

# Red Hat OpenShift Data Foundation

Simplified, consistent, and dynamic persistent storage for OpenShift applications

## Key benefits

- Simplified access and hybrid cloud data management for faster innovation
- Consistent experience for both users and developers across clouds for faster insights
- Dynamic scale, allowing organizations to start small and automate rapid data services growth

Red Hat OpenShift allows organizations to experience the breakthrough capabilities of Kubernetes from a consistent and supported enterprise platform—open to any app, team, or infrastructure.



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

redhat.com

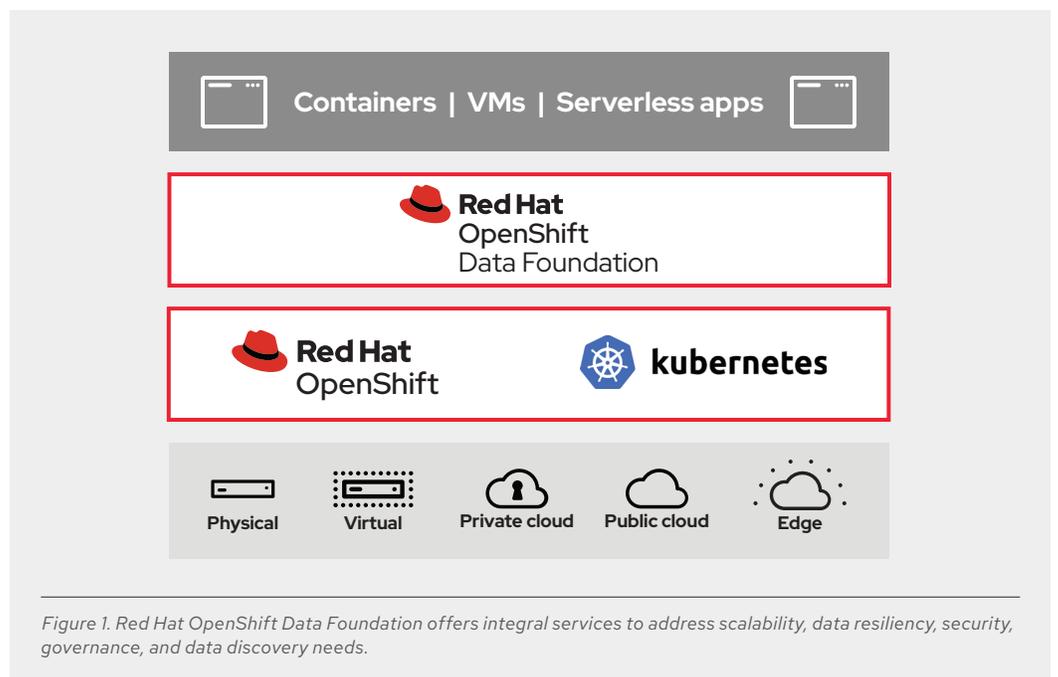
## Red Hat Data Services

Red Hat Data Services is a portfolio of solutions that includes persistent software-defined storage and data services integrated with and optimized for Red Hat® OpenShift® Container Platform. As part of the Red Hat Data Services portfolio, Red Hat OpenShift Data Foundation delivers resilient and persistent software-defined storage based on Ceph® technology. OpenShift Data Foundation runs anywhere Red Hat OpenShift does: on-premise or in the public cloud. The platform provides file, block, and object storage classes, enabling a wide range of data modalities and workloads, including:

- Data at rest, such as databases and data warehouses.
- Data in motion, automating data pipelines.
- Data in action, providing services for continuous deployment models, analytics, artificial intelligence (AI), and machine learning (ML).

## Red Hat OpenShift Data Foundation

Running as a Kubernetes service, OpenShift Data Foundation is engineered, tested, and qualified to provide data services for Red Hat OpenShift Container Platform on any infrastructure (Figure 1). OpenShift Data Foundation can also be decoupled and managed as a separate, independently scalable data store, delivering data for one or many Red Hat OpenShift Container Platform clusters.



Deployed, consumed, and managed through the Red Hat OpenShift administrator console, the platform is built upon Ceph, Noobaa, and Rook technologies, offering tightly integrated, persistent data services for Red Hat OpenShift in hybrid and multicloud environments. This level of integration removes the guesswork from running Kubernetes applications across multiple platforms and provides the data storage functionality, data services, and data protection that enterprises require. Dynamic, stateful, and highly available container-native storage can be provisioned and deprovisioned on demand.

### **Simplified access**

To increase agility, organizations need to reduce complexity for cloud-based applications and data, allowing rapid and flexible deployment of application data in any cloud. OpenShift Data Foundation ignites cross-team collaboration by providing simplified access to a consistent data platform experience. Self-service access from different user, developer, and administrator roles delivers storage on-demand with a click, not with a ticket. Easy to install, intuitive to monitor, and simple to start, organizations spend less time managing separate platforms and more time delivering value, allowing them to:

- Provide accessible data and support for all their Red Hat OpenShift apps.
- Simplify data management across hybrid cloud while providing easy access, faster innovation, and instant insights.

### **Consistent user experience**

Because OpenShift Data Foundation runs as a service in Red Hat OpenShift, users benefit from a consistent user experience independent of where the data resides. Whether you are building cloud-native apps with Kubernetes or just beginning your journey to cloud technology, your team gets a consistent hybrid cloud environment with intuitive data services across any infrastructure. Developers can build once and deploy anywhere. Data scientists get the ability to generate insights faster. Administrators can manage with consistency across cloud footprints, allowing them to:

- Deliver a unified end-user experience for data services for consistency across clouds and confidence across teams.
- Provide an all-in-one view of Red Hat OpenShift data with a comprehensive on-demand data platform.

