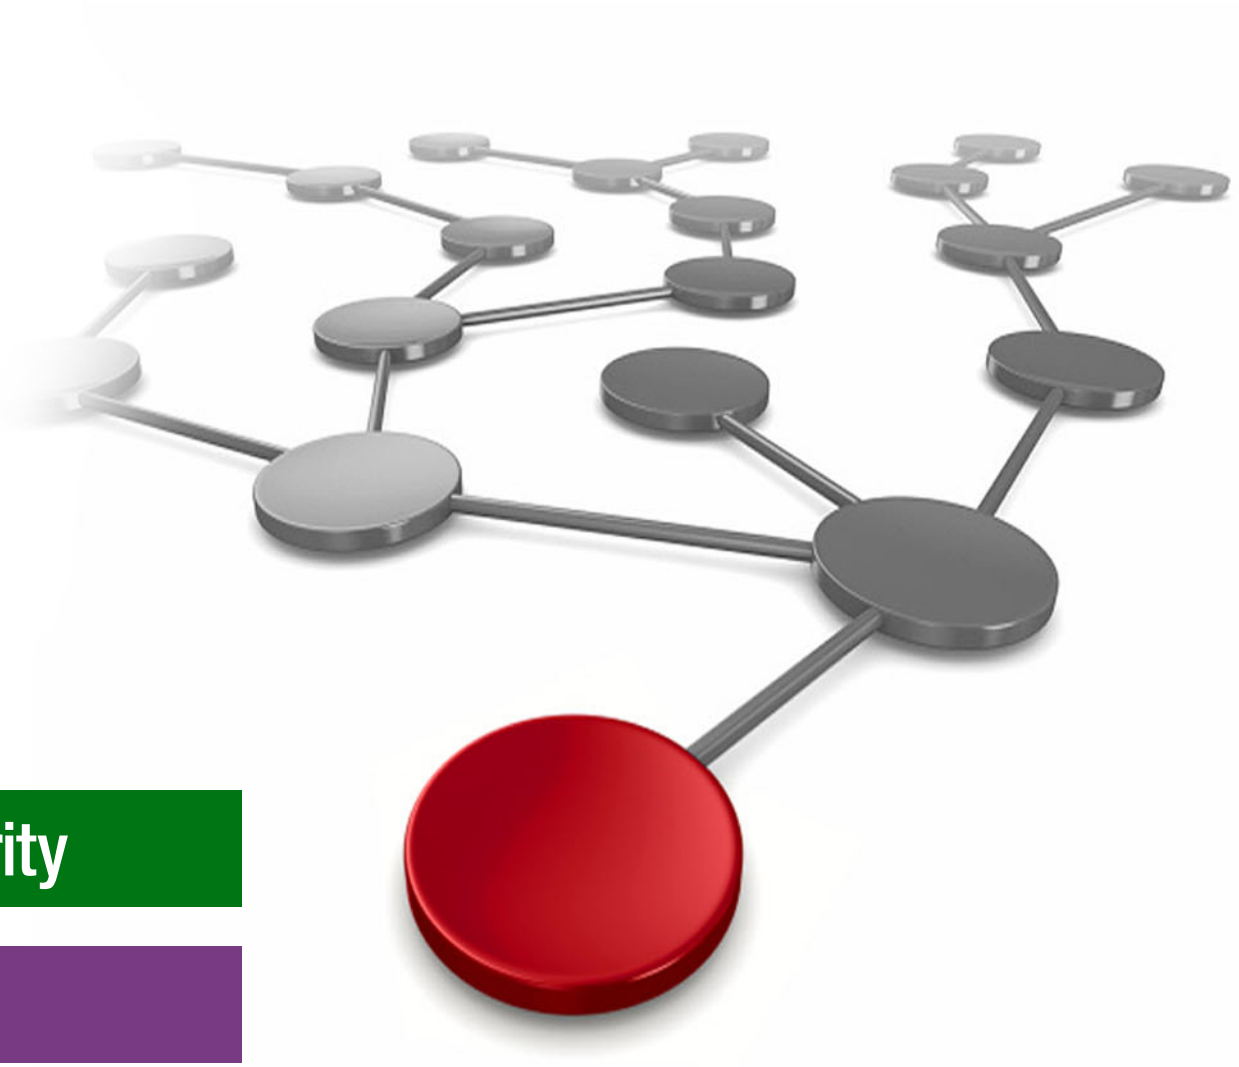
Redbooks
ibm.com/redbooks

IBM Storage Defender: IBM Data Management Service and IBM Data Protect

Christian Burns
Paul Conway
Phillip Gerrard
Gary Graham
Richard Hurst
Juan Carlos Jimenez
James Morassutti
Jack Tedjai
Dan Thompson
Christopher Vollmar



Security

Storage

IBM

Redbooks



IBM Redbooks

IBM Storage Defender: IBM Data Management Service and IBM Data Protect

March 2025

SG24-8554-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (March 2025)

This edition applies to IBM Storage Defender Data Protect (DP) 7.1.1 and 7.1.2.

© Copyright International Business Machines Corporation 2025. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices v
Trademarks vi

Preface vii
Authors vii
Now you can become a published author, too! viii
Comments welcome ix
Stay connected to IBM Redbooks ix

Chapter 1. Introducing IBM Storage Defender 1
1.1 Overview of IBM Storage Defender 2
1.2 Overview of IBM Storage Defender Data Protect 3
1.3 Overview of IBM Storage Defender Data Management Service 4

Chapter 2. IBM Storage Defender: Data Management Service 5
2.1 Data Management Service overview 6
2.2 IBM Storage Defender DMS dashboard 6
2.2.1 Security 8
2.2.2 System 11
2.2.3 Reporting 14
2.2.4 System View 15

Chapter 3. IBM Storage Defender Data Protect 17
3.1 Solution overview 18
3.1.1 Node-based scale-out architecture 18
3.2 What is a cluster and how does it work 18
3.3 Physical, virtual, and ROBO implementations 19
3.4 Cluster setup 20
3.4.1 Configuring an IBM Storage Defender DP Storage Domain 21
3.4.2 Configuring an IBM Storage Defender DP Protection Group 25
3.4.3 Monitoring the IBM Storage Defender DP Protection Group 27
3.5 IBM Storage Ready Nodes 28
3.6 IBM Storage Defender Ready Nodes 29
3.7 Registering with IBM Data Management Service 30
3.8 VMware workload use case 31

Chapter 4. Protecting VMware workloads with IBM Data Protect 33
4.1 IBM Storage Defender DP overview 34
4.2 Configuring IBM Storage Defender DP sources 34
4.2.1 Configuring an IBM Storage Defender DP policy 35
4.3 Recovering data with IBM Storage Defender DP 37
4.3.1 Recovering from a IBM Storage Defender DP cluster 37
4.3.2 Instant Volume Mount recovery from an IBM Storage Defender DP cluster 41
4.4 Replicating workloads with IBM Storage Defender DP clusters 44
4.4.1 VMware replication workflow one-to-one model 44
4.5 Monitoring the Protection job list 48
4.6 Recovering from replicated copies 49
4.6.1 Preparing a failover back to the original source cluster 51
4.7 Archiving workloads with CloudArchive and IBM Storage Defender DP cluster 53

© Copyright IBM Corp. 2025.

4.7.1 VMware CloudArchive workflow S3 to IBM Cloud Object Storage 54
4.7.2 VMware CloudArchive restore from IBM S3 IBM Cloud Object Storage. 57

Chapter 5. Rapid restore of virtual machines. 61
5.1 Virtual machine recovery. 62
5.1.1 Recovering VMs 62
5.1.2 Instant Mass Restore 68
5.2 File or folder recovery 69
5.2.1 Indexing and recovery. 70
5.2.2 Recovering files or folders. 70
5.3 Virtual disk recovery 77
5.3.1 Recovering virtual disks 78
5.4 Instant Volume Mount 82
5.4.1 Using Instant Volume Mount to mount a VMware virtual disk. 83
5.4.2 Tearing down a mounted volume 85

Chapter 6. Integrating IBM Storage Defender DP with IBM Storage Protect. 87
6.1 IBM Storage Defender DP overview 88
6.2 Deployment exercise overview 89
6.3 Creating an S3 bucket on an IBM Storage Protect server 90
6.3.1 How to set up an IBM Storage Protect server to receive IBM Storage Defender DP archives 90
6.4 Capturing access credentials on an IBM Storage Protect server 93
6.5 Registering the S3 bucket as an external target on a DP cluster 93
6.6 Creating policies to archive data to the S3 bucket 96
6.6.1 Creating a Data Protection policy for archiving to the S3 bucket 96
6.6.2 Configuring the Data Protection Group to use the Data Protection policy 98
6.7 Monitoring the Data Protection status 102
6.8 Retrieving data from S3 external storage 104
6.8.1 Retrieving data by using Global Search 104
6.8.2 Retrieving data by using the Recoveries window 106
6.9 Additional process monitoring examples. 108

Abbreviations and acronyms 111

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, US

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.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© Copyright IBM Corp. 2025.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at https://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM®	IBM Spectrum®
Db2®	IBM Cloud®	QRadar®
Enterprise Design Thinking®	IBM FlashSystem®	Redbooks®
Guardium®	IBM Security®	Redbooks (logo)  ®

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat, Ceph, OpenShift, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

VMware, VMware vSphere, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM Redbooks® publication provides an overview of IBM Storage Defender Data Management Service (DMS) and IBM Storage Defender Data Protect (DP), which includes information to better understand how IBM Storage Defender protects many different workloads and integrates with IBM Storage Protect for offload and archival purposes. Configuration and usage examples are provided to further explore defining and configuring protection policies.

This document is intended for use by system administrators and anyone wanting to learn more about implementing IBM Storage Defender DMS and IBM Storage Defender DP.

Authors

This book was produced by a team of specialists from around the world working with IBM Redbooks.

Christian Burns is a Principal Worldwide Storage Data Resiliency Architect and IBM Redbooks Platinum Author who is based in New Jersey. As a member of the Worldwide Storage Technical Sales Team at IBM®, he works with clients, IBM Business Partners, and IBM employees around the globe, designing and implementing solutions that address the rapidly evolving cyber resiliency and data resiliency challenges facing enterprises today. He has decades of industry experience in the areas of sales engineering, solution design, and software development. Christian holds a BA degree in Physics and Computer Science from Rutgers College.

Paul Conway has over 36 years of industry experience in Storage Management and Data Protection. As an IBM customer, Paul protected large Mainframe and Open Systems Storage environments by using IBM data protection software, including IBM Data Facility Storage Management Subsystem (DFSMS), Data Facility Data Set Services (DFDSS), and IBM Storage Protect. Paul has been with IBM 4 years. He is responsible for working with and guiding IBM Storage Protect customers as an IBM Storage Protect and IBM Storage Defender Technical Advisor.

Phillip Gerrard is a Project Leader for the International Technical Support Organization who is based in Beaverton, Oregon. As part of IBM for over 15 years, he has authored and contributed to hundreds of technical documents to IBM.com and worked directly with IBM's largest customers to resolve critical situations. As a team lead and SME for the IBM Spectrum® Protect support team, he is experienced in leading and growing international teams of talented IBM employees, developing and implementing team processes, and creating and delivering education. Phillip holds a degree in computer science and business administration from Oregon State University.

Gary Graham is a Brand Technical Specialist for IBM Storage Defender Data Resiliency Service solutions for IBM customers in the Southeastern US. Gary has presented at the IBM Edge conference and worked on multiple IBM Professional Certification exam teams. He has a Distinguished Technical Specialist certification from The Open Group.

© Copyright IBM Corp. 2025.

Richard Hurst is a cyber resiliency and storage consultant with the IBM Expert Labs team. With 20+ years experience across multiple disciplines, Richard has provided installation and configuration services and support for many different name brand products. Richard has worked with multiple IBM products, including IBM Storage Defender DP, IBM Storage Ceph, and IBM FlashSystem® servers. He consults for the IBM Cybervault Workshop. He continues to broaden his skill set in Red Hat OpenShift and IBM Fusion, and cybersecurity.

Juan Carlos Jimenez is the Worldwide Data Resiliency Product Manager who is based in Dallas, Texas. He is focused on defining roadmaps, initiatives, and strategies within the various data resiliency software products that he manages. Juan Carlos brings an end-to-end view to cyber resilience, and leverages his expertise in both storage and security. Juan Carlos developed the IBM Cyber Resiliency Assessment Tool, which has been helping numerous enterprises identify and close gaps in their IT environments. He holds a Management Information Systems degree from the University of Arizona.

James Morassutti is a Senior Storage Technical Specialist who is based in Toronto, Canada. His career in IT spans over 20 years in key areas like x86, networking, cyber resiliency, and storage solutions. James is the National Data Protection subject matter expert for Canada. He focuses on helping clients design solutions to support their operational resiliency and protect their critical data in support of their cyber resiliency practices. James is educated in Mechanical Engineering.

Jack Tedjai is an IBM Certified Expert IT Specialist and IBM Systems subject matter expert who work at the Northern Europe Infrastructure Lab Expert Services organization. He joined IBM in 1998, and has more than 25 years of experience in storage, storage virtualization, backup, and cyber resilience services for Open Systems. He is mostly involved in architecture and deployments worldwide for IBM Lab Expert Services, with a focus on IBM Storage Protect, IBM Storage Protect Plus, and IBM Cloud® Object Storage.

Daniel Thompson has been working in IT for more than 40 years. His specialty is data protection (backup and restore, disaster recovery, business continuity, and cyber resiliency). He works at the Advanced Technology Group (ATG), IBM Technology, Americas.

Christopher Vollmar is the Principal Worldwide Storage Data Resiliency Architect. Christopher is an IBM Certified IT Specialist (Level 3 Thought Leader) and Storage Architect. He is focused on helping customers design solutions to support operational and cyber resiliency on primary and backup data to complement their cybersecurity practices. He is an author of several IBM Redbooks publications, an IBM Enterprise Design Thinking® Co-Creator, and a frequent speaker at events like IBM THINK, and TechXchange.

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on LinkedIn:

  https://www.linkedin.com/groups/2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/subscribe

- Stay current on recent Redbooks publications with RSS Feeds:

  https://www.redbooks.ibm.com/rss.html

1

Introducing IBM Storage Defender

A few decades ago, considerations for data resilience were a simpler. If a company lost or damaged an important file or folder, they loaded the previous day's backup tape, retrieved a copy of the missing data, and returned to operating normally.

Today, the volume of data and diverse range of workloads make backup and restore operations more complex. Regardless of their size, industry, or location, every organization must have an active security perimeter to keep out bad actors, plus effective recovery mechanisms to get back up and running quickly when an attack gets through.

Although the current world of IT might seem like a dangerous place with new and creative attempts to exploit vulnerabilities, careful planning and execution of the appropriate data security and data resilience processes can help organizations gracefully recover from otherwise dire situations.

This IBM Redbooks publication provides guidance about one of IBM's solutions that is dedicated to these use cases so that customers can recover rapidly and at scale.

This chapter describes the following topics:

- 1.1, "Overview of IBM Storage Defender" on page 2
- 1.2, "Overview of IBM Storage Defender Data Protect" on page 3
- 1.3, "Overview of IBM Storage Defender Data Management Service" on page 4

© Copyright IBM Corp. 2025.

1.1 Overview of IBM Storage Defender

Organizations need a data resilience strategy that encompasses every aspect of their on-premises and cloud environments. Such a strategy should support all traditional, hybrid cloud, virtualized, and containerized workloads. IBM Storage Defender software is designed to meet that need by offering end-to-end data resilience in modern hybrid multi-cloud IT environments that include virtual machines (VMs), databases, applications, file systems, Software as a Service (SaaS) workloads, and containers.

IBM Storage Defender features a combination of scalability, multiple layers of cyber resilience, broad application support, and cost-saving data reduction technologies by using service-level agreement (SLA) based policies to automate the entire data protection process, including backup, replication, and secure data retention on-premises and in the cloud, across primary and backup storage. Cyber resilience is enhanced by key features like immutability, encryption, and support for logical air gap to object storage (WORM technology), and the ability to physically air gap data to tape.

