


IBM Storage Defender: Protecting Red Hat OpenShift Containers and Virtual Machines

Phillip Gerrard
Marcelo Capile
Reginald D'Souza
Maxwell Ferreira
Gerd Franke
Andre Gaschler
Kavish Ghogale
Pallavi Singh
Prajakta Tendulkar
Jörg Walter
Daniel Wendler



Introduction

Purpose of this document

This document provides practical guidance for implementing data protection and cyber resilience for Red Hat OpenShift environments using IBM Storage Fusion and IBM Storage Defender. It is written for architects, administrators, and operations teams who need to design, deploy, and manage enterprise-grade backup, restore, and recovery orchestration capabilities for containerized applications and Red Hat OpenShift Virtualization virtual machines.

The focus is on implementation rather than theory. Although this guide refers to IBM product documentation where appropriate, it emphasizes the decisions you need to make, the configurations that matter in production, and the operational patterns that help you recover successfully when incidents occur.

Who should read this document

This document is intended for the following readers:

- **Infrastructure architects** designing data protection strategies for Red Hat OpenShift environments
- **Platform administrators** responsible for implementing and managing backup and restore services
- **Site reliability engineers** who need to understand recovery procedures and operational patterns
- **Security teams** evaluating cyber resilience capabilities and governance controls
- **Operations teams** who will execute backup policies and perform recovery operations

You should have working knowledge of Red Hat OpenShift, Kubernetes concepts, and basic storage principles. Familiarity with IBM Storage products is helpful but not required.

Section overview

Introduction to Container Protection with IBM Storage Defender establishes the foundational concepts necessary for understanding data protection in containerized environments. This section introduces container technology, explains how Red Hat OpenShift orchestrates containerized applications, and outlines the unique protection requirements for Red Hat OpenShift environments. You will learn about IBM's comprehensive data protection portfolio, including IBM Storage Protect for Containers, IBM Fusion Backup and Restore, and IBM Storage Defender Data Protect. The section also addresses the critical aspects of protecting the Red Hat OpenShift cluster infrastructure itself, including control plane components, configuration management, and disaster recovery strategies.

IBM Storage Fusion Backup and Recovery covers the planning, architecture, installation, and operation of IBM Storage Fusion Backup and Restore. You will learn how to implement the Hub and Spoke model, configure protection policies, create application-consistent backups using recipes, and execute recovery operations for both containers and virtual machines. The section includes detailed procedures for installation, configuration, and day-to-day operations.

IBM Storage Defender Data Protect explains how to protect Red Hat OpenShift workloads using IBM Storage Defender Data Protect. This section covers environment preparation, pushing software images to a private registry, registering Red Hat OpenShift clusters as backup sources, creating backup policies and protection groups, and performing backup and restore operations for both containerized applications and Red Hat OpenShift Virtualization virtual machines. It also covers application-consistent backups using pre- and post-snapshot hooks, auto-protect for namespaces, single file and folder restores, and the key considerations for protecting KubeVirt-based virtual machines.

IBM Storage Defender DRS Integration describes how to integrate IBM Storage Fusion and IBM Defender Data Protect with the Data Resiliency Service (DRS). You will learn how to configure Connection Manager, establish

backup system integrations for both IBM Fusion and IBM Storage Defender Data Protect, create recovery groups with governance profiles, and leverage DRS capabilities for orchestrated recovery workflows. This section also covers DRS visibility concepts—including applications, recovery groups, governance profiles, clean room profiles, and dependency mapping—and ties together the backup mechanics from earlier chapters with the orchestration and governance layer that DRS provides.

Authors

Phillip Gerrard is a Project Leader for the International Technical Support Organization working out of Beaverton, Oregon. As part of IBM® for over 20 years he has authored and contributed to hundreds of technical documents published to IBM.com and worked directly with IBM's largest customers to implement storage solutions and resolve critical situations. As a team lead and Subject Matter Expert for the IBM Spectrum Protect support team, he is experienced in leading and growing international teams of talented IBMers, developing and implementing team processes, creating and delivering education. Phillip holds a degree in computer science and business administration from Oregon State University.

Marcelo Capile is a hands-on technology professional with over 15 years of experience in IT and a career built within IBM. His background spans critical environments, complex operations, and the kind of challenges that require not only technical knowledge, but also composure, consistency, and sound judgment. He has developed strong expertise in infrastructure, automation, cloud, observability, and platforms such as Linux, AIX, Kubernetes, and Red Hat OpenShift. Throughout his career, he has worked with Unix administration, backup strategies, storage, networking, Ansible automation, and Python and Shell scripting, as well as IBM Cloud, AWS, and GCP. Known for his practical mindset and collaborative approach, he combines technical depth with clear communication and a genuine willingness to help others grow.

Reginald D'Souza also known as Reg, is an APAC Storage Software Subject Matter Expert (SME) based in Australia. With over 20 years of experience in infrastructure, multi-vendor storage, and data protection solutions, Reg has been with IBM since 2008, working in various business units including delivery, managed services, and sales. His expertise spans across different software portfolios, including IBM and Red Hat. Currently, Reg's focus is on software-defined storage and modern data protection solutions, specifically IBM Fusion and IBM Storage Ceph, which aid clients in their hybrid cloud and AI journey. With his extensive work with Red Hat, he has earned a nickname, “Reg Hat”.

Maxwell Ferreira is a Senior IT Analyst at IBM Brazil. As a member of the IBM Storage Protect Support Team, working to support customer service requests from IBM's global accounts he is focused on IBM Storage Protect Plus, IBM Storage Protect for Cloud and IBM Fusion. He has over 20 years of experience in customer service, handling analysis requests, providing solutions for backup, restore, and data protection requests with IBM software. Since 2005, he has worked as an IT analyst, mostly with IBM Storage Protect, implementing and managing backup and restore environments with data storage on disk, tape, and cloud. He earned a degree in IT Management, a postgraduate degree in Cybersecurity, and an MBA in Industry 4.0.

Gerd Franke is a Senior IT Architect in the EMEA Storage team of IBM Client Engineering. He is known as an expert for Storage strategy and architecture, focusing on Cyber Security and Resiliency, Hybrid Cloud Storage and Modern Data Protection. He often leads large and complex client projects and is a regular speaker at IBM conferences. Gerd has worked for IBM in various national and international roles for more than 30 years. He holds an engineer's degree in Electrical and Communications Technology.

Andre Gaschler is an IBM Certified IT Specialist based in Frankfurt, Germany. He has over 20 years of experience with IBM's Data Protection portfolio. As senior technical specialist for IBM Systems Andre helps international clients deploy, modernize, and efficiently manage data resiliency infrastructure in enterprise environments. He is part of the Expert Labs DACH team as Solution Architect working in service development, delivery, and training for various strategic IBM Storage products.

Kavish Ghogale is a Senior Solution Architect with IBM Expert Labs Centre of Excellence (CoE), based in Mumbai, India. Since joining IBM in 2006, he has developed more than two decades of expertise in data protection and storage technologies. He began his IBM career under the Tivoli Software brand while working with IBM Software Labs before

joining IBM Systems Lab Services in 2015. Over the following eight years, he led the design and delivery of IBM Storage and modern data protection solutions while driving client enablement across the ISA, ASEAN, ANZ, and MEA regions. Since 2023, Kavish has been part of the IBM Expert Labs Solutioning team, where he focuses on pre-sales solution architecture, project planning, and strategic customer engagements for IBM clients worldwide.

Pallavi Singh is a Senior Technical Staff Member in IBM Storage, based in Pune, India, with over several years of experience in enterprise storage, observability, and system health management. Pallavi has led key initiatives in diagnostics and platform integration, driving scalable automation and improved customer outcomes. Pallavi provides strong technical leadership across cross-functional teams and currently focuses on advancing AI-driven capabilities such as anomaly detection and predictive monitoring. Pallavi has been a co-author on several IBM Redbooks publications and holds a master's degree in Computer Applications.

Prajakta Tendulkar is a QA Architect for IBM Fusion, with more than five years of hands-on experience working on the Fusion platform. In this role, Prajakta plays a critical part in shaping the overall quality strategy and test architecture for Fusion across releases. Having worked extensively across both Software-Defined Storage (SDS) and Hyperconverged Infrastructure (HCI) paradigms, Prajakta brings deep, end-to-end understanding of Fusion's architecture and deployment models. This includes strong expertise across key Fusion software components such as backup and restore, remote mount, and Data Foundation storage, along with a solid grasp of underlying hardware components, including compute and network infrastructure.

Jörg Walter is an IBM Certified IT Specialist based in Frankfurt, Germany. He joined IBM in 2000 and has held various technical roles throughout his career. Starting in network engineering and design, Joerg shifted his focus to storage, data protection, and retention in 2005. From 2011 to 2022, he was part of IBM's Systems Storage Lab Services team, where he supported clients in planning and implementing hardware and software solutions for data protection and retention. Since 2023, Joerg has been working as a Senior Platform Engineer within IBM Client Engineering EMEA, helping clients pilot advanced storage solutions using IBM's comprehensive storage hardware and software portfolio. He holds a diploma in Electrical Engineering and Computer Science from the University of Applied Sciences in Bingen am Rhein, Germany.

Daniel Wendler is a Client Engineer in IBM Global Sales EMEA, specializing in Cyber Resilient Storage architectures and IBM Storage Defender. He joined IBM in 2005 and spent more than a decade in IBM's worldwide last-level support organization, developing deep expertise across hardware and software storage solutions. He subsequently served four years as an IT Consultant within the European Storage Competence Center (ESCC), delivering IBM Expert Labs engagements to design and implement Modern Data Protection and Retention Solutions for clients across the EMEA region. Since 2023, Daniel has been working in Global Sales EMEA as a Client Engineer, where he focuses on architecting Cyber Resilient Storage solutions with IBM Storage Defender, running proofs of concept to demonstrate the value of IBM's Defender portfolio, and enabling clients to build robust, resilient infrastructures. He holds a Diplom-Informatiker degree from the Rhein-Main University of Applied Sciences, Germany.

Feedback

This document reflects practical experience implementing these technologies in production environments. If you encounter issues not covered here, or if you have suggestions for improving the guidance, your feedback helps make this resource more valuable for the community.

[Contact IBM Redbooks](#)

Container protection with IBM Storage Defender

In today's rapidly evolving IT landscape, containerized applications have become the cornerstone of modern software development and deployment. As organizations increasingly adopt container orchestration platforms like Red Hat OpenShift, the need for robust, enterprise-grade data protection solutions has never been more critical. This document explores how IBM Storage Defender and related IBM data protection technologies provide comprehensive

protection for Red Hat OpenShift environments, ensuring business continuity, disaster recovery, and compliance in containerized infrastructures.

Overview of containers and concepts

Containers represent a lightweight approach to application virtualization that packages an application and its dependencies into a single, portable unit. Unlike traditional virtual machines that require a full operating system, containers share the host system's kernel while maintaining isolated user spaces, making them significantly more efficient in terms of resource utilization and startup time.

Key container concepts include:

- **Container Images:** Read-only templates containing the application code, runtime, libraries, and dependencies needed to run an application
- **Container Runtime:** The software responsible for running containers, such as containerd or CRI-O
- **Namespaces:** Linux kernel features that provide isolation for processes, network, and filesystem resources
- **Control Groups (cgroups):** Mechanisms for limiting and monitoring resource usage (CPU, memory, I/O) for containers
- **Persistent Volumes:** Storage resources that exist beyond the lifecycle of individual containers, enabling stateful applications

Containers enable developers to build once and run anywhere, promoting consistency across development, testing, and production environments while accelerating application deployment cycles.

Overview of container orchestration

As containerized applications scale from single containers to hundreds or thousands of interconnected services, manual management becomes impractical. Container orchestration platforms automate the deployment, scaling, networking, and lifecycle management of containerized applications across clusters of hosts.

Red Hat OpenShift is an enterprise Kubernetes platform that extends the open-source Kubernetes orchestration engine with additional developer and operational tools, enhanced security features, and enterprise support. Red Hat OpenShift provides the following capabilities:

- **Automated Deployment and Scaling:** Declarative configuration and automatic scheduling of containers across cluster nodes
- **Self-Healing:** Automatic restart of failed containers and rescheduling on healthy nodes
- **Service Discovery and Load Balancing:** Built-in networking and routing for containerized services
- **Rolling Updates and Rollbacks:** Zero-downtime application updates with automatic rollback capabilities
- **Resource Management:** Intelligent allocation of compute, memory, and storage resources
- **Multi-Tenancy:** Secure isolation between different teams and projects within the same cluster

Red Hat OpenShift's operator framework further simplifies the management of complex stateful applications by encoding operational knowledge into automated workflows.

Overview of protecting an Red Hat OpenShift environment

Protecting an Red Hat OpenShift environment requires a comprehensive approach that addresses multiple layers of the container stack, from application data to cluster configuration. Traditional backup solutions designed for virtual machines often fall short in containerized environments due to the dynamic, ephemeral nature of containers and the distributed architecture of Kubernetes.

Key protection requirements for Red Hat OpenShift include:

- **Application-Consistent Backups:** Capturing not just container images but also persistent volume data, application state, and configuration
- **Kubernetes Resource Protection:** Backing up cluster resources such as deployments, services, config maps, secrets, and custom resource definitions (CRDs)
- **Namespace-Level Granularity:** Enabling selective backup and restore at the project/namespace level
- **Cross-Cluster Mobility:** Supporting disaster recovery scenarios that require restoring applications to different clusters
- **Compliance and Retention:** Meeting regulatory requirements for data retention and audit trails
- **Minimal Performance Impact:** Protecting production workloads without disrupting application performance

Effective Red Hat OpenShift protection must also account for both stateless applications (which can be easily recreated from images) and stateful applications (which require careful handling of persistent data and application state).

Overview of IBM data protection portfolio for Red Hat OpenShift

IBM offers a comprehensive portfolio of data protection solutions specifically designed to address the unique challenges of protecting containerized environments. These solutions integrate seamlessly with Red Hat OpenShift, providing enterprise-grade backup, recovery, and disaster protection capabilities while maintaining the agility and scalability that containers promise.

The IBM data protection portfolio for Red Hat OpenShift encompasses multiple complementary technologies that work together to provide defense-in-depth protection:

- **Backup and Recovery Solutions:** Traditional backup capabilities adapted for container-native architectures
- **Cyber Resilience and Ransomware Protection:** Advanced threat detection and immutable backup copies
- **Disaster Recovery:** Cross-site replication and failover capabilities for business continuity
- **Data Mobility:** Tools for migrating applications and data between clusters and cloud environments

Overview of IBM Fusion Backup and Restore / IBM Storage Protect for Containers

IBM Storage Protect for Containers provides Kubernetes-native backup and recovery capabilities specifically designed for containerized applications running on Red Hat OpenShift and other Kubernetes distributions.

Key capabilities include:

- **Kubernetes-Native Architecture:** Deployed as containers within the Red Hat OpenShift cluster using operators for simplified management
- **Application-Aware Protection:** Understanding of application dependencies and relationships for consistent backups
- **Persistent Volume Backup:** Integration with Container Storage Interface (CSI) for efficient snapshot-based backups
- **Namespace and Label-Based Selection:** Flexible policies for selecting which applications and resources to protect
- **Point-in-Time Recovery:** Granular recovery options including full namespace restore, individual resource recovery, or selective file restoration
- **Multi-Cluster Management:** Centralized management and monitoring of protection across multiple Red Hat OpenShift clusters
- **Backup to Multiple Targets:** Support for backing up to object storage (S3-compatible), IBM Storage Protect, or other enterprise backup repositories

IBM Fusion Backup and Restore extends these capabilities with enhanced integration into IBM's broader storage ecosystem, providing unified data protection across containers, virtual machines, and traditional workloads. IBM Fusion Backup and Restore offers the following features:

- Unified Management Interface: Single pane of glass for managing protection across hybrid cloud environments
- Policy-Based Automation: Intelligent automation of backup schedules, retention, and lifecycle management
- Rapid Recovery: Optimized restore processes that minimize recovery time objectives (RTO)
- Scalability: Architecture designed to protect thousands of containers across large-scale deployments

Overview of IBM Storage Defender Data Protect

IBM Storage Defender Data Protect represents the next generation of data protection for Red Hat OpenShift environments, combining traditional backup and recovery with advanced cyber resilience capabilities. Built on IBM's FlashCore technology and leveraging AI-powered analytics, Storage Defender provides comprehensive protection against both operational failures and cyber threats.

Distinguished features include:

- Cyber Resilience: Immutable backup copies stored on air-gapped or logically isolated storage to protect against ransomware and malicious deletion
- Threat Detection: AI-driven anomaly detection that identifies suspicious patterns in backup data, potentially indicating ransomware or data corruption
- Rapid Recovery at Scale: Leveraging flash storage technology for fast recovery of large containerized applications
- Copy Data Management: Efficient management of backup copies with deduplication and compression to minimize storage footprint
- Compliance and Governance: Built-in reporting and audit capabilities for regulatory compliance (GDPR, HIPAA, SOX)
- Integration with Security Tools: APIs and integrations with security information and event management (SIEM) systems and security orchestration platforms
- Virtual Machine Protection: Unified protection for both containers and virtual machines, supporting hybrid infrastructure models

Storage Defender's architecture is designed for the modern threat landscape, where data protection must extend beyond traditional backup to include proactive threat detection, forensic analysis capabilities, and guaranteed recovery even in the face of sophisticated cyberattacks.

Protecting the Red Hat OpenShift cluster itself

While protecting application data and workloads is critical, comprehensive Red Hat OpenShift protection must also address the cluster infrastructure itself. The Red Hat OpenShift cluster comprises multiple components that require protection to ensure complete recoverability:

Control plane protection

The Red Hat OpenShift control plane includes master nodes running the API server, scheduler, controller manager, and etcd database. Protecting the control plane involves the following tasks:

- etcd Backup: Regular backups of the etcd key-value store, which contains all cluster state and configuration
- Control Plane Configuration: Backup of master node configurations, certificates, and authentication settings
- Disaster Recovery Planning: Procedures for rebuilding control plane nodes and restoring cluster state

Infrastructure configuration

Beyond the control plane, protecting the cluster infrastructure includes the following components:

- Operator Configurations: Backup of installed operators and their custom resource definitions

- Network Policies: Preservation of network segmentation and security policies
- Storage Classes and Provisioners: Configuration of storage infrastructure and persistent volume provisioners
- Registry and Image Repositories: Protection of internal container image registries
- Monitoring and Logging Infrastructure: Backup of Prometheus, Grafana, and logging stack configurations

Identity and access management

Security and access control configurations must be protected and recoverable:

- Authentication Providers: LDAP, OAuth, and other identity provider configurations
- Role-Based Access Control (RBAC): Cluster roles, role bindings, and service accounts
- Security Context Constraints: Policies governing container security privileges
- Secrets and Certificates: Encrypted backup of sensitive credentials and TLS certificates

Cluster recovery strategies

Effective cluster protection requires documented and tested recovery procedures:

- Cluster Rebuild Procedures: Step-by-step processes for recreating cluster infrastructure
- Configuration as Code: Using GitOps practices to maintain cluster configuration in version control
- Disaster Recovery Testing: Regular validation of recovery procedures through testing and drills
- Multi-Cluster Architectures: Implementing redundant clusters for high availability and disaster recovery

IBM's data protection solutions provide tools and automation to simplify cluster-level protection, including automated etcd backups, configuration drift detection, and integration with infrastructure-as-code workflows. By combining application-level protection with comprehensive cluster infrastructure protection, organizations can achieve true business continuity for their Red Hat OpenShift environments.

Note: See the official Red Hat OpenShift documentation for detailed procedures and supported methods for backing up and restoring the Red Hat OpenShift cluster components described in this section.

This introduction establishes the foundation for understanding how IBM Storage Defender and related technologies provide enterprise-grade protection for Red Hat OpenShift environments. The following chapters will dive deeper into implementation details, best practices, and use cases for protecting containerized workloads and virtual machines in production environments.

IBM Fusion Backup and Recovery

IBM Fusion Backup and Restore provides enterprise-grade data protection for Red Hat OpenShift environments, including both containerized applications and Red Hat OpenShift Virtualization virtual machines. Built on proven open source technologies and integrated with IBM's storage portfolio, Fusion delivers policy-driven backup and restore capabilities that scale across hybrid cloud deployments.

This section covers the planning, architecture, installation, and operational aspects of IBM Fusion Backup and Restore. You will learn how to design a resilient backup infrastructure using the Hub and Spoke model, configure protection policies, implement application-consistent backups, and execute recovery operations for both containers and virtual machines. The focus is on practical implementation guidance that helps you build a production-ready data protection solution for your Red Hat OpenShift workloads.

Planning for IBM Fusion Backup and Recovery

Centralized control plane: Why it matters for Fusion DR, data resources, policies, and backups

In IBM Fusion, the centralized control plane is where DR starts to look like an operating model, not just a backup feature. That shift lets you stop managing protection as disconnected tasks spread across clusters and start treating it as one governed flow. Policies here. Locations there. Visibility somewhere else. The control plane pulls this information together and gives the team one place to see what is protected, what failed, what drifted, and what is simply missing. In practice, this removes a lot of operational noise and makes understanding the recovery posture of the environment much more simple.

For Fusion DR, this matters because recovery is never only about restoring data. It is about restoring the right data resources, with the right policy, in the right order, and proving that the recovered application is usable. While this sounds obvious, it is usually not as straightforward as it initially seems when it comes time to actually implement and manage a recovery plan. A central control plane helps because it gives one control surface for protection logic across clusters. You get policy handling, job visibility, storage locations, application assignments, and service-level recovery context in one path.

The data resources side is where this becomes more than basic administration. IBM Fusion design goes beyond protecting single objects. Protecting one PVC here, one namespace there, one VM somewhere else leads . Incidents do not care about those boundaries. The control plane helps by keeping the workload view tied to the protection view, so the team can work in terms of applications, data sets, versions, policies, and recovery targets rather than individual components.

In practice, the centralized control plane provides:

- One place to manage policy intent across clusters, instead of per-cluster improvisation.
- One view for applications, jobs, versions, locations, and service state. Less guessing.
- Better alignment between workload boundaries and the data resources that actually matter at recovery time.
- Cleaner DR operations for Hosted Control Plane environments, where IBM explicitly supports Backup & Restore and disaster recovery for applications.

Policies are another big part of this story. In IBM Fusion, backup policies are a key component of managing data and, are assigned from the application view. Policy should describe intent: frequency, retention, target location, maybe the broader recovery posture too. Proper planning goes beyond just a schedule checkbox. When that policy layer is managed from a central place, the environment becomes more consistent across clusters. Less tribal knowledge is required to understand, maintain and recover the environment and it helps to avoid having to locate or identify the correct backup location.




Figure: IBM Fusion centralized control plane interface

An important note on service protection

Backups under a centralized control plane are not only about workloads. IBM is very explicit here: The Backup & Restore service can protect its own metadata to an S3 bucket. That metadata is what lets you recover backup policies, locations, assignments, and backup custom resources if the cluster or the metadata storage becomes unavailable. This is not a side feature. It is central to DR design. If the service state is gone, the restore path gets ugly fast.

There is also an important distinction that needs to be made. Service protection is for recovery, not for migration. IBM explicitly calls this out as during service backup recovery, the IBM Fusion namespace must remain the same between backup and restore. Another detail that is often missed if local S3 storage was involved, is that you still need a separate strategy to protect that local storage, because the service recovery brings back metadata saved to cloud storage, not the local backup content itself. Missing this key component can result in the design looking complete on paper, but resulting in real recovery issues during a recovery scenario.

Operational reading: What the control plane really buys you

When the recovery environment is designed well, the centralized control plane helps in four direct ways:

- It centralizes visibility of applications, policies, backup locations, and job outcomes.
- It keeps data resources tied to workload recovery intent instead of leaving them fragmented.
- It standardizes policy assignment and makes drift easier to catch before an incident.
- It supports recovery not only of workloads, but also of the Backup & Restore service metadata that the protection model depends on.

It is however important to note, A centralized control plane improves control and repeatability, but it does not fix weak recovery design by itself. If dependencies are not mapped, if recovery order is unclear, if policies are inconsistent, or if service protection was never configured, the control plane mostly gives you a cleaner view of the mess. While this may be useful, it is not enough to ensure recovery is possible. The value comes when governance, backup policy, service protection, and recovery procedures are wired together early in the design process - before the incident call, not during it.

Policies			
Search			
Name	Backup frequency	Retention	Data resources assigned
aws-daily	1 day	5 days	6
new	1 day	2 days	1
service-protection-policy	1 day	30 days	1
service-protection-policy2	1 day	30 days	1
service-protection-policy3	1 day	30 days	1
test	2 months	2 days	1
testpolicy	1 day	3 days	2
Items per page: 25 ▾ 1-7 of 7 items			1 ▾ 1 of 1 page < >

Figure: IBM Fusion policies and data sources configuration screen

Data resources		
Total data resources : 34		
Search		
☐ Name	Backup status	Backup system
☐ aws-20251219-181208	Without backup	console-ibm-fusion.bnr-vcp-titan-98b7318c91b01bd72490e80cc2328915-0000.us-south.containers.appdomain.clo
☐ aws-20251219-185350	Without backup	console-ibm-fusion.bnr-vcp-titan-98b7318c91b01bd72490e80cc2328915-0000.us-south.containers.appdomain.clo
☐ aws-20251219-190504	Without backup	console-ibm-fusion.bnr-vcp-titan-98b7318c91b01bd72490e80cc2328915-0000.us-south.containers.appdomain.clo
☐ calico-apiserver	Without backup	console-ibm-fusion.bnr-vcp-titan-98b7318c91b01bd72490e80cc2328915-0000.us-south.containers.aodomain.clo

Figure: IBM Fusion data resources backup status list

IBM Fusion hub & spoke

IBM Fusion provides Hub & Spoke as management models, where a single administrator from the Hub cluster can manage other clusters known as Spokes. The management workload is concentrated on the Hub, making it possible to manage all Backup & Restore jobs in all of the spoke clusters from a central location.

In the Hub, you can configure and manage Backup & Restore items such as policies, retentions, schedules, and other items related to Backup & Restore.

In the Spoke you need install Backup & Restore Agent, which receives and responds to the Backup & Restore requests provided from the Hub. During the spoke agent installation process, a safety connection needs to be established with the Hub.

For this connection be established, the Hub sends a connection request snippet that the Spoke must respond to within 60 minutes. Once the spoke is able to answer, a secure connection is establish between the Hub and the Spoke.

Backup & Restore service on IBM Fusion provides some benefits:

- Backup centralized of Red Hat OpenShift applications using customized policies.
- The possibility of using Changed Block Detection (CDB) for Ceph RBD block volumes to optimize backup execution, especially for Red Hat OpenShift virtualization VMs based on RBD block volumes.
- Use a custom and specific recipe to automate the consistency validation of application backups.
- Availability to perform backup and restore across all clusters using the HUB and Spoke topology models.




Figure: IBM Fusion dashboard main page showing overview metrics

Red Hat OpenShift Operator Hub

You can search for the operators of IBM Fusion in the Red Hat OpenShift Operator Hub, these operators can be installed in all clusters to enable backup and restore services/agent and other features.




Figure: Operator Hub displaying the IBM Fusion operators available under the IBM Operator Catalog

Once the operator is selected, the details are displayed on the screen, as well as the option to install the component:



IBM Storage Fusion
 2.12.0 provided by IBM



Install

Channel
 v2.0

Version
 2.12.0

Capability level
☒ Basic Install
☒ Seamless Upgrades
☐ Full Lifecycle
☐ Deep Insights
☐ Auto Pilot

Source
 IBM Operator Catalog

Provider
 IBM

Repository
icr.io/cpopen

Introduction
 IBM Storage Fusion provides container-native data and storage services for your OpenShift workloads. Bring applications to production faster, with data services that are simple, consistent, and strategic.

Details
Highly available storage
 Storage provisioning is automated through the Container Storage Interface. IBM Storage Fusion provides two choices for storage provisioning:

- Red Hat OpenShift® Data Foundation - IBM Storage Fusion provides entitlement for Red Hat OpenShift Data Foundation.
- Remote IBM Storage Scale file system mount - Connects to an external IBM Storage Scale cluster to access an existing file system. Container workloads running in Red Hat OpenShift can access existing data sets stored in the IBM Spectrum Scale file system and provision new persistent volumes backed by storage from the remote IBM Spectrum Scale file system.

Data protection
 Protect your applications with efficient incremental snapshots of stateful data. Use local snapshots for fast recovery and offload copies to object storage to protect against cluster outages.

IBM Storage Fusion puts you in control of how data consistency is achieved with support for both application-consistency and crash-consistency:

Figure: IBM Fusion operator details and installation options

Topology

In the Topology tab you can connect to a new spoke cluster and manage the Hub and Spoke clusters connectivity.

Topology						
Displays all clusters involved in Backup & Restore.						
Connection status: Filter...	Service status: Filter...	Search		Connect cluster		
Name	Type	Connection status	Service status	Version	Backed up apps	Success rate
apps.ocp.hub.ibm.lab	Hub	–	Healthy	2.12.0	0/17	–
apps.ocp.pri.ibm.lab	Spoke	Connected	Healthy	2.12.0	0/1	–
apps.ocp.sec.ibm.lab	Spoke	Connected	Healthy	2.12.0	0/1	–
Items per page: 25		1–3 of 3 items		1 1 of 1 page		

Figure: IBM Fusion GUI Backup and restore Topology view

When you connect a new spoke cluster to the hub, the Backup & Restore Agent service is installed on the spoke. After connecting to a cluster, you can verify and manage the following details:

- The **name** of the configured Hub and Spoke clusters.
- The **type** of the cluster, if is a Hub or a Spoke.
- The **connection status** between the hub and all the spoke clusters.

- The IBM Fusion **service status** is healthy on the Hub and Spoke clusters.
- The **Version** of the IBM Fusion service is installed on each cluster.
- The **Backed-up apps** number deployed for the applications on each cluster.
- The **Success rates** display the percentage of successfully finished jobs.

You also have filters and search options to customize the view.

Hub & spoke architecture diagrams

IBM Fusion Backup & Restore supports two deployment models: a Hub-only architecture for single-cluster deployments, and a Hub and Spoke architecture for managing backups across multiple clusters.

Hub diagram

The workflow diagram below displays Hub model architecture with the local agent for Backup & Restore.




Figure: IBM Fusion Backup and Restore architecture diagram with Hub model

Hub and spoke diagram

The workflow diagram below displays Hub and Spoke models architecture for Backup & Restore management.




Figure: IBM Fusion Backup and Restore architecture diagram with Hub and Spoke models

Backup & restore service architecture and components

This section describes the components involved in Backup and Restore using IBM Fusion with Hub & Spoke models.

IBM Fusion Backup & Restore service

The IBM Fusion user interface is the graphical user interface of IBM Fusion where you can admin the Events, Application, Backup & Restore, Storage, Services, and Settings.

The IBM Fusion user interface is configured on the Hub side, and it is possible to manage the Hub cluster and all the others Spokes clusters in the same user interface under Backup & Restore tab.




Figure: IBM Fusion GUI Overview screen

Application & Data Protection CRs and Application & Fusion DP Operator.

The Backup and Restore service interacts with IBM Fusion CRs (Cluster Resources) in the "Application and Data Protection CRs". In addition, the "Application and Fusion DP Operator" uses this interactivity to enable Backup and Restore features for all managed clusters in IBM Fusion.

IBM Fusion Backup & Restore agent

During the execution of a Job, the Backup and Restore agent manages the tasks within that execution. Provide centralized application backups and restore, using local snapshots for fast recovery and protection on external object storage.

The recommendation is to install IBM Fusion and the Backup & restore Agent Service on the cluster that you want to manage application protection. The Backup & Restore agent can be found in the Services tab on IBM Fusion GUI. After installation the Backup & Restore agent is available on Services tab.