### **Dynamic scale**

Moving to cloud technology cannot come at the expense of application performance. OpenShift Data Foundation supports scalability to multiple petabytes with resiliency and peak performance. Organizations can dynamically provision across any environment—bare metal, virtual machines, or hybrid cloud—all through the Red Hat OpenShift administrator console. Native object support in OpenShift Data Foundation dramatically increases input/output (I/O) performance, yielding performance at scale for every workload. Organizations can:

- Scale data services with the confidence to go anywhere and span any cloud.
- Innovate at scale while supporting all types of Red Hat OpenShift workloads, allowing easy object data-sharing across geographic locations and platforms.

## Tight integration with Red Hat OpenShift and Red Hat OpenShift Virtualization

OpenShift Data Foundation is created for container-based environments and is tightly integrated with Red Hat OpenShift Container Platform.<sup>1</sup> With a supported Red Hat OpenShift operator, the platform is simpler to install and manage as a part of the container-based application life cycle. This innovation allows Red Hat to provide support for the entire container-based environment, including cloud-native container management, scheduling, and orchestration, yielding:

- **Storage for trusted enterprise-class Kubernetes.** OpenShift Data Foundation adds support for diverse workloads and multicloud object gateway functionality, along with critical business-continuity functionality for data services.
- **Data protection and resiliency for Red Hat OpenShift.** High availability for stateful enterprise apps demands robust and highly available data storage with sophisticated capabilities. OpenShift Data Foundation supports important features like replication, allowing application data placement across different availability zones. Native support for data resiliency supports backup and restoration services for Kubernetes applications, including their data, resources, and state.
- **A cloud-like experience, everywhere.** Circumstances are constantly changing, favoring one cloud provider over another or in-house deployment versus public cloud. Organizations often need to move quickly to take advantage of favorable pricing or respond to other business pressures. OpenShift Data Foundation provides software-defined storage that lets organizations deploy their apps and storage as needs dictate and adjust as situations change.
- **Increased developer productivity.** Cloud developers want to innovate without arbitrary limitations. Traditional storage has often been an impediment to cloud development, requiring separate and time-consuming arrangements. OpenShift Data Foundation 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 multiple storage technologies. Not only is this time consuming—complex 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, OpenShift Data Foundation provides container-native storage supporting data at rest, data in motion, and data in action (Figure 2). It supports multiple cloud platforms and works with a wide range of partners and technologies.

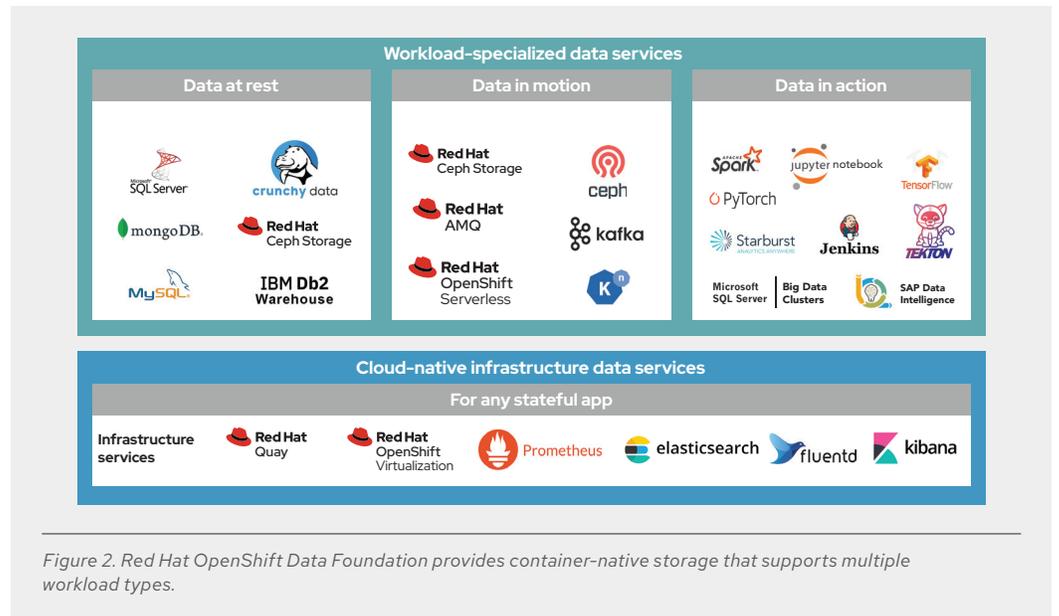
Based on 100% open source technology, Red Hat OpenShift Data Foundation supports file, block, and object storage and employs a single set of Kubernetes operators across all cloud platforms. Organizations can provide software-defined storage for multiple workload types with a single solution, and applications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

Popular use cases for OpenShift Data Foundation include:

- Data repositories and cloud-native application development, including continuous innovation / continuous deployment (CI/CD) models.
- Structured data, including SQL/NoSQL databases and data warehouses.
- Big data workloads such as data analytics and AI/ML workloads.

---

<sup>1</sup> Refer to the latest Red Hat OpenShift Data Foundation [release notes](#) for supported platforms.



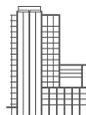
- Container-aware data protection with backup and data resiliency support for both persistent volumes and Kubernetes namespaces.

## Conclusion

Red Hat OpenShift Data Foundation provides simplified access, a consistent experience, and dynamic scale for 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 Data Foundation simplifies data management and allows storage provisioning and deprovisioning 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**  
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

redhat.com  
#F27388\_0321

Copyright © 2021 Red Hat, Inc. Red Hat, OpenShift, Ceph, and the Red Hat logo are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.