Here are some of the key capabilities of IBM Storage Defender:

- Data resilience and compliance

  You can set policies and standards to help ensure resilience and compliance across the data estate. You govern your protections, threat response, and recovery from one auto-updating SaaS console.

- Early threat detection

  IBM Storage Defender is designed to detect threats and anomalies from backup metadata, array snapshots, and other relevant threat indicators by leveraging AI-infused technology. It includes a data resiliency service that enhances existing security systems by including storage-specific malware and anomaly detection, and providing a trust index to help IT leaders decide where to prioritize the allocation of resources.

- Safe and fast recovery at scale

  You can validate, recover, and restore data more quickly. You can restore data at scale from an immutable backup or snapshot for each workload quickly (for example, 1000s of VMs in minutes).

- Flexible licensing

  Licensing is based on resource units, providing a cloud-like, utility-based consumption model for organizations to consume any service within IBM Storage Defender.

IBM Storage Defender is designed to integrate with other IBM Storage and IBM Security® solutions, including IBM QRadar®, IBM Guardium®, IBM FlashSystem, IBM Storage Scale, IBM Storage Ceph, and IBM Fusion. It also includes copy data management tools to manage and orchestrate application-integrated, hardware snapshots by making copies available when and where users need them for instant data recovery or data reuse, automatically cataloging and managing copy data across hybrid cloud infrastructures.

IBM Storage Defender is composed of various components that are designed to meet customers resilience needs. This book takes a deep dive into two of those components: IBM Storage Defender Data Protect (DP) and IBM Storage Defender Data Management Service.

1.2 Overview of IBM Storage Defender Data Protect

IBM Storage Defender DP offers data management and resiliency for the broadest workload support in the industry.

IBM Storage Defender DP supports the following workloads:

- Hypervisors: VMware, M. Hyper V, Nutanix AHV, and Oracle VM (OVM)
- Databases: Oracle, OVM, Oracle RAC, SAP HANA, SAP Oracle, SAP Db2, SAP MS SQL, SAP Sybase ASE, Sybase IQ and ASE, IBM Db2®, MS SQL, Hadoop, IRIS, and Cache (EPIC)
- Modern databases: Cassandra, CouchbaseDB, MySQL, Hbase, MongoDB, PostgreSQL, and Hive
- Cloud-native databases: CockroachDB, AWS RDS, and AWS Aurora
- Cloud applications: AWS virtual machine (VM) (EC2), M365, Exchange Online, Azure VM, and Google Compute
- Physical: Windows, IBM AIX®, Linux, and Solaris
- Containers: Kubernetes, and Tanzu
- File systems: NetApp, IBM Storage Scale, Google EFS, Elastifile, and IBM Pure Storage FlashArray
- Applications: Exchange, Microsoft Active Directory, and Microsoft SharePoint

This book takes a deep dive into how this solution protects the most critical workloads for modern enterprises.

IBM Storage Defender DP has a scale-out architecture that is composed of clusters. These clusters can be deployed virtually, in the cloud, or on-premises through physical nodes. These physical nodes include CPU, memory, storage, a network, the operating system, a file system, and the backup software. An example of these nodes is the IBM Storage Defender Ready Node. By leveraging this cluster and node architecture, IBM Storage Defender DP can run data management operations like backups, cloning, and restores rapidly at scale. You can do upgrades and expansions non-disruptively by adding more nodes to a cluster.

Here are some of the key capabilities of IBM Storage Defender DP that are covered in this document:

- Integrated cybersecurity

  The solution is designed on zero-trust principles to prevent internal attacks and threats. It has ransomware, virus, and vulnerability detection that are built in and can protect data through its immutable architecture, and protect data on immutable targets. Encryption is available both at-rest and in-transit, and integration with SIEM solutions like QRadar, Splunk, and others.

- Instant Mass Restore (IMR)

  By using IMR, you can restore many VMs instantly, for example, 200 VMs in about 10 minutes and 2000 VMs in about 40 minutes. IMR reduces downtime after an incident like a large-scale malware attack or ransomware attack.

- Global actionable search

  Search any data (file, VM, objects, or others) across multiple workloads and across all nodes in a cluster.

- Fast cloning

  Fast cloning of large databases for DevOps, testing, and other development use cases. For example, you can clone a 2 TB SAP HANA database in about 15 seconds.

- Global space efficiency

  The solution offers industry-leading global space efficiency technology through variable length deduplication, compression, and erasure coding (EC). This solution reduces capacity requirements and lowers licensing costs.

- Primary storage integration

  IBM Storage Defender DP integrates with the IBM Storage FlashSystem family to back up volumes to IBM Storage Defender DP through volume snapshots. You can recover by using either on-array snaps or offloaded volume backups. It also coordinates hardware snaps for VM backup to minimize VM stuns.

1.3 Overview of IBM Storage Defender Data Management Service

The IBM Storage Defender Data Management Service (DMS) is the operations center of IBM Storage Defender DP. This SaaS-based GUI enables users to create protection policies, trigger data backups, run fast restores, and complete many other operations.

Users can connect IBM Storage Defender DP clusters, IBM Storage Protect servers, IBM FlashSystem systems, and other assets to the service to drive end-to-end data resiliency operations from a single interface.

Here are some other IBM Storage Defender DMS capabilities:

- Quorum

  This function limits certain actions, including destructive actions, by single users. When user 1 requests an action, a second or third user must approve the request before it runs. You can consider this function as a form of Two-Person Integrity (TPI) checking. This function prevents some destructive attacks and reduces the impact of user errors or intentional damage.

- Security advisor

  The security advisor enables users to view the security posture of your implementation and provides actionable insights so that you can modify the security settings based on best practices and business needs.

- Simulations

  This function offers predictive planning models that can make projections about utilization and storage consumption. This capability is based on historical usage, workloads, and user-defined, what-if scenarios. It empowers users to proactively plan for various situations, such as acquiring new nodes, integrating new workloads, optimizing workloads, and more. You can create simulations with scenarios by using specific clusters and periods to better understand and plan environment changes.

- Reports

  With this function, users can create and view an overall summary of data protection jobs and storage systems. Also, users can analyze data at the granular level by using powerful filtering options. Filter, schedule, email, and download reports to help ensure that users who need detailed information on the environment and its status receive it.

2

IBM Storage Defender: Data Management Service

This chapter contains an overview of the IBM Storage Defender Data Management Service (DMS) and its key concepts.

This chapter describes the following topics:

- 2.1, “Data Management Service overview” on page 6
- 2.2, “IBM Storage Defender DMS dashboard” on page 6

© Copyright IBM Corp. 2025.

2.1 Data Management Service overview

IBM Storage Defender DMS is a Software as a Service (SaaS)-based management platform that provides a single-view, global management interface for all IBM Storage Defender Data Protect (DP) clusters in your environment. Whether the cluster is on-premises, in the cloud, or in IBM Storage Defender Data Protect Virtual Edition, and regardless of cluster size, the IBM Storage Defender DMS dashboard quickly connects clusters to data management, and provides access from anywhere an internet connection is available after logging in with your IBMid credentials.

Data management provides you with the following features:

- Multi-cluster management: Actively manage all your clusters, including multi-cluster monitoring and reporting, from a single dashboard.
- Global actionable search: Search across clusters and act directly from the search results page. For example, search for all unprotected VMs and create jobs to protect them.
- SmartAssist: Provides IBM Storage Defender DMS recommendations based on capacity forecasting and disk failure prediction.
- Security tools: Detect threats and other anomalies across clusters by using the unified Alerts page. Changes that you make directly to the cluster while using data management might not appear in the data management dashboard for 15 minutes.

This chapter provides an overview of IBM Storage Defender Data Management Service.

2.2 IBM Storage Defender DMS dashboard

The IBM Storage Defender DMS dashboard summary provides a comprehensive view of the various aspects that are related to managed IBM Storage Defender DP clusters, which includes cluster health, protection status, and reporting (Figure 2-1).

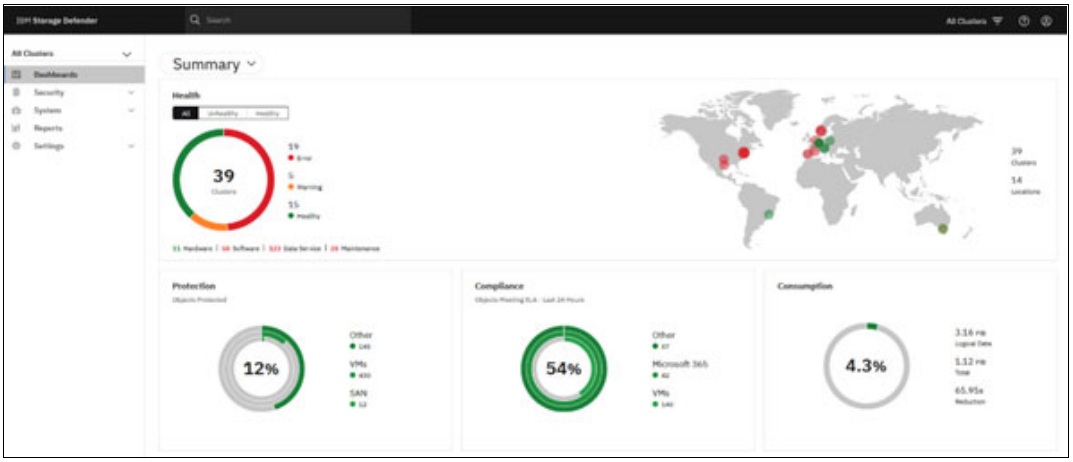


Figure 2-1 IBM Storage Defender DMS dashboard summary

In the summary dashboard, you can view the following key functions of IBM Storage Defender DMS:

- The total number of healthy and unhealthy clusters, an alert summary, and the location of the IBM Storage Defender DP clusters (Figure 2-2).

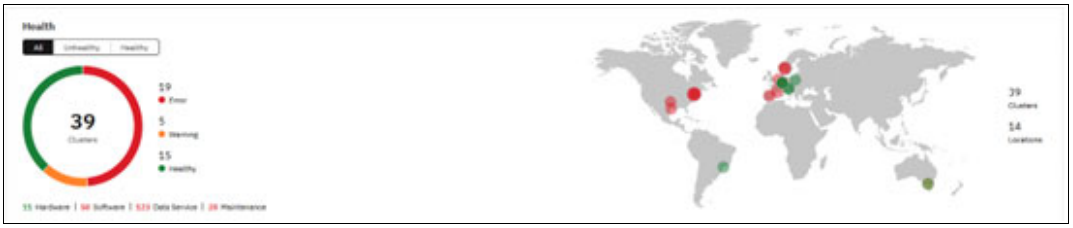


Figure 2-2 IBM Storage Defender DMS cluster health view

- A summary of total objects that are associated with a backup job (Figure 2-3).

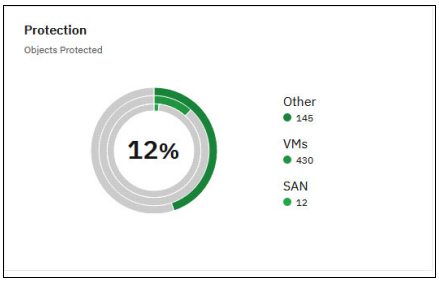


Figure 2-3 IBM Storage Defender DMS backup job protection summary

- Objects that met their SLAs in the last 24 hours (Figure 2-4).

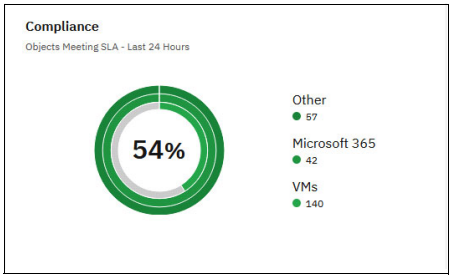


Figure 2-4 IBM Storage Defender DMS backup job status summary

- Storage consumption and data reduction ratios (Figure 2-5).

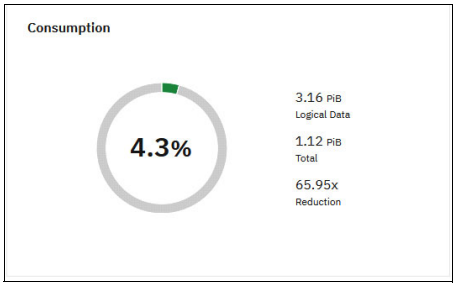


Figure 2-5 IBM Storage Defender DMS storage consumption summary

The Data Protection view provides a topology of the managed IBM Storage Defender DP clusters and their archive and cloud tier targets, and a summary of the recoveries that were made in the last 30 days (Figure 2-6).

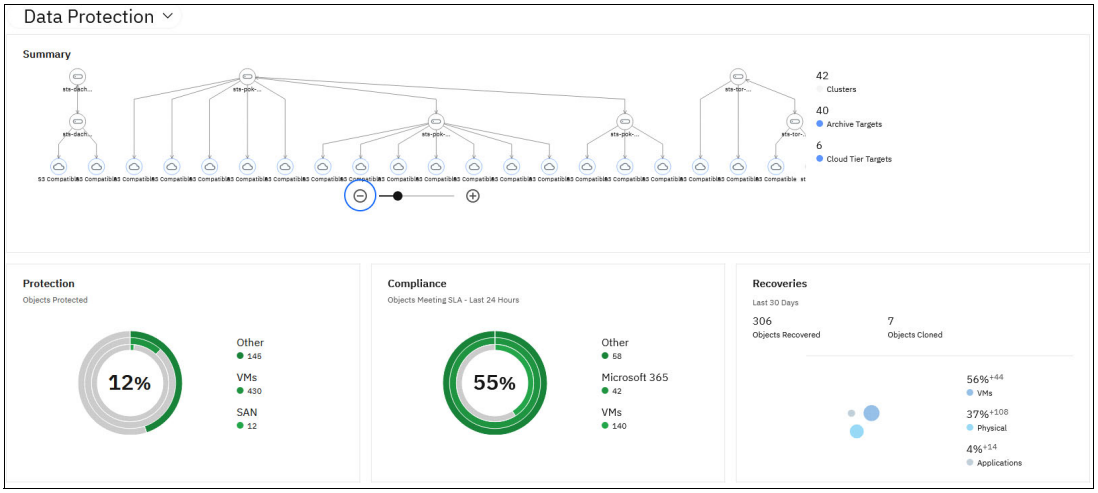


Figure 2-6 IBM Storage Defender DMS data protection topology overview window

The Cloud view provides a summary of cloud activity, including data archived or restored, and consumption by provider (Figure 2-7).



Figure 2-7 IBM Storage Defender DMS storage provider overview window

2.2.1 Security

The Security dashboard provides tools that administrators can use to quickly understand their current security stance. It also provides actionable insights with recommendations that are related to any potential vulnerabilities, which can help improve security in the environment.

The Security Advisor uses a scoring system, evaluating each individual cluster under management to help users understand their current security stance. It provides recommendations to improve scores based on several key attributes. The scores, icons, and other information that is provided in this feature indicate how the organization's deployment and configuration compares with IBM's minimum best practices. These practices are designed to supplement (but are not a substitute for) a robust and comprehensive information security program that is managed by an organization's designated experts.

Figure 2-8 shows the IBM Storage Defender DMS Security overview window.