Figure: IBM Fusion GUI Services tab

For this management, the agent uses three components that act in the execution and control of the tasks:

- **Transaction Manager:** The Transaction Manager handles the task requests during backup or restores jobs. During a job workload, the backup and restore agent is using the Transaction Manager to work on tasks processing, recipe execution, PVC requests for incremental backups, and provides updated status to the server.
- **Data Mover:** The Data Mover is used to mount/read the PVC snapshot during incremental backup jobs. After the PVC snapshot mounted, it will be read and the data compared so that the incremental transfer to the S3 location only the modified data based on the last PVC snapshot that was read.
- **Redis:** Redis works as a mediator between tasks; once a task requests access, Redis performs a lock so that no other task interferes with the execution of the previous task.

Backup & Restore Server

Backup policies

Backup Policies contain the schedules and settings for the backup routines and can be customized according to your needs for each environment.

Policies

The following table contains backup policies created on this cluster. Backup policies contain the schedule and backup location IBM Fusion will use to create recurring backups. [Learn more.](#)

Backup location: All

Q Search

Add policy

Name	Backup location	Frequency	Start time	↑	End time	Retention	Applications
minio.hub.wordpress	minio.hub.wordpress	Every day	4:00 AM Central European Standard Time		8:00 AM Central European Standard Time	14 Days	0
protect1.pri.wordpress	protect1.pri.wordpress	Every day	4:30 AM Central European Standard Time		8:30 AM Central European Standard Time	14 Days	0
protect1.hub.virtualmachines	protect1.hub.virtualmachines	Every day	5:00 AM Central European Standard Time		9:00 AM Central European Standard Time	14 Days	0
protect1.pri.softwarehub	protect1.pri.softwarehub	Every day	5:30 AM Central European Standard Time		9:30 AM Central European Standard Time	14 Days	0

Items per page: 10

1-4 of 4 items

1

1 of 1 page

Figure: IBM Fusion GUI Backup and Restore Policies screen

After configuring a policy for Jobs, the information is displayed in the Policies tab:

- The policy name can be customized to best distinguish between policies in your environment.
- Backup location is where this policy will send the data for storage.

- Frequency is the setting for how often it will be executed, for example, "Every day".
- Start and Finish time is the execution window available for this policy to run.
- Retention is the time setting this data will be stored before it expires from the Location.
- Application is the number of applications associated with each policy.

Backup location

This is where the backup data provided by the workload will be stored. You will need to specify a location during the policy backup creation.




Figure: IBM Fusion GUI Backup and Restore Locations dashboard

Once the locations are configured, you have a dashboard displaying each location where you can check:

- If the location is connected to IBM Fusion and ready to receive workload data.
- The type of storage (IBM COS, Azure, Aws, Minio and so on).
- How much GB/TB is being consumed by policies in that location.
- It is also possible to check how many policies and applications have that location as a backup storage destination.

The S3 object storage providers currently supported by Backup & Restore services in IBM Fusion are:

- IBM Cloud (IBM - Object Storage)
- Azure (Microsoft - Object Storage)
- AWS (Amazon - Object Storage)
- MCG/NooBaa (Red Hat - Object Storage)
- S3 Compliant (Any - Object Storage)
- IBM Storage Protect (IBM - Object Storage and Tape)

Caution: OpenStack Swift is not supported in any version .

Backup service

The IBM Fusion Backup & Restore service ensures protection for all clusters via backup schedules using object storage. Those schedules are responsible for backup of the Backup & Restore data management service to a S3 bucket. If a cluster crashes, you will need to recover the namespace in a new, healthy cluster.

Ensure that the service protection backup has been created and all application backups are performed to avoid data loss. Create schedules with customized retention and frequency for the applications.

The location can be defined to Microsoft Azure, IBM Cloud, S3-compatible object storage, Amazon Web Services (AWS), and IBM Storage Protect. You can immediately start a backup using the "Backup service now" icon. If backups are running, you can monitor the backup progress in the Service Protection table.




Figure: IBM Fusion GUI Backup and Restore Service protection status

The Service backup table displays: **Time**, **Status**, **Location** and **Size**.

- The **Time** indicates when the backup was performed.
- **Status** shows whether the backup was successful or failed.
- **Location** is where the storage is being performed.
- **Size** shows in "Gigabytes" the volume of data that was transferred during this backup.

Note: You also have filters and search options to customize the display information.

The right corner of the screen lists the details of the Backup summary with the following information:

- **Available backups:** Number of available backups that can be used for restoring.
- **Used capacity:** Value in "Gigabytes" showing the amount of space used by the backups.
- **Last successful backup:** Date of the most recent backup among all the available backups.
- **Backup policy:** It indicates which policy is being used and displays the schedule and location details of the policy in use.

Job manager

The Job Manager is responsible for sending the workload to the Backup agents on the clusters. While monitoring the progress of running jobs in the job manager, it also provides the option to cancel workloads if necessary.




Figure: IBM Fusion GUI Backup and Restore Jobs dashboard

The job manager dashboard displays the following information:

- The summary of successful, canceled, and failed for backup and restore job, and you can set the period of this status will be displayed, for example, in the last 24 hours.
- It is also possible to check in the Queue quadrant if there are jobs in progress or are pending, waiting for previous policies to finish.
- Each backup or restore job generates an entry in a task list, where you can check the job name, target cluster, application, status, started/elapsed/finished time, policy name, and if this job is using a recipe.
- In this task list, you can filter the jobs if needs to refine the search for jobs in progress or completed/canceled/failed.

Note: This Job management in a single location brings some advantages, such as avoiding the execution of duplicate workloads in the same application service.

Application service

Applications Service works to provide a registry of all applications that are distributed across all clusters. During backups, the applications in the cluster need to be accessed, and these application registries are queried by the server.

Redis

Like the Redis on the Backup & Restore Agent, we also have Redis on the Backup & Restore Server. Redis acts as a mediator between tasks; once a task requests access, Redis performs a lock so that no other task interferes with the execution of the previous task.

Backup & restore DP operator

The Backup and Restore DP operator will interact with the Data Protection CRs on IBM Fusion to communicate with the Backup and Restore server during backup or restore workloads on the clusters.

Kafka

The communication services for Backup & Restore use Kafka protocol to intermediate the connections. Kafka uses a binary TCP-based protocol and provides scalability, replication, and resilience in these connections between the services during the Backup & Restore jobs demands.

S3 bucket

The S3 bucket is a container for store data on other location with a S3 provider technology. The objects generated during backup processes will be stored, providing scalability, encryption secure and data availability to use the data when restoration is required for a restore job on IBM Fusion. This bucket address will be configured as a Location on IBM Fusion Backup & Restore Service.

The types of S3 object storage currently supported for Backup & Restore services in IBM Fusion are:

- IBM Cloud (IBM - Object Storage)
- Azure (Microsoft - Object Storage)
- AWS (Amazon - Object Storage)
- MCG/NooBaa (Red Hat - Object Storage)
- S3 Compliant (Any - Object Storage)
- IBM Storage Protect (IBM - Object Storage and Tape)

Caution: OpenStack Swift is not supported in any version.

MongoDB

The server uses MongoDB to centralize the metadata generated during the processing of Backup & Restore jobs. MongoDB is a non-relational database that does not use tables but rather flexible documents in "Binary JSON" (JavaScript Object Notation) format. Another characteristic of MongoDB is an open-source database.

Multiple S3 locations on IBM Fusion

IBM Fusion Backup & Restore services can use "S3 Object Storage" as backup location. In addition to the local storage for Red Hat OpenShift application backups, it is recommended to set up one or more additional locations with an S3 destination.




Figure: IBM Fusion Backup and Restore Locations without S3 configuration

The S3 object storage providers currently supported by Backup & Restore services in IBM Fusion are:

- IBM Cloud (IBM - Object Storage)
- Azure (Microsoft - Object Storage)
- AWS (Amazon - Object Storage)
- MCG/NooBaa (Red Hat - Object Storage)
- S3 Compliant (Any - Object Storage)
- IBM Storage Protect (IBM - Object Storage and Tape)

Caution: OpenStack Swift is not supported in any version.

Add a backup location

Backup locations connected to IBM Fusion provide storage for application backups created by backup policies.

Location type

Location name

Choose an object storage type

- MCG/NooBaa**
Red Hat • Object Storage
- Azure**
Microsoft • Object Storage
- IBM Cloud**
IBM • Object Storage
- aws AWS**
Amazon • Object Storage
- S3 Compliant**
Any • Object Storage
- Storage Protect**
IBM • Object Storage and Tape

Cancel Back Next

Figure: IBM Fusion Backup and Restore Add location dialog showing types of object storage

Locations

Backup storage locations are required to set up backup policies. To protect your applications in the event of a cluster failure, consider setting up an additional object storage locations. [Learn more.](#)

Status: All Type: All Search... Add location +

minio.hub.wordpress Connected Used 0 GiB Policies: 0 Applications: 0	protect1.hub.virtualmachine Connected Used 0 GiB Policies: 0 Applications: 0	protect1.pri.softwarehub Connected Used 0 GiB Policies: 0 Applications: 0	protect1.pri.wordpress Connected Used 0 GiB Policies: 0 Applications: 0
---	---	--	--

Figure: IBM Fusion Backup and Restore Locations after S3 storage configured

For your data be safety and protected, the primary productive data and backup data must be in separate locations. Ensure that certain items have been verified to ensure your storage is effective. The IBM Fusion Backup & Restore service cluster must be in a different S3 cluster and namespace. For S3 Compliant, enable the ACL (access control list) with read permission. For MinIO or NooBaa, when using cluster services, the hosts must be in different clusters.

It is not supported to define two backup storage locations that have identical endpoint and bucket names. S3 buckets should not have expiration policies enabled and Buckets should also not have archiving rules defined. If the location is Amazon AWS, then you must also enter the **Region**.

Creating a secret on IBM Fusion

During the configuration of new locations, you can use the option to set a certificate. This setting is applicable only when you create an S3 compliant backup storage location type. In the Certificate settings (optional) section, enter the Secret name for the certificate. On Sample "minio-cert-secret" was used as a secret name for certificate.

The screenshot shows a web form titled "Add a backup location". Below the title is a subtitle: "Backup locations connected to IBM Fusion provide storage for application backups created by backup policies." The form is divided into two main sections. The first section, "Login credentials", contains fields for "Endpoint", "Bucket", "Access key", and "Secret key". The second section, "Certificate settings (optional)", is highlighted with a blue border and contains a text input field for "Secret name for certificate" with the value "minio-cert-secret" entered. At the bottom of the form are three buttons: "Cancel", "Back", and "Add".

Figure: Add a backup location certificate settings with minio-cert-secret

For the certificate to be usable in configuration, you need to create the secret before adding a backup storage location in an S3 provider. You need to use OpenSSL commands to extract the S3 certificate so that it can be used. It is required that the file containing the extracted certificate be named `tls.crt`.

OpenSSL command to extract certified reference:

```
openssl s_client -connect <s3-service-name>-<s3-service-namespace>.apps.<fusion-hostname>.<domainname>:443
```

A sample of OpenSSL command to extract the certificate. The `minio` as service name and the `minio-ns` as namespace an IBM Fusion cluster, using also port 443.

```
echo | openssl s_client -connect minio-minio-ns.apps.myfusionhostname.mydomain:443 2>&1 | sed -n '/-BEGIN
```

Run the `oc` command to create a generic secret in the IBM Fusion namespace using the `tls.crt` file.

```
oc create secret generic <secret-name> --type=opaque --from-file=tls.crt -n <fusion-namespace>
```

A sample, in the default IBM Fusion namespace `ibm-spectrum-fusion-ns`.

```
oc create secret generic minio-cert-secret --type=opaque --from-file=tls.crt -n ibm-spectrum-fusion-ns
```

Note:

- Ensure that the file name must be `tls.crt` on OpenSSL command.
- Save a note of this secret name. You need it to create a backup storage location.

Add location to IBM COS, Azure, AWS and S3 compliant

Procedure for IBM Fusion:

Step 1: Select on **Backup & Restore > Locations**.

Step 2: In the **Locations** page, click **Add location**.

Step 3: The Add backup location wizard page is displayed.

Add a backup location

Backup locations connected to IBM Fusion provide storage for application backups created by backup policies.

Location type

Location name

minio.hub.wordpress

Choose an object storage type

MCG/NooBaa
Red Hat • Object Storage

Azure
Microsoft • Object Storage

IBM Cloud
IBM • Object Storage

aws AWS
Amazon • Object Storage

S3 Compliant
Any • Object Storage

Storage Protect
IBM • Object Storage and Tape

Cancel Back Next

Figure: IBM Fusion add backup location wizard initial screen

Step 4: In the Add backup location, enter the **Location name**.

Step 5: Select the object storage type:

- Azure
- IBM Cloud
- AWS
- S3Compliant
- MCG/NooBaa
- IBM Storage Protect

Note: For MCG/NooBaa and IBM Storage Protect, follow the additional details on below.

Step 6: Click **Next**.

Add a backup location

Backup locations connected to IBM Fusion provide storage for application backups created by backup policies.

Login credentials

Connect IBM Fusion to your backup location.

Endpoint
https://s3.ibm.lab:9000

Bucket
hub.wordpress

Access key
6H6g8Xgd0kp09DHau53y

Secret key
.....

Certificate settings (optional)

SSL secured object storage locations require certificates from authentication.
[Create an OpenShift TLS Secret](#) in namespace `ibm-spectrum-fusion-ns` the certificate for the backup storage location and provide the secret name here.

Secret name for certificate
minio-cert-secret

Cancel Back Add

Figure: IBM Fusion backup location login credentials form

Step 7: In the Login credentials section, enter the following credentials to connect IBM Fusion to your backup location: **Endpoint**, **Bucket**, **Access key** and **Secret key**. For **Azure**, then enter the **Account name** and **Account key** instead of **Access key** and **Secret key**.

Step 8: In the **Certificate settings (optional)** section, enter the **Secret name for the certificate (*)**.

Note: For the procedure to create a secret, see [Creating a secret on IBM Fusion](#)

Step 9: Click **Add**.

A successful message pop-up after adding the location.

Add location for MCG/Noobaa

Create an MCG or Noobaa S3 target backup storage location.

Procedure on Red Hat OpenShift:

Step 1: Log in to Red Hat OpenShift Container Platform.

: Go to Storage > Object Bucket Claims.

Step 3: In the Create ObjectBucketClaim page, create an Object Bucket claim with the following details: capture credential details.

Parameter	Description
ObjectBucketClaim Name	Enter the name of the Object Bucket Claim. If you do not enter a value, a generic name gets generated.
StorageClass	Select openshift-storage.noobaa.io class. It defines the object storage service and the bucket provision.
BucketClass	Select noobaa-default-bucket-class.

Step 4: Click Create.

Step 5: Note down the Object Bucket Claim credentials like Endpoint, Bucket Name, Access Key, Secret key.

Step 6: Extract the certificate from an S3 compliant service to a file. Use the openssl command to extract the certificate into the file tls.crt.

```
openssl s_client -connect <s3-service-name>-<s3-service-namespace>.apps.<fusion-hostname>.<domainname>:443 -showcerts \ | sed -n '/BEGIN CERTIFICATE/,/END CERTIFICATE/'
```

Note: Ensure that the file name must be tls.crt.

Step 7: Run the `oc` command to create a generic secret in the IBM Fusion namespace by using the tls.crt file.

```
oc create secret generic <secret-name> --type=opaque --from-file=tls.crt -n <fusion-namespace>
```

For example, in the default IBM Fusion namespace ibm-spectrum-fusion-ns

```
oc create secret generic bsl-cert --type=opaque --from-file=tls.crt -n ibm-spectrum-fusion-ns
```

Procedure for IBM Fusion:

Step 1: Select on Backup & restore > Locations.

: In the Locations page, click Add location.

Step 3: The Add backup location wizard page is displayed.




Figure: IBM Fusion add MCG Noobaa location wizard

Step 4: In the Add backup location, enter the Location name.

Step 5: Select the type of **MCG/Noobaa** of object storage backup location.

Step 6: click NEXT




Figure: IBM Fusion MCG Noobaa location configuration screen

Step 7: In the Login credentials section, enter the following credentials to connect IBM Fusion to your backup location: **Endpoint, Bucket, Access key and Secret key.**

Step 8: In the **Certificate settings (optional)** section, enter the **Secret name for the certificate(*)** created in the [Creating a secret on IBM Fusion](#) step.

Step 9: Click **Add**.

A successful message pop-up after adding the location.

Additional steps for NooBaa type:

You can add the location as a separate NooBaa type instead of the S3, but you can not do it from the UI. It requires additional steps to create secret credentials and BSL YAML. To create a S3 of NooBaa type, do the following steps:

Step 1: Create a secret using the `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY` from NooBaa.

```

kind: Secret
apiVersion: v1
metadata:
  name: <noobaa bs1 secret>
  namespace: <fusion namespace>
  labels:
    dp.isf.ibm.com/ownedBy: fbsl
    dp.isf.ibm.com/provider-name: isf-backup-restore
data:
  access-key-id: <base64-encoded AWS_ACCESS_KEY_ID>
  secret-access-key: <base64-encoded AWS_SECRET_ACCESS_KEY>
type: Opaque

```

Note: If a YAML is used, encode the values using base64.

: Copy the location of the S3 route to use as the endpoint for the BSL CR.

Use noobaa CLI to display the list of buckets: `noobaa bucket list`.

Step 3: Create a BSL CR with type **noobaa**:

```

apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: BackupStorageLocation
metadata:
  name: <noobaa bs1>
  namespace: <fusion namespace>
  labels:
    dp.isf.ibm.com/provider-name: isf-backup-restore
spec:
  credentialName: <noobaa bs1 secret>
  params:
    bucket: <bucket name>
    endpoint: '<noobaa s3 route>'
    privateCertificateName: <optional noobaa certificate secret>
  provider: isf-backup-restore
  type: noobaa

```

Add location for IBM Storage Protect

Procedure for IBM Fusion:

Step 1: Select on Backup & restore > Locations.

: In the Locations page, click Add location.

Step 3: The Add backup location wizard page is displayed.

Step 4: In the Add backup location, enter the Location name.

Step 5: Select the type of **IBM Storage Protect** of object storage backup location.

Figure: IBM Fusion add IBM Storage Protect location wizard

Step 6: In the Login credentials section, enter the following credentials to connect IBM Fusion to your backup location: **Endpoint**, **Bucket**, **Access key** and **Secret key**.

Step 7: In the **Certificate settings (optional)** section, enter the **Secret name for the certificate(*)** that you created in the [Creating a secret](#) step.

Step 8: Click **Add**.

A successful message pop-up after adding the location.

Installation

This section describes how to install IBM Fusion Backup and Restore. As discussed before in the [planning section](#), Fusion consists of various components:

- A base framework, which includes a Web UI and an Red Hat OpenShift console plugin
- One or more data services that an administrator could select to install

The installation routines and requirements vary slightly, depending on the environment in which the software is to be used, for example, cloud provider managed Red Hat OpenShift, self-managed on-premise Red Hat OpenShift, online or offline (air-gapped). However, the essential installation steps are very similar, so we describe an online installation of IBM Fusion with the Backup & Restore service in a self-managed BareMetal Red Hat OpenShift cluster in this book as an example.

Prerequisites - Prepare for the installation

For a detailed list of environments supported by IBM Fusion and their respective installation methods and prerequisites, see the IBM Fusion [documentation](#).

Review [IBM Fusion services platform support matrix](#) to ensure your environment and platform is supported by IBM Fusion and the data services you plan to deploy. At the time of writing this book, the Fusion Backup and Restore service is supported on all platforms where Fusion can be deployed.

See the [System requirements](#) to understand if your Red Hat OpenShift version is supported by the version of Fusion that you plan to deploy. Furthermore, this page lists the CPU and memory resources required by the various components, for example, Fusion Base, B&R Hub or B&R Agent (Spoke).

Preparing the infrastructure

Network ACLs

If you plan to deploy IBM Fusion on an Red Hat OpenShift Container Platform cluster with a cluster wide proxy configured, then ensure you add these hosts to your allowlist:

- registry.redhat.io
- quay.io
- cp.icr.io
- icr.io
- redhat.com
- registry.connect.redhat.com
- catalog.redhat.com
- access.redhat.com
- cloud.redhat.com
- cdn02.quay.io
- registry.access.redhat.com
- dd0.icr.io
- dd2.icr.io
- dd4.icr.io
- dd6.icr.io

Container Storage Interface (CSI) storage provider

IBM Fusion supports backup and recovery operations on Red Hat OpenShift Container Platform environments with any storage provider that implements the Container Storage Interface (CSI) driver and meets the following prerequisites:

- The Container Storage Interface driver must be v1 or higher (no support for alpha or beta versions).
- The Container Storage Interface driver must support [Volume Snapshot and Restore](#).
- A StorageClass must be created with `volumeBindingMode: Immediate` or `volumeBindingMode: WaitForFirstConsumer` and `allowVolumeExpansion:true`.
- PersistentVolumes (PV) must be created with `volumeMode: Filesystem` or `volumeMode: Block`.
- A VolumeSnapshotClass must be configured for any storage provisioner that is used for volumes that will be backed up by Fusion Backup and Restore. It is also recommended that the VolumeSnapshotClass deletion policy is configured to delete snapshots after the corresponding VolumeSnapshotContent is deleted (deletionPolicy: Delete).

Bastion node

To communicate with the Red Hat OpenShift cluster and to list, add or modify secrets, config maps or other custom resources, a bastion node needs to be available with the following tools installed:

- Red Hat OpenShift CLI (oc) tool The oc tool can be downloaded from the Red Hat OpenShift web interface, for example:

```
https://console-openshift-console.apps.<your_cluster_fqdn>/command-line-tools
```

- JQ JSON processor command line package

The jq command line package is included in most Linux repositories. On Red Hat Enterprise Linux systems, it can be installed using this command:

```
yum install jq
```

Alternatively, the jq tool can be downloaded [here](#)

Obtaining an entitlement key

Entitlement keys determine whether IBM Fusion software operator can automatically pull the required container native images. During installation, image pull failures may occur due to an invalid entitlement key or a key belonging to an account that does not have entitlement to IBM Fusion.

Use the following procedure to obtain your entitlement key:

Step 1: Log in to the [IBM container software library](#) with an IBM id and a password that is associated with the entitled software.

Step 2: In the navigation bar, click **Get entitlement key**.

Step 3: On the **Access your container software** page, click **Copy key** to copy the generated entitlement key.

Step 4: Save the key to a secure location for future use.

Create or update the image pull secret

Update the image pull secret to enable the Red Hat OpenShift cluster to authenticate with the IBM entitled registry. The image pull secret provides the credentials for pulling Docker images from the IBM entitled registry.

The following procedure shows how to add credentials for the IBM entitled registry to an existing pull secret via the Red Hat OpenShift web UI:

Step 1: Log in to the Red Hat OpenShift cluster web UI and go to **Workloads > Secrets**.

Step 2: From the Project pull-down menu, select **openshift-config**.

Step 3: In the context menu of the *pull-secret*, click **Edit Secret**:




Figure: Red Hat OpenShift pull-secret edit dialog

Step 4: Click + Add credentials

Step 5: Fill the form with the following values:

- **Registry server address:** `cp.icr.io`
- **Username:** `cp`
- **Password:** <entitlement key, as described in section [Obtaining an entitlement key](#)>
- **Email:** e-mail address of the entitlement owner (can remain empty)

Step 6: Click the **Save** button.

Alternatively, the IBM container software library entitlement key can be added to the global image pull secret from the bastion node using the `oc` command line tool. This approach is described in detail in the [IBM Fusion Documentation](#)

Adding the IBM operator catalog

The Fusion framework and its services are being installed using an operator that is being provided from the IBM operator catalog. The easiest way to add the IBM operator catalog to your Red Hat OpenShift cluster is to use the Quick create menu from the Red Hat OpenShift GUI:

- Click the plus sign
- Select **Import YAML**
- Select **openshift-marketplace** project and paste the following yaml into the text-input field:

```
apiVersion: operators.coreos.com/v1alpha1
kind: CatalogSource
metadata:
  name: ibm-operator-catalog
  namespace: openshift-marketplace
spec:
  displayName: IBM Operator Catalog
  publisher: IBM
  sourceType: grpc
  image: icr.io/cpopen/ibm-operator-catalog:latest
  updateStrategy:
    registryPoll:
      interval: 45m
```

- Finally, click the **Create** button.

Alternatively, create a file `catalog_source.yaml` with the content above on the bastion node and apply it to the cluster.

Note:

- IBM Operator Catalog is not specific to a version of IBM Fusion.
- If you already have the IBM Maximo software configured in your environment or have plans to configure it, use only the `ibm-operator-catalog` of IBM Maximo software.

Installing the Fusion framework

Once all prerequisites are met and implemented, the Fusion Base framework can be installed.

Install the Fusion operator

If you have not configured the IBM Operator Catalog, then configure it. For the procedure to add IBM Operator catalog, see [Adding the IBM operator catalog](#).

Step 1: Log in to Red Hat OpenShift Container Platform web management console and go to **Operators > OperatorHub**.

Note: This instruction is valid for Red Hat OpenShift version 4.19 and before. Beginning with Red Hat OpenShift v4.20, the OperatorHub is called "Software Catalog". The Fusion Operator will then be found in the Red Hat OpenShift UI under **Ecosystem > Software Catalog**.

Step 2: Under Source, select **IBM Operator Catalog**. It lists all operators that are part of the IBM Operator Catalog including IBM Fusion:




Figure: Red Hat OpenShift OperatorHub showing IBM Fusion operator

Step 3: Click IBM Fusion.

Step 4: The Version, Capability level, Source, and Provider type of IBM Fusion is available:




Figure: IBM Fusion operator details and metadata

Step 5: Click Install. It opens the Install Operator page for IBM Fusion Operator:

Install Operator

Install your Operator by subscribing to one of the update channels to keep the Operator up to date. The strategy determines either manual or automatic updates.

Update channel *
v2.0

Version *
2.12.0

Installation mode *

- ☐ All namespaces on the cluster (default)
This mode is not supported by this Operator
- ☒ A specific namespace on the cluster
Operator will be available in a single Namespace only.

Installed Namespace *

- ☒ Operator recommended Namespace: **PR** `ibm-spectrum-fusion-ns`
- ☐ Select a Namespace

Namespace already exists
Namespace `ibm-spectrum-fusion-ns` already exists and will be used. Other users can already have access to this namespace.

Update approval *

- ☐ Automatic
- ☒ Manual

Manual approval applies to all operators in a namespace
Installing an operator with manual approval causes all operators installed in namespace `ibm-spectrum-fusion-ns` to function as manual approval strategy and will be updated altogether. Install operators into separate namespaces for handling their updates independently. To allow automatic approval, all operators installed in the namespace must use automatic approval strategy.

Install **Cancel**

IBM Storage Fusion
provided by IBM

Provided APIs

- FSD Fusion Service Definition**
FusionServiceDefinition is the Schema for the fusion servicedefinitions API. Documentation for deploying IBM Spectrum Fusion: <https://ibm.biz/Bdbky6>. License information: <https://ibm.biz/BdbumJ>. This product is not deployable standalone and it is part of...
- FSI Fusion Service Instance**
FusionServiceInstance is the Schema for the fusion serviceinstances API. Documentation for deploying IBM Spectrum Fusion: <https://ibm.biz/Bdbky6>. License information: <https://ibm.biz/BdbumJ>. This product is not deployable standalone and it is part of the...
- SC Scale Cluster**
ScaleCluster is the Schema for the scaleclusters API. Documentation for deploying IBM Storage Fusion: <https://ibm.biz/Bdbky6>. License information: <https://ibm.biz/BdbumJ>. This product is not deployable standalone and it is part of the LICCR already...
- C Cluster**
Cluster is the Schema for the clusters API. Documentation for deploying IBM Spectrum Fusion: <https://ibm.biz/Bdbky6>. License information: <https://ibm.biz/BdbumJ>. This product is not deployable standalone and it is part of the LICCR already accepted.
- C Connection**
Connection is the Schema for the connections API. Documentation for deploying IBM Spectrum Fusion: <https://ibm.biz/Bdbky6>. License information: <https://ibm.biz/BdbumJ>. This product is not deployable standalone and it is part of the LICCR already accepted.
- HM Health Manager**
HealthManager is the Schema for the healthmanagers API

Figure: IBM Fusion operator installation configuration

Step 6: Select `v2.0` on the update channel where the current operator is published.

Step 7: In the **Installation mode**, select **A specific namespace on the cluster**. the operator will be available in a single Namespace only.

Step 8: Select `ibm-spectrum-fusion-ns` in the **Operator recommended namespace**. Alternatively, use the **Select a namespace** option to select an existing namespace or create a new namespace.

Step 9: In the **Update approval** section, you can select either **Manual** or **Automatic** strategy. Always set **Update approval** to **Manual** because the **Automatic** option automatically upgrades the operator whenever a new version of the operator is released to the channel. This automatic upgrade might have an impact on your running workloads.

Important: Always set Update approval to Manual as IBM Fusion does not support Automatic.

Step 10: Click **Install**. The installation of the operator begins. As the approval strategy was set to **Manual**, you must manually approve the installation:




Figure: Red Hat OpenShift operator manual approval dialog

Step 11: Wait for the operator to complete the installation. Generally, it completes in few minutes. If it takes more time, check whether all pods are up and running.

After the successful installation of the operator, the following message gets displayed: *Installed operator - operand required*




Figure: IBM Fusion operator installed successfully message

Important: DO NOT click the Create SpectrumFusion button yet. The subsequent installation tasks will be performed in the Fusion Web UI. The IBM Fusion Operator installs a Web UI and a console plugin.

Access the Fusion UI

Step 1: From the Applications menu in the title bar of the Red Hat OpenShift Container Platform UI, click **IBM Fusion**:

Step 2: The IBM Fusion UI is displayed. Click **Accept** to accept the license agreement.




Figure: Red Hat OpenShift applications menu with IBM Fusion link

Step 3: The license agreement is displayed:

License agreement

Please read and accept the IBM Fusion license agreement to continue

[Read our license agreement](#)

[IBM privacy policy](#)

☒ I have read and accept the license agreement

[Continue](#)

Figure: IBM Fusion license agreement acceptance screen

Step 4: Go through the license agreement, click **I have read and accept the license agreement**, and click **Continue**.

Once the license agreement is accepted, the installation continues and when it completes, the Fusion Dashboard is displayed:




Figure: IBM Fusion main dashboard after installation

Browse available Fusion services

In the **IBM Fusion UI**, click **Services** to review the various data services that can be installed:




Figure: IBM Fusion services catalog showing available services

Note: Depending on the license you purchased from IBM, you might not be eligible to install all the services. For example, if you have purchased an IBM Defender Storage Protect for Containers license, you might be allowed to deploy the Fusion Backup & Restore services, but none of other services.

Installing Fusion Backup & Restore on the hub

As discussed before in the [planning section](#), Fusion Backup & Restore can be deployed in two "flavors" - Fusion B&R (HUB) or Fusion B&R Agent (SPOKE). In the following examples, we assume that there are three Red Hat OpenShift clusters in the environment:

- **One OCP cluster is designated as HUB.** This cluster will manage all backup / restore operations in the environment and holds the B&R configuration information, like S3 backup storage locations, policies, applications and policy assignments. The URL of the HUB cluster is <https://console-openshift-console.apps.ocp.hub.ibm.lab>.
- **One OCP cluster is dedicated as the "primary" SPOKE.** This cluster hosts various containerized applications and virtual machines that must be backed up. The URL of the primary SPOKE cluster is <https://console-openshift-console.apps.ocp.pri.ibm.lab>.
- **One OCP cluster is dedicated as the "secondary" SPOKE.** This cluster does not host any applications or VMs, but can be used as a recovery target. The URL of the secondary SPOKE cluster is <https://console-openshift-console.apps.ocp.sec.ibm.lab>.

