulting in lower IT operations costs

- An integrated platform, including container host, Kubernetes, application life-cycle management, and storage, using any infrastructure

- Greater value from operations and development teams, with the ability to easily access data in container-based environments

- Data services that are validated, integrated, tested, and supported with Red Hat OpenShift Container Platform

- Simpler installation and upgrades, leading to faster application development cycles and more frequent software deployments
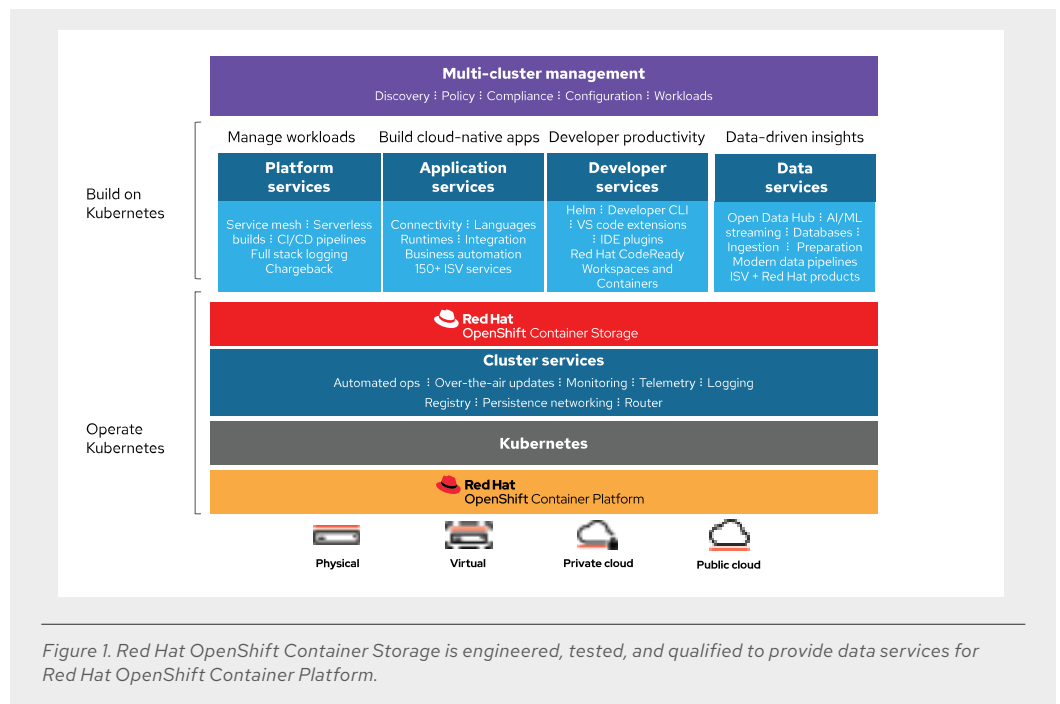
Figure 1. Red Hat OpenShift Container Storage is engineered, tested, and qualified to provide data services for Red Hat OpenShift Container Platform.

Developer-driven agility

To increase agility, organizations need to reduce complexity for cloud-based apps and data, allowing rapid and flexible deployment of application data across any cloud. Red Hat OpenShift Container Platform lets developers move quickly to deliver application programming interface (API)-driven persistent storage. Through container-native virtualization, both containers and virtual machines can be deployed on Red Hat OpenStack® Platform, fully supported by Red Hat OpenShift Container Storage. Red Hat OpenShift operators offer simplified storage management and data portability across public, private, and hybrid cloud.

Deterministic performance at scale

Moving to the cloud cannot come at the expense of application performance. Native object support in Red Hat OpenShift Container Storage dramatically increases input/output (I/O) performance, which increases performance for cloud-based workloads. Red Hat engineers work with partner organizations to ensure the platform provides deterministic performance at scale along with resiliency for a broad range of Red Hat OpenShift workloads, supporting Kubernetes pods, artificial intelligence/machine learning (AI/ML), and other popular workloads.