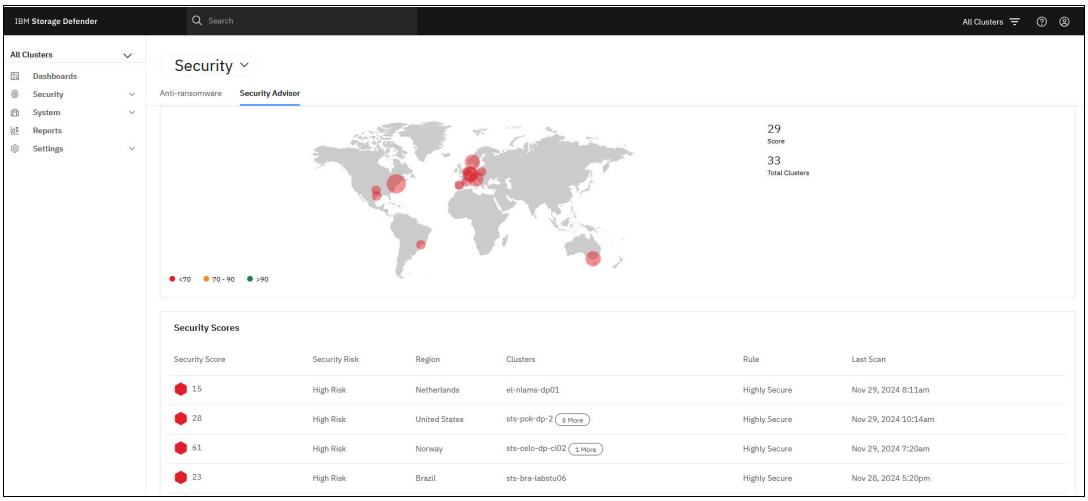


Figure 2-8 IBM Storage Defender DMS Security overview window

The Security view also summarizes threats that are detected across backed up objects and their attributes. This summary is available under the Anit-ransomware header (Figure 2-9).

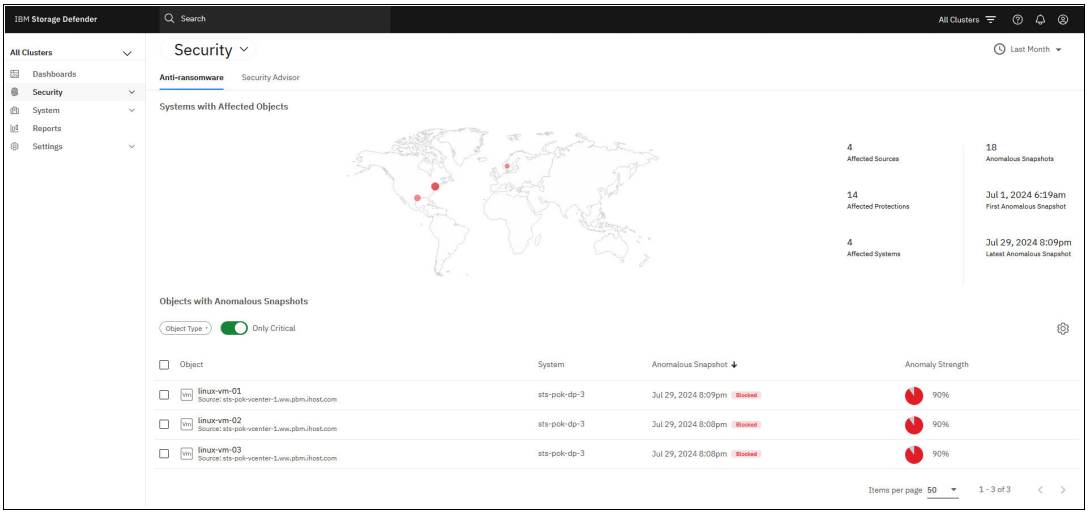


Figure 2-9 IBM Storage Defender DMS Security threat overview window

You can customize the action thresholds, manage notifications, and block recoveries that are tagged as a security risk (Figure 2-10).

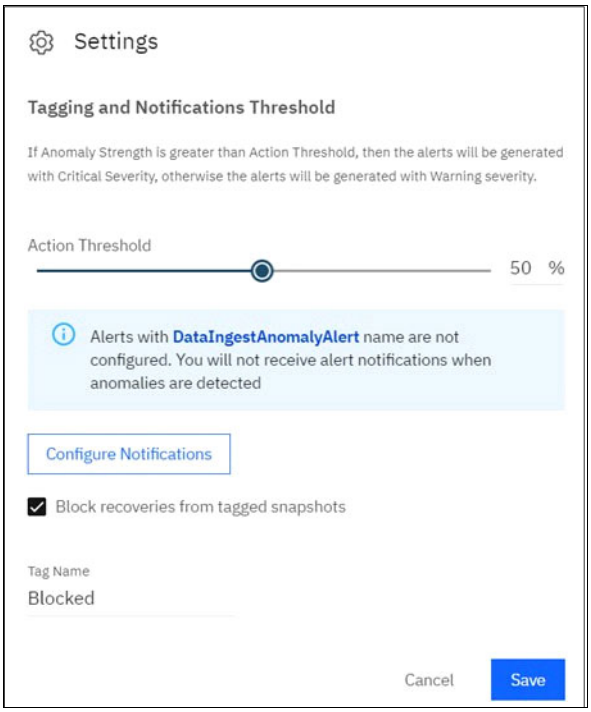


Figure 2-10 IBM Storage Defender DMS notification settings window

Quorum

In the left pane of the Security section of the IBM Storage Defender DMS dashboard is the Quorum function. Much like Two-Person Integrity (TPI), Quorum approvals is a feature of IBM Storage Defender DMS that helps ensure that a pre-defined approver, or group of approvers must approve of actions that are requested by users before those actions occur. Quorum helps eliminate the risk of destructive operations being performed by administrators due to malicious or accidental actions.

In the Quorum dashboard, you can view all quorum requests that are pending approval and current requests (Figure 2-11). Users with the appropriate level of access can also configure quorum groups. Quorum groups are made up of members of the organization that can be assigned as approvers for various operations against selected clusters.

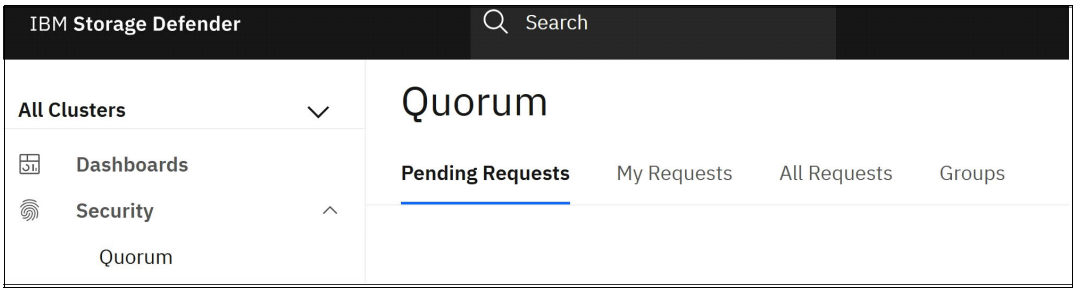


Figure 2-11 Quorum dashboard

Creating a quorum group

By leveraging the Quorum function, you can create a Quorum Group to restrict many pre-designated functions (Operations), which are governed by the group (Figure 2-12). Administrators may select the users that must approve the actions and how many of the group apply, for example, 3 out of 5, or 2 of 4, and so on. You can select the timeline for approval. The request has a certain amount of time to be approved or it is declined automatically. Members of the Quorum Group are notified by email of a pending approval.



Figure 2-12 Create Quorum Group window

2.2.2 System

The **System** tab provides access to the various tools that you use to manage the IBM Storage Defender DP deployment. The **System** tab has several options, which are described in the following sections.

Health Panel

The Health Panel (Figure 2-13) provides administrators with a consolidated list of alerts that are generated by the various clusters in the IBM Storage Defender DP deployment.

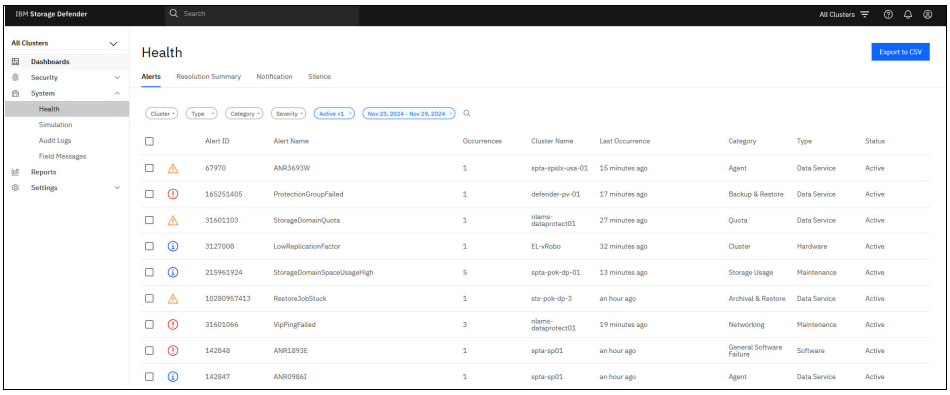


Figure 2-13 Health Panel showing a list of alerts

Alerts

Administrators can select an alert notification to gather more details (Figure 2-14), including the time that the alert was triggered, a detailed description of the severity, the type, and the category of the alert. You can optionally mark an alert as resolved by either tying it to an existing resolution or by creating a resolution description.

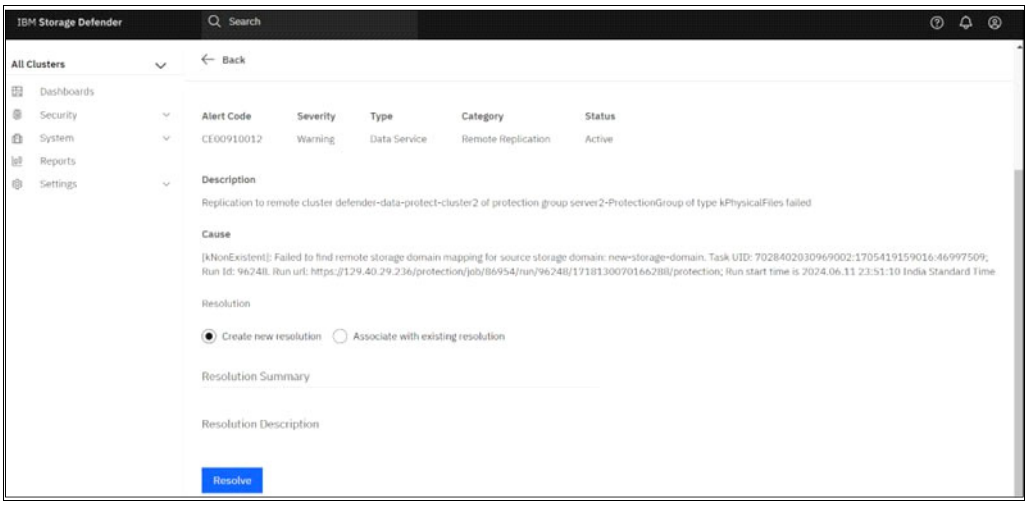


Figure 2-14 Alert details

The Alerts window has the following tabs:

- The **Resolution Summary** tab provides administrators with a consolidated view of all the previously generated resolutions.
- The **Silence** tab enables users to create rules to suppress certain notifications based on the cluster, severity of the alert, category, type, and names.
- The **Notification** tab enables users to customize notification delivery through email or webhooks, which enables customization of alerts and the type of alerts (Figure 2-15 on page 13).

Users may create alerts that are specific to clusters, alert levels, preset alert types, alert categories, and alert names with deep granularity.

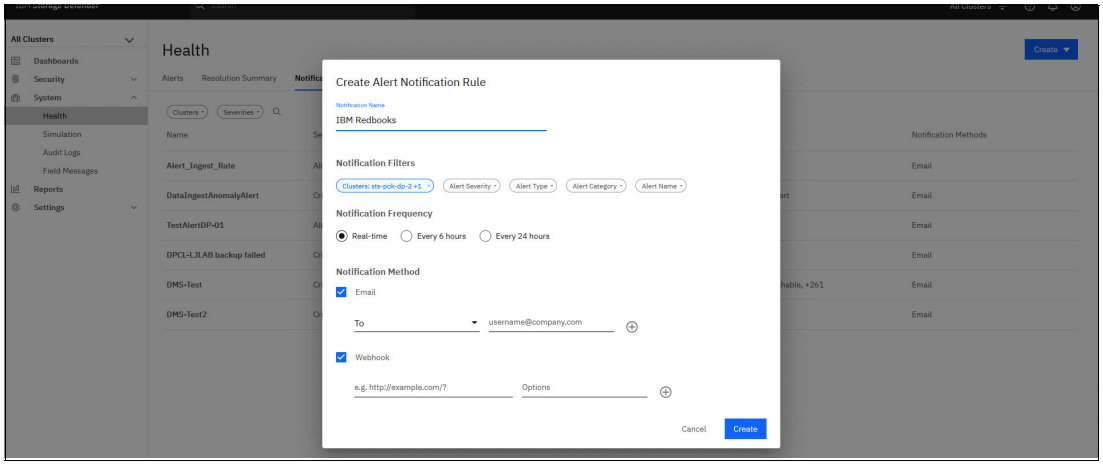


Figure 2-15 Create Alert Notification Rule dialog

Simulation

The **Simulation** tab (Figure 2-16) enables administrators to simulate new workloads against their existing deployment to help predict storage consumption over time. This powerful tool can help administrators proactively plan if and when additional storage capacity might be required to support environment changes.

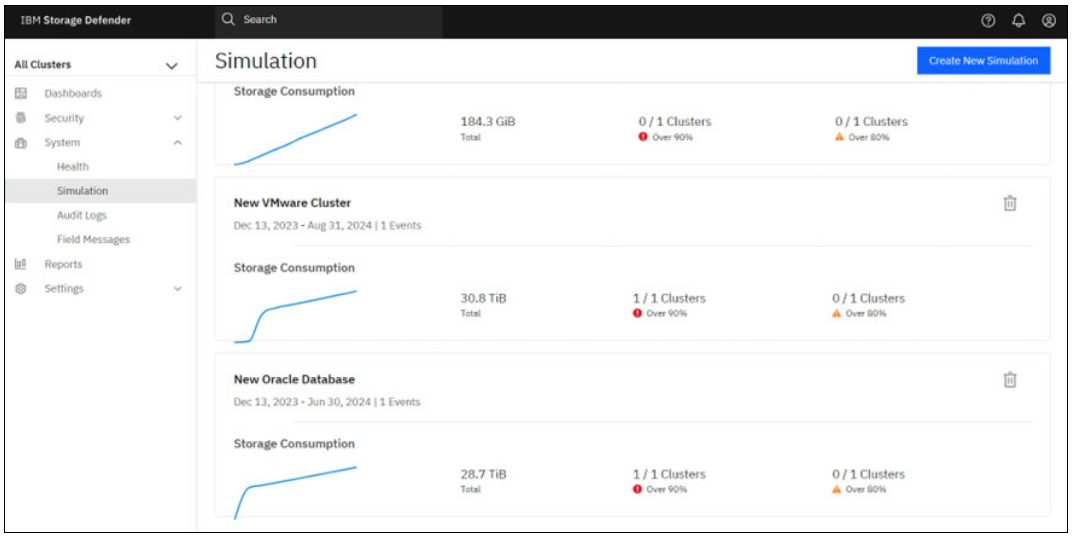


Figure 2-16 Simulation details

To create a simulation, select the blue toggle at the upper right of the window. Then, you can give the simulation a name and an end date for the simulation. The simulation can model the effect of many factors on the IBM Storage Defender DP cluster, such as adding resources, adding new workloads (protection groups), or expanding the protection groups or storage domains.

2.2.3 Reporting

The built-in reporting tool enables users to schedule or create ad hoc reports organization-wide or only for specific clusters or workloads. Reports can be filtered, scheduled, emailed, or downloaded (Figure 2-17).