Note: The HUB service also includes the SPOKE components . Designating a dedicated OCP cluster to be the Backup & Restore control instance and not hosting any other containerized applications or VMs can be a valid approach in very large environments , or if the system shall be under full control of the enterprise backup administrators. However, an OCP cluster that has the Fusion B&R HUB role can also host other applications and VMs and those can be backed up by the B&R HUB running on the same OCP cluster.

Step 1: Begin with the installation of the **Backup & Restore** service on the **HUB**. In the Fusion UI, go to the list of available services and click **Backup & Restore**:




Figure: IBM Fusion services list with Backup and Restore highlighted

Step 2: Review the service description and version and click **Install**.




Figure: IBM Fusion Backup and Restore service details

Step 3: Select the storage class to be used by the B&R HUB service and decide whether to use automatic updates for the service or not:




Figure: IBM Fusion Backup and Restore installation configuration

Note: The Fusion B&R HUB service deploys various PVCs to store persistent information for its own configuration databases. The PVCs use `volumeMode: Filesystem` and `accessMode: ReadWriteOnce (RW0)`.

The installation process might take a few minutes. Once it is completed, review the **Services** in the **IBM Fusion UI** and validate that the Backup & Restore Service is installed:




Figure: IBM Fusion services showing installed Backup and Restore

Note: As the Fusion B&R HUB service also includes the B&R SPOKE role, the Backup Restore Agent service was removed from the list of available services as well.

Step 4: Once the B&R HUB is deployed, a new menu item is inserted into the left-hand side main menu of the IBM Fusion UI. In the **IBM Fusion UI**, go to **Backup & Restore > Topology** and validate that the HUB cluster is properly listed and shown as healthy:

Displays all clusters involved in Backup & Restore.						
Connection status: Filter...		Service status: Filter...		Search		
Name	Type	Connection status	Service status	Version	Backed up apps	Success rate
apps.ocp.hub.ibm.lab	Hub	~	Healthy	2.12.0	0/12	~
Items per page: 25 1-1 of 1 item						

Figure: IBM Fusion Backup and Restore topology showing HUB cluster

If the HUB Red Hat OpenShift cluster is the only cluster in your environment and if no SPOKE clusters exist, you can continue to configure Fusion Backup & Restore as described in [\(LINK to 2.1.3 configuration section\)](#). << where to link? there is no 2.1.3

However, in this example we have two SPOKE clusters running applications and VMs that need to be connected to the HUB.

Installing Fusion Backup & Restore agent on a spoke

Before starting the installation of the Fusion Backup and Restore Agent on a SPOKE cluster, ensure that the prerequisites are met see [Prerequisites – Prepare for the installation](#) and that the Fusion framework was installed using the IBM Fusion operator see [Installing the Fusion framework](#). These preparation tasks are the same for all OCP clusters in an environment, independent of their intended future role – HUB or SPOKE.

Once the Fusion base framework is installed on a spoke, you can access the Fusion UI and start the installation of the Backup & Restore Agent.

Step 1: In the left-hand menu of the **Fusion UI**, click **Services** and select **Backup & Restore Agent**:




Figure: IBM Fusion services with Backup and Restore Agent highlighted

Step 2: Review the service description and version and click **Install**.





Figure: IBM Fusion Backup and Restore Agent service details

As a result, Fusion Backup and Restore Agent will ask for a **Hub connection snippet** and a **Storage class**. The Hub connection snippet needs to be created on the HUB and contains connection information, as well a transient token (valid for 1 hour) that allows a SPOKE to form a connection with the HUB:



Backup & Restore Agent

Install service

Hub Connection Snippet

Storage class to be used to deploy the service.

ocs-storagecluster-ceph-rbd


Storage Class

☒ Automatic updates

Cancel Install

Figure: IBM Fusion Backup and Restore Agent installation form

Step 3: Create a new Connection snippet: On the **HUB** Red Hat OpenShift cluster, open the **Fusion Web UI**, go to **Backup and Restore > Topology** and click **Connect cluster**.



Topology

Displays all clusters involved in Backup & Restore.

Connection status: Filter... Service status: Filter... Search


[Connect cluster](#)

Name	Type	Connection status	Service status	Version	Backed up apps	Success rate
apps.ocp.hub.ibm.lab	Hub	—	Healthy	2.12.0	0/17	—

Items per page: 25 1~1 of 1 item 1 1 of 1 page

Figure: IBM Fusion topology with Connect cluster button

From the following screen, you can select whether to use the Fusion UI to insert the connection snippet on the spoke, or if a YAML or "oc command" output shall be generated, which is useful when automating a Fusion Backup & Restore deployment on multiple clusters. In our scenario, we **select the UI method**:



Connect cluster

Close

Connection method

Create connections to the clusters that you will provide backup services for. If you are only backing up applications in this cluster, you can skip this step.

Use Fusion UI

If the spoke clusters have IBM Fusion installed, you can generate a Code Snippet that can be used while installing the Backup & Restore Spoke service.

→

Automate deployment

Generate a YAML or oc command that can be used to deploy Backup & Restore Spokes into many clusters using automation.

→

Figure: IBM Fusion cluster connection method selection dialog

In the following screen, click Copy snippet to copy the connection snippet into the clipboard, as we need to insert it to the SPOKE later:

Connect cluster

Use the Fusion UI

Go to the cluster that you want to provide backup services for and install the Backup & Restore Spoke service. You will be asked to provide the Connection Snippet, which you can copy from below

During the install of the Backup & Restore service from the Fusion UI, you will be asked to provide the Connection Snippet to establish a secure connection between the spoke and this hub. Copy the Connection Snippet below and paste it into the install form.

Copy Connection Snippet below and paste it into the Spoke cluster.

Connection Snippet

```
{
  "bootstrapToken": "ZX1KaGJHY2lPaUpTVXpJMU5pSXNjbXRwWkNjNk1rWnpMVVZwTVVwNU5GQTJWwFZZWmpsUGN5MMWZUb
  "apiServer": "aHR0cHM6Ly9hcGkub2NwLmh1Yi5pYm0ubGF1OjY0NDM=",
  "clusterName": "apps.ocp.hub.ibm.lab",
  "hubNs": "ibm-spectrum-fusion-ns"
}
```

The token in the Code Snippet is valid for 1 hour.

Back

Copy snippet

Figure: IBM Fusion connection snippet copy dialog

Step 4: Insert the newly created connection snippet to the SPOKE: On the **SPOKE** Red Hat OpenShift cluster, paste the clipboard content to **Fusion Backup & Restore Agent Install service** window and click **Install**:

Backup & Restore Agent

Install service

Hub Connection Snippet

```
5GZUZR3JyLUIwVWVCSm9INTlrYzF0dHVfd2FGZl9LczIxWW8tZzZpdIB3Sno0RVIZM2NwUW9USTBMR2d2U1Jmci16dzliS1E=",
"apiServer": "aHR0cHM6Ly9hcGkub2NwLmh1Yi5pYm0ubGF1OjY0NDM=",
"clusterName": "apps.ocp.hub.ibm.lab",
"hubNs": "ibm-spectrum-fusion-ns"
}
```

Storage class to be used to deploy the service.

ocs-storagecluster-ceph-rbd

Storage Class

Automatic updates

Cancel

Install

Figure: IBM Fusion Backup and Restore Agent with connection snippet

As a result, the Backup & Restore Agent will continue with the installation process.

Verifying the installation

When the Fusion Backup & Restore Agent installation is completed on a **SPOKE** cluster, go to **Fusion UI > Services** and verify that the B&R Agent service is shown as *Running* in the list of installed services. In the following example, the Fusion Data Foundation service was installed before the Backup & Restore service to provide Storage classes that could be used for the B&R agent service. Therefore, two services are shown in the *Installed* section:

Services

Discover and manage available services. [Learn more](#)

Installed



Backup & Restore Agent
IBM • 2.12.0

Provides application-centric backups, using local snapshots for quick recovery and external object storage protection. Enhanced offering with stability...

Data protection

✓ Running



Data Foundation
IBM • 4.19.9

Provides a foundational data layer for applications to function and interact with data in a simplified, consistent and scalable manner.

Storage

✓ Running

Available



Content-Aware Storage
IBM

Provides advanced support for vectorizing data by connecting storage systems to AI pipelines to unlock the business value of GenAI applications.

AI Data Services



Data Cataloging
IBM

Provides rapid automated data discovery and robust metadata capture, curation and enrichment.

AI Data Services



Global Data Platform
IBM

This service provides crash consistent backups, access to data from external storage sources and Container Storage Interfaces (CSI) provisioning of...

Storage

Figure: IBM Fusion services showing installed Backup and Restore Agent

To get an overview of the entire HUB & SPOKE network, review the Topology view on the HUB Red Hat OpenShift cluster. On the **HUB** cluster, go to **Fusion UI > Backup & Restore > Topology** and verify that the HUB, as well as all SPOKES are shown as Healthy and Connected:

Topology

Displays all clusters involved in Backup & Restore.

Connection status: Filter...

Service status: Filter...

Search

Connect cluster

Name	Type	Connection status	Service status	Version	Backed up apps	Success rate
apps.ocp.hub.ibm.lab	Hub	—	✓ Healthy	2.12.0	0/17	—
apps.ocp.pri.ibm.lab	Spoke	✓ Connected	✓ Healthy	2.12.0	0/1	—
apps.ocp.sec.ibm.lab	Spoke	✓ Connected	✓ Healthy	2.12.0	0/1	—

Items per page: 25

1–3 of 3 items

1 of 1 page

Figure: IBM Fusion topology showing HUB and two SPOKE clusters

Uninstalling Fusion Backup & Restore

IBM provides a shell script to uninstall the Fusion Backup and Restore service. The script needs to be run from a bastion node that has the oc and jq tools installed.

The following procedure explains how to uninstall the Fusion B&R service:

Step 1: Copy the [uninstall-backup-restore.sh](#) script to your bastion node and make it executable.

Step 2: Log in to the cluster where the Fusion B&R service shall be uninstalled.

Note: If Fusion B&R shall be uninstalled on all clusters of a HUB<>SPOKE topology, start with the SPOKES and uninstall the HUB last.

Step 3: If the uninstaller script is run on a **Fusion SPOKE** (B&R Agent) and the HUB still exists, run the uninstall script with the `-u` option.

Note: On the Topology page of the Fusion web UI on the HUB , it can take 10 minutes for the Connection Status of the spoke cluster to change from Connected to Unknown . The connection can be removed afterward through the context menu.

Step 4: If the uninstaller script is run on a **Fusion HUB**, use the `-f` option to delete any B&R related custom resources in the Fusion namespace (`ibm-spectrum-fusion-ns`).

Note: If you plan to delete Fusion entirely (incl. the base framework) , use the `-d` option to create deleteBackupRequest CRs to trigger the deletion of any existing backups. As this is also removing/cleaning contents of external object storage buckets, it may take hours to complete.

More details on the Fusion B&R uninstallation procedure can be found in the [IBM Fusion Documentation](#)

Applications and containers

This section describes how to backup and restore Kubernetes applications, persistent volumes, and cluster-specific resources on Red Hat OpenShift using IBM Fusion Backup & Restore.

Prerequisites

Before starting to backup and restore Kubernetes/Red Hat OpenShift applications, ensure that Fusion Backup & Restore has been installed properly and that the configuration is complete.

The examples shown in the following sections of this section assume that Fusion Backup & Restore is installed on multiple Red Hat OpenShift clusters and that the following configuration tasks have been completed:

- One B&R HUB is connected to two B&R SPOKES ("primary" & "secondary")
- The (self) service-protection of the HUB to external object storage is enabled
- At least one backup storage location (*fbsl*) to external object storage is defined
- At least one backup policy is defined (using the *fbsl* as target)

Scope of Fusion Backup and Restore operations

Fusion applications

The scope of a Fusion B&R backup is a Fusion Application (`fapp`). When the Fusion base framework is installed on an Red Hat OpenShift cluster through the IBM Fusion operator, the cluster is being "scanned" for namespaces that do

not belong to the "default set" of Red Hat OpenShift namespaces. For each of those "non-default" namespaces, Fusion creates custom resources (CRs) of type `application.isf.ibm.com`, for example:

- Fusion will create a "wordpress" application CR, if a namespace wordpress exists in the cluster, because wordpress is not an Red Hat OpenShift "default" namespace.
- Fusion will **not** create an "openshift-marketplace" application CR, even if a corresponding namespace exists. This is because "openshift-marketplace" is a "default" namespace in Red Hat OpenShift.

Fusion uses the Application CRs to track parameters and characteristics of Kubernetes applications that relate to features of Fusion services, for example, DR status, backup status, number and size of PVCs, and so on.

In addition, an Application CR allows to specify the `includedNamespaces` parameter - this is important to inform the Fusion Backup & Restore service about containerized applications that span multiple namespaces.

By default, Fusion creates 1:1 mappings of namespaces to applications. These applications can be assigned to a Fusion backuppolicy (`fbp`) later to enable regular backups. The relationship between a Fusion application (`fapp`) and a Fusion backuppolicy creates a Fusion policyassignment (`fpa`) custom resource.

Namespace-scoped resources

As described before, Fusion B&R focuses on Fusion applications, which usually map to a namespace. By default, **all** meta resources of a namespace will be backed up by Fusion Backup & Restore, except those which are explicitly excluded.

Note: Including or excluding resources from backups , as well as specifying custom workflows require the use of Fusion recipes (`frcpe`) , which will be discussed later in this section.

Some examples of namespace-scoped resources are:

- Deployments
- StatefulSets
- Secrets
- Services and Routes
- ServiceAccounts
- Persistent Volume Claims (PVCs)

While PVCs are namespace-scoped resources, Fusion B&R handles them separately from the other resources. Therefore, in a Fusion backup workflow, two different groups of resources are being covered:

- Resource groups (default group name: `rg`)
- Volume groups (default group name: `vg`)

Cluster scoped resources

Cluster-scoped resources are not being backed up by default. However, there are two exceptions:

- A custom Fusion recipe exists that dictates the backup of a cluster-scoped resource that is identified by name or label
- A Persistent Volume (PV) that is being referenced/bound by a claim (PVC) that exists in a namespace that is being backed up

Other cluster-scoped resources, like StorageClasses or ClusterRoles are being ignored.

Persistent volumes

Persistent volumes are storage resources in an Red Hat OpenShift cluster that remain available even if a Pod is deleted or restarted. Therefore, PVs are used to store persistent data of stateful applications. As described before, Fusion detects relationships between (cluster-scoped) PVs and (namespace-scoped) PVCs and ensures that the contents of PVs will be backed up, if the owning PVCs are in scope of a namespace/application backup.

A sample application - WordPress

The architecture of the following example application ("WordPress") can be reviewed for a better understanding of the resources that are in scope or out-of-scope of a Fusion Backup & Restore operation:




Figure: WordPress application architecture diagram showing components

WordPress is a web content management system that allows to publish blogs, host forums, websites and other web content like media galleries. It is written in PHP programming language and consists of two main components, a webserver (front-end) and a MySQL or MariaDB database (back-end).

In the example shown above, all components that make up the WordPress application are grouped into a single namespace ("wordpress").

There are two pods, each of them mounting a PVC to store persistent data:

- The front-end pod mounts a CephFS based volume to store the WordPress binaries and web templates
- The back-end pod mounts a Ceph RBD based volume that stores a MySQL database with configuration data and user entries of a simple guestbook plugin

The front-end service is exposed via a *route*, so it can be accessed by a web-browser.

Note: In this example, the front-end webserver runs on insecure http protocol (TCP port 80), which requires a special configuration in Red Hat OpenShift. The corresponding pod runs under a ServiceAccount "wordpress-sa", which is bound to a very permissive ClusterRole via a

ClusterRoleBinding. We will use this ClusterRoleBinding to demonstrate backup and restore of cluster-scoped resources with the use of Fusion recipes (*frcipe*) later in this section.

The following picture shows the WordPress front-end, accessed via a web-browser (note that it runs on the "primary" SPOKE cluster in our example topology):




Figure: WordPress web interface showing guestbook application

Backup of an application

This section describes how to configure backups for the WordPress sample application using the Fusion web UI.

Note: Application backups can also be configured by making use of Fusion B&R custom resources (CRs) which can be applied as YAML code to the Red Hat OpenShift cluster. This will be discussed later in this section - see section [Backup & restore as code](#).

Create a FusionPolicyAssignment

A Fusion PolicyAssignment (*fpa*) custom resource is created when a backuppolicy is assigned to an application.

Note: If the same backuppolicy is assigned to multiple applications, multiple policyassignment CRs are created as a result. A PolicyAssignment CR is always a 1:1 relation between a backuppolicy and an application.

Step 1: On the HUB, open the Fusion web UI and go to **Backup & Restore > Backed up applications** and click **Protect apps**




Figure: IBM Fusion Backed up applications screen with Protect apps button

Step 2: From the drop-down menu, select the Red Hat OpenShift cluster that **hosts the application** (in this case it is hosted by the "primary" SPOKE), **select the application that should be protected** ("wordpress" in this example) and then click **Next**:




Figure: IBM Fusion application selection dialog for backup

Step 3: From the list of prepared backup policies, **select the policy** that shall be used (in this case it is "protect1.pri.wordpress") and click **Assign**:

Protect applications

Select applications

Assign policies

Assign policies

Select the backup policies you would like to assign.

1 item selected | Select all 4 items

Cancel

☐	Name	Location	Type	Frequency	Start time	End time	Retention	Back up now
☐	minio.hub.wordpress	minio.hub.wordpress	s3	Every day	4:00 AM Central European Standard Time	8:00 AM Central European Standard Time	14 Days	☐
☐	protect1.hub.virtualmach...	protect1.hub.virtualmachines	storageprotect	Every day	5:00 AM Central European Standard Time	9:00 AM Central European Standard Time	14 Days	☐
☐	protect1.pri.softwarehub	protect1.pri.softwarehub	storageprotect	Every day	5:30 AM Central European Standard Time	9:30 AM Central European Standard Time	14 Days	☐
☒	protect1.pri.wordpress	protect1.pri.wordpress	storageprotect	Every day	4:30 AM Central European Standard Time	8:30 AM Central European Standard Time	14 Days	☒

Items per page: 25

1-4 of 4 items

1

1 of 1 page

Cancel

Back

Assign

Figure: IBM Fusion backup policy assignment dialog

Note: Once the backup policy is assigned to an application , an ad hoc backup operation is started , even if the policy is configured with a different time schedule. In some cases, it is not preferable to start an immediate backup, e.g. if the Fusion PolicyAssignment (fpa) needs to be patched with a recipe before the first backup run to achieve application consistency. In this case, turn the "Backup now" toggle-switch off before pressing the Assign button.

Step 4: Go back to **Backup & Restore > Backed up applications** ****and review the newly created PolicyAssignment (fpa):

Backed up applications

Backup status

1 X Filter by cluster...

1 Apps

100%

Completed

In progress

Cancelled

Failed

Cluster: Filter...

Backup status: Filter...

Search

Protect apps +

☐	Name	Cluster	Backup status	Last backup on	Success rate	Policies	Recipes	Backup capacity (GiB)
☐	wordpress	apps.ocp.pri.ibm.lab	Checking inventory		0/1	protect1.pri.wordpress	0	0.00

Items per page: 25

1-1 of 1 item

1

1 of 1 page

Figure: IBM Fusion backed up applications showing WordPress with backup in progress

Step 5: List PolicyAssignments from the CLI: The `policyassignment` (fpa) CRs are created on the HUB cluster. Log in to the HUB cluster from a bastion host and use the following `oc` command to list all PolicyAssignments:

```
oc get policyassignment -n ibm-spectrum-fusion-ns
```

Note: Besides the newly created fpa CR for the wordpress application, you will also find the fpa for the HUB cluster self - service backup with name **ibm-spectrum-fusion-ns-service-protection-policy**. If this fpa is not listed, remember to establish regular self - service backups for the Fusion HUB !

To get more details on a specific `fpa`, list the output in `yaml` format:

```
oc get fpa wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab -n ibm-spectrum-fusion-ns -o yaml

apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: PolicyAssignment
...
spec:
  appCluster: apps.ocp.pri.ibm.lab
  application: wordpress
  backupPolicy: protect1.pri.wordpress
...
status:
  backups: 1
  capacity: 337120150
  lastBackupTimestamp: "2026-02-18T03:30:58Z"
  phase: Assigned
```

Monitoring a backup job

You can review the progress of a backup job from the Fusion web UI and from the CLI:

Step 1: In the Fusion web UI, go to **Backup & Restore > Jobs** and click the backup job that you want to monitor:

IBM Fusion

Jobs /

wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab-202601261036

Completed

Actions

Summary
Last updated: Just now

Summary
Last updated: Just now

List view Log view

Started Jan 26, 2026, 11:36 AM Duration 48 seconds

Inventory

Workloads

Resource	Count
Deployment	1
DeploymentConfig	0
StatefulSet	1
Pod	2

Storage

Resource	Count
PVC	2

Resources

Resource	Count
Route	1
Secret	1
Service	2
ConfigMap	2
VirtualMachine	0

Download resource list

Backup sequence

Resource group	Status
rg	Completed
vg	Completed

Data transfer

Property	Value
Backup location	protect1.pri.wordpress
Backup size	0.31 GiB
Upload speed	15.89 Mbps
Uploading	Complete

Details

Property	Value
Application	wordpress
Size	0.31 GiB
Policy	protect1.pri.wordpress
Recipe	-
Location	protect1.pri.wordpress
Cluster	apps.ocp.pri.ibm.lab
Job ID	de29b6c2-0799-4882-967c-97ae18e5bfe2

Figure: IBM Fusion backup job details showing workflow steps

Note: The backup sequence shows that this job was executed with the built-in default workflow, which runs a backup of all namespace-scoped meta resources (Resource group " rg "), followed by a temporary snapshot of all PVCs in the namespace (Volume group " vg "). When the backup sequence is complete, the Data transfer phase ensures that the VolumeSnapshotContents are copied to external object storage (as the backup policy used in this example is configured to do so) and the temporary PVC snapshots will be removed again.

You can toggle from "List view" to "Log view" to see a more detailed view of the backup job results.

Step 2: From a bastion node, connect to the HUB cluster and issue the following `oc` command to get a list of all Fusion backup (`fbackup`) CRs:

```
oc get backup -n ibm-spectrum-fusion-ns
```

To get more details on a specific backup, list the output in `yaml` format:

```
oc get fbackup wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab-202602180330 -n ibm-spectrum-fusion-n
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: Backup
name: wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab-202602180330
  namespace: ibm-spectrum-fusion-ns
...
spec:
  appCluster: apps.ocp.pri.ibm.lab
  application: wordpress
  backupPolicy: protect1.pri.wordpress
status:
  backupCapacity: 1749921
  endTimestamp: "2026-02-18T03:30:58Z"
  expiration: "2026-03-04T15:30:58Z"
  phase: Completed
...
```

Application consistent backups

In the previous section we demonstrated the backup of a Kubernetes application using Fusion's built-in default workflow. The temporary volume snapshots that are created during the job run are crash-consistent only.

While this might be sufficient for simple applications, there are situations where a more sophisticated backup workflow needs to be applied.

In this section we will examine an IBM Fusion capability called *recipes*, which allows the user or the application platform to define a repeatable, orchestrated backup and recovery procedure to ensure that data can be restored properly anywhere at any time.

Note: Fusion recipes allow to specify orchestrated backup and restore workflows. While custom backup workflows help to create consistent application backups, custom restore workflows help to recover complex applications in the right order (for example, if one app component relies on another component that needs to be restored first).

The default backup (and restore) workflow

The **default backup workflow** that is used by the Backup & Restore service is as follows:

Step 1: Backup all **namespace-scoped resources** that are associated with the application and any dependent cluster-scoped resources.

Step 2: Run any **Velero pre-backup hooks** that are specified through pod annotations in an arbitrary order.

Step 3: Create **snapshots of all PVCs** associated with the application.

Step 4: Run any **Velero post-backup hooks** that are specified through pod annotations in an arbitrary order.

Step 5: If the backup policy specifies backup to an object storage, then **copy the data to object storage**.

The **default restore workflow** that is used by the Backup & Restore service is as follows:

Step 1: Restore all **PVCs**.

Step 2: Restore all **resources** in an arbitrary order.

Step 3: Run **Velero post-restore hooks** that are specified through pod annotations in an arbitrary order.

Recipe overview

The Fusion recipe (`frcpe`) custom resource has three basic specification elements:

- **Groups:** A group defines a set of resources or PVCs that are processed together within a backup or restore. For example, a group can be a subset of specific resources that are specified with an exclude statement or an include statement. Within a group, even more specifications are possible, for example:
 - if a resource is essential (or if a failed backup this resource can be ignored)
 - if a cluster-scoped resource shall be included in the backup
 - if a resource shall be overwritten during a restore, if it already exists
- **Hooks:** Hooks can be used to start external scripts before and after snapshots, scale deployments up and down, or wait for certain conditions (such as pods starting up). The following types of hooks are supported:
 - **exec hooks:** Start arbitrary commands and scripts within target containers.
 - **scale hooks:** Scale workload resources up and down.
 - **check hooks:** Check for conditions on workload resources (such as readiness, replicaCount and so on).
 - **job hooks:** Support to run custom operations during specific events, such as backups, restores or other workflows. It provides a flexible way to run long-running commands, provide dedicated environments for tasks, or perform batch processing and one-time tasks.
 - **label hooks:** Modify resources by applying a new or removing existing labels.
 - **annotation hooks:** Modify resources by applying a new or removing existing annotations.
- **Workflows:** A workflow defines the sequence of steps for a backup or restore operation, specifically the order in which the groups of resources and PVCs must be processed and any hooks that need to be applied.

At least one backup workflow (named `backup`) and one restore workflow (named `restore`) must be specified in the recipe CR.

Note: Recipe CRs can be created in the `ibm-spectrum-fusion-ns` namespace on the cluster that hosts the application, or in any other namespace. If an application developer is responsible to create or modify a recipe, it should be in a namespace that the developer can access (for example, the application namespace that will be protected with the recipe).

More details on creating recipes can be found in the [IBM Fusion documentation](#).

Assigning a recipe

You can assign a recipe to your backup and restore operations. To assign a recipe to backup operations, specify it in the PolicyAssignment (`fpa`) CR. You can specify only a single recipe in each PolicyAssignment CR.

Step 1: Log in to the **HUB** cluster and patch the PolicyAssignment custom resource with the following `oc` command:

```
oc -n ibm-spectrum-fusion-ns patch policyassignment POLICY_ASSIGNMENT_NAME --type merge -p '{"spec":{"recipeName": "RECIPE_NAME", "recipeNamespace": "RECIPE_NAMESPACE"}}
```

where:

- `POLICY_ASSIGNMENT_NAME` is the name as specified in the PolicyAssignment CR.
- `RECIPE_NAME` is the name of the recipe as specified in the Recipe CR.
- `RECIPE_NAMESPACE` is the namespace where the Recipe CR is located.

Step 2: Use the following `oc` command to review the specific PolicyAssignment CR to ensure that it was properly patched (note if there are new recipe-related entries in the spec section):

```
oc get fpa POLICY_ASSIGNMENT_NAME -n ibm-spectrum-fusion-ns -o yaml
```

- The answer to the question of **which recipe is used in a restore** depends on the circumstances:
 - If there is a Recipe CR with the **same name** and **in the same namespace** as the Recipe CR that was used during the backup operation, it will take precedence. It can be helpful in situations when you want to specify a different workflow or different hooks during the restore (compared to what was specified in the recipe used for the backup).
- If there is no Recipe CR with the same name and in the same namespace, the recipe that was used during the backup operation will take precedence. This information is stored as part of the backup metadata and will be automatically applied in the following situations:
 - The original recipe CR does not exist any longer, for example, because it was deleted or contained in the application namespace and the application namespace was deleted (DR restore)
 - The restore is going to a different namespace on the same cluster
 - The restore is going to the same or different namespace on another cluster (cross-cluster restore)
- If the restore should use a recipe with a different name (for example, if a second recipe should be created for recovery testing, or if the restore workflow in the original recipe does not work as expected), a Fusion restore (frestore) custom resource must be created by applying YAML code to the HUB cluster. This frestore CR can refer to the changed/new recipe. Restores triggered from the Fusion web UI will always refer to the original recipe, as it was used during the backup. For an example, see section Backup & restore as code later in this section.

More details on assigning recipes to backup & restore operations can be found in the [IBM Fusion documentation](#).

Example: Create and assign a simple recipe for WordPress

In section Backup of an application we discussed the protection of a sample WordPress application using the default workflow. In this section we will demonstrate how to enhance the backup workflow with a custom recipe.

Step 1: Develop the custom recipe ("wordpress-recipe") to orchestrate backup and restore the WordPress sample application:

```

apiVersion: spp-data-protection.isf.ibm.com/v1alpha1
kind: Recipe
metadata:
  name: wordpress-recipe
  namespace: wordpress
spec:
  appType: wordpress
  groups:
    - name: data_volumes
      type: volume
      labelSelector: 'exclude-pv-from-backup notin (true)'
    - name: meta_resources
      type: resource
    - name: namespace_resources
      backupRef: meta_resources
      type: resource
    - name: cluster-role-binding
      backupRef: meta_resources
      type: resource
      labelSelector: backup-crb-with-fusion=true
      includeClusterResources: true
      restoreOverwriteResources: true
  hooks:
    - name: mysqlscalehook
      type: scale
      namespace: ${GROUP.meta_resources.namespace}
      selectResource: statefulset
      nameSelector: wordpress-mysql
  workflows:
    - name: backup
      sequence:
        - group: meta_resources
        - hook: mysqlscalehook/down
        - hook: mysqlscalehook/sync
        - group: data_volumes
        - hook: mysqlscalehook/up
    - name: restore
      sequence:
        - group: data_volumes
        - group: cluster-role-binding
        - group: namespace_resources

```

As described in "Recipe overview", a recipe has three basic specification elements that will be discussed in more detail:

- **Groups:**

The "wordpress-recipe" defines four groups - **one volume group** ("data_volumes") and **three resource groups** ("meta_resources", "namespace_resources" and "cluster-role-binding").

While we can cover the backup of **all** etcd meta resources within a single resource group ("meta_resources"), the reason for having additional resource group definitions in the recipe is the restore process, where we want to ensure that the ClusterRoleBinding is restored first and all remaining namespace-scoped resources are restored afterward.

In this example, the resource group "meta_resources" is being referenced by the other resource groups via the `backupRef` parameter. This will improve backup performance, as the Fusion B&R service needs to "read" the Red Hat OpenShift `etcd` database only once.

- **Hooks:**

In the WordPress sample application, the backend MySQL database is managed by a StatefulSet. To ensure data consistency, we implemented a **scale hook** "mysqlscalehook", that scales the "wordpress-mysql" StatefulSet up or down.

The StatefulSet is being identified by its name but could also be identified by a label (if this would have been created by the application admins before).

A **scale hook** has three built-in default operations, that can be called from a backup or restore workflow later:

- **down**: scales the StatefulSet down
- **sync**: wait until the desired state (up or down) was reached
- **up**: scales the StatefulSet up

- **Workflows:**

The "wordpress-recipe" implements the following **backup workflow**:

- Backup all meta resources that reference the `meta_resources` group. Besides the namespace-scoped resources of the wordpress application, this will include the ClusterRoleBinding that is being labelled with the `backup-crb-with-fusion=true` label.
- Scale down the MySQL database
- Wait until the StatefulSet is really scaled down
- Snapshot all PVCs, as defined in the `data_volumes` volume group
- Scale up the MySQL database

Note: When the backup workflow of the recipe is completed, Fusion starts to clone the PVC snapshots to new volumes, mounts those to velero DataUploader pods, copies the contents to external object storage and finally removes the clone volumes and PVC snapshots. The recipe itself does not need to take care of the upload process.