Consistent user experience

Because OpenShift Container Storage runs as a service in Red Hat OpenShift, users can benefit from a consistent user experience regardless of where the data resides. Whether data is being aggregated at edge, accessed in a core data center, in hybrid cloud, or in multiple public clouds, OpenShift Container Storage provides policy-based administration for scale and automation, providing deployment flexibility. Organizations can unlock the power of data in new and impactful ways while making data accessible, resilient, and actionable to applications. Following this practice, organizations can create a consistent open hybrid cloud application environment with intuitive services for intelligently moving, storing, transforming, responding to, and learning from enterprise data.

Benefit	Features and details
Agility	Fully integrates with Red Hat OpenShift Container Platform for Day 1 and Day 2 installation and management. A single unified platform supports: ▸ Block storage for databases and messaging. ▸ Shared file storage for continuous integration and data aggregation. ▸ Object storage for application data, data lakes, archival, backup, and media storage. Storage nodes are members of the Red Hat OpenShift cluster. OpenShift Container Storage nodes are managed through the Red Hat OpenShift administrator console.

Benefit	Features and details
Scalability	Supports all types of OpenShift workloads, allowing easy object data-sharing across geographic locations and platforms, and scales to orders of magnitude more persistent volumes (PVs) per OpenShift Container Storage cluster than previous releases.
Consistency	Delivers a unified end-user experience for data services across Red Hat platforms in the hybrid cloud, along with hybrid and multicloud data protection for enterprise applications. Consistent Red Hat OpenShift management tools work across environments, whether on-premise or in public cloud. The multicloud object gateway provides data federation across multiple private and public clouds.

Tight integration with Red Hat OpenShift Container Platform

Red Hat OpenShift Container Storage 4 is created for container-based environments and is tightly integrated with Red Hat OpenShift Container Platform.[1] With a supported Red Hat OpenShift operator, OpenShift Container Storage is simpler to install and manage as a part of the container-based application life cycle. With this innovation, Red Hat can provide support for the entire container-based environment, including cloud-native container management, scheduling, and orchestration, yielding:

▸ **Data resiliency for Red Hat OpenShift.** Enterprise applications require storage with enterprise capabilities. For a stateful app to exhibit high availability, its data must first be highly available. Red Hat OpenShift Container Platform supports important features like replication, allowing application data to be placed across different availability zones.

▸ **A cloud-like experience, everywhere.** Circumstances are constantly changing, favoring one cloud provider over another or in-house deployment versus public cloud. Organizations need the ability to move quickly to take advantage of favorable pricing or respond to other business pressures. Red Hat OpenShift Container Storage provides software-defined storage that lets organizations deploy their apps and storage as needs dictate and adjust as they move forward.

▸ **Increased developer productivity.** Cloud developers want to innovate without arbitrary limitations. Traditional storage often has been an impediment to cloud development, requiring separate and time-consuming arrangements. Red Hat OpenShift Container Storage provides common functionality across all cloud platforms, simplifying processes for developers.

Embracing diverse workloads

Most cloud providers support data storage for diverse workloads, but they typically do so with different storage technologies. Not only is this complexity time-consuming to understand and manage, but it can lock applications into a given cloud vendor because other vendors offer a different mix of storage technologies and capabilities. In contrast, Red Hat OpenShift Container Storage provides container-native storage that supports multiple workload types (Figure 2), and it does so across multiple cloud platforms.

---

1 *Refer to the latest Red Hat OpenShift Container Storage 4 release notes for supported platforms.*