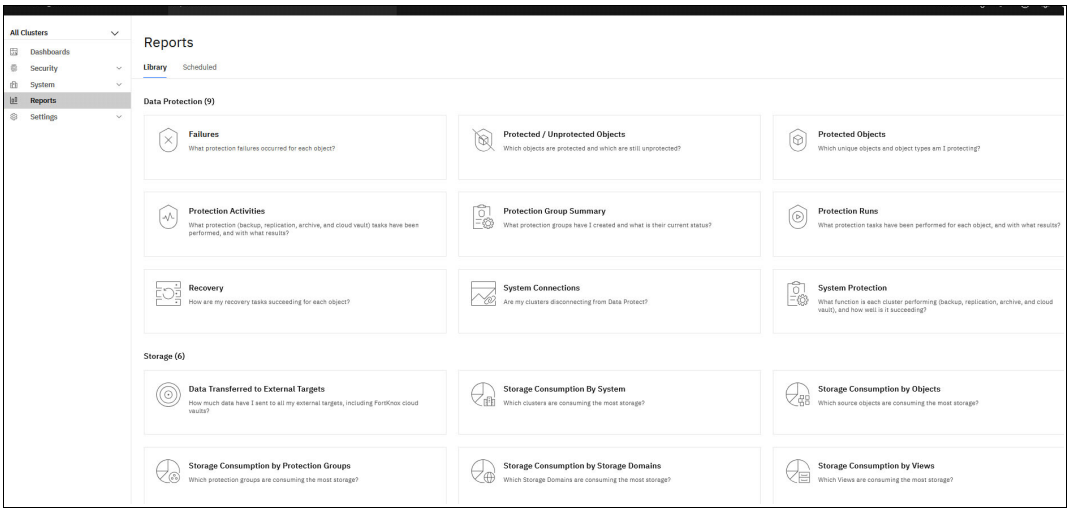


Figure 2-17 IBM Storage Defender DMS Reports

IBM Storage Defender DMS provides a comprehensive list of built-in reports to help users gather detailed information about the environment (Figure 2-18 on page 15). At the time of writing, IBM Storage Defender DMS provides the following built-in reports:

- Data Protection:
  - Failures
  - Protected/Unprotected Objects
  - Protected Objects
  - Protection Activities
  - Protection Group Summary
  - Protection Runs
  - Recovery
  - System Connection
  - System Protection
- Storage:
  - Data transferred to external targets
  - Storage Consumption by System
  - Storage Consumption by Objects
  - Storage Consumption by Protection Groups
  - Storage Consumption by Storage Domains
  - Storage Consumption by Views

These reports provide a comprehensive view of the specific category for the environment. You can further customize the results be selecting from several filtered options that are available within each report.

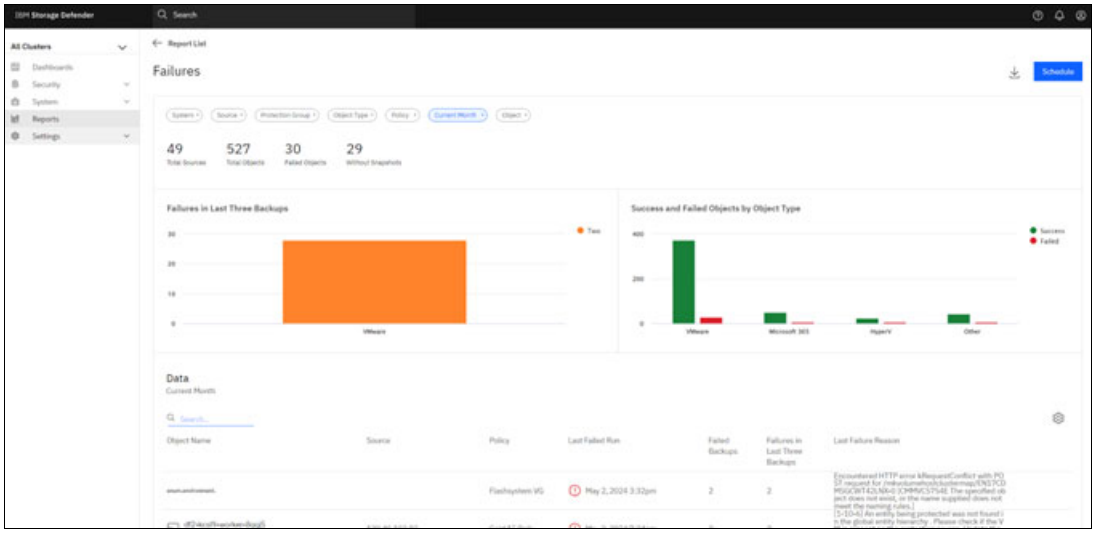


Figure 2-18 IBM Storage Defender DMS report example

Reports can be downloaded locally by selecting the download icon and choosing the file format, or you can schedule a delivery in several formats (Figure 2-19).

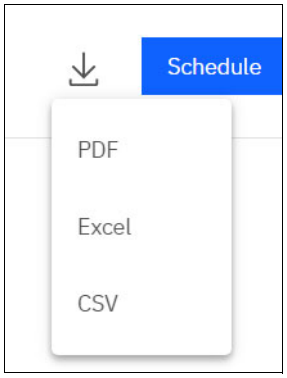


Figure 2-19 Downloading report formats

2.2.4 System View

The IBM Storage Defender DMS consolidates system information of all the managed IBM Storage Defender DP clusters to provide a single view of the health of the clusters organization-wide. The IBM Storage Defender DMS dashboard provides the following information:

- Health information of all clusters, which includes healthy clusters that have warnings or errors.
- A protection view, which reports on the total number of objects that are under protection.
- A compliance view, which reports on the total number of objects that met their SLAs over the last 24 hours.
- A consumption view, which consolidates the total storage that was consumed and the available capacity and overall data reduction ratios.

Figure 2-20 shows the IBM Storage Defender DMS System View window.

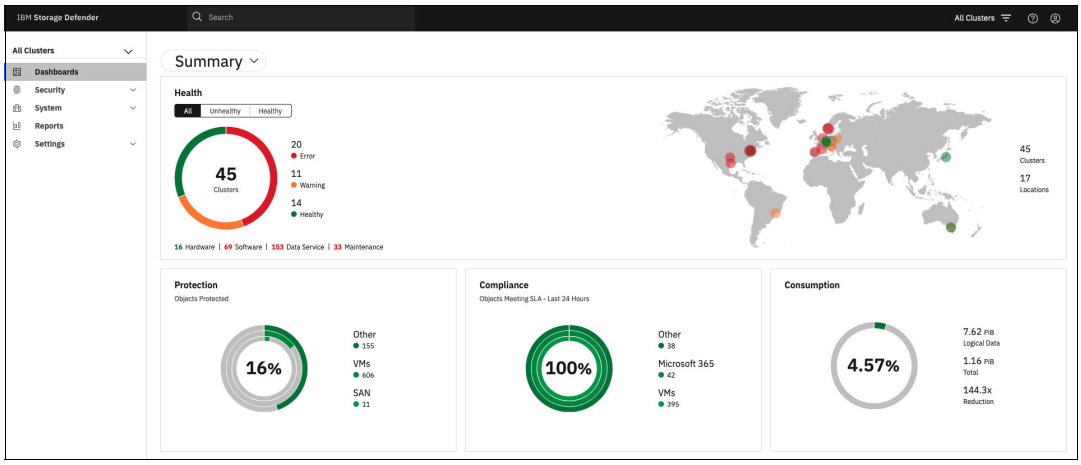


Figure 2-20 IBM Storage Defender DMS dashboard: System View

3

IBM Storage Defender Data Protect

The IBM Storage Defender Data Protect (DP) solution extends the data protection and cyber resiliency capabilities of the IBM Storage Portfolio by adding new architecture options for data resiliency by using clustered, scale-out data protection services that can also provide rapid recovery capabilities. These IBM Storage Defender DP clusters are managed globally by the IBM Storage Defender Data Management Service (DMS), which is provided as a Software as a Service (SaaS) subscription. The new capabilities in IBM Storage Defender DP are licensed by using the new Resource Unit license model.

This chapter introduces the new architecture and key components, and then explores several deployment options to build a IBM Storage Defender DP and replica solution to service Instant Mass Restore (IMR) capabilities for hundreds of VMs at once.

This chapter describes the following topics:

- 3.1, “Solution overview” on page 18
- 3.2, “What is a cluster and how does it work” on page 18
- 3.3, “Physical, virtual, and ROBO implementations” on page 19
- 3.4, “Cluster setup” on page 20
- 3.5, “IBM Storage Ready Nodes” on page 28
- 3.6, “IBM Storage Defender Ready Nodes” on page 29
- 3.7, “Registering with IBM Data Management Service” on page 30
- 3.8, “VMware workload use case” on page 31

© Copyright IBM Corp. 2025.

3.1 Solution overview

Figure 3-1 shows a IBM Storage Defender DP environment with a management plane that is called the Data Management Service (DMS). The DMS is cloud-based and licensed as a SaaS subscription. The Data Protection Resources below the DMS can be scaled out to fulfill the services that are requested by the IBM Storage Defender DMS.

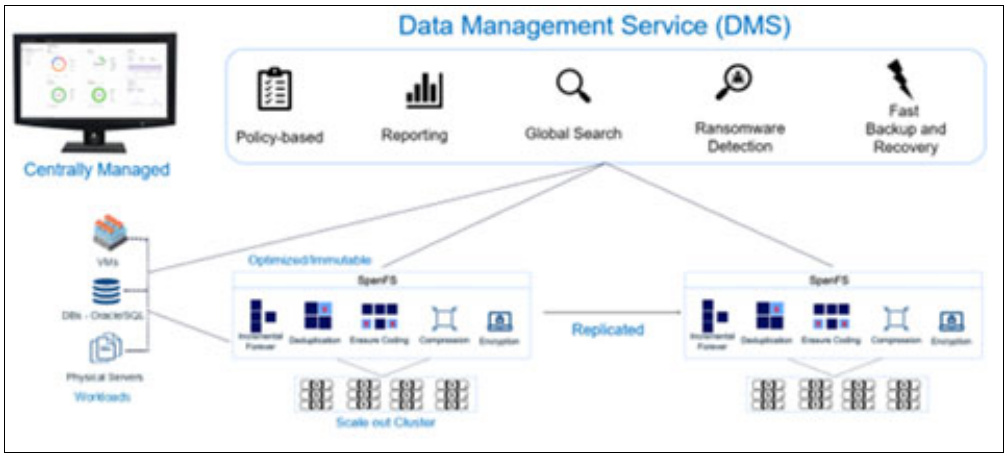


Figure 3-1 IBM Storage Defender DP architecture overview

3.1.1 Node-based scale-out architecture

IBM Storage Defender DP was designed with scalability in mind, so the architecture is described as “scale-out” rather than “scale-up”. It scales out by adding more processor nodes to make a more powerful server component. You can create a cluster on a physical node or a virtual appliance that runs the IBM Storage Defender DP software. Each node has CPU, memory, network, and storage capacity resources that are managed by the IBM Storage Defender DP software. Nodes can operate as a single-node cluster (remote and branch office (ROBO) node or an IBM Storage Defender Data Protect Virtual Edition node) or in groups of nodes that are called clusters. A minimum of four nodes are required to form an IBM Storage Defender DP cluster. Three-node clusters are possible, but n+2 node redundancy requires four nodes so that the cluster can run with up to two nodes being unavailable.

3.2 What is a cluster and how does it work

The IBM Storage Defender DP scale-out architecture is made with building blocks that are called *nodes*. The nodes are server resources that are arranged in clusters to share network, CPU, and storage capacity resources. The clusters have a powerful local cluster GUI for cluster-specific management and reporting tasks. The clusters roll up their statuses and health to the IBM Storage Defender DMS. Each cluster has a disk capacity that is provided as an integrated immutable file system that is called *SpanFS* to provide always-on availability, non-disruptive upgrades, and a pay-as-you-grow consumption model. SpanFS supports all the key capabilities that are required for data protection use cases, including globally distributed NFS, SMB, and S3 storage, unlimited snapshots, global deduplication, encryption, replication, global indexing and search, and the performance that is needed for both sequential and random I/O.

SpanFS is a part of IBM Storage Defender DP that provides file services for the data protection cluster, spreading the I/O across the nodes in the cluster to maximize throughput and performance. The solution offers industry-leading, global space efficiency technology through variable length deduplication, compression, and erasure coding (EC). It reduces capacity requirements and lowers software licensing costs. IBM Storage Defender DP provides immutable storage, data encryption, and multi-site replication to increase data resiliency.

Immutability helps ensure that WORM-protected data cannot be changed or modified by any process or user, including cluster administrators. Only the backup service running on the IBM Storage DP cluster can write to the file system through trusted APIs.

Data at rest encryption that uses FIPS-validated cryptography supports external key management.

Deduplication is performed by using a unique, variable-length data deduplication technology that spans an entire cluster, resulting in savings across a customer's entire storage footprint. SpanFS creates variable-length chunks of data, which optimize the level of deduplication no matter the type of file.

The inline deduplication process deduplicates data as it is written to the cluster and the post-process deduplication deduplicates data after it is written to the cluster.

Compression by default is an always-on capability, but it can be turned off in select storage domains. IBM Storage Defender DP uses inline ZSTD compression. Organizations can achieve enterprise-level resiliency with site-to-site replication between IBM Storage Defender DP clusters. All data on a single IBM Storage Defender DP cluster can be replicated to one or more clusters through Protection Groups. A Protection Group specifies the objects (like databases, physical servers, VMs, views, and others) to be backed up and replicated. Full and partial failover functions are fully automated by IBM Storage Defender DP at either the cluster level or the IBM Storage Defender Data DMS level for simple, seamless recovery.

3.3 Physical, virtual, and ROBO implementations

A cluster can be built to run in traditional data centers (like production and disaster recovery (DR)); on the edge of the enterprise at a ROBO by using physical or virtual nodes; or out in the public cloud by using IBM Storage Defender Data Protect Virtual Edition.

Figure 3-2 shows multiple deployment options for virtual and physical nodes, and some general suggestions and deployment examples. In the figure, the Front End TB (FETB) is the amount of data in a workload that is protected.

Remote Site Considerations	Deployment Type	Positioning	Things To know (Cons)
	Virtual Edition Node (Single Node)	< 12TB FETB	• Limited Compute • Must Replicate • Hypervisor Vulnerable • Disruptive Upgrade
	Virtual Edition Cluster (Multi-Node)	12TB - 100TB FETB	• Limited Compute • Must Replicate • Hypervisor Vulnerable • 24 Node Max
	Physical ROBO Node (Single Node)	< 24 TB FETB	• Limited Capacity • Must Replicate • No Node Redundancy • Disruptive Upgrade
	Physical Cluster (Multi-Node)	All Deployments	• Dense (Compute and storage) Node Form Factor Important • Highest price point

Figure 3-2 A comparison of virtual and physical deployment options

IBM Storage Defender Data Protect Virtual Edition is used for a single-node IBM Storage Defender DP cluster or multi-node IBM Storage Defender DP cluster that is hosted on a virtual machine (VM) on a VMware vCenter server, a Microsoft Hyper-V server. or a Nutanix AHV. You can use single-node IBM Storage Defender Data Protect Virtual Edition if your environment has smaller offices with reduced workloads that do not require the full computing power of a IBM Storage Defender DP cluster running on multiple nodes. In addition, IBM Storage Defender Data Protection Clustered Virtual Edition is a multiple-node IBM Storage Defender DP cluster that can support larger workflows. It is hosted on multiple VMs in a VMware vCenter server.

Physical deployment options include the ROBO single node that offers more compute power than the single virtual node, and limited capacity. Because this node is a single node, there is no node redundancy and there is limited disk capacity. To protect the ROBO node, set up replication to an IBM Storage Defender DP cluster to protect the backup data and provide an additional copy of the backup data at a centrally managed site.

Although IBM Storage Defender DP is designed with IBM Storage Ready Nodes in mind, it does not require the usage of a particular vendor's hardware offerings. Several different vendor platforms were tested and approved for usage, including Dell, Cisco, HP, and others. For more information about vendor platforms, see the IBM Storage Defender Data Protect: Third-Party Hardware Support Matrix (requires login with IBMid).

Physical clusters are the best options for scaling performance and capacity to meet the needs of production backup windows and replication to a DR site. The cluster resources are also available for cyber resiliency design considerations like IMR and recovery assurance testing.

3.4 Cluster setup

The details for setting up the cluster vary depending on the choice of hardware or virtual servers that are used for the cluster nodes. For more information, see IBM Storage Defender (requires login with an IBMid).

In this book, we summarize the cluster setup process by using IBM Storage Ready Nodes.

Before you bring the services team to your site, complete the following tasks:

1. Rack the cluster.
2. Cable the cluster.
3. Verify the prerequisites for cluster setup.
4. Configure the BIOS on IBM Storage Ready Nodes.