The "wordpress-recipe" implements the following **restore workflow**:

- Restore all PVCs, as defined in the `data_volumes` volume group
- Restore the cluster-scoped ClusterRoleBinding, as specified in resource group `cluster-role-binding`
- Restore all remaining meta resources that belong to namespace `wordpress`

Step 2: Apply the recipe to the Red Hat OpenShift cluster:

In this example, the recipe CR will be created in the wordpress namespace on the **"primary" SPOKE cluster**. Copy the recipe YAML code to a file and apply it to the cluster:

```
[root@ocp-bastion]# oc project wordpress
Now using project "wordpress" on server
"https://api.ocp.**pri**.ibm.lab:6443".
```

```
[root@ocp-bastion]# oc apply -f wordpress-recipe.yaml
recipe.spp-data-protection.isf.ibm.com/wordpress-recipe created
```

```
[root@ocp-bastion]# oc get frcpe -n wordpress
NAME                                AGE                PARENT RECIPE
wordpress-recipe                   2026-01-26T11:17:01Z
```

Step 3: Patch the Fusion PolicyAssignment custom resource:

As all PolicyAssignment CRs exist on the **Fusion HUB**, log in to the HUB cluster and patch the `fpa` that connects the wordpress application to its backuppolicy:

```
[root@ocp-bastion]# oc project ibm-spectrum-fusion-ns
Now using project "ibm-spectrum-fusion-ns" on server "https://api.ocp.hub.ibm.lab:6443".
```

```
[root@ocp-bastion]# oc -n ibm-spectrum-fusion-ns patch policyassignment wordpress-protect1.pri.wordpress-a
policyassignment.data-protection.isf.ibm.com/wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab patched
```

Step 4: Review the specific PolicyAssignment CR to ensure that it was properly patched (note the new recipe-related entries in the spec section):

```
[root@ocp-bastion]# oc get fpa wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab -n ibm-spectrum-fusio
```

```
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: PolicyAssignment
...
spec:
  appCluster: apps.ocp.pri.ibm.lab
  application: wordpress
  backupPolicy: protect1.pri.wordpress
** recipe:
**   apiVersion: spp-data-protection.isf.ibm.com/v1alpha1
**   name: wordpress-recipe
**   namespace: wordpress
...
```

Review the backup job that was using the assigned recipe

When the next scheduled or ad hoc backup is completed, open the Fusion web UI, go to **Backup & restore > Jobs** and click the latest wordpress application backup:




Figure: IBM Fusion backup job showing custom recipe workflow

Note that the Backup sequence shows the new workflow, as specified in the *recipe* CR.

IBM provided recipes for complex apps (e.g. IBM Software Hub)

As Kubernetes applications can become much more complex than the simple WordPress demo that we discussed earlier in this document, there will be also much more sophisticated backup & restore workflows required to always ensure backup consistency and safe recovery of those apps.

While it is hardly possible for backup administrators to determine the best backup approach for a containerized application that they don't know in detail, IBM is shipping customized, [ready-to-use recipes with its most popular Red Hat OpenShift applications](#)

One example of this is IBM Software Hub. IBM Software Hub serves as a universal platform for on-premises deployments, supporting diverse workloads across IBM's broader Data and AI portfolio, including watsonx and IBM Cloud Pak for Data.

The process of enabling backup & restore for IBM Software Hub with Fusion is using approaches and [tools](#) that a Cloud Pak application administrator will be familiar with.

The following list shows a summary of the tasks required to enable backup & restore for IBM Software Hub:

- Prepare a shell script to [set environment variables](#) that help to run sample commands as-is, allowing to do cut & paste from the online documentation. The same approach is being used when deploying an instance of IBM Software Hub – there is just a few more [Fusion specific environments](#) to be added.
- [Patch the velero component of Fusion](#) to increase resource limits
- Patch Fusion with additional plugins for swhub, db2u or Cloud Pak Foundational Services (cpfs)
- Patch the Software Hub operator application CR (`fapp`) the to cover the Software Hub operands namespace(s) in the `includedNamespaces` section
- Leverage the cpd-cli oadp install tool to [implement a Cloud Pak for Data "side-car" container](#) (cpdbr-tenant pod)
- Install a "Parent" recipe "ibmcpd-tenant" in the Software Hub operator namespace
- Dependent on the installed Software Hub services (for example, db2u, db2wh, and so on), deploy a set of "Child" recipes in the Software Hub operator namespace

Once the preparation tasks are complete, the remaining tasks are the same as for all other applications:

- Create a Fusion BackupStorageLocation (`fbsl`) on external object storage)
- Create a Fusion BackupPolicy (`fbp`)
- Assign the BackupPolicy to the Software Hub operator application (do not start an ad hoc backup yet)
- Patch the resulting PolicyAssignment (`fpa`) to use the ibmcpd-tenant recipe
- Start a backup job and monitor the proper execution

The following figure shows the backup sequence of the pre-defined recipe shipped with IBM Software Hub:




Figure: IBM Software Hub backup workflow using ibmcpd-tenant recipe

A detailed description on how to backup and restore IBM Software Hub with IBM Fusion can be found in the [online documentation](#)

GitHub repository for recipes

As application or backup administrators might oversee the development of other, complex recipes (for example, for non-IBM applications), IBM has established a GitHub repository with a collection of recipe examples for a bunch of popular applications and databases.

This repository can be leveraged to get ideas, or to combine fragments from various recipes into your own recipe.

As an example, you'll find a recipe to protect a WordPress application that uses a different approach to achieve application consistency for the backend database. Instead of scaling the MySQL pod (as we do in the sample recipe described earlier) the GitHub example proposes to use a MySQL utility ("FLUSH TABLES WITH READ LOCK") to prepare the database for a subsequent PVC snapshot. The recipe collection can be found here: <https://github.com/IBM/storage-fusion/tree/master/backup-restore/recipes>

Restoring an application

A restore operation of a K8s/Red Hat OpenShift application can be triggered either by the Fusion web UI, or by applying a Fusion restore CR (`frestore`) to the HUB cluster. The restore action allows you to perform the following operations:

- Select the application to be restored
- In a HUB/SPOKE topology, select the target cluster for the restore
- Select whether to restore to the original or to an alternate namespace
- Select the backup image (PIT) to be restored
- Select to change the target StorageClass for persistent volumes

Note:

- Some applications do not allow to be restored to an alternate namespace.
- Some applications that use "hardcoded" namespaces or storage classes might not be able to leverage the Fusion built-in mechanism to change these parameters to alternate values . In this case, a Transform CR could be used.

In the following sections we will demonstrate the restore of a small WordPress and a more complex application (IBM Software Hub).

Cross-cluster restore of the WordPress sample application

This section describes how to restore the sample WordPress application to the same namespace, but on another Red Hat OpenShift cluster.

Assume the original application on the "primary" SPOKE was hit by a cyberattack, so it needs to be recovered to a different cluster that is not impacted. The cyberattack happened on February 2, 2026, at 6:09 PM:




Figure: WordPress application showing cyber-attack message

Step 1: In the Fusion web UI on the HUB cluster, go to **Backup & restore > Backed up applications**, open the **context-menu** of the wordpress application and click **Restore**:




Figure: IBM Fusion backed up applications with restore context menu

Step 2: Select the **"secondary" SPOKE** from the **Target cluster drop-down menu** and select **Use same project** for the **Project destination**:




Figure: IBM Fusion restore target cluster and project selection

Step 3: Select the backup PIT to restore (ensure the backup was taken before the cyberattack happened):

Select a backup						
Q Search ⚙						
Backup time	Location	Policy	Status	Size	PVs	
☒ Feb 2, 2026, 5:28 PM	protect1.pri.wordpress	protect1.pri.wordpress	✓ Completed	< 0.01 GiB	2	
☐ Feb 2, 2026, 5:09 PM	protect1.pri.wordpress	protect1.pri.wordpress	✓ Completed	< 0.01 GiB	2	
☐ Feb 2, 2026, 4:30 AM	protect1.pri.wordpress	protect1.pri.wordpress	✓ Completed	< 0.01 GiB	2	
☐ Feb 1, 2026, 4:30 AM	protect1.pri.wordpress	protect1.pri.wordpress	✓ Completed	< 0.01 GiB	2	
☐ Jan 30, 2026, 4:30 AM	protect1.pri.wordpress	protect1.pri.wordpress	✓ Completed	< 0.01 GiB	2	
☐ Jan 27, 2026, 9:19 AM	protect1.pri.wordpress	protect1.pri.wordpress	✓ Completed	< 0.01 GiB	2	

Figure: IBM Fusion backup version selection for restore

Step 4: Decide whether the target StorageClass should be changed or not:

Target storage class

By default, the restore occurs to the same storage class name as the backup PVCs. If the storage class name is not found on the target cluster, the default storage class of target cluster is used instead. If you wish to use a different storage class for the restore, specify it below. [Learn more](#)

☐ Change target storage class ⓘ

Figure: IBM Fusion restore StorageClass mapping configuration

Step 5: Review the summary and confirm it to start the restore.

Step 6: In the Fusion web UI, go to **Backup & restore > Jobs** and **click the Restore tab** to list the history of completed and running restore operations:

Backups Restores									
Q Search ⌵ ⚙									
Name	Target cluster	Application	Status	Started	Elapsed time	Finished	Policy	Recipe	
restore-wordpress-202602021646	apps.ocp.sec.ibm.lab	wordpress	Restore snapshot in progress	Feb 2, 2026, 5:46 PM	0:00:31	–	protect1.pri.wordpress	–	

Figure: IBM Fusion jobs screen showing restore operations

Step 7: Click the wordpress restore item and review the job details.




Figure: IBM Fusion restore job details and workflow

Note: that the restore operation uses the custom workflow as defined in the wordpress-recipe, even if the target namespace (and thus the *frcpe* CR) did not exist on the target cluster before. This is because the recipe is part of the backup meta data and can be "recalled" by the Fusion B&R service during a restore process.

Step 8: Perform post-restore tasks.

Dependent to whether the application was restored to the original or to a different target cluster and to the method how services are exposed to the outside world, it might be necessary to perform post-restore tasks - either within the restored application itself, or within remote or external systems that have a connection to the restored Red Hat OpenShift application.

In our sample WordPress application, the external route (the URL of the website) is being created dynamically and will change if the application is being restored to a cluster with a different domain. WordPress stores its own URL inside the MySQL backend-database, so these database entries must be modified after the restore - we do this with a small script:

```
[root@ocp-bastion wordpress]# ./url-fix.sh
Login successful.

## Replace wordpress home and siteurl in wp_options table with proper## route from restore target cluster

## Display home and siteurl as it was restored:
option_id    option_name    option_value    autoload
2            siteurl        http://wordpress-wordpress.apps.ocp.pri.ibm.lab on
3            home          http://wordpress-wordpress.apps.ocp.pri.ibm.lab on

## Display wordpress home and siteurl after change:
option_id    option_name    option_value    autoload
2            siteurl        http://wordpress-wordpress.apps.ocp.sec.ibm.lab on
3            home          http://wordpress-wordpress.apps.ocp.sec.ibm.lab on

## Restart wordpress frontend and backend pods to activate the changes
pod "wordpress-mysql-0" deleted
pod "wordpress-6f7dcf8fcf-pb874" deleted
```

Note: Changes within the restored application could also be performed as part of the restore workflow of a custom recipe.

Step 9: Check that the wordpress application was successfully restored to a PIT before the cyberattack:




Figure: WordPress application successfully restored on secondary SPOKE

Note: that the application URL now uses the domain of the "secondary" SPOKE.

Cross-cluster restore of IBM Software Hub

In short, the restore process of IBM Software Hub can be described as an orchestrated application deployment from scratch, but with injection of persistent configuration and user data.

While the restore process re-deploys the application, the target cluster needs to be prepared in the same way as if the installation would be performed manually.

The following list shows a summary of the [tasks required to prepare a target cluster](#) for a restore of IBM Software Hub:

- [Set environment variables](#) Fusion specific environments and add the [Fusion specific environments](#) that match the TARGET cluster
- Log in to the target cluster using the `oc` tool
- [Patch the velero component of Fusion](#)
- Patch Fusion with additional plugins for swhub, db2u or Cloud Pak Foundational Services (cpfs)
- Leverage the cpd-cli oadp install tool to implement a Software Hub specific ClusterRole (*cpdbr-tenant-service-clusterrole*) and ClusterRoleBinding (*cpdbr-tenant-service-crb*)
- Install the prerequisite software - IBM License Service and the Certificate Manager (if not already present on the target cluster)
- Optional (if this is not the first restore): [Clean up the target cluster](#) before starting another restore

Once the preparation tasks are complete, the remaining tasks are the same as for all other application restores:

- Install the Fusion base framework and the Fusion B&R agent, if the restore is directed to a new or re-built SPOKE cluster
- Start a restore job and monitor the proper execution
- [Complete any post-restore tasks](#), if required (IBM Software Hub can provide various features and services, where some of them require special tasks after a restoration).

Understanding of job details - "how" and "what"

As described before, you can review the backup or restore job logs via the Fusion web UI, or even generate detailed yaml outputs of *fbackup* or *frestore* CRs to understand if a backup or restore job was successful.

However, in some situations it might be helpful to understand the detailed workflows and the contents of a backup or restore operation (for example, when troubleshooting a failed recovery).

IBM offers two scripts that help to understand the B&R workflows and to obtain a list of resources that have been in scope of an operation.

Analyzing the workflow of a job that used a recipe

The script [getRecipeWorkflow.sh](#) can be run from a bastion node to list the details of a job workflow. The following example shows the abbreviated output of a backup job for the wordpress sample application (which uses a recipe for backup):

```
[root@ocp-bastion HUB]# ./getRecipeWorkflow.sh backup -n wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm
{
  "uid": "b0f0d3df-3d50-43aa-a662-cc015b891735",
  "type": "backup",
  "logsType": "RecipeWorkflow",
}
'2026-02-23 14:14:56 [INFO]: === Backup & recipe validation ===
2026-02-23 14:14:56 [INFO]: App namespace: ibm-spectrum-fusion-ns name: wordpress
2026-02-23 14:14:57 [INFO]: Recipe found with labels = dp.isf.ibm.com/parent-recipe:wordpress-recipe and d
2026-02-23 14:14:57 [INFO]: Evaluating volume group data_volumes with 2 pvc candidates...
2026-02-23 14:14:59 [INFO]: Starting workflow "backup" of recipe "wordpress:wordpress-recipe"
2026-02-23 14:14:59 [INFO]: Start execution sequence "group/meta_resources" ...
2026-02-23 14:15:09 [INFO]: The backup operation of resources from namespace ['\''wordpress'\'''] completed
2026-02-23 14:15:09 [INFO]: Executing scale hook/operation mysqlscalehook/down
2026-02-23 14:15:10 [INFO]: Scale hook op mysqlscalehook: found statefulset: wordpress-mysql, replicas: 1
2026-02-23 14:15:10 [INFO]: statefulset: name: wordpress-mysql scaled down successfully, current replica c
2026-02-23 14:15:10 [INFO]: Start execution sequence "hook/mysqlscalehook/sync" ...
2026-02-23 14:15:10 [INFO]: Executing scale hook/operation mysqlscalehook/sync
2026-02-23 14:15:10 [INFO]: Scale hook op mysqlscalehook: found statefulset: wordpress-mysql, replicas: 0
2026-02-23 14:15:15 [INFO]: Scale hook/operation mysqlscalehook/sync completed successfully.
2026-02-23 14:15:15 [INFO]: Start execution sequence "group/data_volumes" ...
2026-02-23 14:15:17 [INFO]: Reevaluation complete: 1 volume groups
2026-02-23 14:15:17 [INFO]: Vols in vg: 2
2026-02-23 14:15:17 [INFO]: Volume from inventory: wordpress:mysql-pvc
2026-02-23 14:15:17 [INFO]: Volume from inventory: wordpress:wp-pvc
2026-02-23 14:15:17 [INFO]: Executing VolumeGroup data_volumes with workflow type backup...
2026-02-23 14:15:17 [INFO]: Starting CSI snapshot of PVC mysql-pvc
2026-02-23 14:15:17 [INFO]: Starting CSI snapshot of PVC wp-pvc
2026-02-23 14:15:21 [INFO]: End execution sequence "group/data_volumes" completed successfully.
2026-02-23 14:15:21 [INFO]: Executing scale hook/operation mysqlscalehook/up
2026-02-23 14:15:21 [INFO]: Scale hook op mysqlscalehook: found statefulset: wordpress-mysql, replicas: 0
2026-02-23 14:15:21 [INFO]: statefulset: name: wordpress-mysql scaled up successfully to 1 replicas, curre
2026-02-23 14:15:21 [INFO]: Scale hook/operation mysqlscalehook/up completed successfully.
2026-02-23 14:15:21 [INFO]: Execution of workflow backup completed. Number of failed commands: 0, 0 are es
2026-02-23 14:15:21 [INFO]: Execution of workflow "backup" of recipe "wordpress:wordpress-recipe" complete
2026-02-23 14:15:24 [INFO]: **=== Post recipe execution ===**
2026-02-23 14:15:38 [INFO]: Data upload is in progress. Successfully transferred 88308639.
2026-02-23 14:15:40 [INFO]: Data upload is in progress. Successfully transferred 321773514.
2026-02-23 14:15:45 [INFO]: Cleanup the snapshots created during the backup.
2026-02-23 14:15:45 [INFO]: Deleting snapshot 1b3905e2-6eb9-44d7-9b0a-5f7a87253a8f-1771856117.626644 in ns
2026-02-23 14:15:45 [INFO]: Deleting snapshot 0dd1f2cf-61d5-42ab-8f13-0cd96d58c150-1771856117.6330936 in n
2026-02-23 14:15:47 [INFO]: Copy backup job completed successfully.
```

Analyzing the resources that have been backed up or restored

The script [getResources.sh](#) can be run from a bastion node to list the details of all re-sources that have been contained in a backup or restore job. The following example shows the abbreviated output of a backup job for the wordpress sample application:

```
[root@ocp-bastion HUB]# ./getResources.sh backup -n wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab-
{
  "uid": "b0f0d3df-3d50-43aa-a662-cc015b891735",
  "type": "backup",
  "logsType": "ResourceDetails",
  "logs":
===== CSV format =====
'namespace,name,group,version,kind,count,recipeGroup
,pvc-29df9d1f-a79e-4067-abb7-757f72e36b82,,v1,PersistentVolume,1,data_volumes
,pvc-fb8a6529-64a2-4482-a463-4e9a2516981a,,v1,PersistentVolume,1,data_volumes
wordpress,mysql-pvc,,v1,PersistentVolumeClaim,1,data_volumes
wordpress,wp-pvc,,v1,PersistentVolumeClaim,1,data_volumes
,clusterserviceversions.operators.coreos.com,apiextensions.k8s.io,v1,CustomResourceDefinition
,recipes.spp-data-protection.isf.ibm.com,apiextensions.k8s.io,v1,CustomResourceDefinition,1,meta_resources
wordpress,wordpress-mysql-69f4c7655d,apps,v1,ControllerRevision,1,meta_resources
wordpress,wordpress,apps,v1,Deployment,1,meta_resources
wordpress,wordpress-6f7dcf8fcf,apps,v1,ReplicaSet,1,meta_resources
wordpress,wordpress-mysql,apps,v1,StatefulSet,1,meta_resources
...
,system:openshift:scc:anyuid,rbac.authorization.k8s.io,v1,ClusterRole,1,meta_resources
,system:openshift:scc:anyuid,rbac.authorization.k8s.io,v1,ClusterRoleBinding,1,meta_resources
...
wordpress,wordpress,route.openshift.io,v1,Route,1,meta_resources
wordpress,wordpress-recipe,spp-data-protection.isf.ibm.com,v1alpha1,Recipe,1,meta_resources
wordpress,kube-root-ca.crt,,v1,ConfigMap,1,meta_resources
wordpress,openshift-service-ca.crt,,v1,ConfigMap,1,meta_resources
...
,wordpress,,v1,Namespace,1,meta_resources
wordpress,wordpress-6f7dcf8fcf-pb874,,v1,Pod,1,meta_resources
wordpress,wordpress-mysql-0,,v1,Pod,1,meta_resources
wordpress,mysql-pass-5m579c2th2,,v1,Secret,1,meta_resources
wordpress,wordpress,,v1,Service,1,meta_resources
wordpress,wordpress-mysql,,v1,Service,1,meta_resources
...
wordpress,wordpress-sa,,v1,ServiceAccount,1,meta_resources
```

Self-service backup & restore

In a default Fusion Backup & Restore setup, it is common practice to grant backup administrators with cluster-admin permissions on the Fusion B&R HUB. Cluster-admin permissions are granted for the following two reasons:

- Cluster-admin permissions are mandatory for accessing the Fusion web UI.
- All B&R relevant custom resources are in the *ibm-spectrum-fusion-ns* namespace on the HUB cluster. Thus, a user that manages the Fusion B&R service configuration and performs backup and restore operations must have admin permissions over the *ibm-spectrum-fusion-ns* namespace on the HUB.

Whenever an application owner or developer wants to configure backup for a new application or needs to trigger a restore, a backup- or cluster-administrator needs to be involved ("help-desk approach").

However, it is also possible to allow non-admin users to view, edit, create or delete Fusion B&R custom resources, so that they can trigger backups & restores, or even create new backup storage locations or policies on their own ("self-service approach").

Self-service capabilities for non-cluster-admins can be achieved in two ways:

- Create self-service Backup & Restore CRs in the application namespace:

In this case, all self-service Backup & Restore CRs, including scheduled Backup CRs exist in the namespace of the application.

To create Backup & Restore CRs in the application's namespace, you need to be an application namespace administrator, an application user, or a service account with granted permission to the namespace. Application namespace administrators can directly create Backup & Restore CRs within their application namespace.

Non-admin application users and service accounts need the namespace administrator to first grant IBM Fusion Backup & Restore permissions through standard Red Hat OpenShift RoleBindings. These RoleBindings grant permissions for the user or service account in the application's namespace with the specific Backup & Restore roles.

There are roles that grant a non-admin user with granular permission to the following backup & restore CRs:

- Backup Policy (*fbp*)
- Backup (*fbackup*)
- Backup Storage Location (*fbsl*)
- Delete Backup Request (*fdbr*)
- Policy Assignment (*fpa*)
- Recipe (*frcpe*)
- Restore (*frestore*)
- Recover Backup CRs (*frbkpcrs*)

Note: The "recoverbackupcrs" custom resource was specifically introduced to support self-service for non-cluster admins. It allows a user to request the re-creation of B&R custom resources that existed in an application namespace which was accidentally deleted.

The roles that can be assigned to a non-admin user via RoleBinding are of type "view", "crdview", "edit" or "admin" and adhere to the following naming scheme:

```
<Fusion backup & restore CR>-v1alpha-<role>
```

Example: If a non-admin user shall be granted permission to edit backup policies in a namespace, the user needs to be bound to the following role:

```
backuppolicies.data-protection.isf.ibm.com-v1alpha-edit
```

A more detailed example on configuring Fusion B&R custom resources by a non-cluster-admin can be found in the 'Backup & restore as code' section .

Custom Backup & Restore administrator

Another mode of operation for managing the Backup & Restore service is to grant a non-admin user the admin of the Backup & Restore service. To do this, create the RoleBindings of the user in the IBM Fusion namespace (*ibm-spectrum-fusion-ns*) with the Backup & Restore admin roles.

This user does not have IBM Fusion web UI access but can manage Backup & Restore activity of **all** applications on both hub and spoke clusters through Backup & Restore CRs within the IBM Fusion namespace.

The following limitations apply to self-service B&R for both HUB and SPOKE:

- Only command line support is available.
- Cross-cluster restore is not available.

- Self-service jobs are not shown in the IBM Fusion web UI.
- Spoke self-service policies and backup storage locations are not shown in the IBM Fusion web UI.
- Applications with only self-service policies or completed self-service backup or restore jobs do not appear in the backed-up applications table.
- Applications configured with self-service policies do not appear in the IBM Fusion web UI: **Backed Up Applications > Protect Apps**.
- On the Topology table, the backed-up apps ratio and success rate may also be inaccurate as it includes self-service applications and self-service backup jobs.

More details on Self-service backup and restore can be found in the [IBM Fusion documentation](#).

Backup & restore as code

Besides managing backup and restore through the Fusion web UI, it is also possible to make use of the backup and restore services through custom resources (CRs). One example use case for this would be to integrate B&R functionalities with the application code. This allows to manage application backups by using a declarative state and to maintain the backup posture for an application alongside the application in Git. The following are examples where using B&R as code is mandatory:

- Use a new or modified recipe during a restore operation which differs from the recipe that was used during backup
- Manage B&R operations for non-cluster-admins

Example 1: change the recipe used for restoring an application

As discussed in section "Application consistent backups", a recipe defines the workflows to be used for backup and restore. There might be situations where a restore is not working properly, because:

- The backups have been created without using a custom recipe (using the default workflow) and it was noticed later that the default restore workflow is not feasible.
- The backups have been created with a recipe, but it was noticed during recovery testing that the prepared custom restore workflow does not work as expected and needs to be changed.

As it is not possible to do an "on-demand" change of the recipe during a restore operation when triggered via the Fusion web UI, a restore CR can be applied to the cluster via yaml code.

In the following example, we assume that a modified recipe "wordpress-restore-recipe" was created in the wordpress namespace and should be applied during a cross-cluster restore of the sample wordpress application to the "secondary" SPOKE.

Step 1: To restore as code using a different recipe, on a bastion node, create a `restore-with-modified-recipe.yaml` file with the following content:

```
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: Restore
metadata:
  annotations:
    dp.isf.ibm.com/provider-name: isf-backup-restore
  name: wordpress-restore-with-modified-recipe-01
  namespace: ibm-spectrum-fusion-ns
spec:
  backup: wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab-202602231412
  targetCluster: apps.ocp.sec.ibm.lab
  recipe:
    name: wordpress-restore-recipe
    namespace: wordpress
```

Step 2: Log in to the Fusion hub and apply the yaml to the cluster:

```
oc apply -f ./restore-with-modified-recipe.yaml
```

As a result, a restore operation will start that uses a recipe called `wordpress-restore-recipe` to perform a restoration of backup image `wordpress-protect1.pri.wordpress-apps.ocp.pri.ibm.lab-202602231412` to target cluster `apps.ocp.sec.ibm.lab`. Note that the recipe CR needs to be created in advance so that it can be referenced by the restore CR later.

Example 2: self-service backup & restore for a non-cluster-admin

In this example, a user named **redbook** is responsible to manage backup and restore for a project called **self-service**.

Note: User redbook is a project administrator for namespace self-service, but does not have the cluster-admin role. Thus, user redbook cannot use the Fusion web UI but must create Fusion B&R custom resources "as code" within the application namespace.

The following steps demonstrate the self-service operations a non-cluster-admin user can perform:

- Create a Fusion backup storage location
- Create a Fusion backup policy using the previously defined location
- Trigger an ad hoc backup using the previously defined policy
- Delete the Fusion backup custom resource again within a specific project using yaml code

Step 1: From the bastion node, log in to the Red Hat OpenShift cluster as a non-cluster-admin user (project-admin user) "redbook":

```
[root@ocp-bastion users]# **oc login** --token=sha256~F9y4FuraqlPzvikG-CzJGn4eFqNAckxvvdvZDjslbc --server
**Logged into "https://api.ocp.hub.ibm.lab:6443" as "redbook"** using the token provided.
You have access to the following projects and can switch between them with 'oc project <projectname>':
  self-service
Using project "".

[root@ocp-bastion self-service]# **oc project self-service**
Now using project "self-service" on server "https://api.ocp.hub.ibm.lab:6443".
```

Step 2: Create a Fusion backup storage location (*fbsl*) in the application namespace. Create a file `self-service-fbsl.yaml` with the following sample content:

```

apiVersion: v1
data:
  access-key-id: <base64-encoded access-key>
  secret-access-key: <base64-encoded secret-access-key>
kind: Secret
metadata:
  name: minio.hub.self-service-secret-0
  namespace: self-service
---
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: BackupStorageLocation
metadata:
  name: minio.hub.self-service
  namespace: self-service
spec:
  credentialName: minio.hub.self-service-secret-0
  params:
    bucket: hub.self-service
    endpoint: 'https://s3.ibm.lab:9000'
  privateCertificateName: minio-cert-secret
  provider: isf-backup-restore
  type: s3

```

to Create the BackupStorageLocation using code, apply the YAML to the cluster. Note that the resulting *fbsl* resource was created in the application namespace (*self-service*) and not in the Fusion namespace:

```

[root@ocp-bastion self-service]# **oc apply -f self-service-fbsl.yaml**
secret/minio.hub.self-service-secret-0 created
backupstoragelocation.data-protection.isf.ibm.com/minio.hub.self-service created

[root@ocp-bastion self-service]# **oc get fbsl -n self-service**
NAME                                PHASE      STORAGETYPE
minio.hub.self-service      Connected   s3

```

Step 3: Create a Fusion backup policy custom resource (*fbp*) in the application namespace Create a file `self-service-fbp.yaml` with the following sample content:

```

apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: BackupPolicy
metadata:
  name: minio.hub.self-service
  namespace: self-service
spec:
  provider: isf-backup-restore
  backupStorageLocation: minio.hub.self-service
  retention:
    number: 3
    unit: days
  schedule:
    cron: '00 4 * * * '
    timezone: Europe/Berlin
    windowEndTime: "08:30"

```

Apply the YAML to the cluster to create the BackupPolicy using code. Note that the resulting *fbp* resource was created in the application namespace (*self-service*) and not in the Fusion namespace:

```
[root@ocp-bastion self-service]# oc apply -f self-service-fbp.yaml
backuppolicy.data-protection.isf.ibm.com/minio.hub.self-service created

[root@ocp-bastion self-service]# oc get fbp -n self-service
```

NAME	BACKUPSTORAGELOCATION	SCHEDULE	RETENTION	RETENTIONUNIT
minio.hub.self-service	minio.hub.self-service	00 4 * * *	3 days	

Step 4: Trigger an ad hoc backup using the policy (*fbbackup*) in the application namespace Create a file `self-service-fbackup.yaml` with the following sample content:

```
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: Backup
metadata:
  name: self-service-backup-01
  namespace: self-service
spec:
  appCluster: apps.ocp.hub.ibm.lab
  application: self-service
  backupPolicy: self-service/minio.hub.self-service
```

Apply the YAML to the cluster. Note that the resulting *fbbackup* resource was created in the application namespace (*self-service*) and not in the Fusion namespace:

```
[root@ocp-bastion self-service]# oc apply -f self-service-fbackup.yaml
backup.data-protection.isf.ibm.com/self-service-backup-01 created

[root@ocp-bastion self-service]# oc get fbackup -n self-service
```

NAME	CLUSTER	APPLICATION	PHASE
self-service-backup-01	apps.ocp.hub.ibm.lab	self-service	Completed

Step 5: Finally, delete the backup again using a Fusion DeleteBackupRequest (*fdbr*): Create a file `self-service-fdbr.yaml` with the following content:

```
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: DeleteBackupRequest
metadata:
  name: self-service-backup-01-delete-request
  namespace: self-service
spec:
  backup: self-service-backup-01
```

Apply the YAML to the cluster. Note that the resulting *fdbr* resource will start to delete the selected backup image, including data that was copied to external object storage. It might take a while until the *fbbackup* resource is removed from the self-service namespace again:

```
[root@ocp-bastion self-service]# oc apply -f self-service-fdbr.yaml
deletebackuprequest.data-protection.isf.ibm.com/self-service-backup-01-delete-request created

[root@ocp-bastion self-service]# oc get fbackup -n self-service
No resources found in self-service namespace.
```

More examples and details on backup & restore as code can be found in the [IBM Fusion Documentation](#)

Virtual machines

This section describes how to backup and restore virtual machines running on Red Hat OpenShift Virtualization (OCP-V) using the IBM Fusion Backup & Restore service.

Red Hat OpenShift Virtualization basics

Red Hat OpenShift Virtualization allows you to run virtual machines (VMs) alongside containers on an Red Hat OpenShift cluster. It is powered by KubeVirt, an open-source extension that adds virtualization capabilities to Kubernetes, treating VMs as native Kubernetes resources.

