

Data management for Red Hat OpenShift services and Red Hat Quay

Highlights:

- ▶ Provide highly available data services with storage designed for Red Hat OpenShift.
- ▶ Deploy block, file, and object storage with Red Hat OpenShift Data Foundation to serve common application needs and simplify operations.
- ▶ Launch Red Hat OpenShift services in-house, in the public cloud, or in hybrid environments, without changing your software-defined storage platform.
- ▶ Support Red Hat Quay for enhanced enterprise registry services, backed by highly available OpenShift Data Foundation block and object storage.

Persistent and highly available storage

Red Hat® OpenShift® services have become essential as organizations embrace containers and Kubernetes orchestration for essential applications. However, making Red Hat OpenShift services highly available requires cluster data management that provides persistent data storage, scalable performance, and resiliency to protect critical internal data from failures and outages. Red Hat OpenShift Platform Plus includes [Red Hat OpenShift Data Foundation](#), a persistent, software-defined storage platform integrated with and optimized for Red Hat OpenShift. OpenShift Data Foundation is ideal for all types of applications. It is also an ideal persistent storage platform for internal Red Hat OpenShift services, regardless of whether they are deployed in-house, in a cloud environment, or in hybrid environments.

Serving critical Red Hat OpenShift services

Various internal Red Hat OpenShift services maintain state, creating a need for persistent storage. Although each service has a default data storage method, alternative storage solutions can offer additional benefits. For example, some software-defined storage solutions do not support bare-metal or public cloud deployment options, artificially limiting flexibility. Moreover, while default storage solutions rightly focus on allowing the service to function, they may offer only ephemeral (non-persistent) data storage. Even with persistent storage, data services may not be resilient, causing potential issues during outages or failover scenarios. In a worst-case scenario, data service failure could result in the loss of state for a critical service. Essential internal Red Hat OpenShift services include:

- ▶ **Logging.** Logging is an important service of the Red Hat OpenShift monitoring stack. Default emptyDir volumes do not provide persistence for logging information. Red Hat recommends data persistence for the logging service to retain the logs of all cluster components and processes. Whether deployed in-house or in a public cloud environment, persistent storage can aggregate the logs from your Red Hat OpenShift cluster, including node system logs and application container logs.
- ▶ **Metrics.** Red Hat OpenShift collects metrics for almost everything that happens within an installation. Saving these metrics to persistent storage allows the retention of this vital information no matter what happens inside the cluster. The Prometheus stack deployed by default to collect metrics does not offer data persistence.
- ▶ **Registry.** The primary purpose of the built-in Red Hat OpenShift registry is to store and serve container images. A registry service requires back-end storage for this purpose. By default, Amazon Web Services (AWS) uses Amazon Simple Storage Service (S3). For all other infrastructure situations, a ReadWriteMany (RWX) persistent volume is the recommended and Red Hat-supported way to connect the registry to persistent storage. Alternatively, Red Hat Quay backed by OpenShift Data Foundation can also be used to add a suite of additional registry functionality for enterprise needs.

In most cases, OpenShift Data Foundation presents an ideal persistent data storage solution for Red Hat OpenShift services. Combining Ceph® storage, the Rook operator, and NooBaa Multicloud Object Gateway technology, the platform offers tightly integrated persistent data services for Red Hat OpenShift within any hybrid and multicloud environment. It supports:

- ▶ **A choice of storage types.** OpenShift Data Foundation supports all common types of storage needed for container-based applications, including block storage (ReadWriteOnce-RWO), file storage (RWX), and S3-compatible object storage.
- ▶ **Maximum deployment flexibility.** OpenShift Data Foundation can be deployed on-premise, in a public cloud environment, and for edge computing.
- ▶ **Resilience for business continuity.** OpenShift Data Foundation offers replication across multiple cloud provider availability zones, allowing for failover of Red Hat OpenShift services for enhanced business continuity for critical applications.

With Red Hat OpenShift, organizations can choose whether they want to use installer-provisioned infrastructure (IPI) or user-provisioned infrastructure (UPI). Table 1 lists default and recommended storage technology for several deployment infrastructure options.