3.4.1 Configuring an IBM Storage Defender DP Storage Domain

To configure an IBM Storage Defender DP Storage Domain, complete the following steps:

1. Expand **Settings**.
2. Select Summary.
3. Click **Storage Domains**.
4. Click **Create Storage Domain**.

   Figure 3-3 shows steps 1 - 4.

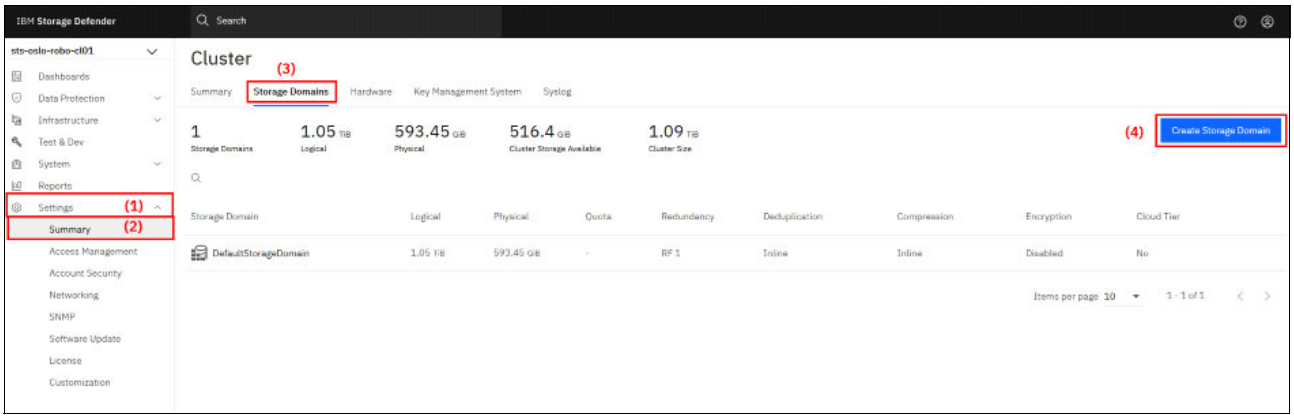


Figure 3-3 Storage Domain configuration: Steps 1 - 4

5. Input a meaningful name for the storage domain.
6. Select the type of authentication provider (if required).
7. Alter the deduplication and compression setting if your requirements differ from the default.

   Figure 3-4 shows steps 5 - 7.

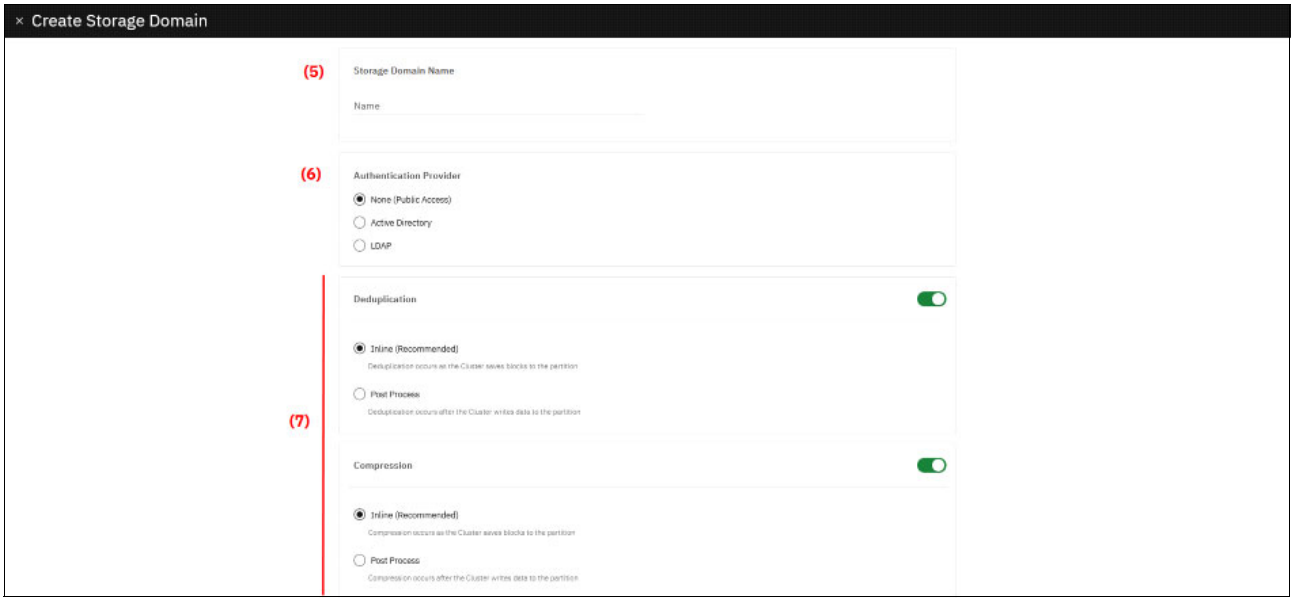


Figure 3-4 Storage Domain configuration: Steps 5 - 7

Note: Encryption is either on or off. It cannot be modified after creation. If encryption must be disabled, delete the Storage Domain and re-create it as needed or create another Storage Domain with the same settings but include encryption.

8. Enable encryption if required.
9. Expand the more icon.
10. Enable the quota and alert options, if required.
11. Set the quota limit for each Storage Domain.

    Figure 3-5 on page 23 shows steps 8 - 11.

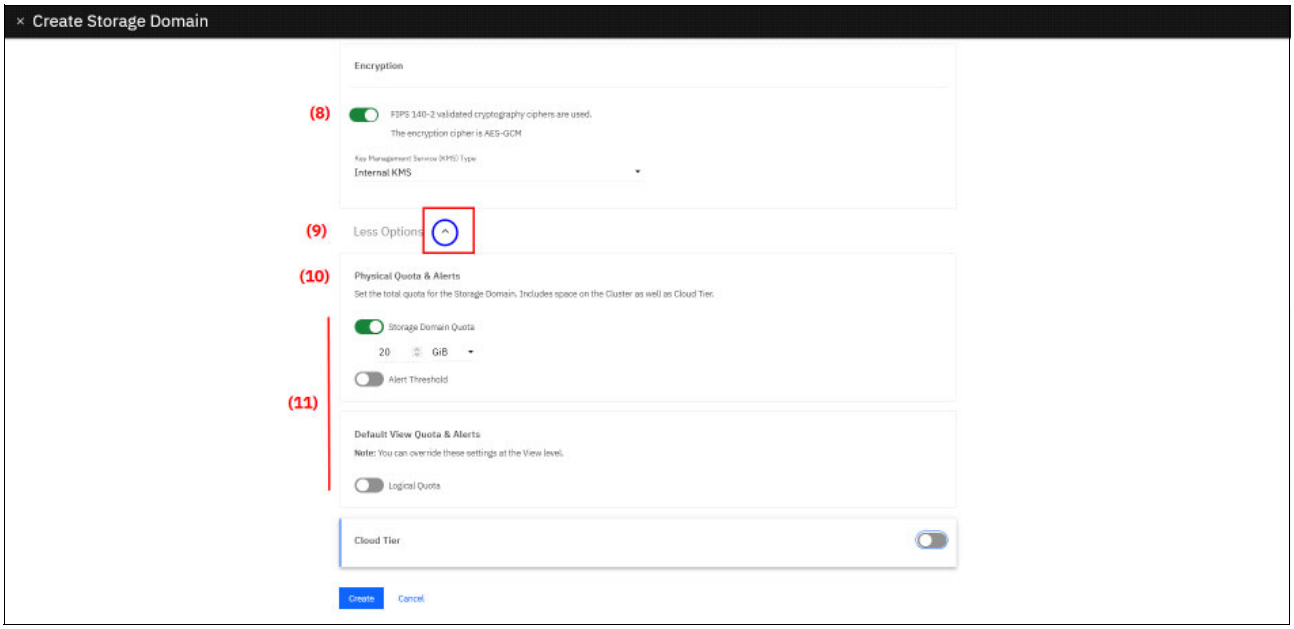


Figure 3-5 Storage Domain configuration: Steps 8 - 11

12. Enable **Cloud Tier** if you are using an external S3 target.
13. Click **Register External Target** and select the target.
14. Specify the Storage Type, Storage Class, and other options. Click Register.

    Figure 3-6 shows steps 12 - 14.

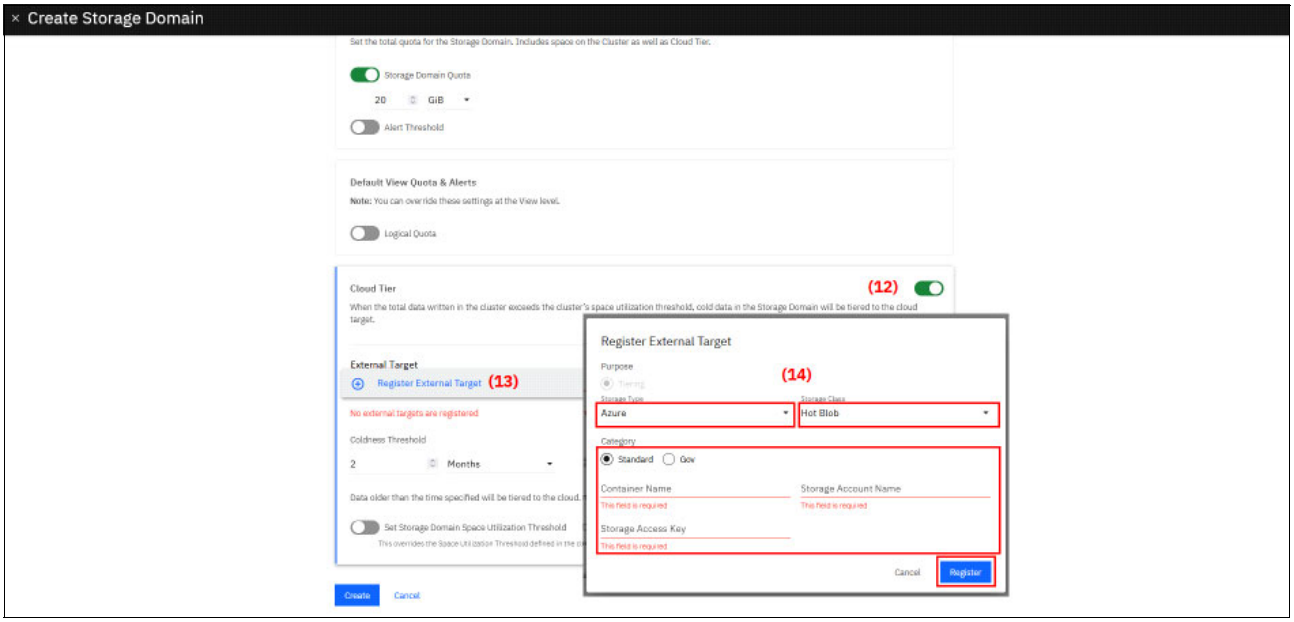


Figure 3-6 Storage Domain configuration: Steps 12 - 14

15. Set the Configure Threshold for moving data to the cloud tier.
16. Select the required fault tolerance and redundancy requirements.
17. If required, enable **Erasure Coding**.
18. Select the Erasure Coding configuration from the drop-down menu.
19. Click **Create**.

Figure 3-7 shows steps 15 - 19.

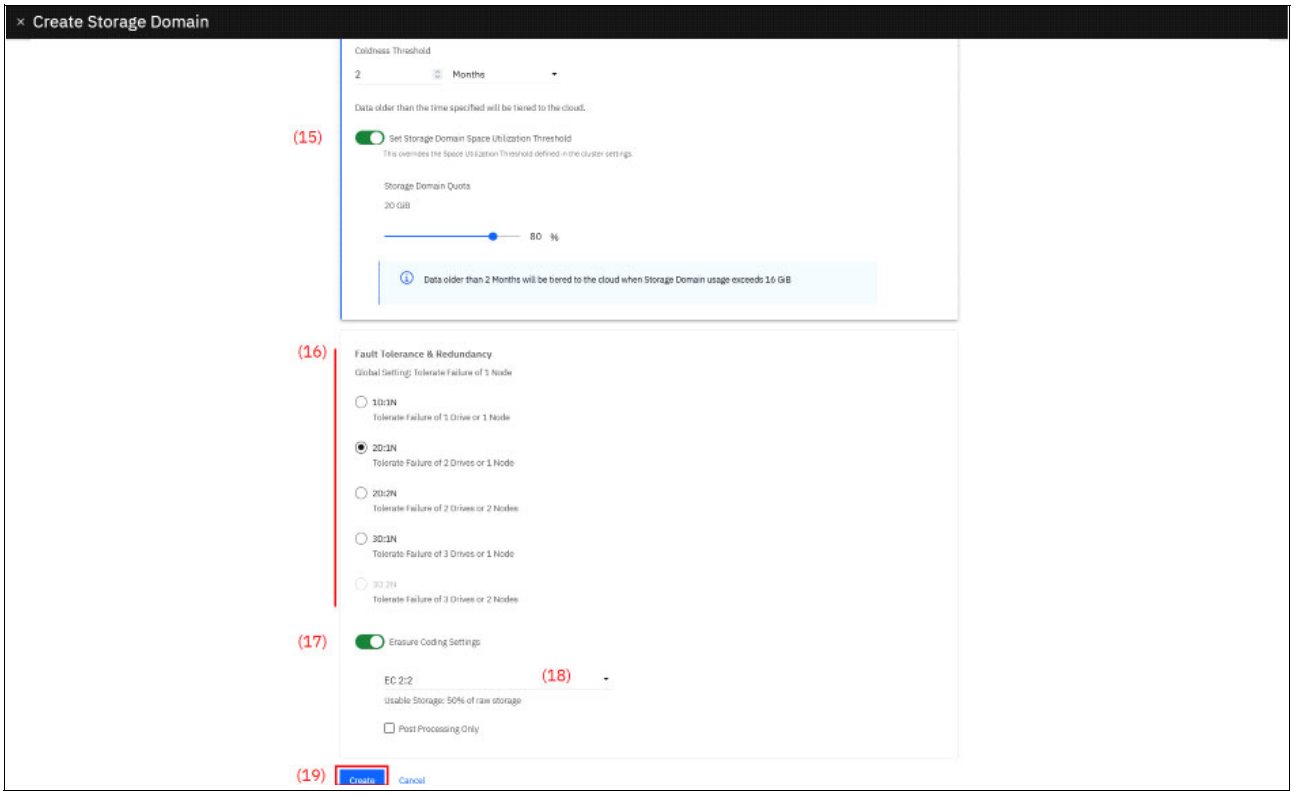


Figure 3-7 Storage Domain configuration: Steps 15 - 19

A note on fault tolerance and redundancy

IBM Storage Defender DP supports the configuration of fault tolerance at the cluster level and the storage domain level. When deploying a cluster, cluster-level fault tolerance is selected. This fault tolerance setting becomes the default value for any new storage domains that are created on the cluster. Optionally, a different fault tolerance setting can be selected when you create a storage domain or edit an existing domain. Modifying storage domain fault tolerance is a non-disruptive action and can be performed without an outage.

This feature is invaluable in environments that grow over time. As the amount of physical hardware within an IBM Storage Defender DP cluster grows to meet the increasing storage demands or changes are made to the protected workloads, the likelihood of a hardware change increases too. IBM Storage Defender DP can adapt to provide increased resilience in such scenarios without user disruption.

When the new settings are applied, all new writes to the system are protected according to the newly chosen resiliency model, and a background process handles re-protecting the existing data in the storage domain according to the new model settings.

The available Erasure Coding options in step 18 on page 24 depend directly on the Fault Tolerance value that was selected in step 16 on page 24.

3.4.2 Configuring an IBM Storage Defender DP Protection Group

This section demonstrates how to create a Protection Group, combining all configured base components to protect a VMware cluster or ESXi host. Complete the following steps:

1. Expand **Data Protection**.
2. Select **Protection**.
3. Click **Protect**.
4. Click **Virtual machine**.

   Figure 3-8 shows steps 1 - 4.