Red Hat OpenShift users can define VMs declaratively, and the platform orchestrates them using the same tools, APIs, scheduling, networking, and storage mechanisms that any other containers use.

KubeVirt: The virtualization engine inside Red Hat OpenShift

KubeVirt extends Kubernetes by adding:

- Custom Resource Definitions (CRDs) such as `VirtualMachine` (VM) and `VirtualMachineInstance` (VMI) to represent VM workloads.
- Controllers and daemons that manage VM lifecycle tasks.
- A virtualization stack using `libvirt` and QEMU/KVM to run VMs.

VMs run in a special pod called a **virt-launcher** pod. This pod includes QEMU/KVM binaries and the environment required to run the VM. Inside the pod, the VM runs as a QEMU-KVM process. These processes use the underlying node's Linux kernel KVM module for hardware acceleration.

If you inspect such a virt-launcher pod, you will find the QEMU-KVM processes and associated tools, like `qemu-img`, `virsh`, and so on.

When a user applies a VM or VMI YAML definition, the `virt-controller` detects it and initiates the creation and scheduling of a pod that will ultimately host the VM.

Because virtualization happens inside a pod, the VM inherits all Kubernetes pod behaviors, like:

- Scheduling rules (affinity/anti-affinity, taints, tolerations)
- HA and restart semantics
- Resource cgroups and namespaces

KubeVirt integrates with host resources:

- Networking: via Multus or SR-IOV
- Storage: via standard Kubernetes StorageClasses (PVCs)
- Security: via RBAC and SELinux

The key components of KubeVirt

The KubeVirt system consists of three main components:

- `virt-controller` (the control-plane orchestrator)
 - Watches VM/VMI objects and ensures the correct pod (virt-launcher) is created.
 - Maintains the desired state of VMs (start/stop/restart).
- `virt-handler` (node-level agent)
 - Runs on every OCP cluster node capable of hosting VMs.
 - Manages local VM operations and interacts with KubeVirt's runtime (libvirt/QEMU).
 - Ensures VMs launch and run properly on a node.

- `virt-launcher` (per-VM pod)
 - A pod specifically created to host a single VirtualMachineInstance (VMI).
 - Contains QEMU/KVM and runs the VM process.
 - Provides isolation and integrates the VM with Kubernetes networking/storage.

How Fusion B&R handles Virtual Machines

General VM backup & restore considerations

As described in section "Scope of Fusion Backup and Restore operations", the Fusion B&R backup policies are applied to Fusion Application CRs (`fapp`), which usually map 1:1 to a namespace.

A namespace can contain multiple VMs (and thus multiple `virt-launcher` pods). If a backup policy is assigned to an application that refers to a "multi-VM" namespace, all VMs in that namespace will be backed up by the resulting backup jobs. It is not possible to include or exclude VMs from a backup or restore job for a given namespace.

Important: If the intended level of granularity for a backup or restore job is on a per VM basis, ensure to have a dedicated namespace/project per VM, otherwise all VMs inside of the namespace will be backed up / restored at the same time.

While a VirtualMachineInstance (VMI) is hosted by a `virt-launcher` pod that belongs to a namespace, the procedure to enable backup and restore for a virtual machine in Red Hat OpenShift is essentially identical to the procedure for containerized applications. However, there are some differences in Fusion Backup & Restore workflow if a namespace contains one or more VMs.

VM backup workflow

The Fusion B&R **backup workflow** for namespaces that contain VMs is as follows:

Step 1: Backup all namespace-scoped resources that are associated with the application and any dependent cluster-scoped resources.

Step 2: Run any Velero pre-backup hooks that are specified through pod annotations in an arbitrary order.

Step 3: Create snapshots of all PVCs associated with applications in the namespace. **IMPORTANT: PVCs that belong to VM disks are EXCLUDED here!**

Step 4: Run any Velero post-backup hooks that are specified through pod annotations in an arbitrary order.

Step 5: For all VMs in the namespace, create a VirtualMachineSnapshot (VMS) custom resource.

Note: As part of this operation, snapshots will be created for all PVCs that belong to VM disks.

Step 6: If the backup policy specifies backup to an external object storage, clone the PVC snapshots in to new, temporary PVCs and start Velero DataUploader (DU) pods to copy the data to object storage.

Note: While the workflow shown above is similar to the workflow for backing up a non-VM application, it is important to understand that the VM - disks are snapshotted at a later point in time, leveraging the VirtualMachineSnapshot CR.

VM restore workflow

The Fusion B&R restore workflow for namespaces that contain VMs is as follows:

Step 1: DELETE the restore target namespace

Step 2: Restore all **PVCs**.

Step 3: Restore all namespace-scoped **resources** in an arbitrary order.

Step 4: Run **Velero post-restore hooks** that are specified through pod annotations in an arbitrary order.

Note: While the workflow shown above is similar to the workflow for restoring a non-VM application, it is important to understand that **the target namespace is being deleted before the restore begins**. If a single VM shall be restored to a namespace that already exists and that contains other VMs or even non-VM containerized applications, restore the VM to a different (new) namespace and migrate it to the target namespace later.

The detailed procedure for backing up and restoring an OCP-V virtual machine is described in [End-2-End Example: Creating an application consistent VM backup](#).

Implementing consistency for apps running inside VMs

For containerized (non-VM) applications, a Fusion Backup & Restore recipe CR (`frcpe`) can be used to ensure that a backup is consistent, and data can be restored properly anywhere at any time.

However, the snapshot of VM disks happens at a later phase of the Fusion backup workflow, when the "recipe phase" has finished already (see "VM BACKUP workflow"). Therefore, a recipe cannot be used to achieve consistency for applications running inside OCP-V virtual machines. Standard KVM/QEMU freeze and unfreeze scripts can be used instead. << x ref >>

The QEMU Guest Agent

The **qemu-guest-agent** is a helper daemon installed inside a virtual machine (VM) that enables active communication between the host hypervisor (for example, KVM, QEMU) and the guest operating system. The agent provides the following capabilities:

- Ensures graceful, clean VM shutdowns and reboots
- Allows for accurate backups via filesystem freezing
- Allows the execution of commands inside the VM (permissions required)
- Allows information exchange, for example, reporting of IP or MAC addresses
- Synchronizes time between the VM and the host
- Is available for most Linux and Windows operating systems

Important: If not installed automatically during the VM deployment, ensure that the `qemu-guest-agent` is installed and activated in a virtual machine. Otherwise, online backups of the VM will be crash-consistent only. More details on installing the QEMU guest agent can be found [here](#).

To achieve filesystem consistency when triggering snapshots of VM disks, the `qemu-guest-agent` issues freeze and unfreeze operations. This happens in the following situations:

- As part of the namespace-scoped resources backup, if the VM's `virt-launcher` pod has the following annotations assigned
 - `pre.hook.backup.velero.io/command = ["/usr/bin/virt-freezer", "--freeze" ...]`
 - `post.hook.backup.velero.io/command = ["/usr/bin/virt-freezer", "--unfreeze" ...]`
- When VM disk-snapshots are being created as part of a `VirtualMachineSnapshot (VMS)` operation.

Note: The `[pre/post].hook.backup.velero.io` annotations are automatically assigned to the virt-launcher pod, when the pod is created (or re-created after a VM reboot) . **These annotations should be removed**, as the freeze and unfreeze activities are useless during the phase of the namespace meta data backup (see also <https://access.redhat.com/solutions/7137659> for more details). The freeze and unfreeze operations that are triggered with the VirtualMachineSnapshot operation are scheduled properly before and after the VM disk snapshots.

Pre- and Post-Snapshot actions with fsfreeze-hook scripts

In some cases, it may not be sufficient to trigger filesystem freeze and unfreeze operations together with the VM disk snapshots, for example, if a database shall suspend writes before the fsfreeze and resume normal operations afterward. For VMs running the Linux operating system, the QEMU guest agent offers the ability to execute custom scripts together with the VM disk snapshots. The fsfreeze-hook mechanism works as follows:

- The main script `/etc/qemu/fsfreeze-hook` is called from the qemu-guest-agent whenever a filesystem freeze or thaw action is triggered
- The main script inspects folder `/etc/qemu/fsfreeze-hook.d/` and executes **all** custom user scripts that are found there
- The custom user scripts are called with parameters `freeze` (before snapshot) and `thaw` (after the snapshot)

The following sample script could be used to suspend an IBM Db2 database before freezing the filesystems and to resume the database after thawing the filesystems:

```
#!/bin/bash
#####
# Example: Db2 write suspend/resume when doing fs freeze/thaw
#####

case "$1" in
    **freeze**)
        # Activity BEFORE the snapshot
        export INSTHOME=/db2home/db2inst1
        export INSTUSER=db2inst1
        su - $INSTUSER -c "($INSTHOME/sqlllib/bin/db2 connect to mydb); \
($INSTHOME/sqlllib/bin/**db2 set write suspend** for database); \
($INSTHOME/sqlllib/bin/db2 commit); \
($INSTHOME/sqlllib/bin/db2 disconnect mydb)"
        exit 0
        ;;
    **thaw**)
        # Activity AFTER the snapshot
        export INSTHOME=/db2home/db2inst1
        export INSTUSER=db2inst1
        su - $INSTUSER -c "($INSTHOME/sqlllib/bin/db2 connect to mydb); \
($INSTHOME/sqlllib/bin/**db2 set write ****resume** for database); \
($INSTHOME/sqlllib/bin/db2 commit); \
($INSTHOME/sqlllib/bin/db2 disconnect mydb)"
        exit 0
        ;;
    *)
        echo "Usage: $0 {freeze|thaw}"
        exit 1
        ;;
esac
```

Note: Scripts in / etc / qemu / fsfreeze-hook.d / must be made executable (`chmod +x`).

End-2-End example: Creating an application consistent VM backup

This example demonstrates how to backup and restore an OCP-V virtual machine and how to implement a fsfreeze-hook script to achieve application consistency.

Prerequisites

Before starting to backup and restore OCP-V virtual machines, ensure that Fusion Backup & Restore has been installed properly (see "[Installation](#)") and that the configuration is complete.

The examples shown in the following sections of this section assume that Fusion Backup & Restore is installed on multiple Red Hat OpenShift clusters and that the following configuration tasks have been completed:

- One B&R HUB is connected to two B&R SPOKES ("primary" & "secondary")
- The (self) service-protection of the HUB to external object storage is enabled
- At least one backup storage location (*fbsl*) to external object storage is defined
- At least one backup policy is defined (using the *fbsl* as target)

The sample webserver VM

For demonstration purposes, a virtual machine "vm-websrv" was prepared with the following properties:

- RHEL 9 operating system on a 30GB DataVolume (*DV*)
- Apache HTTP server with sample web page
- QEMU guest agent up and running
- Virtual network adapter using the Red Hat OpenShift Pod network (NAT)
- MetalLB load balancer services expose ssh and http via an external IP address

The following figure shows the VM details in the Red Hat OpenShift web UI (note that it was created in a dedicated namespace to allow backup & restore granularity on a per VM basis):




Figure: Red Hat OpenShift virtual machine details for vm-websrv

The VM uses a DataVolume (DV) based on a PVC that leverages a StorageClass (SC) with CSI "snapshot" and "create volume from snapshot" capabilities. This is a general prerequisite to create application and VM backups in Red Hat OpenShift using the Fusion Backup and Restore service:

Disks

Add disk

Filter

Search by name...

Name	Source	Size	Drive	Interface	Storage class
rootdisk bootable	DV vm-websrv	30 GiB	Disk	virtio	ocs-storagecluster- ceph-rbd

Figure: Virtual machine storage showing DataVolume and PVC

The virtual network interface is connected to the Red Hat OpenShift Pod network:

Network interfaces

Add network interface

Filter

Name

Search by name...

Name	Model	Network	State	Type	MAC address
pod-network	virtio	Pod networking		Masquerade	02:85:57:00:00:00

Figure: Virtual machine network interface configuration

MetallLB load balancer services have been created to expose the ssh and httpd ports, so that they can be accessed from outside the Red Hat OpenShift cluster:

Services (3)

Name	Labels	Pod selector	Location
headless	No labels	network.kubevirt.io/headlessService=headless	None
vm-websrv-httpd	No labels	kubevirt.io/domain=vm-websrv	10.760.40
vm-websrv-ssh	No labels	kubevirt.io/domain=vm-websrv	10.760.40

Figure: Red Hat OpenShift services showing MetalLB load balancer configuration

The Apache web server hosts a sample HTML-page that lists the contents of its own html-data directory (empty now):



Figure: Apache web server page showing empty directory listing

Verify the QEMU guest agent is installed and active:

```
[root@vm-websrv ~]# systemctl status qemu-guest-agent
● qemu-guest-agent.service - QEMU Guest Agent
   Loaded: loaded (/usr/lib/systemd/system/qemu-guest-agent.service; enabled; preset: enabled)
   Active: active (running) since Wed 2026-02-04 16:17:45 CET; 1min 12s ago
     Main PID: 2141 (qemu-ga)
        Tasks: 2 (limit: 10300)
       Memory: 2.7M
          CPU: 68ms
      CGroup: /system.slice/qemu-guest-agent.service
              └─2141 /usr/bin/qemu-ga --method=virtio-serial --path=/dev/virtio-ports/org.qemu.guest_agent.

Feb 04 16:17:45 vm-websrv systemd[1]: Started QEMU Guest Agent.
```

Remove the default Velero pre- and post-backup annotations

To prevent unnecessary filesystem freeze and unfreeze operations while Fusion B&R is doing the backup of meta data resources in the VM's namespace, the Velero pre- and post-backup annotations must be removed from the virt-launcher pod that contains the VM instance. Locate the virt-launcher pod in the Red Hat OpenShift web UI:

- Go to **Workloads > Pods > virt-launcher-[vm_name]-.....**, or
- Go to **Virtualization > VirtualMachines** and review the VM's "General" information tab:




Figure: Red Hat OpenShift virtual machine general information tab

Note: The VM must be in running state to have a virt -launcher pod available.

In the YAML view of the virt-launcher pod, identify and **remove the velero pre- and post-hook command annotations**:




Figure: Virt-launcher pod YAML showing Velero hook annotations

Prepare a fsfreeze-hook script

For demonstration purposes, a fsfreeze-hook script was created that writes log files into the Apache web server html-data directory. The files contain a timestamp and an indication whether the log was created before the freeze or after the thaw operation:

```
[root@vm-websrv ~]# cat /etc/qemu-ga/fsfreeze-hook.d/pre-post-script.sh

#!/bin/sh
#####
# This script is called the fsfreeze --freeze and --thaw actions
# It runs pre- and post- tasks around freezing the VM's filesystems
#####

case "$1" in
  **freeze**)
    #-----
    # "Suspend" part - called with freeze
    #-----
    **touch "/var/www/html/$(date +%****Y%m%d****_%H.%M.%S").PRE-APSHOT"**
    exit 0
    ;;
  **thaw**)
    #-----
    # "Resume" part - called with thaw
    #-----
    **touch "/var/www/html/$(date +%****Y%m%d****_%H.%M.%S").POST-SAPSHOT"**
    exit 0
    ;;
esac
```

Perform a VM backup

As discussed in section (link to 2.2.2.3 Backup of an application), a Fusion PolicyAssignment (*fpa*) must be created to assign a backup policy to the application that represents the VM namespace.

Step 1: On the HUB, open the Fusion web UI and go to **Backup & Restore > Backed up applications** and click **Protect apps**:




Figure: IBM Fusion Backed up applications with Protect apps button

Step 2: From the drop-down menu, **select the Red Hat OpenShift cluster that hosts the virtual machine** (in this case it is hosted by the "primary" SPOKE), **select the application that should be protected** ("vm-websrv" in this example) and then click **Next**:




Figure: IBM Fusion application selection showing vm-websrv

Step 3: From the list of prepared backup policies, select the policy that shall be used (in this case it is "protect1.pri.virtualmachines") and click **Assign**:




Figure: IBM Fusion policy assignment for virtual machine backup

Note: Once the backup policy is assigned to an application, an ad hoc backup operation is started , even if the policy is configured with a different time schedule. In some cases, it is not preferable to start an immediate backup, for example, if the VM needs to be prepared with fsfreeze -hook scripts to achieve application consistency. In this case, turn the "Backup now" toggle-switch off before pressing the Assign button , otherwise let it run an ad hoc backup immediately.

Step 4: Go back to **Backup & Restore > Backed up applications** and review the newly created PolicyAssignment (*fpa*):




Figure: IBM Fusion backed up applications showing vm-websrv backup

Step 5: In the Fusion web UI, go to **Backup & Restore > Jobs** and monitor the backup job:




Figure: IBM Fusion VM backup job workflow details

Step 6: When the backup job for the VM is completed, we would expect to have two files logged in the `html-data` directory of the Apache web server, which have been created as part of the `fsfreeze-hook` mechanism. After the backup job is completed, the sample web page looks like this:




Figure: Apache web page showing pre and post snapshot log files

Perform a VM restore

This section describes how to restore the sample VM “vm-websrv” to a different namespace on the same Red Hat OpenShift cluster.

Step 1: In the Fusion web UI on the HUB cluster, go to **Backup & restore > Backed up applications**, open the **context-menu** of the vm-websrv application and click **Restore**.

1 item selected Select all 3 items								
Back up now Manage policies Cancel								
	Name	Cluster	Backup status ↑	Last backup on	Success rate	Policies	Recipes	Backup capacity (GiB)
☐	wordpress	apps.ocp.pri.ibm.lab	Completed	Feb 6, 2026, 4:32 AM	15/15	protect1.pri.wordpress	0	0.31
☒	vm-websrv	apps.ocp.pri.ibm.lab	Completed	Feb 6, 2026, 10:07 PM	15/17	protect1.pri.virtualmachines	0	5.89
☐	vm-new	apps.ocp.pri.ibm.lab	Completed	Feb 6, 2026, 5:04 PM	1/1	0	0	30.00
Items per page: 25 1-3 of 3 items								

View details
Back up now
Manage policies
Restore

Figure: IBM Fusion backed up applications with VM restore menu

Step 2: Select the "primary" SPOKE from the Target cluster drop-down menu and select **Create new project** for the Project destination. Specify the name of the target project for the VM:

Restore vm-websrv

Restore applications to an earlier point in time.

Cluster destination

Choose where the application will be restored.

Target cluster

apps.ocp.pri.ibm.lab

Project destination

Choose which project the application data should be placed in.

Use same project

Use the same project the application is already using.

Choose existing project

Choose a different existing project for this restore.

Create new project

Create a new project specifically for this restore.

Project name

vm-restore

Summary

The following changes will be made:

Before restore

Cluster apps.ocp.pri.ibm.lab

Type Spoke

Project vm-websrv

Recipe -

After restore

Cluster apps.ocp.pri.ibm.lab

Type Spoke

Project vm-restore

PVCs -

Recipe -

Restore point

-

Cancel

Back

Next

Figure: IBM Fusion VM restore target and new project configuration

Step 3: Select the backup PIT to restore:

Restore vm-websrv

Restore applications to an earlier point in time.

Select a backup

Search

Backup time	Location	Policy	Status	Size	PVs
☒ Feb 6, 2026, 10:07 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	0 GiB	1
☐ Feb 6, 2026, 8:49 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	0 GiB	1
☐ Feb 6, 2026, 6:36 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	< 0.01 GiB	1
☐ Feb 6, 2026, 6:15 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	0.3 GiB	1
☐ Feb 6, 2026, 5:55 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	< 0.01 GiB	1
☐ Feb 6, 2026, 5:53 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	0.32 GiB	1
☐ Feb 4, 2026, 8:15 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	< 0.01 GiB	1
☐ Feb 4, 2026, 7:50 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	< 0.01 GiB	1
☐ Feb 4, 2026, 7:32 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	< 0.01 GiB	1
☐ Feb 4, 2026, 6:57 PM	protect1.pri.virtualmachines	protect1.pri.virtualmachines	Completed	0.17 GiB	1

Items per page: 10 1-10 of 15 items

1 1 of 2 pages

Summary

The following changes will be made:

Before restore

Cluster apps.ocp.pri.ibm.lab

Type Spoke

Project vm-websrv

Recipe -

After restore

Cluster apps.ocp.pri.ibm.lab

Type Spoke

Project vm-restore

PVCs 1

Recipe -

Restore point

Feb 6, 2026, 10:07 PM

Cancel

Back

Next

Figure: IBM Fusion VM backup version selection for restore

87

Step 4: Decide whether the target StorageClass should be changed or not. Furthermore, decide whether all PVCs or only a subset of them should be restored:

Restore vm-websrv

Restore applications to an earlier point in time.[Learn more](#)

Target storage class

By default, the restore occurs to the same storage class name as the backup PVCs. If the storage class name is not found on the target cluster, the default storage class of target cluster is used instead. If you wish to use a different storage class for the restore, specify it below. [Learn more](#)

☐ Change target storage class ⓘ

Missing etcd resources

Restore any etcd resources that are not present in the existing OpenShift project.

☒ Include missing etcd resources

PVCs

By default all PVCs in the backup will be restored. You can select individual PVCs to restore if you only need to restore certain volumes.

ⓘ The restore will not overwrite existing PVCs. New PVCs will be named with a numerical count suffix such as <PVC Name>-<1>. Modify your application deployment to use the new PVCs.

☐ Only restore a subset of selected PVCs

Cancel

Back

Restore

Summary

The following changes will be made:

Before restore

Cluster apps.ocp.pri.ibm.lab

Type Spoke

Project vm-websrv

Recipe -

After restore

Cluster apps.ocp.pri.ibm.lab

Type Spoke

Project vm-restore

PVCs 1

Recipe -

Restore point

Feb 6, 2026, 10:07 PM

Figure: IBM Fusion VM restore StorageClass and PVC selection

Step 5: Review the summary and confirm it to start the VM restoration.

Step 6: In the Fusion web UI, go to **Backup & restore > Jobs** and **click the Restore tab** to list the history of completed and running restore operations

Cluster: Filter...		Backup status: Filter...		Q Search		Protect apps +		
☐	Name	Cluster	Backup status ↑	Last backup on	Success rate	Policies	Recipes	Backup capacity (GiB)
☐	wordpress	apps.ocp.pri.ibm.lab	Completed	Feb 6, 2026, 4:32 AM	15/15	protect1.pri.wordpress	0	0.31
☐	vm-new	apps.ocp.pri.ibm.lab	Completed	Feb 6, 2026, 5:04 PM	1/1	0	0	30.00
☐	vm-websrv	apps.ocp.pri.ibm.lab	Restoring	Feb 6, 2026, 10:07 PM	15/17	protect1.pri.virtualmachines	0	5.89

Figure: IBM Fusion jobs screen showing VM restore operations

Step 7: Wait until the restore job is completed.

Step 8: Go to **Virtualization > VirtualMachines** and review the details of the restored VM. The name of the virtual machine is unchanged ("vm-websrv"), but it resides in a new namespace ("vm-restore"). The "Services" section should show MetalLB services for ssh and httpd, as these have been restored as part of the namespace-scoped resource data. Note that the services have been assigned a new IP address from the MetalLB IP address pool:

88

Name ↓	Labels ↓	Pod selector ↓	Location ↓	
 headless	velero.io/backup-name=backu... velero.io/restore-name=backup...	 network.kubevirt.io/headlessService=headless	None	⋮
 vm-websrv-httpd	velero.io/backup-name=backu... velero.io/restore-name=backup...	 kubevirt.io/domain=vm-websrv	10.7.60.41	⋮
 vm-websrv-ssh	velero.io/backup-name=backu... velero.io/restore-name=backup...	 kubevirt.io/domain=vm-websrv	10.7.60.41	⋮

Figure: Red Hat OpenShift restored VM showing new MetalLB service IPs

Step 9: Review the sample web page that is provided by the restored VM:

←

→

↺

Not Secure

http://10.7.60.41



Redbooks - IBM Storage Defender: Protecting OpenShift Containers and Virtual Machines

Name	Last modified	Size
 20260206_22.05.23.PRE-SNAPSHOT	2026-02-06 22:05	0

Figure: Apache web page of restored VM showing pre-snapshot log

Note: As the post-snapshot log file is missing after the restore. This is an expected condition, because that log file was created only after the VM disk snapshot was taken.

With VMs hosting more complex applications or databases, post-restore activities might need to be performed. If, for example, an IBM Db2 database is in write-suspend mode during the backup, it will also be in write-suspend mode after the restore and must therefore be manually resumed.

IBM Storage Defender Data Protect

IBM Storage Defender Data Protect extends enterprise data protection capabilities to Red Hat OpenShift environments through integration with IBM Fusion. Built on the Data Resiliency Service framework, Defender provides centralized visibility, policy management, and recovery orchestration for containerized applications and Red Hat OpenShift Virtualization workloads across distributed clusters.

This chapter describes how IBM Defender DataProtect protects Kubernetes workloads with a focus on Red Hat OpenShift and Red Hat OpenShift Virtualization. The main goal is to protect namespaces, PVCs, cluster resources, metadata, and Red Hat OpenShift VMs, so they can be restored later to the same cluster or to another registered Red Hat OpenShift cluster.

For Red Hat OpenShift, Defender DataProtect deploys components into the Red Hat OpenShift cluster, including a Data Mover, Velero, and a DataProtect Plugin/OADP-related components. These components are used to discover, snapshot, copy, and restore Red Hat OpenShift resources. Defender DataProtect can protect the following resources:

- Kubernetes namespaces
- Namespace metadata
- Kubernetes cluster resources
- PVCs
- Red Hat OpenShift VMs
- VM-related PVCs

For normal Red Hat OpenShift workloads, Defender DataProtect backs up PVCs, Kubernetes cluster resources, and namespace metadata. The Data Mover pod is deployed inside the Red Hat OpenShift cluster and is used to back up PVC data.

For Red Hat OpenShift Virtualization, Defender DataProtect additionally supports VM backup. It protects PVCs, VMs, Red Hat OpenShift cluster resources, and namespace metadata.

The goal of this chapter is not to replicate the DataProtect User Guide, but to provide a quick start guide for protecting PVCs and Red Hat OpenShift Virtualization workloads. It also outlines selected use cases and useful customizations.

Applications: protecting Red Hat OpenShift containers and Red Hat OpenShift Virtualization VMs

In the Data Resiliency Service, Applications represent a fundamental shift in how you approach data protection. Rather than managing individual objects in isolation—one virtual machine here, one persistent volume claim there—Applications allow you to define and protect complete workload boundaries. This approach is critical because incidents do not respect artificial boundaries between components. When recovery is needed, you need to restore complete, functional applications, not scattered pieces.

In the user interface, IBM places Applications next to Recovery Groups, which makes sense from a recovery perspective. Recovery follows a logical sequence: data must be restored first, then the platform, and finally the services. If you ignore these dependencies, you can end up with systems that appear to be restored but remain non-functional. Applications define the workload boundary and ensure that all necessary components are recovered together in the correct order.

Defining application boundaries

An Application in the Data Resiliency Service ties the right recovery groups together to prevent partial restores. You define the recovery order explicitly—typically data before the application, though there are exceptions. You also document what "done" means: the validation checks, smoke tests, and health indicators that must be green before you can consider the recovery complete.

An Application typically encompasses one or more namespaces, often more than you initially expect. Within that boundary, you will find the following components:

- Deployments, ConfigMaps, Secrets, Services, and Routes
- Persistent data backed by Persistent Volume Claims (the components you absolutely cannot lose)
- Optionally, Red Hat OpenShift Virtualization virtual machines in the same namespace, which share the same recovery intent but have different mechanics and failure modes

During an incident, this Application view is what teams keep open because it matches how you actually rebuild systems. It may not be the most visually appealing interface, but it is accurate and actionable.

Application inventory and metadata

From the IBM Fusion inventory that the Data Resiliency Service consumes, an Application displays the following comprehensive information:

- Services, Pods, Secrets, Deployments, and ConfigMaps
- Routes (which provide external access—these are easy to overlook but critical for recovery)
- Virtual machine names in the namespace
- Backup versions with start time, expiration, policy identifier, status, captured metadata, and size in object storage

This information provides proof that protection is working and gives you the audit trail needed during recovery operations.




Figure: IBM Storage Defender Data Protect protection view showing Remote Adapter and workload types

Operational visibility and governance

In daily operations, this Application view helps keep conversations focused on outcomes rather than technical details. You can quickly assess the following criteria:

- Application X backs up every N hours, retains Y days, and stores data in location Z
- Application Y has no policy attached (yes, this happens more often than it should)
- Application Z has backup versions, but the last one failed, expired, or falls outside the service level agreement

This is what true resiliency looks like: understanding workload outcomes, not just confirming that "a copy exists somewhere."

The typical workflow from end to end works as follows: IBM Fusion protects the application according to policy and schedule. Connection Manager collects inventory from clusters, applications, backups, policies, and locations. The Data Resiliency Service displays this information in a unified view. You then create recovery groups for IBM Fusion Applications and bind governance profiles to them. This completes the protection and governance loop.

Container and virtual machine recovery considerations

A key point that is often missed: containers and virtual machines can exist under the same application boundary. The grouping approach is the same, but the execution differs significantly.

Container recovery workflow

For containers, recovery involves rebuilding the namespace and restoring the data, then validating the result. There are no shortcuts:

Step 1. Restore namespace objects (Deployments, Services, Routes or Ingress, ConfigMaps, Secrets). Scope and rules determine what gets restored.

Step 2. Restore persistent data (Persistent Volume Claim content).

Step 3. Reconcile and validate until the application is healthy. Do not skip this step—problems that are ignored during recovery will cause issues later.

Virtual machine recovery workflow

For virtual machines, recovery focuses on the virtual machine definition and disks, along with infrastructure dependencies that frequently cause problems:

Step 1. Restore virtual machine resource definitions (KubeVirt) and virtual machine disks (typically backed by Persistent Volume Claims).

Step 2. Address dependencies: networks, storage classes, node placement, and availability. This is where recovery operations become sensitive and require careful attention.

Virtual machine recovery comes down to one fundamental choice: speed versus clean restore.

- Fast bring-up: Get the virtual machine running immediately and copy data in the background later.
- Copy-first: Restore the complete data set before starting the virtual machine.

Choose your approach based on your Recovery Time Objective (RTO) versus your restore-first discipline. Do not make this decision during an incident—decide it earlier when you have time to think clearly and document the rationale.

Policy-based protection: defining intent clearly

IBM frames policy-based protection as the service level agreement path, encompassing cadence, retention, consistency, immutable targets, anomaly detection, and integration with IBM Storage Protect Vault. While this context is useful, in practice you should focus on writing clear intent rather than trying to implement every feature mentioned in product literature.

When reviewing configurations, focus on these key elements:

- Policy: Intent (frequency, retention, location, security posture).
- Application: Boundary (namespaces, resources, virtual machines) where that intent applies.
- Versions: Proof that protection ran (timestamp, status, expiry).
- Data Resiliency Service inventory: The audit trail. If something is not in inventory, it is not real. Period.

Network connectivity requirements

Protection only works if the control plane can actually see the inventory. The Connection Manager requires access to outbound port 443 to reach the Defender cloud service endpoints. Using a proxy is common and should be planned for from the beginning.

Treat this as a hard requirement. Otherwise you will encounter the classic pattern: backups appear to run successfully, but governance and visibility stop updating. While this might work initially and show backups completing, the problem will be noticed later when other components do not accurately report data. By then, you may have gaps in your protection that you were not aware of, which can be discovered at the worst possible time—during an actual recovery scenario.

Preparing the environment

Prerequisites

Before starting to backup and restore Red Hat OpenShift / Kubernetes applications, ensure that IBM Defender DataProtect and Red Hat OpenShift Systems are fulfilling the requirements and check the support matrix for the supported code levels.