Table 1. Supported Red Hat OpenShift infrastructure and service storage repository recommendations

	Supported Red Hat OpenShift infrastructure	Logging	Metrics	Registry
AWS	Default (IPI)	AWS Elastic Block Store gp2 storage class	Ephemeral	AWS S3
	Recommended	OpenShift Data Foundation (RWO)	OpenShift Data Foundation (RWO)	AWS S3
VMware	Default (UPI)	Virtual machine disk (VMDK)	Ephemeral	VMDK
	Recommended	OpenShift Data Foundation (RWO)	OpenShift Data Foundation (RWO)	OpenShift Data Foundation (RWX)
Bare metal	Default (UPI)	Local disk	Ephemeral	Local disk
	Recommended	OpenShift Data Foundation (RWO)	OpenShift Data Foundation (RWO)	OpenShift Data Foundation (RWX)

Using Red Hat Quay for the Red Hat OpenShift registry

The embedded image registry provided with Red Hat OpenShift offers basic private-registry functionality. Red Hat OpenShift Platform Plus includes [Red Hat Quay](#) and offers [many additional enterprise features](#) such as automation, security capabilities, high availability, speed, content distribution, and application programming interface (API) integrations. Red Hat Quay requires both a back-end store for container images and a database to catalog the available images.

- ▶ **Back-end store.** An image repository acts as a container image store. Red Hat Quay additionally offers support for artifacts and Helm charts. As such, Quay requires a storage back end to provide more durable and reliable storage for container images. The back-end storage also needs to be accessible in a shared manner to provide access to multiple cluster resources simultaneously. Object storage is mandatory for Quay's geo-replication capabilities. S3-compatible storage is the most common way to fulfill Quay's back-end storage needs.
- ▶ **Database.** Quay additionally employs an external database that must be highly available. The preferred database back end is PostgreSQL, which can also serve Clair, Quay's image security scanning component. Providing high availability for PostgreSQL can be accomplished through PostgreSQL directly or through replication at the storage layer using OpenShift Data Foundation.

Because OpenShift Data Foundation supports block, file, and object storage, it is an ideal back end for enterprise applications like Red Hat Quay that require multiple types of software-defined storage (Figure 1). It is important to note that the same OpenShift Data Foundation cluster can provide all three types of storage for Red Hat OpenShift simultaneously. In the context of Red Hat Quay, OpenShift Data Foundation offers:

- ▶ **S3-compatible object storage.** Multicloud Object Gateway (MCG) functionality provides an AWS S3-compatible interface for Red Hat Quay container image storage. MCG uses NooBaa technology that offers a rich suite of capabilities in addition to basic S3 functionality.
- ▶ **Highly available block storage.** PostgreSQL requires persistent RWO volumes as a storage resource. OpenShift Data Foundation supports this requirement and can supply high availability through internal replication that can extend across AWS Availability Zones for business continuity.