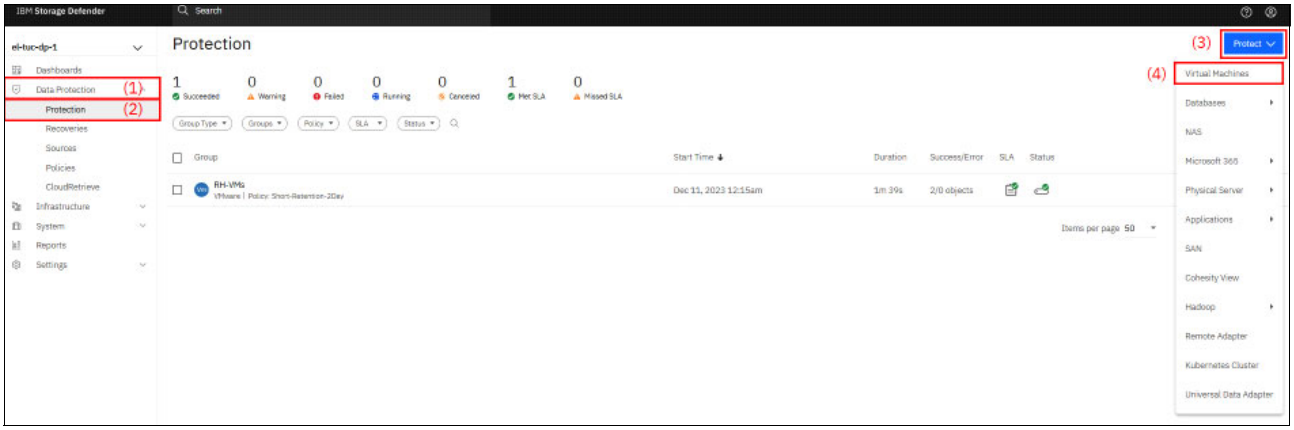


Figure 3-8 Data Protection policy for virtual machines configuration: Steps 1 - 4

A new Protection wizard window opens to complete the configuration.

5. Click **Add Objects**.
6. Select your registered VMware source.
7. Search for and select the VMs or hosts that you want to protect.
8. Click **Continue**.

Figure 3-9 shows steps 5 - 8.

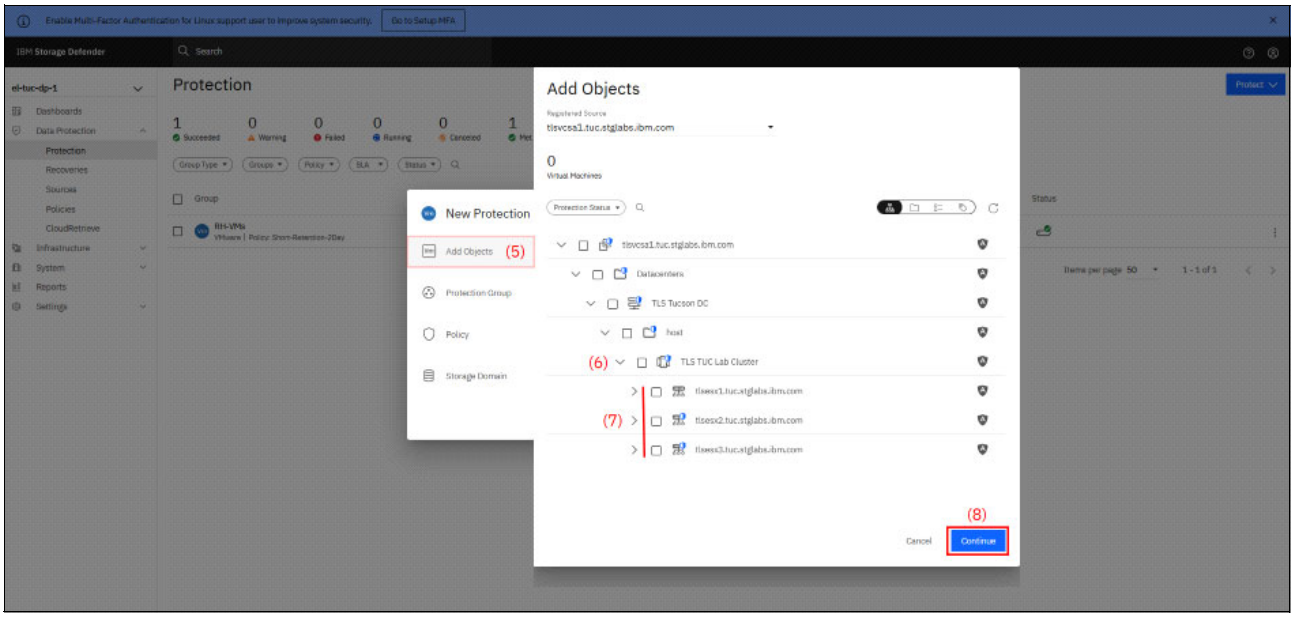


Figure 3-9 Data Protection policy for virtual machines configuration: Steps 5 - 8

Note: It is possible to view and select VMs and objects by the folder list of tag views, but it is a best practice to always use VMware vSphere tags support when selecting objects to help ensure consistent data management practices. To use tags, select the **Tags** icon (1), expand the tags list (2), and then select the VMware tag (3), as shown in Figure 3-10.

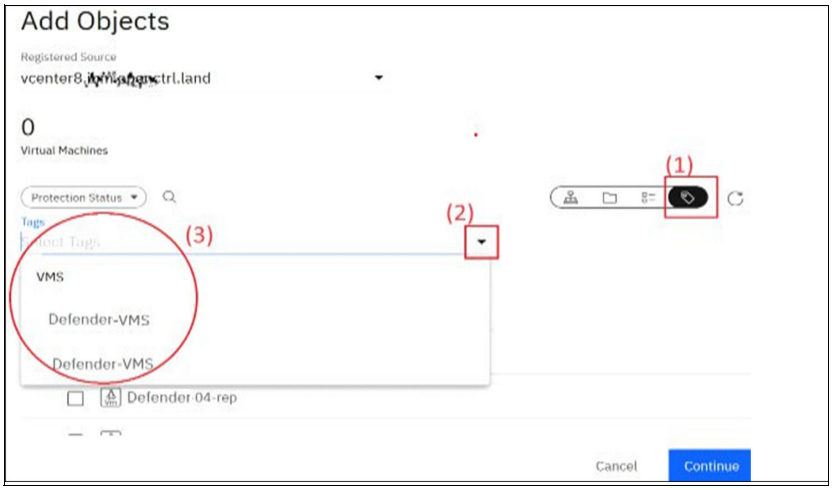


Figure 3-10 Virtual Machine Data Protection Policy configuration: vSphere tags

9. Give the protection group a meaningful name.
10. Select a policy for the protection group.
11. Select your storage domain.
12. Click **Protect** to complete the configuration.

    Figure 3-11 shows steps 9 - 12.

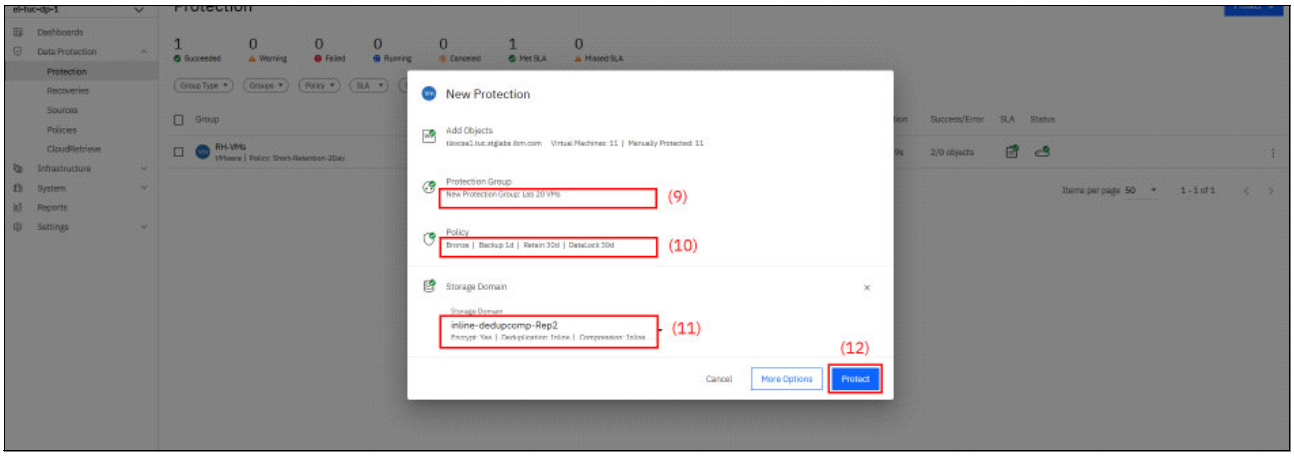


Figure 3-11 Data Protection policy for virtual machines configuration: Steps 9 - 12

The Protection Group is ready to run based on the Protection Policy schedule.

3.4.3 Monitoring the IBM Storage Defender DP Protection Group

This section demonstrates how to monitor a Protection Group task. Complete the following steps:

1. Expand **Data Protection**.
2. Select **Protection**.
3. Use the filter option and type the VMware Tags “VMtags”.

   Figure 3-12 shows steps 1 - 3.



Figure 3-12 Monitoring a Protection Group: Steps 1-3

4. In Figure 3-13, click the Protection Group **VMtags** and select a date and time to view which VMs are in the Protection Group backup set.



Figure 3-13 Monitoring the list of protected VMs in a Protection Groups backup set

5. Select the VM to view detailed task events. The window that is shown in Figure 3-14 opens.

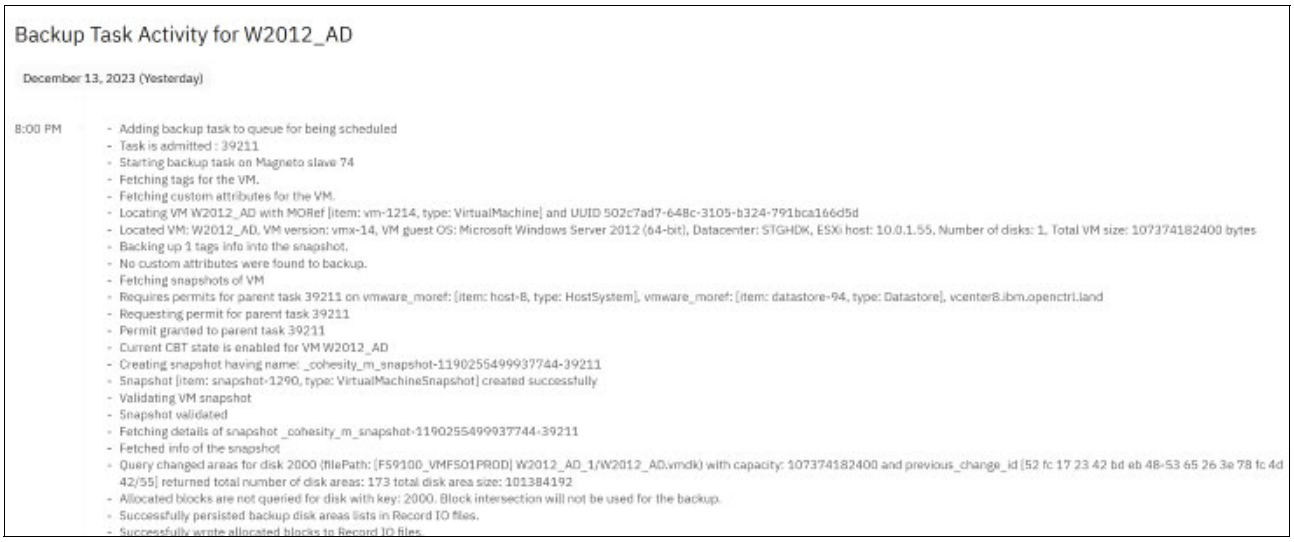


Figure 3-14 Protection Group backup task activity details

3.5 IBM Storage Ready Nodes

IBM Storage Defender DP software installation and configuration services are available from IBM Professional Services or an approved IBM Business Partner.

The following steps are typically done by the IBM Professional Services team, who follow the network requirements and data resiliency parameters that are used in the sizing:

1. Setting up nodes
2. Creating the initial cluster
3. Setting up data sources and protection policies

Note: The ISO software installation must be done on each node. Downloading the ISO image is a time saver, and having multiple USB devices with the ISO image can save additional time when there are many nodes to install.

3.6 IBM Storage Defender Ready Nodes

The IBM Storage Ready Nodes platform can be used as the foundation for an advanced, scale-out cluster that is highly secure, resilient, and provides exceptional performance by using IBM Storage Defender. For detailed advice about designing the ideal solution for your enterprise data resiliency needs, contact your IBM seller or IBM Business Partner. Together, they review the workload types, the size, retention, and change rate of your data to suggest the necessary scale-out cluster architecture that you need to meet your enterprise data resiliency goals.

A minimal cluster can be made by using four nodes (8U total rack space).

For information about the latest IBM Storage Defender supported hardware, see the IBM Storage Defender Data Protect: Third-Party Hardware Support Matrix.

Figure 3-15 shows the IBM Storage Ready Nodes storage bays and USB port access from the front of the cabinet.

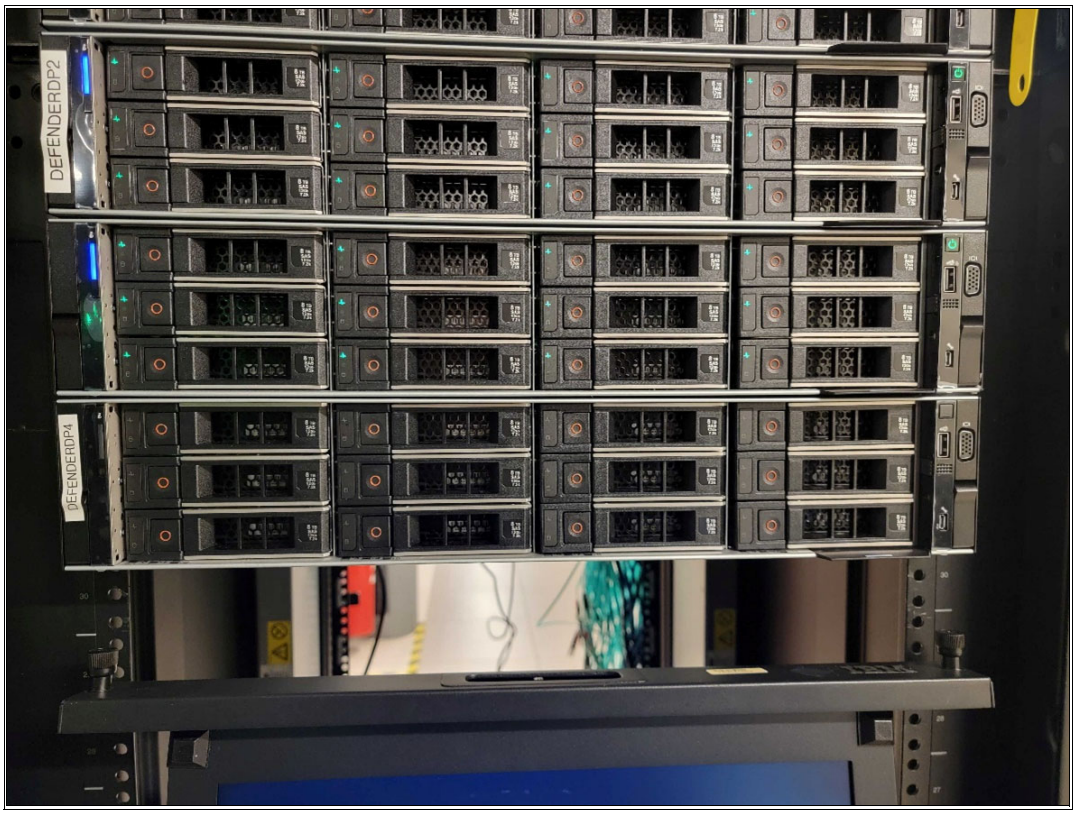


Figure 3-15 IBM Storage Ready Nodes storage bays and USB port access through the front of the cabinet

Figure 3-16 shows IBM Storage Ready Nodes management and network and optional Fibre Channel adapter access from the rear of the cabinet.