The following tasks provide an overview of the capabilities:

- Import required plugins into an Red Hat OpenShift registry

- Prepare the Red Hat OpenShift Cluster: Service Accounts, Cluster Role Binding, Access Token
- Configure the Red Hat OpenShift cluster as a Source in Defender DataProtect
- Configure DataProtect Policies and Protection Groups
- Customize object-level backup parameters such as PVC or VM includes/excludes, CSI Snapshot configuration, pre- and post-backup scripts and autoprotection by labels

These instructions are based on the following versions:

- IBM Storage Defender DataProtect version 7.4
- Red Hat OpenShift version 4.20.17

Prepare the Defender DataProtect software images

Obtain the required plugins

This example assumes that the Red Hat OpenShift cluster has internet connectivity and as such has access to online repositories outside of the corporate network. However, the following packages need to be downloaded and pushed to your own private registry:

- DataProtect Datamover: Download from the IBM Passport Advantage download page.
- DataProtect Plugin: Download from the IBM Passport Advantage download page. This component is only required if you need to backup Red Hat OpenShift Virtualization virtual machines and need to select VMs granularly from a namespace. If you plan to protect the namespace with all VMs, this plugin is not required.

Push the software images into your private registry

The following example demonstrates the scenario with a Red Hat Quay registry. All components are tagged and pushed, whereas the DataProtect Plugin is only required for Red Hat OpenShift Virtualization and only if a subset of VMs within a namespace needs to be protected.

```
# Load the images into local podman image store
podman load -i cohesity-dataprotect-plugin-7.4.tar
podman load -i cohesity-datamover-7.4.tar

# Tag locally
podman tag localhost/cohesity-dataprotect-plugin:7.4 \
    localhost/ibm/defender/cohesity-dataprotect-plugin:7.4

podman tag localhost/cohesity-datamover:7.4 \
    localhost/ibm/defender/cohesity-datamover:7.4

# Tag for Quay registry
podman tag localhost/ibm/defender/cohesity-dataprotect-plugin:7.4 \
    quay-enterprise.apps.ocp.ddp.ibm.lab/ibm/defender/cohesity-dataprotect-plugin:7.4

podman tag localhost/ibm/defender/cohesity-datamover:7.4 \
    quay-enterprise.apps.ocp.ddp.ibm.lab/ibm/defender/cohesity-datamover:7.4

# Login to Quay registry
podman login quay-enterprise.apps.ocp.ddp.ibm.lab

# Push to your Quay registry
podman push quay-enterprise.apps.ocp.ddp.ibm.lab/ibm/defender/cohesity-dataprotect-plugin:7.4
podman push quay-enterprise.apps.ocp.ddp.ibm.lab/ibm/defender/cohesity-datamover:7.4
```

Validate that the images are uploaded successfully with the correct tags as shown in the follow figure.




Figure: Red Hat Quay Registry with Required Images

Prepare the Red Hat OpenShift cluster

To register the Red Hat OpenShift cluster with the DataProtect cluster, meet the following requirements before starting configuration and management:

- Have a valid certificate with populated SAN fields.
- Ensure Red Hat OpenShift Cluster permissions to create service accounts, credentials, and Cluster Role Bindings.
- Ensure versions of Red Hat OpenShift, Velero, and DataProtect are supported.
- Ensure firewall rules are in place.

Valid certificate with Subject Alternative Name (SAN) field

Install a valid HTTPS certificate on the DataProtect cluster that includes a SAN field. The SAN field defines additional hostnames (e.g., DNS names, IP addresses) covered by the certificate, enabling multi-domain support. See the DataProtect documentation for guidance on creating and updating custom CA-signed certificates with multiple DNS and IP entries.

Create an Red Hat OpenShift service account for DataProtect

To register the Red Hat OpenShift cluster as a source in the DataProtect cluster, provide the bearer token of an Red Hat OpenShift service account. The service account must have the cluster-admin role on the cluster.

The following three-step example describes how to create the service account, assign a cluster role binding, and extract the bearer token:

```
oc create namespace cohesity-dataprotect-ns

oc project cohesity-dataprotect-ns

oc create serviceaccount cohesity-dataprotect-sa -n cohesity-dataprotect-ns

oc create clusterrolebinding cohesity-dataprotect-crb \
  --clusterrole=cluster-admin \
  --serviceaccount=cohesity-dataprotect-ns:cohesity-dataprotect-sa

oc apply -n cohesity-dataprotect-ns -f - <<EOF
apiVersion: v1
kind: Secret
metadata:
  name: cohesity-dataprotect-token
  namespace: cohesity-dataprotect-ns
  annotations:
    kubernetes.io/service-account.name: cohesity-dataprotect-sa
type: kubernetes.io/service-account-token
EOF

oc describe secrets -n cohesity-dataprotect-ns cohesity-dataprotect-token | grep token: | awk '{print $2}'
```

Prepare the DataProtect cluster

Register the Red Hat OpenShift cluster as backup source

Before you can protect Red Hat OpenShift namespaces, you first need to register the cluster as a Kubernetes source in Defender DataProtect. The following figures show how to launch the Kubernetes Registration Wizard from the Defender Data Management Service (DMS) user interface.




Figure: Step 1 — Launch the Red Hat OpenShift Registration Wizard

Register Kubernetes Source

Basic Information

Kubernetes Distribution

Red Hat OpenShift

Kubernetes Cluster Endpoint

https://api.ocp.ddp.ibm.lab:6443

Eg IP Address:Port Number

To ensure you're using a supported version of kubernetes, we recommend checking our official [documentation](#).

Password-Bearer Token

3Iif251pgDr1M-_YGnjLoxfDxF31AdUMfRVwL8TNoKvw962NlYk9Axx8hsVEYzYB3EBRtQdV0Y

Every Service Account has a Secret with valid Bearer Token that can be used to log in to Dashboard

Cohesity uses Data Mover pods and other components to protect your environment's data. These are automatically downloaded when internet connectivity is available. If not, you may need to configure the settings in the 'Optional' section below. [Learn More](#)

Optional

Cancel

Save

Figure: Step 2 — Provide OCP Cluster Details and Credentials

Register Kubernetes Source

Optional ^

Datamover Connection Type

NodePort

Velero OpenShift Plugin Image Path

Optional

Velero Image Path

Velero Image Path

To process metadata related to Kubernetes resources

Datamover Image Path

quay-enterprise.apps.ocp.ddp.ibm.lab/ibm/defender/cohesity-datamover:7.4

To process data during backup and recovery

Cohesity DataProtect Plugin Image Path

quay-enterprise.apps.ocp.ddp.ibm.lab/ibm/defender/cohesity-dataprotect-plugin:7.4

To interface with any kubernetes distribution e.g. OpenShift, Mainline etc

Velero AWS Plugin Image Path

To access Cohesity storage through S3 protocol

Cancel

Save

Figure: Step 3 — Provide OCP Cluster Image Paths

After saving the configuration in the DataProtect registration wizard, a new namespace is created in the Red Hat OpenShift cluster. In addition to the `cohesity-dataprotect-ns` namespace (which includes service account, secret, and cluster role binding), another namespace is created. This second namespace has the DataProtect cluster ID as a suffix, e.g. `cohesity-3606920535574357`, as shown below.

97




Figure: OCP Namespaces Created for DataProtect Backups

Display the Red Hat OpenShift namespaces

Upon successful registration of the Red Hat OpenShift source the inventory gets pulled and can be displayed in the DataProtect GUI. The following figure lists the registered Red Hat OpenShift Clusters that can be protected with the DataProtect cluster. The types of sources are filtered by the Kubernetes tag and the last inventory was fetched a few seconds ago.




Figure: Status of Registered Kubernetes Sources

The Red Hat OpenShift namespace and project inventory is displayed by selecting the source name. The following figure shows a list of all namespaces and projects available for protection in DataProtect. Currently, 24 objects are not

protected by any backup job, and 1 object is protected. An object refers to a namespace or a project with one or more pods or virtual machines.

Note: Default Red Hat OpenShift namespaces are not shown in this view.




Figure: Red Hat OpenShift Namespace Inventory Listing

Create backup policies and protection groups

In IBM Storage Defender DataProtect, Protection Policies and Protection Groups are two closely related configuration objects. Together, they define how data protection is applied to workloads.

A *Protection Policy* defines the backup rules. It describes when backups run, how often they run, how long recovery points are retained, and whether additional operations such as replication or archival are performed. In simple terms, the policy defines the schedule, retention, and lifecycle behavior of the backup data. A policy is reusable and can be assigned to multiple Protection Groups.

A *Protection Group* defines what is protected. It contains the selected source and workload objects, such as virtual machines, EC2 instances, databases, file shares, Kubernetes namespaces, or other supported workloads. The Protection Group also contains workload-specific settings, such as selected objects, exclusions, start time, priority, storage domain, and advanced application or platform options.

The relationship between the two is straightforward: the Protection Group selects the workloads, and the Protection Policy defines the protection behavior applied to those workloads. In other words, the Protection Group answers the question "What should be protected?", while the Protection Policy answers the question "How and when should it be protected?"

For example, a Protection Policy may define that backups run every night at 22:00, local snapshots are retained for 30 days, and monthly copies are archived to object storage. A Protection Group can then use this policy to protect a defined set of namespaces or virtual machines. If the same schedule and retention requirements apply to another workload, the same policy can be reused by another Protection Group.

For the example configuration, the following protection policies and protection groups are defined for Red Hat OpenShift application workloads and Red Hat OpenShift virtual machines:

- **OCP-Bronze**: Basic protection level. Suitable for workloads where daily recovery points are sufficient and where no additional secondary copy is required.
- **OCP-Silver**: Increased protection level. Data Lock is enabled and replication provides an additional copy of the protected data, improving resilience if the primary cluster or site is unavailable.
- **OCP-Gold**: Highest protection level in this example. Includes replication to a second Defender DataProtect cluster and archive to an S3 tape target. Suitable for workloads that require more frequent recovery points, an additional replicated copy, and longer-term retention on external object storage.

Policy name	Backup frequency	Retention	Data Lock	Replication	Archive
OCP-Bronze	Once per day	One week	Disabled	Not configured	Not configured
OCP-Silver	Once per day	Two weeks	Enabled	Enabled	Not configured
OCP-Gold	Every six hours	Four weeks	Enabled	Enabled	Enabled to an S3 tape target

To create a new Protection Policy, log in to the IBM Defender Data Management Service and select the required Defender DataProtect cluster. From the cluster view, navigate to **Policies** and create a new policy named **OCP-Gold**.

As shown in the following figure, the policy is configured to run every six hours each day and to retain backup snapshots for two weeks. Data Lock is enabled for this policy. This helps prevent protected recovery points from being deleted or modified before the configured retention period expires.

×

Create Protection Policy

Build

Summary

Policy Name

OCP-Gold

DataLock ⓘ

Backup

Backup every

6

Hours

Primary Copy

Keep on

Local

Retain for

2

Weeks

Lock

2

Weeks

Figure: Create Policy OCP-Gold — Schedule and Retention Settings

The policy also includes replication to a second Defender DataProtect cluster and an additional archive copy to IBM Storage Protect with tape attachment.

× Create Protection Policy

Build Summary

Replication

Replicate to

Remote Cluster ▼

Replication Target

CE-EMEA-DDP-ROBO1 ▼

Every

Run ▼

Retain for

1 ▼

Month ▼

Lock

14 ▼

Days ▼

Archive

Archive to

Storage-Protect-S3-Glacier ▼

Every

1 ▼

Week ▼

Retain for

3 ▼

Months ▼

[+ Add DataLock](#)

☐ Archive only fully successful runs

☒ Full Backup Only

Figure: Create Policy OCP-Gold — Replication and Archive Settings

The following figure shows the policy overview filtered by policy name, filtered by filter value `OCP`.

Policy ↑	Protection Groups	Protected Objects	Backup	Replication	Archive	CloudSpin	Cloud Vault
OCP-Bronze Backup 1d Retain 1w	0	0					
OCP-Gold Backup 6h Retain 2w DataLock 2w Replicate Archive	0	0					
OCP-Silver Backup 1d Retain 2w DataLock 2w Replicate	0	0					

Figure: Policy Overview Filtered by OCP

Protecting Red Hat OpenShift applications

To protect Red Hat OpenShift workloads, a Protection Group is required. Protection Groups define which workloads use the defined policies. During the different use cases, a set of Protection Groups is created. `OCP-App-PG1` is an application Protection Group that protects a selected set of Red Hat OpenShift application objects.

`OCP-App-autoprotect-PG2` is also an application Protection Group, but it uses an auto-protect approach. Instead of manually selecting every object, auto-protect rules can be used to automatically include matching namespaces, applications, or other supported objects.

`OCP-VM-PG3` is used for Red Hat OpenShift virtual machines.

Backup of a sample WordPress application

Protection Groups are created as part of the process of selecting the workloads that require protection. Like the chapter that describes application protection with IBM Fusion Backup and Restore, this section describes how to protect a simple WordPress application by using a Protection Group and assigning the `OCP-Gold` policy.

To protect the WordPress application, navigate to **Data Protection > Sources** and select the registered Red Hat OpenShift cluster source. From there, you may filter the namespaces (1), select the relevant application workload (2) and select **Protect** (3).

IBM Storage Defender

← Sources

api.ocp.ddp.ibm.lab:6443

Objects Settings

Unprotected		Protected	
24	45.1 GB	1	107.4 GB
Objects	Data	Objects	Data

Protection Status

3 [-] [Protect] [Recover]

✓ [-] [📁] https://api.ocp.ddp.ibm.lab:6443 -

☐	🔒 restored-wordpress	🚫	13 days ago
2 ☒	🔒 wordpress	🚫	12 days ago

Unprotected

Figure: Select the WordPress Application

The next step prompts you to create a new Protection Group. During the Protection Group configuration, assign the `OCP-Gold` policy so that the WordPress application is protected according to the defined Gold service level. Select **more options** to review and define the required parameters.

× New Protection



Source

Registered Source

<https://api.ocp.ddp.ibm.lab:6443>

Objects

1

Namespaces

1

Manually Protected

Protection Group

 New Group

☐ Existing Group

Name

OCP-App-PG1

Policy

OCP-Gold ☐

Backup

Every 6 hours | Retain 2 weeks |
DataLock 2 weeks



Retry Options

Retry 3 times on error 5 minutes apart

Native Source Snapshot

Replicate to CE-EMEA-DDP-ROBO1

Every run | Retain 1 month | DataLock 2 weeks

Archive to ddp-archive-bronze

Every week | Retain 3 months

Figure: Create Protection Group with OCP-Gold Policy Assignment

Settings

Start Time

i Start time is unnecessary since runs will occur on regular intervals

Storage Domain

DefaultStorageDomain
Encrypt: AES-256-GCM | Deduplication: Inline | Compression: Inline

Leverage CSI snapshot

1

Leverage CSI snapshot

Additional Settings

2

Protect

Cancel

Figure: Enable CSI Snapshots for Protection Group

Monitoring of application backup jobs

After the Protection Group has been created and one or more applications have been assigned, the protection job starts immediately. Additional workloads can be added to the same Protection Group at any time. These newly added workloads will be protected during the next scheduled run, based on the assigned Protection Policy.

To get an overview of running and completed protection runs, select the required Defender DataProtect cluster and navigate to **Data Protection > Protection**. The protection run overview is shown here.

IBM Storage Defender

DataProtect

CE-EMEA-DDP-ROB02

Dashboards

Data Protection

Protection

Recoveries

Sources

Policies

CloudRetrieve

Infrastructure

SmartFiles

Test & Dev

Marketplace

System

Reports

Settings

Protection

7 Succeeded

1 Warning

0 Failed

1 Running

0 Canceled

7 Met SLA

0 Missed SLA

Group Type

Groups

Policy

SLA

Status

☐	Group	Start Time	Duration	Success/Error	SLA	Status
☐	K8 OCP-App-PG1 Kubernetes Policy: OCP-Gold 1	Jun 24, 2026 7:00pm	4m 17s	0/0 objects		
☐	Db Oracle-RMAN-Adapter Oracle Policy: Databases	Jun 24, 2026 4:49pm	6m	1/0 objects		
☐	K8 K8S-DDP-PG2 Kubernetes Policy: OpenShift_app1_policy	Jun 24, 2026 4:17pm	11m 25s	1/0 objects		
☐	Db MSSQL-Adapter-VolumeBased Microsoft SQL Policy: Databases	Jun 24, 2026 3:17pm	21s	5/0 objects		
☐	Vm SLES-VM HyperV Policy: Bronze	Jun 24, 2026 2:00pm	3m 16s	1/0 objects		
☐	Vm RHEL-VM HyperV Policy: Bronze	Jun 24, 2026 1:00pm	1m 10s	2/0 objects		
☐	Ap ActiveDirectory-Adapter Active Directory Policy: Bronze	Jun 24, 2026 11:00am	1m 53s	1/0 objects		

Figure: Protection Run Overview

To view more details for the new protection job, collect error information, or review the replication and archive status, select the required protection group (1) in the protection overview. The detailed job view shows the individual protection runs, their status, start and end times, and any warnings or errors that occurred during execution.

From this view, administrators can validate whether the local backup completed successfully and whether the configured replication and archive tasks were triggered and completed as expected.




Figure: Protection Run Job Details

The figure above shows that the protection run has been completed successfully. The local backup task finished without errors (1), the replication to the second Defender DataProtect cluster completed successfully (2), and the archive copy to the S3 target also completed successfully (3).

To view even more details for a specific job run, select the job by clicking the desired timestamp. Select the workload of interest (1) and review the **Backup Task Activity** (2) for the WordPress application.




Figure: Backup Task Activity Details

Customize the backup scope for an application

Red Hat OpenShift namespace protection can be customized at several levels:

- **Namespace selection and Auto Protect:** Administrators can manually select namespaces for backup or use Auto Protect at the Kubernetes cluster level or based on namespace labels. This allows newly created matching namespaces to be protected automatically.
- **PVC include/exclude:** For each namespace, Defender DataProtect can include or exclude specific Persistent Volume Claims. These settings can be configured at namespace level or Protection Group level. Namespace-level PVC include/exclude settings take precedence over the Protection Group settings.
- **Resource include/exclude:** Administrators can include or exclude Kubernetes resources and custom resources. There is also an option to back up only PVCs and related resources, which skips other resource types such as Secrets.
- **Quiescing and pre-/post-snapshot hooks:** For stateful workloads, Defender DataProtect supports quiescing before PVC snapshots are taken. Pre-snapshot scripts can be executed inside selected pods and containers before the snapshot, and post-snapshot scripts can unquiesce the application afterwards. This helps create more application-consistent snapshots.
- **Label-based script targeting:** Quiescing rules can target pods by label selectors. The scripts run at container level, and a specific container can be selected if the pod has multiple containers. Each script has a timeout, with a default of 60 seconds.
- **Quiesce execution modes:** There are three quiesce modes: together, independent, and sequential. Together mode executes rules in parallel before a common snapshot step. Independent mode uses multiple volume groups and provides the fastest backup behavior. Sequential mode runs rules one after another and is slower, but provides more controlled ordering.
- **Script failure handling:** For pre- and post-snapshot scripts, administrators can define whether the backup should continue if a script fails. This allows either strict application consistency, where the backup fails on script failure, or a more tolerant mode where the PVC backup continues.

In the previous section, the namespace `wordpress` was selected for protection. The WordPress namespace consists of two pods: one pod running the WordPress application and another pod running the database. Both pods must be included in the backup to protect the complete application stack.

In the following section, the focus is on creating an application-consistent backup by configuring pre- and post-snapshot actions. These actions are used to prepare the application before the CSI snapshots of the Persistent Volume Claims are created.

The pre-snapshot action places the application or database into a consistent state before the snapshot operation starts. After the snapshot has been created, the post-snapshot action releases the application again and allows normal operations to continue.

Prepare the Red Hat OpenShift pods with the pre- and post-scripts

Before the Red Hat OpenShift namespace protection is configured in Defender DataProtect, the required application containers must be prepared with pre- and post-snapshot scripts.

For application-consistent backups, every container that requires a specific action before or after the PVC snapshot must contain the required shell scripts. Defender DataProtect can execute these scripts during the backup run as pre-snapshot and post-snapshot hooks. The scripts are executed inside the selected container and can be used to place the application into a consistent state before the snapshot is taken and to resume normal operation afterwards.

Important: The required command sequence and logging logic must be placed into a shell script inside the target container. Passing multiple or complex commands directly from the backup application to the container as shell execution parameters is not supported.

For this reason, the pre- and post-snapshot logic is implemented as scripts inside the WordPress and MySQL containers. Defender DataProtect only triggers the script path during the backup run, while the

script itself contains the required application commands, error handling, return code handling, and log file creation.

In this example, the WordPress namespace contains two relevant containers: the WordPress application container and the MySQL database container. The WordPress application is placed into maintenance mode before the snapshot and taken out of maintenance mode after the snapshot. The MySQL database is set into a read-only state and table locks are flushed before the snapshot. After the snapshot has completed, the database is returned to normal read/write operation.

The following four scripts are created inside the respective containers. When they are triggered by Defender DataProtect during a backup run, each script also writes a log file to the `/tmp` directory. This makes it easier to validate whether the pre- and post-snapshot actions were executed successfully.

Script location	Script name	Purpose
WordPress container	pre_script_wp.sh	Pre-snapshot script: This script enables WordPress maintenance mode before the snapshot is created.
WordPress container	post_script_wp.sh	Post-snapshot script: This script disables WordPress maintenance mode after the snapshot has completed.
WordPress container	/tmp/pre_post_script_wp.log	Logging of pre- and post-execution runtimes and return codes
MySQL container	pre_script_mysql.sh	Pre-snapshot script: This script prepares the database for the snapshot by enabling read-only mode, enabling super read-only mode, and flushing tables with a read lock.
MySQL container	post_script_mysql.sh	Post-snapshot script: This script resumes normal database operation by disabling super read-only mode, disabling read-only mode, and unlocking the tables.
MySQL container	pre_post_script_mysql.log	Logging of pre- and post-execution runtimes and return codes

Below is the source code of the four scripts used for the WordPress and MySQL pre- and post-snapshot actions. Log in to the WordPress pod and create the following pre- and post-scripts:

WordPress pre-snapshot script

```
mkdir /var/www/html/wp-cli/pre-post-backup

cat > /var/www/html/wp-cli/pre-post-backup/pre_script_wp.sh <<'EOF'
#!/bin/sh

LOG_FILE="/var/www/html/wp-cli/pre-post-backup/pre_post_script_wp.log"
WP_CLI="/var/www/html/wp-cli/wp"

echo "$(date): Starting WordPress maintenance mode ENABLE" >> "$LOG_FILE"

if "$WP_CLI" maintenance-mode activate --allow-root >> "$LOG_FILE" 2>&1; then
    echo "$(date): Maintenance mode ENABLE completed successfully" >> "$LOG_FILE"
else
    echo "$(date): ERROR - Maintenance mode ENABLE failed" >> "$LOG_FILE"
    exit 1
fi

exit 0
EOF

chmod +x /var/www/html/wp-cli/pre-post-backup/pre_script_wp.sh
```

WordPress post-snapshot script

```
cat > /var/www/html/wp-cli/pre-post-backup/post_script_wp.sh <<'EOF'
#!/bin/sh

LOG_FILE="/var/www/html/wp-cli/pre-post-backup/pre_post_script_wp.log"
WP_CLI="/var/www/html/wp-cli/wp"

echo "$(date): Starting WordPress maintenance mode DISABLE" >> "$LOG_FILE"

if "$WP_CLI" maintenance-mode deactivate --allow-root >> "$LOG_FILE" 2>&1; then
    echo "$(date): Maintenance mode DISABLE completed successfully" >> "$LOG_FILE"
else
    echo "$(date): ERROR - Maintenance mode DISABLE failed" >> "$LOG_FILE"
    exit 1
fi

exit 0
EOF

chmod +x /var/www/html/wp-cli/pre-post-backup/post_script_wp.sh
```

MySQL pre-snapshot script

```
mkdir -p /var/lib/mysql/pre-post-backup

cat > /var/lib/mysql/pre-post-backup/pre_script_mysql.sh <<'EOF'
#!/bin/sh

LOG_FILE="/var/lib/mysql/pre-post-backup/pre_post_script_mysql.log"

echo "$(date): MySQL mode: freezing IO" >> "$LOG_FILE"

if MYSQL_PWD="$MYSQL_ROOT_PASSWORD" mysql -u "$MYSQL_USER" "$DB" \
    -e "SET GLOBAL read_only = ON; SET GLOBAL super_read_only = ON; FLUSH TABLES WITH READ LOCK;" \
    >> "$LOG_FILE" 2>&1; then

    echo "$(date): MySQL freeze completed successfully" >> "$LOG_FILE"
    exit 0
else
    RC=$?
    echo "$(date): ERROR - MySQL freeze failed with return code $RC" >> "$LOG_FILE"
    exit "$RC"
fi
EOF

chmod +x /var/lib/mysql/pre-post-backup/pre_script_mysql.sh
```

MySQL post-snapshot script

```
cat > /var/lib/mysql/pre-post-backup/post_script_mysql.sh <<'EOF'
#!/bin/sh

LOG_FILE="/var/lib/mysql/pre-post-backup/pre_post_script_mysql.log"

echo "$(date): MySQL mode: resuming IO" >> "$LOG_FILE"

if MYSQL_PWD="$MYSQL_ROOT_PASSWORD" mysql -u "$MYSQL_USER" "$DB" \
    -e "SET GLOBAL super_read_only = OFF; SET GLOBAL read_only = OFF; UNLOCK TABLES;" \
    >> "$LOG_FILE" 2>&1; then

    echo "$(date): MySQL resume completed successfully" >> "$LOG_FILE"
    exit 0
else
    RC=$?
    echo "$(date): ERROR - MySQL resume failed with return code $RC" >> "$LOG_FILE"
    exit "$RC"
fi
EOF

chmod +x /var/lib/mysql/pre-post-backup/post_script_mysql.sh
```

Sample log files

```
cat /tmp/pre_post_script_wp.log

Thu Jun 25 16:55:19 UTC 2026: Starting WordPress maintenance mode ENABLE
Enabling Maintenance mode...
Success: Activated Maintenance mode.
Thu Jun 25 16:55:21 UTC 2026: Maintenance mode ENABLE completed successfully

Thu Jun 25 17:00:55 UTC 2026: Starting WordPress maintenance mode DISABLE
Disabling Maintenance mode...
Success: Deactivated Maintenance mode.
Thu Jun 25 17:00:57 UTC 2026: Maintenance mode DISABLE completed successfully

cat /var/lib/mysql/pre-post-backup/pre_post_script_mysql.log

Thu Jun 25 16:56:20 UTC 2026: MySQL mode: freezing IO
Thu Jun 25 16:56:20 UTC 2026: MySQL freeze completed successfully
Thu Jun 25 16:59:55 UTC 2026: MySQL mode: resuming IO
Thu Jun 25 16:59:55 UTC 2026: MySQL resume completed successfully
```

The log files confirm that the pre- and post-snapshot actions were executed in the correct sequence: WordPress was paused first, then the database was paused, afterwards the database was resumed, and finally WordPress was resumed.

Configure the DataProtect protection group to use pre- and post-snapshot hooks

1. Open the protection group.

In Defender DataProtect, select the required cluster and navigate to **Data Protection > Protection**. Open and edit the Protection Group that protects the Red Hat OpenShift namespace `wordpress`. In the Protection Group, go to the object selection section and edit the `wordpress` namespace as shown in the following figure.




Figure: Customize the Namespace

2. Configure quiesce, pre- and post-snapshot hooks.

In the namespace settings, switch to the **Scripts** tab and select **Quiesce Rules**. This is where the pre- and post-snapshot actions are configured. Choose how the rules should be executed. There are three modes: apply rules together, apply rules independently, or apply rules sequentially. For a WordPress and database stack where order matters, sequential execution is usually the safest option. Quiescing is enabled by using **Quiesce Rules**. The quiesce mode is configured to **Apply following rules sequentially**, which means that the configured rules are executed one after another in a defined order.

Options for wordpress

Resources

Scripts

☒ Fail Backups on Hook Failure

Quiesce Mode

Apply following rules together

+ Add Rule

Cancel

Save

Figure: Define Snapshot Sequence and Hooks

Two rules are configured: Rule-1 and Rule-2. These rules are used to define the pre- and post-snapshot actions for the application components involved. The following figures show the detailed configuration of both rules:

Rule-1

▼

Pod Selector Labels

Key	Value
tier	frontend

+ Add Label

Pre Snapshot Scripts

Script-1

Container

wordpress

Optional Command

sh -c /var/www/html/wp-cli/pre_script_wp.sh

Timeout (Seconds)

▼

Optional

+ Add Script

Post Snapshot Scripts

Script-1

Container

wordpress

Optional Command

sh -c /var/www/html/wp-cli/post_script_wp.sh

Timeout (Seconds)

▼

Optional

+ Add Script

Figure: Rule-1: WordPress Pod Pre- and Post-Snapshot Scripts

Rule-1

Rule-2

Pod Selector Labels

Key	Value
tier	mysql

+ Add Label

Pre Snapshot Scripts

Script-1

Container

Optional Command

sh -c /pre_script_mysql.sh

Timeout (Seconds)

Optional

+ Add Script

Post Snapshot Scripts

Script-1

Container

Optional Command

sh -c /post_script_mysql.sh

Timeout (Seconds)

Optional

+ Add Script

Figure: Rule-2: MySQL Pod Pre- and Post-Snapshot Scripts

The commands below show how to capture the labels of the pods within the protected namespace. Each quiesce rule requires a pod selector label. The pod selector label is used to assign the rule to the correct pod.

The definition of the container within the pod is optional. If the selected pod contains only one container, Defender DataProtect automatically selects the correct container. If the pod contains multiple containers, the container name should be specified explicitly.

```
oc project wordpress
Now using project "wordpress" on server "https://api.ocp.ddp.ibm.lab:6443".

oc label pod wordpress-5c97bdbb5b-jtlv4 ddpPodLabel=wordpress
pod/wordpress-5c97bdbb5b-jtlv4 labeled

oc label pod wordpress-mysql-0 ddpPodLabel=mysql
pod/wordpress-mysql-0 labeled

oc get pods --show-labels
```

NAME	READY	STATUS	RESTARTS	AGE	LABELS
wordpress-5c97bdbb5b-jtlv4	1/1	Running	0	2d13h	app=wordpress,ddpPodLabel=wordpress
wordpress-mysql-0	1/1	Running	0	2d13h	app=wordpress,apps.kubernetes.io/po

3. Start a backup and observe the custom backup sequences.

After the quiesce rules have been configured, start a manual backup run to validate the complete sequence. Select the Protection Group and start the job by using **Run Now**. After the job has started, open the latest job run by clicking the corresponding timestamp. Then select the WordPress application and review the **Backup Task Activity** log. The activity log shows the full backup workflow. Below is a summary of the relevant backup activity:

```
The metadata backup is then started through Velero and completes successfully:
Starting backup job on Velero. Metadata backup job submitted successfully. Metadata backup job complete

After the metadata backup, the configured quiesce rules are executed.
The following pod selector label for quiescing: tier=frontend; passed.
The following pod selector label for quiescing: tier=mysql; passed.

After the application components have been quiesced, the CSI snapshots of both PVCs are created:
[volume wp-pvc] Taking snapshot: cohesity-snap-wp-pvc-49602.
[volume mysql-pvc] Taking snapshot: cohesity-snap-mysql-pvc-49602.
[volume wp-pvc] Snapshot completed successfully.
[volume mysql-pvc] Snapshot completed successfully.

After the snapshots are completed, the unquiesce sequence starts:
The following pod selector label for unquiescing: tier=mysql; passed.
The following pod selector label for unquiescing: tier=frontend; passed.

The PVC backup then completes successfully for both volumes:
Backup of volume: mysql-pvc completed successfully.
Backup of volume: wp-pvc completed successfully.
PVC Backup Summary: Total Volumes: 2 Successful Backups: 2 Failed Backups: 0 Skipped Backups: 0

The summary also confirms that CSI snapshots were used for both PVCs:
PVC: Name - wp-pvc, Data Read: 33B, CSI snapshot used: Yes, Volume Mode: Filesystem
PVC: Name - mysql-pvc, Data Read: 92.3 MiB, CSI snapshot used: Yes, Volume Mode: Filesystem
```

This confirms that the custom backup sequence was executed in the expected order: the WordPress application was quiesced first, then the MySQL database was quiesced, the PVC snapshots were created, the MySQL database was resumed, and finally the WordPress application was resumed.