Based on 100% open source technology, Red Hat OpenShift Container Storage supports file, block, and object storage and employs a single set of Kubernetes Operators across all cloud platforms. Organizations can support multiple workload types with a single software-defined storage solution, and applications can move easily between cloud platforms with compatibility.
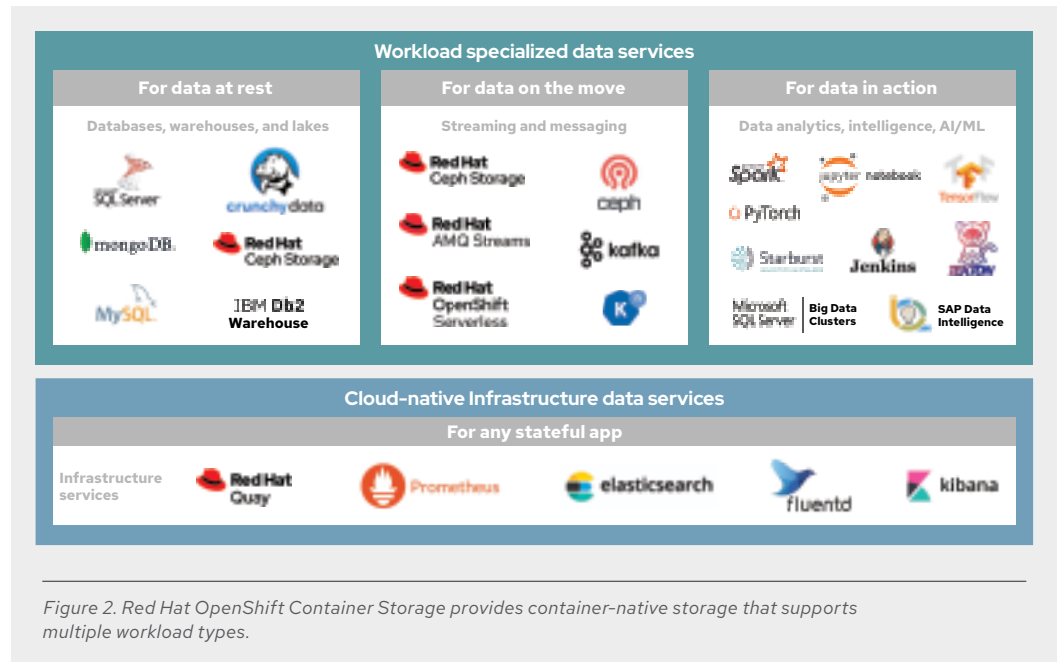

Figure 2. Red Hat OpenShift Container Storage provides container-native storage that supports multiple workload types.

Accelerating application development

Developer productivity depends on agile continuous integration/continuous delivery (CI/CD) pipelines and responsive infrastructure. To perform effectively, developers require self-service storage deployment access, freeing them from having to think about storage provisioning as a separate activity that often causes delays. With comprehensive support for Kubernetes, Red Hat OpenShift Container Storage automates the provisioning of storage alongside the provisioning of application resources, all within the OpenShift administrator console. Container-based tools like Jenkins can run on Red Hat OpenShift Container Storage to automate the CI/CD pipeline with improved performance as build dependencies don't need to become retrieved with each build cycle.

Supporting databases with deterministic performance

As databases have moved to container-based environments, the amount of stored data has grown, which has created an urgent need for performant container-based storage. Databases require persistent storage volumes to support data availability and failover scenarios. If one or more Red Hat OpenShift pods representing a database service move from one node to another, the data they use must remain persistent and available for access after the pods move. Red Hat OpenShift Container Storage provides persistent block storage for databases that supports the databases' availability needs and provides simplified portability across multiple cloud platforms.

Simplifying storage for data analytics

Data analytics is evolving rapidly, both in terms of traditional static data analysis as well as dynamic AI/ML environments. Like developers, data engineers and analysts don't have time to become expert storage managers or wait on time-consuming storage allocation from other groups within the enterprise. Red Hat OpenShift Container Storage lets data scientists, and those who support them, deploy and manage cloud-portable storage on demand. Data scientists can deploy the platforms they need without thinking about how data is stored or what is required to move datasets to other platforms.

Conclusion

Red Hat OpenShift Container Storage storage provides agile, scalable, and persistent data services anywhere that Red Hat OpenShift runs — across on-premise infrastructure, public cloud, or hybrid cloud. The platform offers tightly integrated persistent data services that can serve a broad range of workload types. Deploying Red Hat OpenShift Container Storage simplifies data management and allows storage to be provisioned and deprovisioned on demand as an integral part of orchestrated, container-based environments.

About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers integrate new and existing IT applications, develop cloud-native applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500. As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.

facebook.com/redhatinc
@RedHat
linkedin.com/company/red-hat

North America	**Europe, Middle East, and Africa**	**Asia Pacific**	**Latin America**
1 888 REDHAT1	00800 7334 2835	+65 6490 4200	+54 11 4329 7300
www.redhat.com	europe@redhat.com	apac@redhat.com	info-latam@redhat.com

redhat.com
#F25244_0920