Figure 3-16 IBM Storage Ready Nodes management and network and optional Fibre Channel adapter access through the rear of the cabinet

3.7 Registering with IBM Data Management Service

IBM Storage Defender provide a convenient way to manage the status of IBM Storage Protect servers along with VMs and other protected workloads that are registered to the IBM Storage Defender DMS GUI. To register an IBM Storage Protect server with the IBM Storage Defender DP DMS, go to the DMS login and select **Connection Agents** from the left pane. You see an option to download the connection agent for the IBM Storage Protect server. Use the directions to download and extract the connection agent onto the IBM Storage Protect server:

1. Log in to the IBM Storage Defender DMS by going to the Data Management window and clicking the **Launch Data Management**.
2. From the Data Management Service window, select the hamburger icon in the upper left and select **Settings** → **Access Management**.
3. Select the **Tokens** tab.
4. Select the **Create** checkbox in the upper right. A Create Token input window opens with two input fields: Name and Type. The Name field is a user-provided name for this token. Change the Type field to `IBM`. Select **Create** to create the token.
5. Before closing the Create Token window, copy the token by selecting the copy icon.
6. Register the claim by running the following command:

   ```
   /opt/ibm/defender/bin/dcli claim register --token=<token copied from previous step>
   ```

7. To start the Connection Agent, run the following command:

```
/opt/ibm/defender/bin/dcli eagle-agent start
```

IBM Storage Defender DP 7.1.0 and later can be integrated with snapshots from the following SAN storage arrays:

- IBM FlashSystem (includes immutable IBM Safeguarded Copy snapshots)
- Nimble Storage, HPE Alletra 5000, or HPE Alletra 6000
- IBM Pure Storage FlashArray

Registering VM Hypervisors

When you create a Protection Group, you can select an existing source, policy, or storage domain. You can also create them while creating the Protection Group. However, you might find it more efficient to create them before you create the Protection Group, as described in Chapter 6, “Integrating IBM Storage Defender DP with IBM Storage Protect” on page 87.

3.8 VMware workload use case

One of the key capabilities of IBM Storage Defender DP is that it can perform IMR to drastically reduce downtime after an incident like a large-scale malware attack or ransomware attack. For more information, see Cohesity.

IBM Storage Defender DP can recover Protected Objects (such as VMs) from a snapshot that was created by a Protection Group. You can recover VMs from a cluster or a registered archive. You can recover VMs to their original location or a new location. The scale-out architecture of the IBM Storage Defender DP cluster provides the resources that are needed to do large-scale recovery of hundreds of VMs at a time, which y improves the RTO for groups of servers and applications.

For more information, see Chapter 6, “Integrating IBM Storage Defender DP with IBM Storage Protect” on page 87. It explores the VM backup process and the IMR capability by using a lab environment.

4

Protecting VMware workloads with IBM Data Protect

This chapter explores key concepts and the different functions and features of IBM Storage Defender Data Protect (DP), and how to configure agentless backup workloads for a VMware environment through vCenter or a single ESXi host. Backups are snapshot-based, incremental forever backups that can recover a virtual machine (VM) or restore a specific file or folder within a VM.

This chapter describes the following topics:

- 4.1, “IBM Storage Defender DP overview” on page 34
- 4.2, “Configuring IBM Storage Defender DP sources” on page 34
- 4.3, “Recovering data with IBM Storage Defender DP” on page 37
- 4.4, “Replicating workloads with IBM Storage Defender DP clusters” on page 44
- 4.5, “Monitoring the Protection job list” on page 48
- 4.6, “Recovering from replicated copies” on page 49
- 4.7, “Archiving workloads with CloudArchive and IBM Storage Defender DP cluster” on page 53

© Copyright IBM Corp. 2025.

4.1 IBM Storage Defender DP overview

IBM Storage Defender DP is a powerful tool in the hands of an administrator. An administrator can benefit from understanding its many features and functions that help an organization realize the full potential of IBM Storage Defender Protect to secure data from potential threats.

Here is a list of functions and definitions to help you understand IBM Storage Defender DP:

Protection	Used to define the parameters around a specific workload, such as a type of workload, objects to backup and protect, a policy for retention and frequency, and the storage domain to which the datasets are stored.
Sources	Used to register (attach) a source host.
Policies	Used to define parameters for retention, frequency, location to where the backup is being stored, and designating replication, archive location, and cloud destination.
Storage domains	Used to define how the data is processed (deduplicated or compressed) or stored. You can set quotas and define Fault tolerance and redundancy.
Remote cluster	Used to configure replication targets for replicated protection.
External targets	Used to configure archival or tiering targets for additional copies of data sets.

The following sections describe how to configure each feature, starting with registering sources to configuring a complete protection policy. The goal is to complete a protection policy for VMware environments, and configure additional copies for replication or archival.

4.2 Configuring IBM Storage Defender DP sources

When you configure IBM Storage Defender DP sources, *sources* refer to a host or hosts that are designated for protection. The focus of this book and its examples use sources that contain VMs, specifically VMware.

Figure 4-1 on page 35 shows the steps to register new VMs and configure a new source:

1. Expand the **Data Protection** section in the left pane of the IBM Storage Defender GUI.
2. Click **Sources**, and then select **Register**.
3. Select **Virtual Machines** for registration, and then select the source type from the drop-down menu.
4. Select VMware: **vCenter** or **ESXi host**.
5. Complete the fields for the IP and hostname, username, and password for the designated VMware cluster or host.
6. Enable the listed options that you want IBM Storage Defender to use with the VMware environment.
7. Click **Register** to complete the registration process.

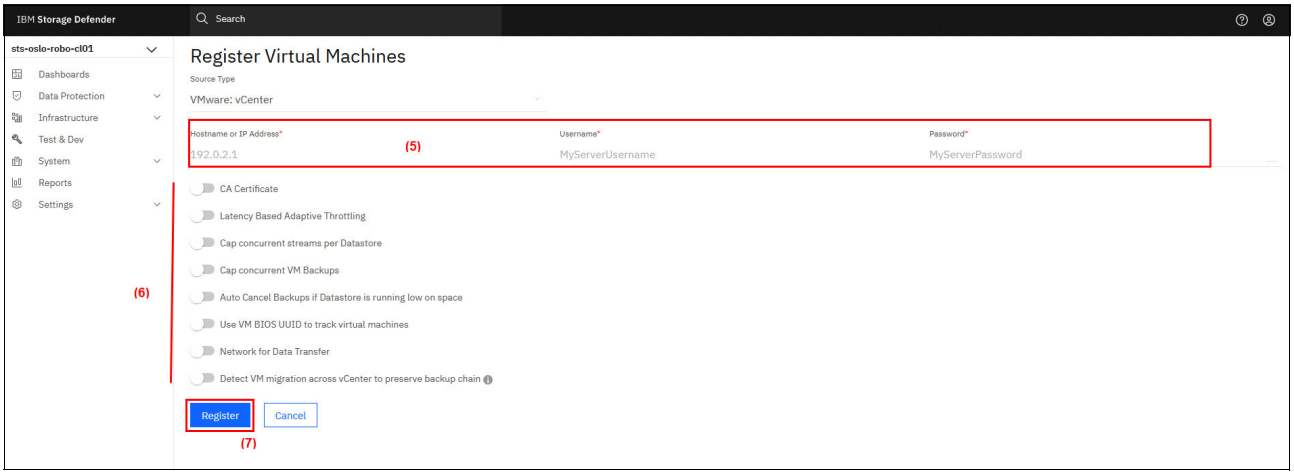


Figure 4-1 Register Virtual Machines: Steps 5 - 7

Note: As a best practice, have a minimum of 10 GbE links or greater between the IBM Storage Defender DP cluster and the VMware hosts to achieve the most effective RPO and RTO.

4.2.1 Configuring an IBM Storage Defender DP policy

Policies configure the retention and frequency of recovery points, the location where the backup is stored, a replication target, an archive location, and a cloud destination.

The following example shows the steps to configure a IBM Storage Defender DP policy for the previously configured VMware backup:

1. Expand **Data Protection**.
2. Click **Policies**.
3. Click **Create Policy**.
4. Give the policy a name.
5. Set the frequency of the backup (Backup).
6. Set the retention of each copy (Primary Copy).
7. Click **More options**.

Figure 4-2 shows steps 1 on page 35 - 7 on page 35.

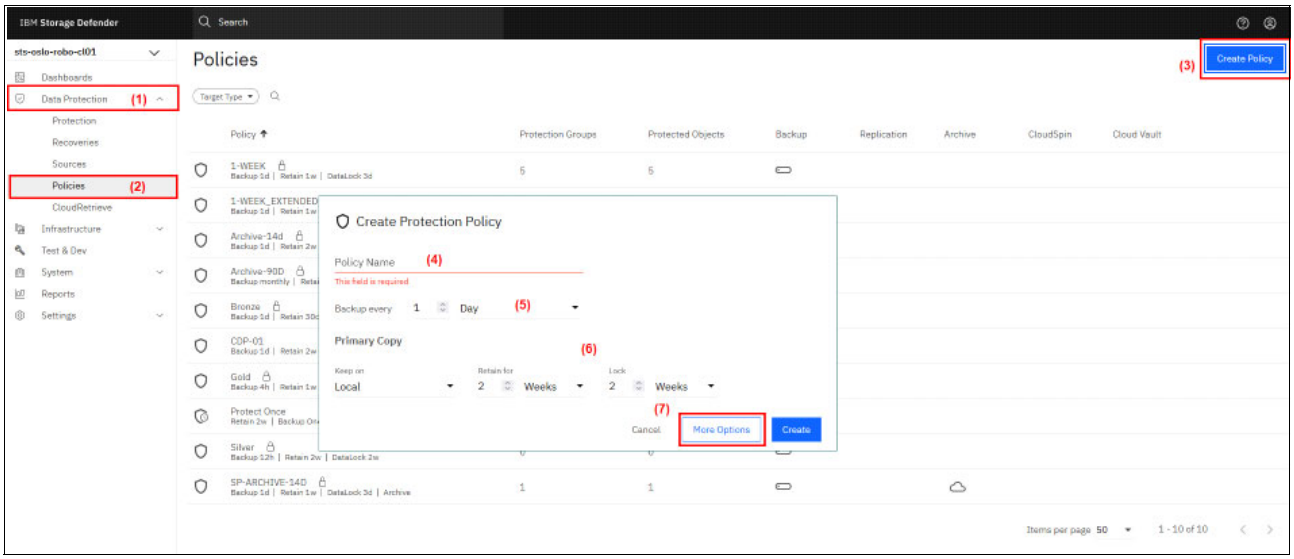


Figure 4-2 Policy sources configuration: Steps 1 - 7

8. Complete the Replication, Archive, or CloudSpin fields as needed for the backup settings.
9. Click **Create**.

Figure 4-3 shows steps 8 and 9.

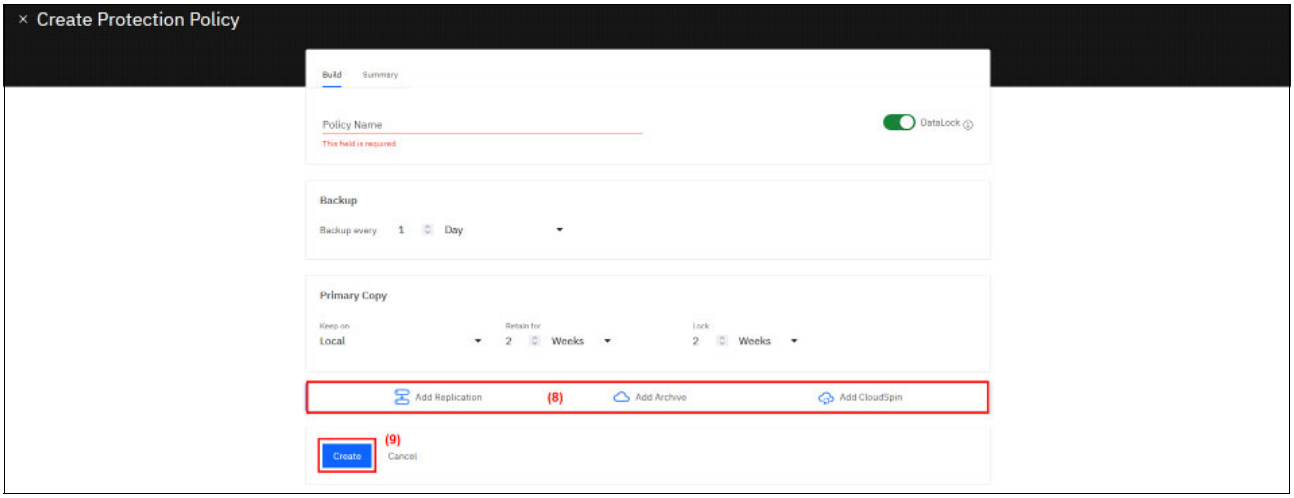


Figure 4-3 Policy sources configuration: Steps 8 and 9

Note: The Replication, Archive, or Cloudspin features can be applied only to a policy that is attached to a data protection group. The Replication, Archive, and Cloudspin features have their own frequency and retention.

To help meet legal obligations, the following additional backup options are available:

- Quality of Services (QoS): Determines the type of storage that is used and the latency factor in writing data to the view.
- Periodic Full Backup: Scheduled for Day, Week, Month, or Year. Schedules can also be made for multiple periodic full backups with different frequencies and dates.
- Continuous Data Protection: Defines the Continuous Data Protection schedule for capturing logs from the VMs, which enables the restore of VM data to a point-in-time rather than relying on the periodic snapshots that are available with regular backups.
- Quiet Times: Defines periods when new protection runs are not started.
- Customize Retires: By default, IBM Storage Defender DP attempts to capture snapshots three times before the protection run fails. The default time between retries is 5 minutes. You can customize the number of retries and how long to wait between each attempt.
- BMR Backup: Provides a backup schedule and retention period for Bare Machine Recovery (BMR) system data on a physical server.
- Log Backup: Adds a database logs schedule if there are plans to restore databases to a specific point in time between two full database server backups.

4.3 Recovering data with IBM Storage Defender DP

This section outlines the recovery types that are available, the tasks that are related to recovery, and the steps that are required when recovering data for a protected IBM Storage Defender DP cluster.

4.3.1 Recovering from a IBM Storage Defender DP cluster

When recovering from an IBM Storage Defender DP cluster, the following tasks and operations are created and run:

1. A recovery task is created.
2. The IBM Storage Defender DP cluster communicates with ESXi or vCenter (based on the source registration) to validate the current inventory and the selected recovery task settings.
3. The IBM Storage Defender DP cluster creates an internal View, clones the VM snapshot (which contains VM files), and mounts the View on the target ESXi host as an NFS data store.
4. The cloned VM snapshot is migrated to the primary data store by Storage vMotion.
5. IBM Storage Defender DP unmounts and detaches the IBM Storage Defender DP NFS data store and the View from the ESXi host once the Storage vMotion is complete.

IBM Storage Defender DP Cluster supports two methods of recovery:

- Instant recovery

  Instant recovery enables the recovery of, and access to VMs within minutes, with access to the VMs available while the restore is ongoing. Once the recovery is initiated, the VMs are instantly available after the recovery to the target location begins. On completion of the recovery, the VM data is moved to the target storage location. The VM can be accessed when the data is copied to the VM from the IBM Storage Defender DP cluster. In these cases, the VM's performance might be slow while the Storage vMotion is still in progress. To perform instant recovery, specify a recovery point (by selecting a backup or backup copy) and a target location to where the recovered VM will be located.

- Copy recovery

  The VMs are available in the target location only after all data is copied to the target storage from the source location (IBM Storage Defender DP cluster or cloud). Once the restore is complete, the VM is available to use.

Note: If you chose to keep the VM in a powered-off state when initiating instant recovery, wait for the Storage vMotion to complete before powering on the VM. As a best practice for better performance, choose **Copy Recovery** for powered-off VMs.