The successful execution of the custom sequence can also be verified by reviewing the log files inside the containers, as shown in the figure below. The timestamps in the log files confirm that the actions were executed in the expected order: WordPress was paused first, then the database was paused, afterwards the database was resumed, and finally WordPress was resumed.

```
cat /tmp/pre_post_script_wp.log
```

```
Thu Jun 25 16:55:19 UTC 2026: Starting WordPress maintenance mode ENABLE
```

```
Enabling Maintenance mode...
```

```
Success: Activated Maintenance mode.
```

```
Thu Jun 25 16:55:21 UTC 2026: Maintenance mode ENABLE completed successfully
```

```
Thu Jun 25 17:00:55 UTC 2026: Starting WordPress maintenance mode DISABLE
```

```
Disabling Maintenance mode...
```

```
Success: Deactivated Maintenance mode.
```

```
Thu Jun 25 17:00:57 UTC 2026: Maintenance mode DISABLE completed successfully
```

```
cat /tmp/pre_post_script_mysql.log
```

```
Thu Jun 25 16:56:20 UTC 2026: MySQL mode: freezing IO
```

```
Thu Jun 25 16:56:20 UTC 2026: MySQL freeze completed successfully
```

```
Thu Jun 25 16:59:55 UTC 2026: MySQL mode: resuming IO
```

```
Thu Jun 25 16:59:55 UTC 2026: MySQL resume completed successfully
```

Figure: Pre- and Post-Snapshot Log File Validation

Auto protect Red Hat OpenShift namespaces

Auto Protect is used to automatically protect new Kubernetes or Red Hat OpenShift namespaces without manually editing the Protection Group each time a new namespace is created.

Auto Protect can be enabled at two levels. At the cluster level, all new namespaces that are added to the Red Hat OpenShift cluster are automatically added to the Protection Group. At the label level, only namespaces that have the selected label are automatically added to the Protection Group.

When Auto Protect is enabled, newly created namespaces are automatically added to the Protection Group and are protected from the next scheduled protection run. If a namespace is removed from the cluster, or if the selected label is removed from the namespace, the namespace is automatically removed from the Protection Group. Existing backups of that namespace are retained until the configured snapshot retention expires.

For cluster-level Auto Protect, enable Auto Protect at the cluster level and save the selection. With this setting, all future namespaces in the cluster are added automatically to the Protection Group.

To configure Auto Protect, navigate to **Data Protection > Protection** and select the Protection Group that protects the Red Hat OpenShift cluster. Open the actions menu for the Protection Group and select **Edit**. In the object selection area, click the edit icon and then enable the Auto protect function by clicking the A-sign that will toggle its color upon activation.




Figure: Auto Protect the Whole Cluster

For label-based Auto Protect, the `ocp-silver` DataProtect policy is assigned to the Protection Group. This Protection Group automatically protects all Red Hat OpenShift namespaces that are labeled with `ddp-policy=ocp-silver`.

Therefore, selected namespaces are first labeled with `ddp-policy=ocp-silver`:

```
oc label namespace openshift-ovirt-infra ddp-policy=ocp-silver
oc label namespace openshift-cloud-platform-infra ddp-policy=ocp-silver
oc label namespace openshift-infra ddp-policy=ocp-silver

oc get namespaces -l ddp-policy=ocp-silver
NAME                                STATUS    AGE
openshift-cloud-platform-infra      Active    71d
openshift-infra                     Active    71d
openshift-ovirt-infra               Active    71d
```

After the labels are applied it may be required to manually trigger an inventory refresh of the Red Hat OpenShift Source to retrieve the new label information. To configure label-based Auto Protect, navigate to **Data Protection > Sources** and select the registered Red Hat OpenShift cluster source. In the source view, select the Label Selection option (1). The label selection view shows the available namespace labels discovered from the Red Hat OpenShift cluster (2). Select the label **ddp-policy=ocp-silver** (3). This label is used to identify the namespaces that should be automatically protected by the new Protection Group. After selecting the label, continue to protect in step (5).

← Sources

api.ocp.ddp.ibm.lab:6443

Objects

Settings

⚠️ Unprotected

23

Objects

36 GiB

Data

✅ Protected

2

Objects

106 GiB

Data

Protection Status ▾

🔍

Labels

ddp-policy:ocp-silver ×

Select Labels

2

1

📦

☰

🔍

☐ Object

Status

Latest Snapshot

▼ ☐ 🏷️ ddp-policy:ocp-silver 3	-	
4 Auto Protect This Label cloud-platform-infra	🚫	-
Select All Child Objects fra	🚫	-
☐ ⚙️ openshift-ovirt-infra	🚫	-

Figure: Auto Protect Label Selection




Figure: Commit Auto Protect Label Selection

The following figure shows the assignment of the Protection Group **OCP-App-autoprotect-PG2** and the OCP-Silver policy. After the configuration is saved, Auto Protect is applied from the next protection run. Usually, the protection job starts immediately after the Protection Group is created.

▲ Unprotected

23
Objects

36 GiB
Data

✔ Protected

2
Objects

106 GiB
Data

Protection Status ▾

🔍

Labels

ddp-policy:ocp-silver ✕

Select Labels ▾

☐

Protect

5

▼			ddp-policy:ocp-silver	-
			openshift-cloud-platform-infra	 -
			openshift-infra	 -
			openshift-ovirt-infra	 -

Figure: Assign Protection Group and Policy

To observe the backup progress, navigate to **Data Protection > Protection** and select the new Protection Group OCP-App-autoprotect-PG2. The protection view shows the current job status and allows you to monitor the progress of the running backup.

New Protection

Add Objects

Source: <https://api.ocp.ddp.ibm.lab:6443> | Namespaces: 3 | Auto Protected: 1

Protection Group

Name
OCP-App-autoprotect-PG2

Policy

OCP-Silver
Backup 1d | Retain 2w | DataLock 2w | Replicate

Cancel

More Options

Protect

Figure: Auto Protect Protection Group Backup Run

Restoring an Red Hat OpenShift application

IBM Defender DataProtect can recover Red Hat OpenShift namespace data from protected recovery points. A namespace recovery can restore the PVCs and the metadata of the protected namespace. The recovery can be performed to the original location or to an alternate location, such as another registered Red Hat OpenShift cluster.

Restore an Red Hat OpenShift application namespace

When restoring an Red Hat OpenShift namespace back to its original location, Defender DataProtect brings back objects that are no longer present in the namespace. Resources that still exist are left unchanged, so the restore does not replace existing namespace content.

A namespace can also be restored to another registered Red Hat OpenShift cluster. In this case, the recovery workflow allows the administrator to choose the target cluster and map the original storage classes to storage classes that are available on the target environment. The restore scope can be limited when only selected data is required.

To restore an Red Hat OpenShift namespace, start the recovery workflow from the required recovery point. In the DataProtect UI select **Data Protection > Recoveries > Recover > Kubernetes Cluster > Namespace**.

Projects

★

Create Project

Filter

Name

wordpress

Name

wordpress

[Clear all filters](#)

Name	Display name	Status	Requester	Memory	CPU	Created
PR redbook-restore-wordpress	No display name	Active	kube:admin	475.7 MiB	0.003 cores	5 minutes ago
PR wordpress	No display name	Active	kube:admin	658.8 MiB	0.005 cores	Apr 17, 2026, 9:25 AM

Figure: Restore Kubernetes Namespace — Launch Recovery Workflow

Search and filter for the namespace name (1), select the namespace (2) and pick the date and source of the snapshot from either the local DataProtect Cluster or an archive target (3).

Recovery Options

Rename

Add Prefix

redbook-restore-

Add Suffix

Task Name

Recover_Kubernetes_Namespaces_Jul_1_2026_7_09_PM

Namespace Resources

 wordpress

All Resources



Recover PVCs Only

No

Include or Exclude Labels

Disabled

Recover

Cancel

Figure: Namespace Recovery — Select Namespace and Pick Snapshot

The namespace can be restored to the original location or to an alternate Red Hat OpenShift cluster. In this step, the restore scope can also be adjusted by selecting which PVCs should be included or excluded.

Additional options allow a prefix or suffix to be added to the restored namespace name. The recovery can also be further customized by including or excluding namespace resources, filtering resources by labels, or restoring only the PVCs without the additional namespace resources.




Figure: Restore Namespace — Select PVCs




Figure: Restore Namespace — Custom Settings

Review the restore results in the DataProtect Restore Job report. After the restore job has completed, also verify the Red Hat OpenShift cluster and confirm that the restored namespace was created successfully.

×

New Recovery

K8

Kubernetes Namespaces

wordpress

Namespaces

Jun 24, 2026 10:56pm

Snapshot

Local

Location

Recover To

☒

Original Location

☐

New Location

Granular Recovery

☒

Enable Granular Recovery

PVCs

☒

PVCs

☒

wp-pvc

☒

mysql-pvc

Figure: Restored Namespace Verified in Red Hat OpenShift

After the restore completed, the recovered namespace was verified in the Red Hat OpenShift console. The MySQL pod started successfully, but the WordPress pod did not start and entered the `CrashLoopBackOff` state.

The pod logs show that the WordPress webserver cannot bind to port 80 because of a permission issue. The error message indicates that no listening socket can be opened and the container shuts down. This means that the WordPress container was restored, but the runtime security configuration does not allow the application process to bind to the required port. This behavior is not a backup data problem. It indicates that additional Red Hat OpenShift runtime settings must be reviewed and adjusted before the restored WordPress application can start successfully.




Figure: WordPress Pod in CrashLoopBackOff State After Restore — Post-Restore Fix Required

Restore single files or folders

For Red Hat OpenShift workloads, Defender DataProtect can restore individual files and folders from protected PVC backups. This is useful when only a small part of the application data must be recovered, for example after accidental deletion or file corruption.

Instead of restoring the complete PVC, administrators can browse the backup content and select only the required files or directories. This reduces recovery time and avoids replacing data that is still valid in the running environment. For file and folder level recovery, administrators can either save the selected items outside of the cluster or restore them back through a recovery workflow. When the download option is used, a target download location can be defined.

The first example shows how to restore a single file back to the original PVC. For demonstration purposes, the pre- and post-backup hook log file was deleted from the WordPress application.




Figure: Single File Restore — WordPress PVC Content After File Deletion

In the DataProtect UI, navigate to **Data Protection > Recoveries > Recover > Files and Folders**. To reduce the search result list, set the object name to **WordPress (1)** and enter the name of the deleted log file as the search value (2).

The search result shows two matching log files: one from the WordPress container and one from the MySQL database container. Select the file `pre_post_script_wp.log` (3) and choose the point-in-time snapshot that should be used for the restore (4).




Figure: Single File Restore — Search and Select File

Continue and specify the recovery options and click **Recover**.

FI

Files

1

Files

Jul 2, 2026 09:26

Snapshot

Local

Location

wordpress

Namespace

Recover To

☒ Original Location
 ☐ New Location

Recovery Options

Overwrite Existing File/Folder

☒ Overwrite

Preserve File/Folder Attributes

☐ Preserve File/Folder Attributes

Continue on Error

Yes

Cluster Interface

Auto Select

Task Name

Recover_Files_Jul_3_2026_7_56_PM

Recover

Cancel

Figure: Single File Restore — Recover to Original Location

After successful recovery, the directory listing in the WordPress PVC lists the restored file.




Figure: Single File Restore — WordPress PVC Listing After File Restore

The second example shows how a single file can be downloaded from a selected point-in-time backup of the WordPress PVC. This time, the file is not located by searching for its name. Instead, the backup snapshot is opened and the required file is selected by browsing through the protected file system.

Navigate to **Data Protection > Recoveries > Recover > Files and Folders**. Instead of using the file search option, select **Browse** (1) and enter `wordpress` (2) in the search field. Press Enter to list the available backup snapshots for the WordPress application.

Select the required snapshot and click **Browse** (3) to open the file browser for the WordPress backup. From there, navigate through the directory structure and select the file that should be downloaded or restored.




Figure: Download a File From a Namespace Backup Snapshot

In the file browser, select the required PVC. In this example, the namespace contains one PVC for the WordPress application and one PVC for the MySQL database.

Browse through the directory structure of the selected PVC and locate the file that should be downloaded. The pre-and post-backup hook log file is selected (1) and downloaded to the local workstation for review.

After selecting the file, choose the required snapshot time (2), (3), confirm with **Apply** (4) and click the **Download** button (5).




Figure: Download File From PVC Snapshot

After the download request is submitted, DataProtect creates a new recovery task. The progress of this task can be monitored under **Data Protection > Recoveries > Recovery Task**.

12 Succeeded	0 Warning	0 Failed	0 Running	2 Canceled
Recovered From ▾	Recovery Type ▾	Status ▾	Past 7 Days ▾	🔍
Recovery Task	Start Time ↓	Status	Duration	
 Download_Files_Jul_3_2026_8_21_PM 1 Object	Jul 3, 2026 20:21	✓ Succeeded	1s	
 Recover_Files_Jul_3_2026_7_56_PM 1 Object	Jul 3, 2026 20:02	✓ Succeeded	1s	
 Resubmit_Resubmit_Recover_Files_Jul_3_2026_7_25_PM 1 Object	Jul 3, 2026 19:36	✓ Succeeded	1s	
 Resubmit_Recover_Files_Jul_3_2026_7_25_PM 1 Object	Jul 3, 2026 19:35	✓ Succeeded	1s	
 Recover_Files_Jul_3_2026_7_25_PM 1 Object	Jul 3, 2026 19:25	✓ Succeeded	1s	
 Recover_Kubernetes_Namespaces_Jul_3_2026_6_49_PM 1 Object	Jul 3, 2026 18:51	✓ Succeeded	18m 14s	

Figure: Recovery Task List

To download the recovered file, open the recovery task and select the download option. DataProtect downloads the selected file as a compressed ZIP file to the local workstation.

If multiple files, folders, or a complete PVC were selected, the selected content is also packaged into a single ZIP file before it is downloaded.

 **Files**

☐ Files and Folders
 ☒ Browse 1

WordPress 2

Object Type ▾

Source ▾

Protection Group ▾

Storage Domain ▾ 3

 **wordpress**
 Cluster: https://api.ocp.ddp.ibm.iab:6443 | Protection Group: OCP-App-PG1 | Last Backup: a day ago | Size: 6 GiB

Browse

Figure: Download the File From Backup Snapshot

If both PVCs of the WordPress application are selected instead, the downloaded ZIP file contains the data from both volumes. After the ZIP file is extracted on the local workstation, the directory structure shows the content of each selected PVC.




Figure: Downloaded ZIP Containing Full Namespace Content

Protecting OCP-V virtual machines

IBM Defender DataProtect can also protect Red Hat OpenShift Virtualization workloads. Red Hat OpenShift Virtualization allows virtual machines to run inside Red Hat OpenShift by using KubeVirt. From a protection point of view, virtual machines are part of Kubernetes namespaces and use PVCs as their virtual disks. More details on KubeVirt and Red Hat OpenShift Virtualization can be found in [Red Hat OpenShift Virtualization Basics](#).

Defender DataProtect supports the backup of Red Hat OpenShift Virtualization VMs, PVCs, Kubernetes namespace metadata, and Red Hat OpenShift cluster resources. PVCs with Filesystem and Block volume mode are supported. For VM backups, Defender DataProtect leverages the `VirtualMachineSnapshot` (VMS) custom resource, which can create one snapshot for all volumes that belong to a VM in a single backup operation. PVCs in the same namespace that are not attached to a VM can still be protected by using individual CSI volume snapshots.

During a backup, the metadata of all resources in the namespace is backed up. However, only the volumes of the selected VMs are included in the VM backup. During recovery, if all VMs in the namespace are recovered, the namespace metadata is also recovered. If only selected VMs are recovered, only the metadata related to these selected VMs and their volumes are restored.

Important requirements and considerations

Review the following requirements and considerations before configuring protection for Red Hat OpenShift Virtualization:

- Red Hat OpenShift Virtualization protection is supported for Red Hat OpenShift on-premises environments only.
- CSI snapshot support should be enabled and used for crash-consistent backups.
- The Data Mover must be available in the Red Hat OpenShift cluster, and the latest Data Mover version should be used for Block mode VM protection.
- The Data Protect Plugin and KubeVirt/Velero related images must be available in an image registry that the Red Hat OpenShift cluster can access.
- OADP support depends on the Red Hat OpenShift version. For Red Hat OpenShift versions earlier than 4.17, deployment of the OADP operator is not supported.
- Incremental backup is not supported for Red Hat OpenShift Virtualization at the time of this writing. This may change when Red Hat supports change block tracking in ODF (see Red Hat tech preview).
- Granular file and folder recovery inside VMs is not supported.
- VM quiescing during backup is not supported, but the qemu-guest-agent (if installed in a VM) can execute pre- and post-snapshot scripts.
- If SELinux is enabled inside the VM, the required permissions must be set for qemu-agent to allow guest filesystem freeze commands.
- Recovery depends on the availability of the original storage class. If the storage class used during backup no longer exists and is also not available on the target cluster, recovery is not supported.
- The required VM machine type must exist on the original or alternate target cluster. Otherwise, the recovery process may complete, but the VM may not run correctly.

Configure protection for Red Hat OpenShift Virtualization namespaces

After the Red Hat OpenShift Virtualization requirements have been reviewed, the next step is to create a Protection Group for the namespace that contains the virtual machines.

To configure the protection, log in to the IBM Defender Data Management Service and select the required Defender DataProtect cluster. Navigate to **Data Protection > Sources** and select the registered Red Hat OpenShift cluster source. In the Red Hat OpenShift source view, filter the namespace (1) that contains the Red Hat OpenShift Virtualization virtual machines (2).




Figure: KubeVirt Namespace Protection Selection

When the namespace has been selected, continue with the wizard by clicking the **Protect** button (3). Enter the Protection Group name, for example `ocp-vm-pg3`, and assign the required policy. In this example, the `ocp-gold` policy is being assigned, which implements a frequent backup with a backup data retention of two weeks, data lock, replication, and archive.

Monitoring the backup progress of kubevirt namespaces

After the Protection Group has been created, the first backup usually starts immediately. To review the backup progress, navigate to **Data Protection > Protection** and select the Protection Group `ocp-vm-pg3`. The protection view shows the current job status and allows administrators to monitor the progress of the running backup.




Figure: OCP VM Protection Group Job Overview

To view more details for the new protection job, collect error information, or review the replication and archive status, select the required Protection Group `OCP-VM-PG3` in the protection overview. The detailed job view shows the individual protection runs, their status, start and end times, and any warnings or errors that occurred during execution. The following figure shows the Backup Task activities for the kubevirt namespace `vm-websrv`. The namespace consists of two virtual machines and — in this configuration — all VMs and all PVCs are protected by default.




Figure: Backup Task Activities for KubeVirt Namespace

Customize the backup scope for Virtualization namespace

The backup scope can be customized by using include or exclude settings. These settings can be used to include or exclude VMs or PVCs. This allows administrators to define which virtual machines or persistent volumes should be part of the Protection Group and which objects should be skipped.

When customizing the backup scope for an Red Hat OpenShift Virtualization namespace, some options are mutually exclusive. The Protection Group can be configured to use include or exclude rules for VMs, include or exclude rules for PVCs, or resource exclusions. These options should be selected depending on the required backup scope.

In addition, the script configuration for quiescing cannot be selected when VM inclusion/exclusion is enabled.

Note: Including or excluding specific VMs within an Red Hat OpenShift Virtualization namespace requires the Data Protect Plugin to be installed in the Red Hat OpenShift cluster. This prerequisite is covered in "Prepare the Red Hat OpenShift cluster".

This section describes the option to exclude VMs from a namespace backup.

In the Defender DataProtect UI navigate to **Data Protection > Protection**. Open and edit the Protection Group **OCP-VM-PG3** that protects the Red Hat OpenShift Virtual Machine namespace **vm-websrv**. In the Protection Group, go to the object selection section and edit the `vm-websrv` namespace as shown in the following figures.

Edit Protection

K8 Kubernetes

Source

Registered Source

https://api.ocp.ddp.ibm.lab:6443

Objects

1

Namespaces

1

Manually Protected

Protection Status ▼

vm-websrv



https://api.ocp.ddp.ibm.lab:6443



vm-websrv



Figure: Customize VM Namespace Backup

Options for vm-websrv

Resources

Scripts

☒ VM Inclusion/Exclusion

☒ Include
 ☐ Exclude

Include VMs

vm-websrv

vm-websrv-clone

vm-websrv

☐ Persistent Volume Claim(PVC) Inclusion/Exclusion

This overrides inclusion or exclusion settings made for Protection Group level in Additional Settings

☐ Resource Inclusion/Exclusion

Cancel

Save

Figure: Selecting a Dedicated VM to be Included in Namespace Backup

Implementing consistency for apps running inside VMs

Defender DataProtect does not support quiescing a VM during the backup process, but the qemu-guest-agent (if installed in a VM) can execute pre- and post-snapshot scripts.

For more information, see "The QEMU Guest Agent" and "Pre- and Post-Snapshot actions with fsfreeze-hook scripts" in chapter 2 of this publication.

Important: If not installed automatically during the VM deployment, ensure that the qemu-guest-agent is installed and activated in a virtual machine. Otherwise, online backups of the VM will be crash-consistent only. More details on installing the QEMU guest agent can be found in the product documentation.

Note: The `[pre/post].hook.backup.velero.io` annotations are automatically assigned to the virt-launcher pod when the pod is created (or re-created after a VM reboot). These annotations should be removed, as the freeze and unfreeze activities are unnecessary during the namespace metadata backup phase (see also [Red Hat Solution 7137659](#) for more details). The freeze and unfreeze

operations that are triggered with the VirtualMachineSnapshot (VMS) operation are scheduled properly before and after the VM disk snapshots.

Sample fsfreeze-hook script


The VM (`vm-websrv`) which was backed up in the previous example includes an Apache webserver. For demonstration purposes, an fsfreeze-hook script was created that writes log files into the Apache web server html-data directory. The files contain a timestamp and an indication whether the log was created before the freeze or after the thaw operation:

```
[root@vm-websrv ~]# cat /etc/qemu-ga/fsfreeze-hook.d/pre-post-script.sh

#!/bin/sh
#####
# This script is called by the fsfreeze --freeze and --thaw actions
# It runs pre- and post- tasks around freezing the VM's filesystems
#####

case "$1" in
    freeze)
        #-----
        # "Suspend" part - called with freeze
        #-----
        touch "/var/www/html/$(date +"%Y%m%d_%H.%M.%S").PRE-SNAPSHOT"
        exit 0
        ;;
    thaw)
        #-----
        # "Resume" part - called with thaw
        #-----
        touch "/var/www/html/$(date +"%Y%m%d_%H.%M.%S").POST-SNAPSHOT"
        exit 0
        ;;
esac
```

When a backup job for the VM is completed, we would expect to have two files logged in the html-data directory of the Apache web server, which have been created as part of the fsfreeze-hook mechanism. After the backup job is completed, the sample web page looks like this:



The screenshot shows a web browser window with the address bar displaying "http://10.7.60.60". The page title is "Redbooks - IBM Storage Defender: Protecting OpenShift Containers and Virtual Machines". Below the title is a table listing files in the html-data directory. The table has three columns: "Name", "Last modified", and "Size". There are two entries in the table, both with a size of 0.

Name	Last modified	Size
20260706_11.15.24.POST-SNAPSHOT	2026-07-06 11:15	0
20260706_11.15.20.PRE-SNAPSHOT	2026-07-06 11:15	0

Figure: The fsfreeze-hook Script Has Logged the Freeze and Thaw Operations

Perform a VM restore

This section describes how to restore the sample VM `vm-websrv` to a different namespace on the same Red Hat OpenShift cluster using Defender DataProtect.

1. In the Defender DataProtect GUI, click **Data Protection > Recoveries > Recover** and select **Kubernetes Cluster > Namespace** from the drop-down menu.




Figure: The Recoveries Menu in the DataProtect UI

2. In the **New Recovery** screen, search for available backup items (in this example, use the search string `OCP`, as this string is part of all Protection Group names that cover Red Hat OpenShift resources), select the namespace `vm-websrv`, select the PIT to be restored (leave it at the default, which is the latest backup) and click **Next: Recover Options**.




Figure: Select the Namespace and PIT to Restore

Note: The Recover Options presented on the next screen vary depending on the following recovery scopes:

- If you select an entire Protection Group to be restored (e.g. `OCP-VM-PG3`), it is not possible to include or exclude specific cluster- or namespace resources later.
- If you select a dedicated namespace (`vm-websrv` in this example), it is possible to select dedicated VMs (if multiple VMs reside in the same namespace).

3. In the Recovery Options menu, select **Recover to Original Location** (to perform the restore to the original Red Hat OpenShift cluster) and add a suffix `-copy` to redirect the restore to a new namespace.

Kubernetes Namespaces

vm-websrv
Namespaces

Latest
Snapshot

Local
Location

Recover To

☒ Original Location ☐ New Location

Recovery Options

Rename

Add Prefix

Add Suffix
-copy

Figure: Select the Restore Location and Target Namespace Suffix

4. Click **Namespace Resources > VM Inclusion/Exclusion** and select the VM `vm-websrv`.

Options for vm-websrv

Resources Storage Class

☒ VM Inclusion/Exclusion

☒ Include ☐ Exclude

Include VMs

vm-websrv

☐ Persistent Volume Claim(PVC) Inclusion/Exclusion

☐ Resource Inclusion/Exclusion

Cancel Apply

Figure: Select a Dedicated VM 'vm-websrv' to be Included in the Restore

- Click the **Recover** button to start the restore of the VM `vm-websrv` to a new target namespace `vm-websrv-copy` on the same Red Hat OpenShift cluster.
- In the Defender DataProtect UI, click **Data Protection > Recoveries** to get an overview of recent recovery tasks (active and completed).




Figure: Overview of Recovery Tasks

- Click the recovery task that was just started to see more details.




Figure: Recovery Activity Details

- In the Red Hat OpenShift UI, go to **Virtualization > VirtualMachines** and validate that the new namespace was created and contains the restored VM.




Figure: Validate in Red Hat OpenShift That the VM Was Properly Restored

- Note that the original namespace contains two VMs, but the restore process just recovered a single VM (as expected).

Note: When only a subset of VMs is being recovered from a namespace, the restore process only takes care of the selected VMs and their related PVCs, but not of other namespace-scoped resources (such as services).

- The VM restore did not recover the LoadBalancer services from the original namespace, as the recovery options used a VM include/exclude configuration. To allow SSH and HTTP access to the restored VM, create these services manually by applying the following YAML to the Red Hat OpenShift cluster:

```

apiVersion: v1
kind: Service
metadata:
  name: vm-websrv-ssh
  namespace: vm-websrv-copy
  annotations:
    metallb.universe.tf/loadBalancerIPs: '10.7.60.61'
    metallb.universe.tf/allow-shared-ip: "key-to-share-10.7.60.61"
spec:
  ports:
    - name: ssh
      protocol: TCP
      port: 22
  type: LoadBalancer
  selector:
    kubevirt.io/domain: vm-websrv
---
apiVersion: v1
kind: Service
metadata:
  name: vm-websrv-httpd
  namespace: vm-websrv-copy
  annotations:
    metallb.universe.tf/loadBalancerIPs: '10.7.60.61'
    metallb.universe.tf/allow-shared-ip: "key-to-share-10.7.60.61"
spec:
  ports:
    - name: httpd
      protocol: TCP
      port: 80
  type: LoadBalancer
  selector:
    kubevirt.io/domain: vm-websrv

```

Figure: LoadBalancer services for SSH and HTTP access to the restored VM

1. Review the sample web page that is provided by the restored VM.




Figure: Status of the Restored VM

Note that the post-snapshot log file is missing after the restore. This is an expected condition, because that log file was created only after the VM disk snapshot was taken.

With VMs hosting more complex applications or databases, post-restore activities might need to be performed. If, for example, an IBM Db2 database is in write-suspend mode during the backup, it will also be in write-suspend mode after the restore and must therefore be manually resumed.

Summary

This chapter covered how IBM Storage Defender DataProtect protects Red Hat OpenShift workloads, from environment preparation through to recovery. After registering an Red Hat OpenShift cluster as a source and defining Protection Policies and Protection Groups, the chapter demonstrated backing up a sample WordPress application, customising backup scope with pre- and post-snapshot quiesce hooks, and using Auto Protect to extend coverage to new namespaces automatically. Recovery scenarios included full namespace restore, single file and folder recovery, and downloading backup content directly to a workstation. The chapter also addressed Red Hat OpenShift Virtualization, showing how to protect KubeVirt namespaces, select individual VMs for backup, implement filesystem consistency using the fsfreeze-hook mechanism, and restore a virtual machine to a new namespace on the same cluster.

IBM Storage Defender DRS integration

Data protection without visibility is just expensive storage. You can have perfect backups running on schedule, immutable copies in multiple locations, and policies that look good on paper, but if you cannot see what is actually protected, what failed, or what drifted from policy, you are operating blind.

In the previous chapters, we covered how IBM Fusion provides backup and restore capabilities for Red Hat OpenShift containers and virtual machines. IBM Storage Defender Data Resiliency Service (DRS) provides the orchestration layer above these backup systems. It gives you centralized visibility, policy governance, and orchestrated recovery workflows across IBM Fusion, IBM Defender Data Protect, and other supported platforms.

This section explains how to integrate these backup systems with DRS, how to use recovery groups and governance profiles to ensure protection meets business requirements, and how to leverage DRS for faster, more reliable recovery operations.

IBM Storage Defender Data Resilience Service overview

The information that is included in this section is intentionally kept brief. For more information about this topic, see the [IBM Storage Defender documentation](#) or see the IBM Redbooks publication [IBM Storage Defender: Data Resiliency Service \(REDP5744\)](#).

IBM Storage Defender Data Resiliency Services (DRS) is a SaaS-based orchestration platform designed to unify and automate cyber-resilience workflows across heterogeneous, hybrid cloud storage environments. By integrating data protection, proactive monitoring, and rapid recovery orchestration, DRS addresses the critical gap between detection and restoration in the event of a ransomware or cyber-destructive event.

The primary targets for IBM Storage Defender DRS are enterprise-scale hybrid cloud infrastructures that necessitate high-availability and rapid recovery for mission-critical workloads. Specifically, it is engineered for the following workload types:

- **Mission-Critical Applications:** Workloads requiring stringent Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO).
- **Cyber Vault Architectures:** Environments utilizing air-gapped or immutable storage tiers, such as IBM Safeguarded Copy.
- **Heterogeneous Ecosystems:** Organizations managing a mix of IBM and third-party primary and secondary storage assets.




Figure: IBM Storage Defender DRS dashboard showing overview

Core technical features

- **Orchestrated Recovery Workflows:** DRS provides automated, policy-driven restoration of complex multi-tier applications. This eliminates manual intervention, reducing the risk of human error during high-stress recovery scenarios.
- **Clean Room Recovery:** The platform facilitates the automated provisioning of isolated "Clean Room" environments. These sandboxes allow for the forensic analysis and validation of recovery points to ensure data is free of latent malware before being reintroduced to the production environment.
- **Centralized Control Plane:** A unified, SaaS-based dashboard provides global visibility into the resiliency posture of the entire data estate, encompassing both on-premises and cloud-resident assets.
- **Cyber Resilience Integration:** DRS integrates natively with SIEM (Security Information and Event Management) and SOAR (Security Orchestration, Automation, and Response) platforms, such as IBM QRadar. This allows for automated recovery triggers based on high-fidelity security alerts.
- **Immutable Snapshot Management:** The service manages the lifecycle of immutable snapshots and logical air-gaps, ensuring that recovery candidates remain protected from encryption or deletion by unauthorized actors.