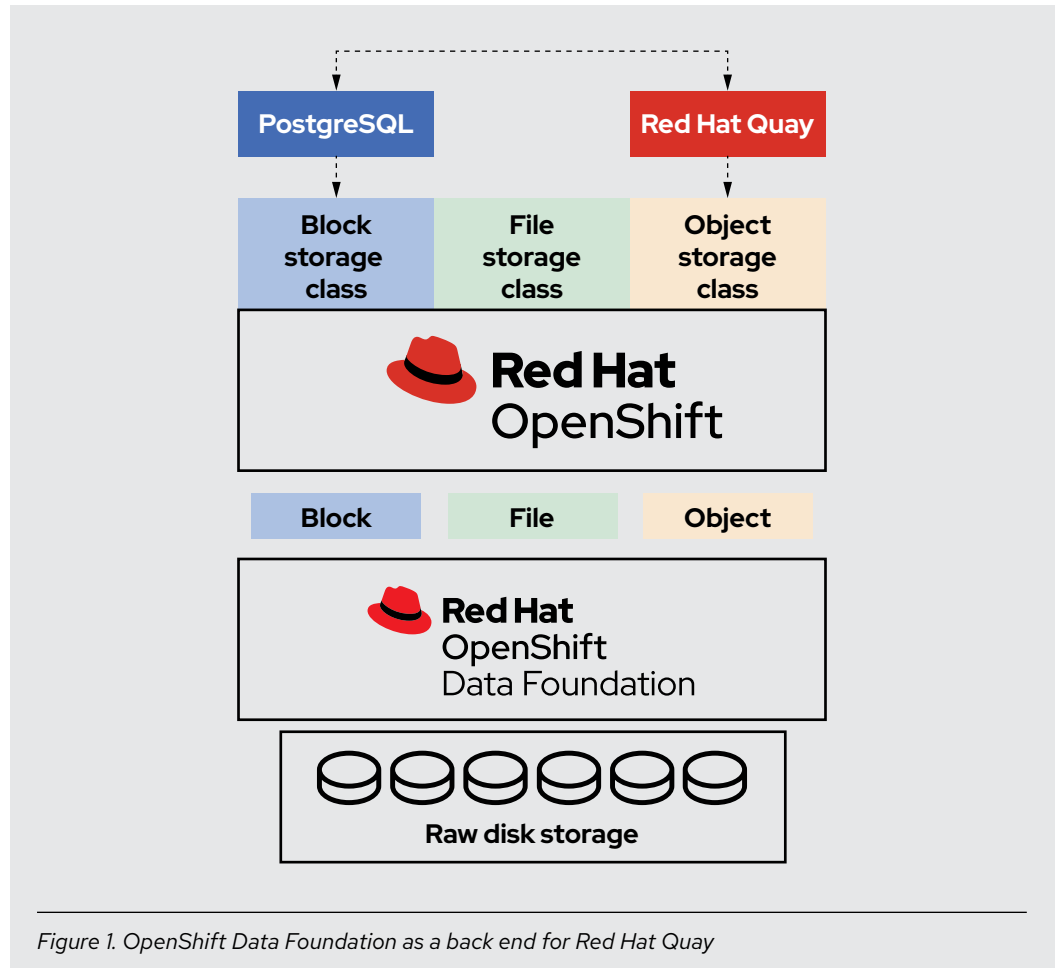


Figure 1. OpenShift Data Foundation as a back end for Red Hat Quay

In the configuration depicted, Red Hat Quay and PostgreSQL both run as containerized workloads inside Red Hat OpenShift pods. The PostgreSQL database used by Red Hat Quay requires a high-availability architecture, independent from Quay. PostgreSQL connects to a storage class for block storage (RWO) with data protection provided through OpenShift Data Foundation replication.

As shown, Red Hat Quay connects to an S3-compatible storage class for object storage, provided by OpenShift Data Foundation Multicloud Object Gateway, using NooBaa technology. Because OpenShift Data Foundation runs as a containerized service inside Red Hat OpenShift, it provides RWO and RWX persistent volumes (PVs), as well as object storage from a single solution.

Conclusion

Red Hat OpenShift Platform Plus includes OpenShift Container Platform, OpenShift Data Foundation, and Red Hat Quay, in addition to Red Hat Advanced Cluster Management for Kubernetes and Red Hat Advanced Cluster Security for Kubernetes, thus providing a full value proposition for enterprise Red Hat OpenShift deployments.

With its support for block, file, and object storage, Red Hat OpenShift Data Foundation provides software-defined storage for on-premise, public cloud, and hybrid cloud environments. This flexibility makes it an ideal highly available, persistent storage solution for Red Hat OpenShift logging, metrics, and registry services as well as for all other Red Hat OpenShift workloads. Red Hat OpenShift Data Foundation is a fully container-native solution that can simultaneously provide persistent and highly available storage for Red Hat OpenShift services, even as it supports a wide range of applications with more security capabilities, resiliency, and performance optimized storage.



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 develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. [A trusted adviser to the Fortune 500](#), Red Hat provides [award-winning](#) support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

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

redhat.com
#160770_1222

North America

1 888 REDHAT1
www.redhat.com

Europe, Middle East, and Africa

00800 7334 2835
europe@redhat.com

Asia Pacific

+65 6490 4200
apac@redhat.com

Latin America

+54 11 4329 7300
info-latam@redhat.com