Note: IBM Storage Defender DRS is architected to operate as a modular component of the broader IBM Storage Defender suite, allowing for seamless scaling as data volumes and security requirements evolve.

Connection Manager

A key architectural component of IBM DRS is the Connection Manager, a VM-based application that needs to be installed on-premises within each data center that DRS needs to acquire information from. It acts as the information broker between the data sources within the data center and DRS itself.




Figure: Connection Manager architecture diagram within DRS

Data sources can be any of the following types:

- Storage Systems
- Backup Systems
- Virtual machines which either host production workloads, or which can be used as clean rooms

Each data source needs to allow data to be collected, either using an API token, or a user/password combination. An overview of all data that will be collected by Connection Manager can be found in the following appendix: [Inventory metadata collected from Data Resiliency Service](#).

Defender Data Resiliency Service integration: Visibility into what is being protected

In IBM Storage Defender, the Data Resiliency Service (DRS) is the place where the environments “protection reality” becomes visible. The goal is not only to say that backups exist, but to make it clear what assets are connected, what copies exist, and how the current protection posture aligns with what the business expects. DRS does this by combining inventory metadata, copy awareness, and governance alignment into a single operational view in a Defender architecture. This visibility is not a “nice dashboard feature”, it is the key mechanism that allows you to answer, with evidence, these basic questions that matter when protecting production resources:

- What is protected right now?
- What is missing (and no one noticed yet)?
- What is protected, but outside the service level agreement. Or off policy?
- If we had to recover today, what versions do we actually have (what versions, where are they and how old)?

The Data Resiliency Service is the governance and orchestration layer that ties actions, resources, and integrations together. Access points are how IBM keeps permissions cleanly mapped to your enterprise role-based access control (RBAC) and separation of duties: you scope who can see items, who can run actions, and where. This level of control and organization becomes important as the same tenant can have multiple clusters, multiple backup domains, and multiple recovery locations. With complex environments it's even more crucial to still be able to have a controlled model of who can 'view' versus who can execute recovery actions.

Where visibility comes from

In production, asset visibility should never be an "open bar" where everyone choose what they want. Access points are the guardrails, a managed bundle of actions/resources/integrations treated as a unit that can be assigned to specific users. You decide what a team can see versus what it can do and provide access organized by environment, by location, by domain. This makes administration simple to understand and saves you later, when an audit asks who touched what.

From a protection view, DRS does not "guess" or assume what is protected or being protected. It consumes inventory metadata coming from the on-prem integration layer (connection manager) and from the systems that actually manage protection (for example Data Protect clusters and other registered sources). IBM explicitly describes that inventory metadata is securely stored in the DRS SaaS instance and retained until the connection manager sending it is removed. This is a subtle but important operational point: visibility is not only real-time; it is also designed to survive short outages and provide consistent governance data.

Connection Manager: The bridge

The Data Resiliency Service is inventory-driven. Connection Manager pulls inventory metadata from your connected sources and pushes it to the software-as-a-service endpoint. The cloud service uses that metadata to build the dashboards and views. If Connection Manager is down or removed, inventory stops updating and the dashboard becomes stale (and can mislead you). No Connection Manager, no trustworthy visibility.

In practice, you deploy Connection Manager as a virtual machine in vCenter, or on bare metal using a bootable ISO image. Under the hood it is Red Hat Enterprise Linux with a small setup wizard. The fact that it has a simple deployment is why people often underestimate the importance and key role of the connection manager.

Name	Type	Connection managers	Last inventory	Last inventory time	Duration	Credential
[Redacted]	IBM Fusion	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	8 seconds	[Redacted]
[Redacted]	IBM Fusion	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	8 seconds	[Redacted]
[Redacted]	IBM Fusion	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	8 seconds	[Redacted]
[Redacted]	IBM Data Protect	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	3 seconds	[Redacted]
[Redacted]	IBM Data Protect	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	10 seconds	[Redacted]
[Redacted]	IBM FlashSystem	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	0 seconds	[Redacted]
[Redacted]	IBM FlashSystem	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	0 seconds	[Redacted]
[Redacted]	IBM Storage Protect	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	2 seconds	[Redacted]
[Redacted]	IBM Storage Protect	CE4S-EMEA-ConMan1	Success	Jun 8, 2026, 9:16 AM	2 seconds	[Redacted]

Figure: IBM Storage Defender Connection Manager connections panel

For more information, see [the following IBM documentation](#).

From a network perspective, Connection Manager is your controlled egress point. It initiates outbound connections to the cloud service and as such, outbound internet access is required. If you use a proxy, set it up as you normally would. If there is no proxy, make sure Domain Name System resolution works and your certificate inspection rules do not break Transport Layer Security. This is most often where the configuration runs into issues during deployment. Usually there is no inbound traffic required.

IBM calls out the use of TCP port 443 to domains under `storage-defender.ibm.com` and `cloudfront.net`. WebSocket connections in the [DRS network requirements](#) as required for Data Resiliency Service traffic. HTTPS is used for content delivery and updates (CloudFront).

Note: A firewall or proxy rule issue, or deep inspection that resets the session is the most common type of issue seen during a deployment.

Another common issue is that onboarding is driven by an application programming interface key. This key expires after a short amount of time once generated. Because of this it is suggested to generate the key and use it right away to ensure there is not an issue with the key being expired and the connection being rejected.

Once Connection Manager is in place, DRS organizes visibility into a few core angles:

- Resources and assets: what systems exist, and what the Data Resiliency Service can actually protect and recover. IBM's own definition is clear: a "resource" is the unit you protect and later recover, and it is typically managed under Storage systems (volumes) and Backup systems (data resources).
- Recovery groups: your 'recover these together' bucket. Some are created automatically; others you create to match how the business runs. This becomes the bridge between raw inventory and real recovery workflows.
- Posture, copies, and readiness: DRS uses the collected inventory to expose recovery posture views (dashboards/graphs) and to show "available resources" and "available copies" (the practical confirmation that backups and retention are actually there). This is explicitly mentioned in IBM's Fusion integration flow as well: clusters, applications, backups, policies, and backup locations are collected and then shown in posture/copies views
- Integrations: forward posture and events to Splunk or QRadar.




Figure: Data Resiliency Service integrations tab showing SIEM and notifications

What visibility means for Red Hat OpenShift containers and virtual machines

On Red Hat OpenShift, 'protected' only matters at the application level - the namespace, labels, and whatever groups your app into a recoverable unit. Cluster level matters too, because you might have multiple clusters (hub and spoke) and multiple environments. With IBM Fusion feeding the Data Resiliency Service, you see application-level inventory and copies, not just a cluster sitting there as 'registered'. Here is what you can verify, quickly:

- Cluster identity (hub and spoke)
- Which cluster you are looking at
- Application identity, and the protection policy attached
- Kubernetes objects in the app namespace such as Services, Pods, Secrets, Deployments, ConfigMaps
- Virtual machine names that exist inside the application namespace (relevant for Red Hat OpenShift Virtualization workloads)
- Backup versions with timestamps, expiry, status, and backup size
- Policy metadata (interval, retention) and backup locations (S3/COS style)

Put it together and you get a chain that can be defended in audit: application to objects to versions to policy to location.

This is the “visibility story” you want. Defender/DRS is not only telling you “a cluster is registered” but provides a governance-grade inventory that ties together: application → objects → versions → policies → backup location. Visibility connects directly to operational KPIs (health + SLA + capacity).

From the Data Management perspective, you can get a consolidated view across clusters including cluster health, protection status, SLA compliance, and capacity available. That matters because visibility is not only a list — it becomes the basis for an operator’s daily controls (“what is drifting”, “what is failing”, “what is at risk”)

Control plane

IBM Storage Defender Data Resiliency Service (DRS) acts as a centralized control plane by seamlessly integrating IBM Defender Data Protect and IBM Fusion as backup sources, enabling a unified approach to ransomware-resilient data protection. By aggregating these systems, DRS gives backup and security administrators a single, trusted view of cyber-recovery readiness, including:

- **Centralized Policy Visibility:** Unified insight into protection policies defined across backup platforms, ensuring ransomware-resilient configurations—such as immutability, retention, and isolation—are consistently applied.
- **Backup Posture & Risk Awareness:** Consolidated backup statistics that help administrators quickly assess overall protection coverage and identify gaps that could be exposed during a ransomware attack.
- **Protection Health & Freshness Monitoring:** Clear visibility into which resources are protected, whether backups exist, and whether they are current—helping detect stale or missing backups that increase recovery risk during an incident.
- **Workload-Level Protection Coverage:** Transparency into the types of resources being protected (e.g., VMware, containers, Kubernetes applications), ensuring critical workloads are included in ransomware recovery plans.
- **Policy Enforcement & Validation:** Visibility into which backup and resiliency policies are applied to each workload, simplifying validation that ransomware-ready controls are in place.
- **Resiliency & Governance Posture:** Insight into whether governance profiles are configured and whether protected resources meet key ransomware-resilience attributes, such as immutable backups, compliance with RPO/RTO targets, and secure recovery readiness.

By unifying policy governance, backup health, and resiliency posture across IBM Data Protect and IBM Fusion, DRS enables faster attack impact assessment, trusted recovery decisions, and reduced risk of reinfection—transforming backups from a passive safeguard into an active cyber-resilience strategy.

The following sections provide the steps to onboard and integrate IBM Fusion and IBM Defender Data Protect as backup systems on the DRS platform.

[Integration of Fusion data protect in DRS](#)

[Integration of IBM Defender Data protect in DRS](#)

Orchestrated recovery workflows

IBM Storage Defender Data Resiliency Service (DRS) manages the entire flow of recovery while IBM Defender Data Protect and IBM Fusion execute the workflows.

The recovery workflow used by DRS is completed by going through the following steps: Detect -> Assess & Validate -> Orchestrate -> Recover -> Verify

Detection

Recovery can be triggered by the following events:

- Security incidents or anomalies detected by Defender sensors
- Operational failures or data corruption
- Manual trigger of recovery workflow by security administrators in case of a cyber incident

DRS evaluates the context of the event, with the help of governance profiles, including affected workloads, protection status, and recovery objectives, before initiating any restore action.

Assessment & validation

DRS, with the help of the control plane and governance profiles, assesses the following items:

- Availability of recovery points (backups, snapshots, replicas)
- Freshness and immutability of backups
- Compliance with defined RPO and RTO targets
- Governance controls required for ransomware recovery

Orchestration

DRS orchestrates the following multi-step recovery workflows with the help of governance profiles across IBM Data Protect and IBM Fusion:

- Selection of the most appropriate recovery point
- Controlled restore sequencing (application → data → dependencies)
- Cross-platform coordination:
 - IBM Data Protect for backup-based restores
 - IBM Fusion for Kubernetes-native and snapshot-based recovery
- Automated execution of recovery runbooks

Workflows allow for consistent, repeatable, and policy-driven orchestration.

Recover

For ransomware scenarios, DRS supports secure recovery models, such as:

- Recovery to isolated or clean environments
- Controlled restore authorization
- Prevention of overwrite on compromised systems
- Validation-first recovery approaches

This ensures recovered workloads are not immediately re-exposed to threats.

Verify

During and after recovery, DRS provides:

- Real-time recovery progress visibility
- Success/failure validation
- Post-recovery compliance confirmation
- Audit logging of every recovery action

For more information about how to set up governance profiles for IBM Data Protect & IBM Fusion, see [grouping and governance profiles in DRS](#).

Integration of IBM Fusion in DRS

Within the DRS context, IBM Fusion is treated as a backup system. The Connection Manager is used to establish the connection with DRS and collects relevant metadata.

Prerequisites

The following prerequisites are required for integrating IBM Fusion with DRS:

- Access to the DRS tenant where the new backup system will be integrated.
- Access to the IBM Fusion cluster that is planned to be integrated with DRS. It needs to be the Fusion cluster which runs the Backup - Restore service as the hub role.
- Credentials that the Connection Manager will use to access the Fusion cluster. The credential can either be a user / password combination, or an API key.
- A Connection Manager which is linked to the DRS tenant. The Connection Manager needs to have local network access to the IBM Fusion cluster.

Finding the Fusion cluster URL

For the DRS integration, you must find the correct URL which will be used by the Connection Manager for collecting metadata. First, identify the Fusion cluster which runs the Backup - Restore service as the hub role:

Step 1: Log in to the RHOS cluster GUI where the Fusion SW is installed

Step 2: Go to 'Storage / Backup & Restore' and check the 'Clusters' Tab. It should look similar to the screenshot below. The first entry shows the cluster name of the backup / restore hub (in this example it is 'apps.ocp.hub.ibm.lab'):




Figure: Red Hat OpenShift Backup and Restore clusters tab showing hub cluster

Step 3: Go to Networking / Routes. Select the project named ibm-spectrum-fusion-ns. Find the route named console and save the location entry to the clipboard or a text file (in this example it is '<https://console-ibm-spectrum-fusion-ns.apps.ocp.hub.ibm.lab/>')




Figure: IBM DRS Connection Manager selection dialog

Step 3: Now the backup system source types which are supported by DRS are displayed. Select the source type 'IBM Fusion' and click 'Next':




Figure: IBM DRS backup system source type selection showing IBM Fusion

Step 4: In the connection details panel, enter the Fusion cluster URL which you have saved in [Finding the Fusion cluster URL](#)




Figure: IBM DRS Fusion cluster connection details

Step 5: Click 'Check certificate' to verify the certificate details and if the connection is available. When the certificate check is successful, click 'Next'.




Figure: IBM DRS Fusion cluster certificate verification

Step 6: For allowing access to the Fusion cluster, either a valid user/password combination or a valid API key is required.

At this step, you have 3 options:

- A. A credential entry for the Fusion cluster already exists, e.g. from an earlier setup: select the correct credential entry from the list and click 'Next'




Figure: IBM DRS Fusion cluster credential selection dialog

- B. You need to set up a new credential entry with a user / password combination: Click 'Create credential', select Access type 'Password' and provide the user and password which have been created in the Fusion cluster .

Create credential

You can use a credential to access multiple sources. [Learn more](#)

Name

credential for console-ibm-spectrum-fusion-ns.apps.ocp.hub...

Description

Describe the application (optional)

Access type

☒ Password
 ☐ API key

Username

Enter username

Password

Enter password

Cancel

Create

Figure: IBM DRS create credential dialog with password option

- C. You need to set up a new credential leveraging an API key: Click 'Create credential', select Access type 'API key' and provide the API key which has been created in the Fusion cluster.

Create credential

credential for console-ibm-spectrum-fusion-ns.apps.ocp.hub...

Describe the application (optional)

Access type

☐ Password
☒ API key

API key

Enter API key

Cancel

Create

Figure: IBM DRS create credential dialog with API key option

Step 7: You will now see a new line in the list of Backup Systems showing your new Fusion cluster. The 'Last scan' column may first show an 'installing agent' message. This is a connection agent that is automatically deployed to the Connection Manager. Next, it will show 'Inventorying' while the Connection Manager is collecting the first set of metadata. When the scan is ready, it will show 'Success' and your DDP cluster is successfully integrated with DRS. In case a 'Failed' message is shown, you can click on the 'Failed' link and review the Job details for further analysis.

IBM Storage Defender

Data Resiliency / Connection managers /

CE4S-EMEA-ConMan1 Connected

Overview

Jobs

Sources

Sensor control nodes

Sensors

Search sources

Name

Type

Connection managers

Last inventory

Last inventory time

Duration

Credential

console-ibm-spectrum-fusion-ns.apps.ocp.hub...

IBM Fusion

CE4S-EMEA-ConMan1

Inventorying

Apr 28, 2026, 12:23 PM

13 seconds

Password for console-ibm-spectrum-fusion-ns.apps.ocp.hub...

dataprotec-robo...

Data Protect

CE4S-EMEA-ConMan1

Success

Apr 28, 2026, 12:23 PM

4 seconds

API for dataprotec-robo...

dataprotec-robo...

Data Protect

CE4S-EMEA-ConMan1

Success

Apr 28, 2026, 12:23 PM

10 seconds

API for dataprotec-robo...

fs5300-rm10...

IBM FlashSystem

CE4S-EMEA-ConMan1

Success

Apr 28, 2026, 12:23 PM

0 seconds

CE4S FlashSystem Grid at

fs5300-rm11...

IBM FlashSystem

CE4S-EMEA-ConMan1

Failed

Feb 6, 2026, 2:15 PM

0 seconds

CE4S FlashSystem Grid at

fs9110-cleanroom...

IBM FlashSystem

CE4S-EMEA-ConMan1

Failed

Feb 6, 2026, 3:15 PM

0 seconds

CE4S FlashSystem Grid at

protect1...

IBM Storage Protect

CE4S-EMEA-ConMan1

Success

Apr 28, 2026, 12:23 PM

2 seconds

Credentials for protect1...

protect2...

IBM Storage Protect

CE4S-EMEA-ConMan1

Success

Apr 28, 2026, 12:23 PM

2 seconds

Credentials for protect2...

Items per page: 10

1-8 of 8 items

1 1 of 1 page

Figure: IBM DRS Fusion cluster first scan image

Step 8: Now you can click on the name of your Fusion cluster and review its state regarding data resources, policies, backup statistics, and backup resilience posture:




Figure: IBM DRS Fusion cluster dashboard showing resources and statistics

Integration of IBM Defender Data Protect in DRS

Within the DRS context, IBM Defender Data Protect (IBM DDP) is treated as a backup system. The Connection Manager is used to establish the connection with DRS and collects relevant metadata.

Prerequisites

The following prerequisites are required for integrating IBM DDP with DRS:

- Access to the DRS tenant where the new backup system will be integrated.
- Access to the IBM DDP cluster that is planned to be integrated with DRS.
- Credentials that the Connection Manager will use to access the DDP cluster. The credential can either be a user / password combination, or an API key.
- A Connection Manager which is linked to the DRS tenant. The Connection Manager needs to have local network access to the IBM DDP cluster.

Integration steps for IBM Defender Data Protect

Perform the following steps for integrating IBM DDP in DRS:

Step 1: From the DRS dashboard, go to Backup systems/Backup systems and click 'Add backup system':




Figure: IBM DRS backup systems page with Add backup system button

Step 2: Next, a list of available Connection Managers will be displayed. Select the Connection Manager which has local network access to the IBM DDP cluster and click 'Next':




Figure: IBM DRS Connection Manager selection for Data Protect

Step 3: Now the backup system source types which are supported by DRS are displayed. Select the source type 'Data Protect' and click 'Next':




Figure: IBM DRS backup system source type showing Data Protect

Step 4: In the connection details panel, enter the IP address or fully qualified domain name (FQDN) of the IBM DDP cluster. Click 'Check certificate' to verify the certificate details and if the connection is available. When the certificate check is successful, click 'Next'.




Figure: IBM DRS Data Protect connection details and certificate verification

Step 5: For allowing access to the DDP cluster, either a valid user/password combination or a valid API key is required.

You have 3 options:

- A. A credential entry for the DDP cluster already exists, e.g. from an earlier setup: select the correct credential entry and click 'Next'.




Figure: IBM DRS Data Protect cluster credential selection

- B. You need to set up a new credential entry with a user / password combination: Click 'Create credential', select Access type 'Password' and provide the user and password which have been created in the DMS cluster




Figure: IBM DRS create credential with password for Data Protect

- C. You need to set up a new credential leveraging an API key: Click 'Create credential', select Access type 'API key' and provide the API key which has been created in the DMS cluster.

Step 6: You will now see a new line in the list of Backup Systems showing your new DDP cluster. The 'Last scan' column may first show an 'installing agent' message, next it will show 'Scanning' while the CM is collecting the first set of metadata. When the scan is ready, it will show 'Success' and your DDP cluster is successfully integrated with DRS. In case a 'Failed' message is shown, you can click on the 'Failed' link and review the Job details for further analysis.

Step 7: Now you can click on the name of your DDP cluster and review its state regarding data resources, policies, backup statistics, and backup resilience posture:




Figure: IBM DRS Data Protect cluster dashboard with statistics

Grouping and governance profiles in DRS

The information that is included in this section is intentionally kept brief. For more information about this topic, see the [IBM Storage Defender documentation](#) or see the IBM Redbooks publication [IBM Storage Defender: Data Resiliency Service \(REDP5744\)](#).

Defender DRS establishes a strong concept of groupings. The idea of grouping follows the way applications are managed and protected. A DRS application reflects the business-level orchestration, while the DRS recovery group shows data-level protection.

For example, an application can consist of multiple interconnected components, such as web servers, application servers, databases, and others. The different elements of an application may have different backup / restore mechanisms which are collected within a recovery group.

Applications

An application is a logical collection of recovery groups that jointly represent a complete application. It ensures that all associated recovery groups are restored in the correct order during the recovery process, helping to minimize downtime and reduce the risk of recovery failure.

The DRS application dashboard provides a centralized view of the application's health and status.




Figure: IBM DRS application dashboard showing health and status

Furthermore, in large enterprise ecosystems, many applications work together to support different business functions. These applications are often interdependent and complement each other to deliver essential services. To illustrate and manage these interdependencies, the Data Resiliency Service supports the concept of application dependency mapping.

To protect an application, each element and its resources need to be backed up in a consistent manner. The backup and restore processes may be different between the components, for example, while a database is often protected by establishing a consistent backup copy of the database itself and the database logs, for a simple web server a backup copy of its configuration may be sufficient.

Recovering a complete application often requires its elements to be restored in a specific order, for example, before a database itself can be restored, the container(s) running the database application needs to be recovered first.

Recovery groups

A recovery group is the atomic management entity within DRS. It is a logical collection of related resources (for example, virtual machines, storage volumes, or databases) that share identical protection requirements and must be recovered to the same point-in-time.

As part of the recovery group, a governance profile defines the key recovery expectations, for example, the RPO targets, for the data sources. In addition, a clean room profile can be defined as an alternate recovery location that can be used for recovery testing or data validation.

If all requirements set for a recovery group are met, the resources monitored by this group are recoverable, and the recovery group shows as 'Ready' in the Recovery Group overview. Otherwise, a 'Not compliant' status is shown if any of the requirements is violated.

Name	Status	Governance profile	Resources	Resource type	Workload type	Connection manager	Application group	Last updated	Labels
ce4s-core-shopping	⚠ Not compliant	ce4s-emea-pilot-...	3	VM	VMware	sts-dach-dcm-1	ce4s-shopping	Apr 30, 2026, 3:51 PM	User defined
ce4s-ActiveDirectory	✅ Ready		0	Placeholder	Microsoft Active Directory		ce4s-shopping	Apr 30, 2026, 3:21 PM	User defined
ce4s-web-server	✅ Ready		0	Placeholder	VMware		ce4s-shopping	Apr 30, 2026, 3:24 PM	User defined
ce4s-contentmgmt	✅ Ready		0	Placeholder	Oracle Database		ce4s-shopping	Apr 30, 2026, 3:25 PM	User defined
ce4s-webcontent	✅ Ready		0	Placeholder	MySQL		ce4s-shopping	Apr 30, 2026, 3:27 PM	User defined
ce4s-payments	✅ Ready		0	Placeholder	VMware		ce4s-shopping	Apr 30, 2026, 3:21 PM	User defined

Figure: IBM DRS recovery groups list showing Ready and Not compliant status

Resources

A recovery group knows two types of resources:

- Production systems: (for example, virtual machines, or database servers). By analyzing the metadata collected by the connection manager, the recovery group is automatically aware of available backups or snapshots from these systems.
- Storage systems: where the volumes hosted by the storage systems are not protected as part of a production system as described above. In this case, the recovery group will be aware of available volume backups or volume snapshots.

Governance profiles

The governance profile defines the recovery expectations for the recovery group:

- For Storage Systems: Snapshot frequency and retention days, in addition attributes like threat detection, encryption, HA, DR, and immutability can be evaluated.
- For Backup Systems: Backup frequency, number of copies, backup retention time. In addition, attributes like locations, and media types, air gap, and backup encryption can be evaluated.

Recovery Testing Frequency

If all expectations defined in the governance profile are met, the recovery group will be shown as 'Ready' in the dashboard. If any of the defined expectations is violated, for example, due to a failed backup, a forgotten recovery test, or incorrect implementation of retention times, the recovery group will be shown as 'non-compliant'.




Figure: IBM DRS recovery group details showing non-compliant status

Clean room profiles

A clean room is important for customers who want to ensure their data can be recovered without impact on their production. Instead of restoring data into a production environment where original production data will be overwritten, the clean room provides an alternate recovery target which can be used for recovery testing, malware scanning, and data validation.

DRS can orchestrate recovery testing and can perform recovery into a VMware-based clean room with a single mouse click. Upon successful data recovery the user is asked to perform data validation. Data validation must be handled manually, depending on the application and the customer's requirements. The basic options for data validation are:

- Starting a VM and checking whether it operates as expected.
- Performing a deep scan of the data, using, for example, IBM Storage Defender Sentinel, to ensure there is no corrupted data or other signs of malware activity.
- Leveraging application specific tooling, for example, a database integrity check, to ensure the application can operate on the data that has been recovered.

In addition, it provides the option to mark the result of validation testing within DRS. In the list of recovery points available, those which have undergone successful validation testing are marked appropriately. This allows us to select the best recovery point without additional delay, speeding up recovery after, for example, a cyber incident.

Dependency mapping in IBM Storage Defender DRS

In IBM Storage Defender DRS, dependency maps provide a sophisticated topological visualization of the interconnections between disparate infrastructure components, ranging from virtualized hosts to underlying storage volumes and network dependencies.

These maps are dynamically generated through the analysis of metadata and telemetry, leveraging heuristic algorithms and AI-driven auto-grouping to identify functional relationships that might otherwise remain opaque in complex hybrid environments. By surfacing these critical interdependencies, DRS enables administrators to proactively identify gaps in the resiliency posture, visualize the potential blast radius of a cyber-attack, and architect Phased Recovery plans that honor strict boot-order requirements. This automated discovery ensures that the orchestrated recovery of an application tier is logically synchronized with its respective database and infrastructure prerequisites, thereby minimizing Recovery Time Objectives (RTO) and eliminating the risk of service failure due to "orphan" or missing resources during a restoration event.




Figure: IBM DRS application dependency map visualization

Abbreviations

Abbreviation	Expansion
ACL	Access Control List
AI	Artificial Intelligence
API	Application Programming Interface
AWS	Amazon Web Services
B&R	Backup and Restore
CBD	Changed Block Detection
CLI	Command Line Interface
COS	Cloud Object Storage
CPU	Central Processing Unit
CR	Custom Resource
CRB	Cluster Role Binding
CRD	Custom Resource Definition
CSI	Container Storage Interface
DACH	Germany, Austria, and Switzerland (Deutsch, Austria, Confoederatio Helvetica)
DDP	Defender Data Protect
DMS	Data Management Service
DNS	Domain Name System
DR	Disaster Recovery
DRS	Data Resiliency Service
EMEA	Europe, Middle East, and Africa
fapp	Fusion Application (custom resource)
fbackup	Fusion Backup (custom resource)
fbp	Fusion Backup Policy (custom resource)

Abbreviation	Expansion
fbsl	Fusion Backup Storage Location (custom resource)
fpa	Fusion Policy Assignment (custom resource)
frcpe	Fusion Recipe (custom resource)
frestore	Fusion Restore (custom resource)
FQDN	Fully Qualified Domain Name
GDPR	General Data Protection Regulation
GUI	Graphical User Interface
HA	High Availability
HIPAA	Health Insurance Portability and Accountability Act
HTTPS	Hypertext Transfer Protocol Secure
HUB	Centralized IBM Storage Fusion Backup and Restore management cluster
I/O	Input/Output
IBM	International Business Machines
ISO	International Organization for Standardization
ITSO	International Technical Support Organization
JSON	JavaScript Object Notation
LDAP	Lightweight Directory Access Protocol
MCG	Multicloud Gateway
OCP	Red Hat OpenShift Container Platform
OCP-V	Red Hat OpenShift Virtualization
OADP	OpenShift API for Data Protection
OAuth	Open Authorization
PIT	Point In Time

Abbreviation	Expansion
PHP	PHP: Hypertext Preprocessor
PV	Persistent Volume
PVC	Persistent Volume Claim
QEMU	Quick Emulator
QRadar	IBM QRadar Security Information and Event Management platform
RBAC	Role-Based Access Control
RBD	RADOS Block Device
RPO	Recovery Point Objective
RHOS	Red Hat OpenShift (cluster)
RTO	Recovery Time Objective
S3	Simple Storage Service
SaaS	Software as a Service
SAN	Subject Alternative Name
SELinux	Security-Enhanced Linux
SIEM	Security Information and Event Management
SLA	Service Level Agreement
SME	Subject Matter Expert
SOAR	Security Orchestration, Automation, and Response
SOX	Sarbanes-Oxley Act
SPOKE	Remote IBM Storage Fusion Backup and Restore agent cluster
SQL	Structured Query Language
TCP	Transmission Control Protocol
TLS	Transport Layer Security

Abbreviation	Expansion
UI	User Interface
URL	Uniform Resource Locator
VM	Virtual Machine
VMS	VirtualMachineSnapshot (custom resource)
YAML	YAML Ain't Markup Language

Additional Resources

This document complements but does not replace IBM product documentation. For detailed API references, complete configuration options, and the latest updates, consult the following additional resources:

IBM Redbooks Publications:

- IBM Storage Defender Data Resiliency Service (REDP-5744): <https://www.redbooks.ibm.com/abstracts/redp5744.html>
- IBM Storage Defender: IBM Data Management Service and IBM Data Protect (SG24-8554): <https://www.redbooks.ibm.com/abstracts/sg248554.html>
- IBM Storage Fusion Backup and Restore for IBM Cloud Pak for Data (REDP-5706): <https://www.redbooks.ibm.com/abstracts/redp5706.html>

IBM Documentation:

- IBM Storage Fusion documentation: <https://www.ibm.com/docs/en/fusion-software>
- IBM Storage Fusion Github repository: <https://github.com/IBM/storage-fusion>
- IBM Storage Defender documentation: <https://www.ibm.com/docs/en/storage-defender>

Other related Documentation:

- Red Hat OpenShift documentation: <https://docs.openshift.com>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM® in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property

right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 United States of America

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Trademarks

The following terms are trademarks of the International Business Machines Corp.

AI Attribution

This work was primarily human-created. AI was used to make stylistic edits, such as changes to structure, wording, and clarity. AI was used to make content edits, such as changes to scope, information, and ideas. AI was prompted for its contributions, or AI assistance was enabled. AI-generated content was reviewed and approved. The following model(s) or application(s) were used: IBM Bob, Microsoft CoPilot and Google Gemini.

Version History

Version 1.0 - Initial Publication (July 2026)

Document Number: MD260014

Title: IBM Storage Defender: Protecting OpenShift Containers and Virtual Machines

Authors:

- Phillip Gerrard
- Marcelo Capile
- Reginald D'Souza
- Maxwell Ferreira
- Gerd Franke

- Andre Gaschler
- Kavish Ghogale
- Pallavi Singh
- Prajakta Tendulkar
- Jörg Walter
- Daniel Wendler

Initial Release Content:

- Container Protection with IBM Storage Defender
- IBM Storage Fusion Backup and Recovery
- IBM Storage Defender Data Protect
- IBM Storage Defender DRS Integration

Publication Date: July 2026

Summary: Initial publication providing practical guidance for implementing data protection and cyber resilience for Red Hat OpenShift environments using IBM Storage Fusion and IBM Storage Defender. This document covers backup and restore operations for containerized applications and OpenShift Virtualization virtual machines, including container concepts, Hub and Spoke architecture, policy-based protection, application-consistent backups using recipes, and integration with the Data Resiliency Service (DRS) for orchestrated recovery workflows across both containerized applications and OpenShift Virtualization virtual machines.