To perform a VM recovery, complete the following steps:

1. Expand **Data Protection**.
2. Select **Recoveries**.
3. Click **Recover**.
4. Select **Virtual Machines**.
5. Select **VMs**.

   Figure 4-4 shows steps 4 and 5.

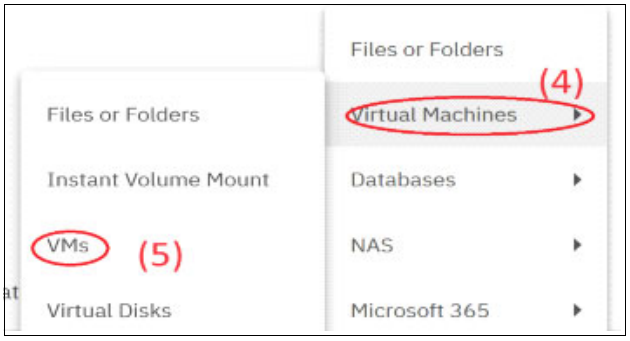


Figure 4-4 Recovering a VM with IBM Storage Defender DP: Steps 4 and 5

6. Filter by Protection Group (VMtags).
7. Search by words (W2012_AD).

   Figure 4-5 on page 39 shows steps 6 and 7.

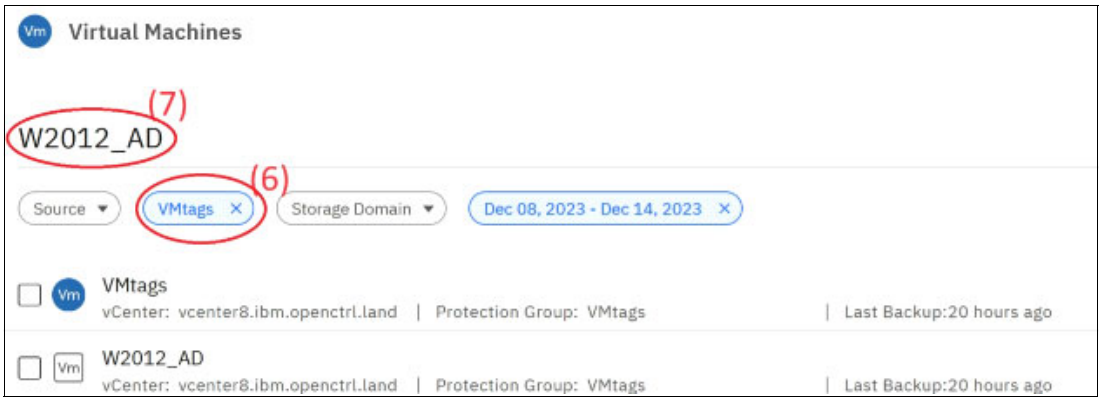


Figure 4-5 Recovering a VM with IBM Storage Defender DP: Steps 6 and 7

8. Select the VM to recover and click **Next**. For the recovery method, select **Instant Recovery**.
9. Select one of the following options for Existing VM Handling:

None	The recovery fails if there is an existing VM. Recovers the VM as a new VM and retains the original VM. In the Rename field that is available in the Recovery Options, specify the name of the recovered VM by adding prefix or suffix strings to the name of the original VM. By default, the recovered VM is named `copy-<original_VM_name>`.
Overwrite Existing VM	Recovers the VM by deleting the original VM. The recovered VM has the original VM name. When choosing this option, the original VM is deleted before the recovery. A recovery failure results in the loss of the original VM.
Attempt Differential Recovery	Attempts to recover the VM by overwriting only the difference between the original VM and the snapshot that is selected for recovery. Any newly added data in the original VM is deleted. If you want to reclaim space for thin-provisioned disks, do not to attempt differential recovery, and perform a recovery only by using **Overwrite Existing VM**. Differential recovery reduces the amount of data that is transferred during the recovery process. In the pulse log, you can view the amount of data transfer that is reduced by selecting **Differential Recovery**.
Keep Existing VM	Recover the VM as a new VM with the original VM name. The original VM is retained as powered off. The original VM is renamed to `obsolete_<time_stamp><Orginal_VM_Name>`.

Figure 4-6 shows steps 8 on page 39 and 9 on page 39, including the options for step 9 on page 39.

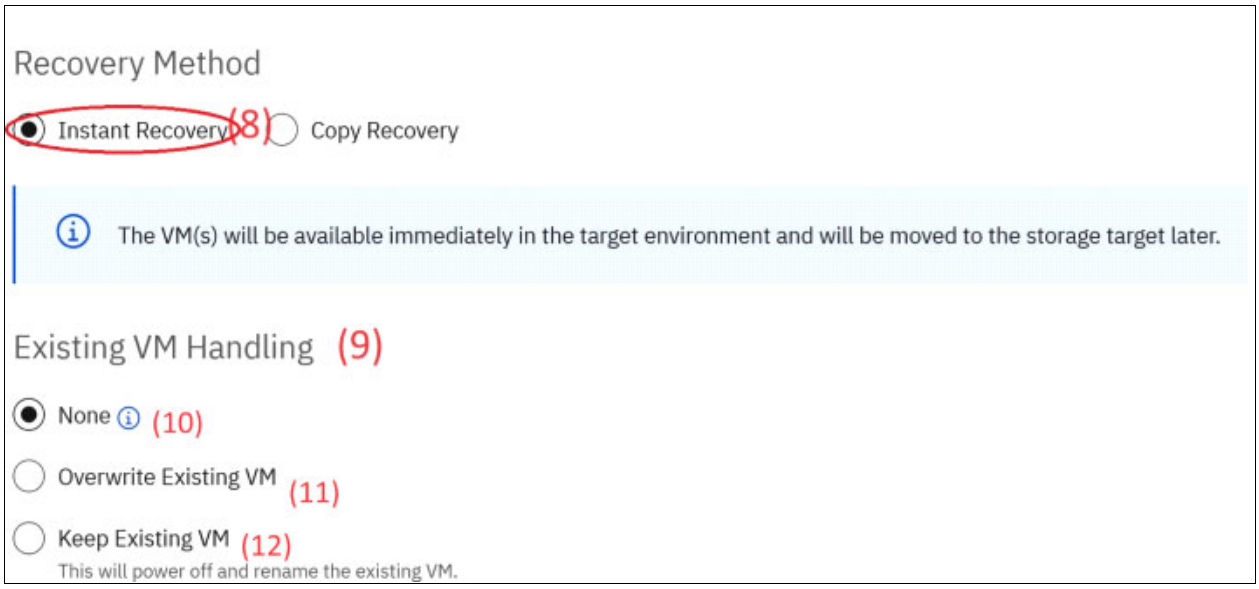


Figure 4-6 Recovering VM with IBM Storage Defender DP: Steps 8 and 9 with options

10. Finish the recover task and monitor the workflow (Figure 4-7). If needed, review the recovery activity log that is associated with the recovery process (Figure 4-8 on page 41).

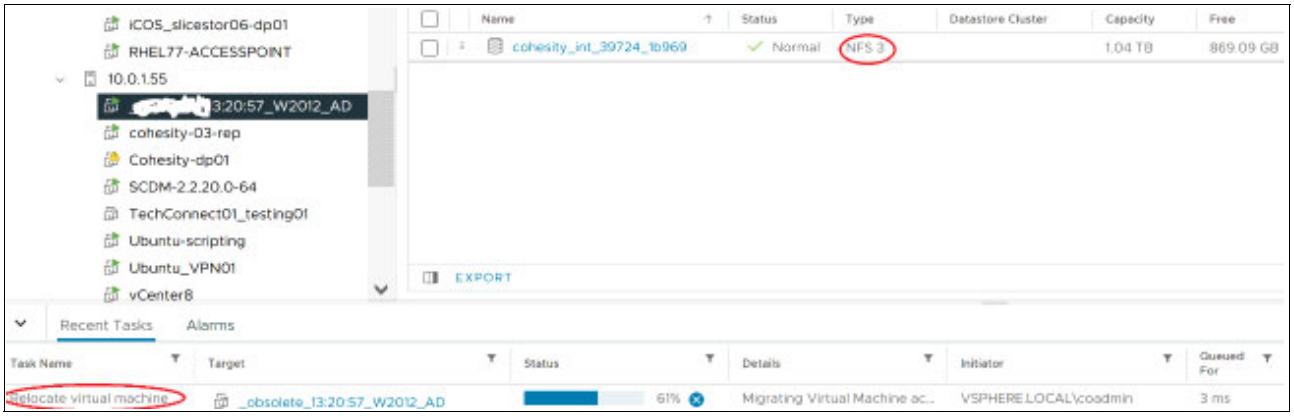


Figure 4-7 IBM Storage Defender DP: Instance Recovery vMotion from NFS to production

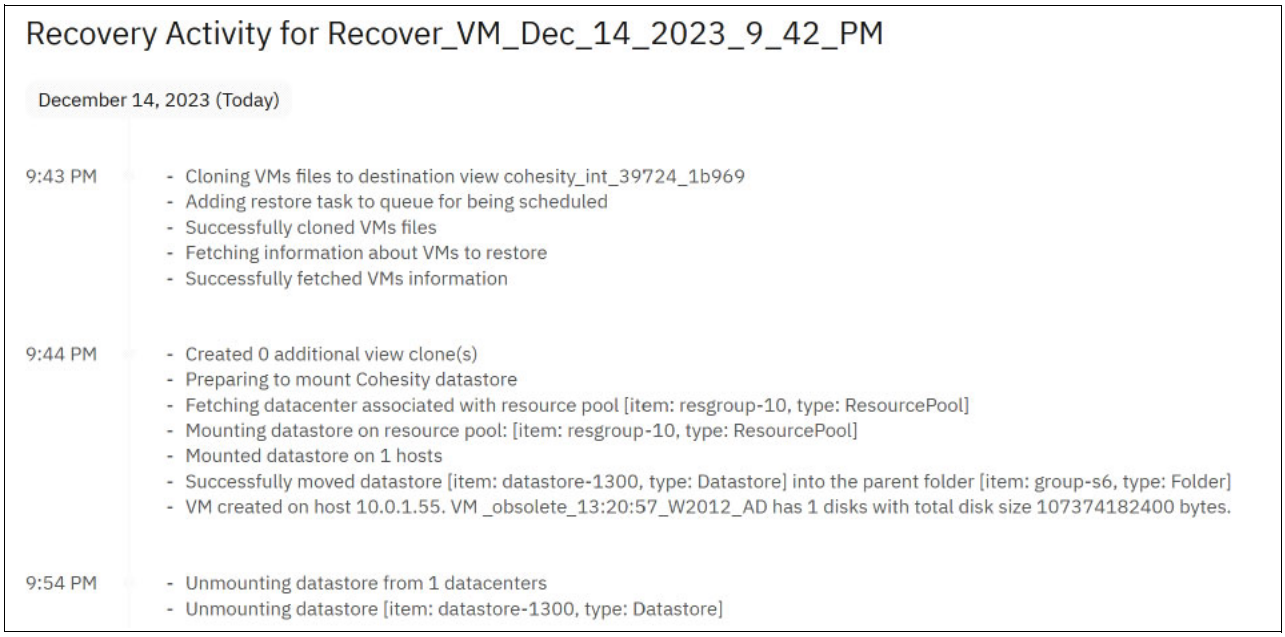


Figure 4-8 IBM Storage Defender DP: Recover Activity log

4.3.2 Instant Volume Mount recovery from an IBM Storage Defender DP cluster

The Instant Volume Mount (IVM) operation restores an entire volume to a running VM. It is the best way to restore data when a large data set must be recovered. This restoration method mounts the restored Virtual Machine Disk (VMDK) to a Windows VM to enable access to copy the required data wherever it is needed.

To recover a VMDK by using IVM, complete the following steps:

1. Expand **Data Protection**.
2. Select **Recoveries**.
3. Click **Recover**.
4. Select **Virtual Machines**.
5. Select **Instant Volume Mount**.

   Figure 4-9 shows steps 4 and 5.

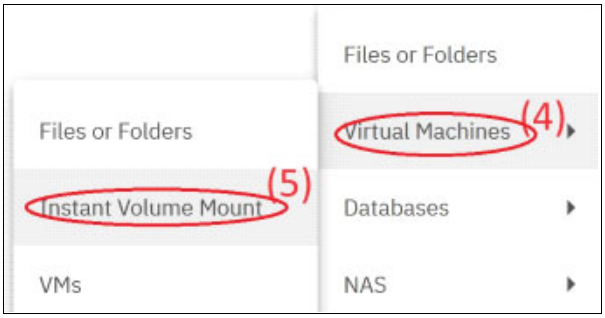


Figure 4-9 Instant Volume Mount recovery from IBM Storage Defender DP: Steps 4 and 5

6. Filter by Protection Group (VMtags).
7. Search by string (W2012_AD).

Figure 4-10 shows steps 6 on page 41 and 7 on page 41.

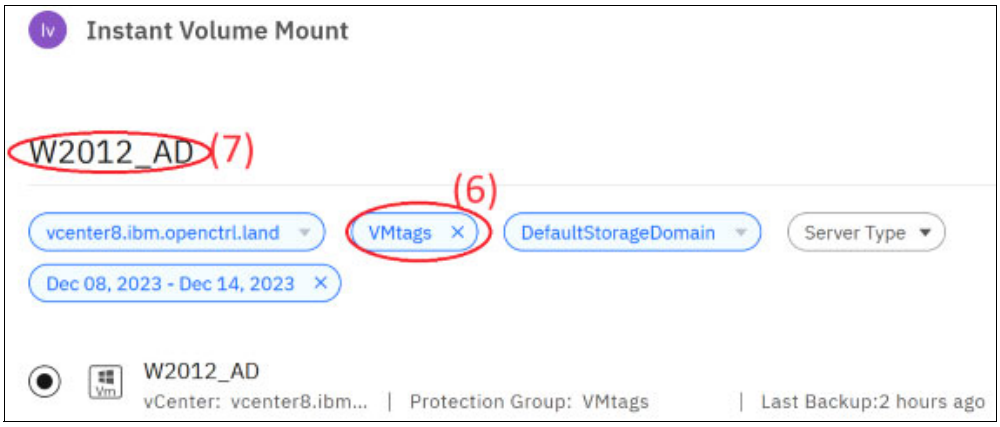


Figure 4-10 Instant Volume Mount recovery from IBM Storage Defender DP: Steps 6 and 7

8. Select the recover options and the volume to recover.
9. The default recover option is the Original Location (Figure 4-11). The VMDK must mount to another VM target. To do so, select **New Location** and browse to the vCenter and the VM where the VMDK should be mounted (Figure 4-12 on page 43).

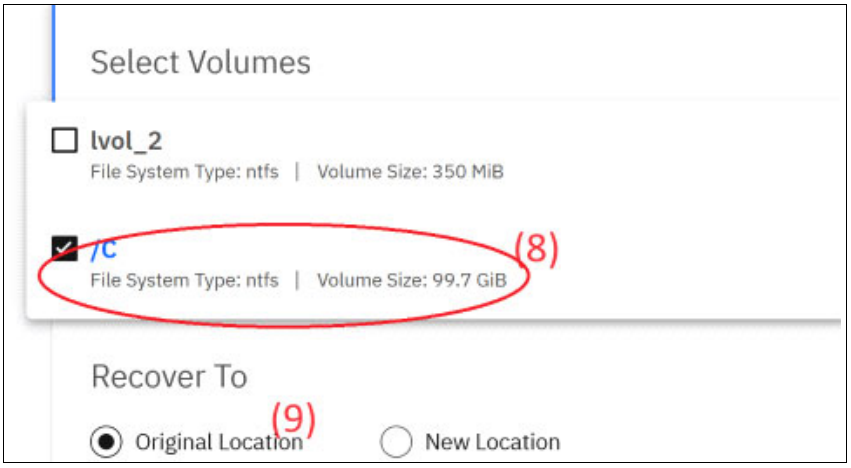


Figure 4-11 Instant Volume Mount recovery from IBM Storage Defender DP: Steps 8 and 9

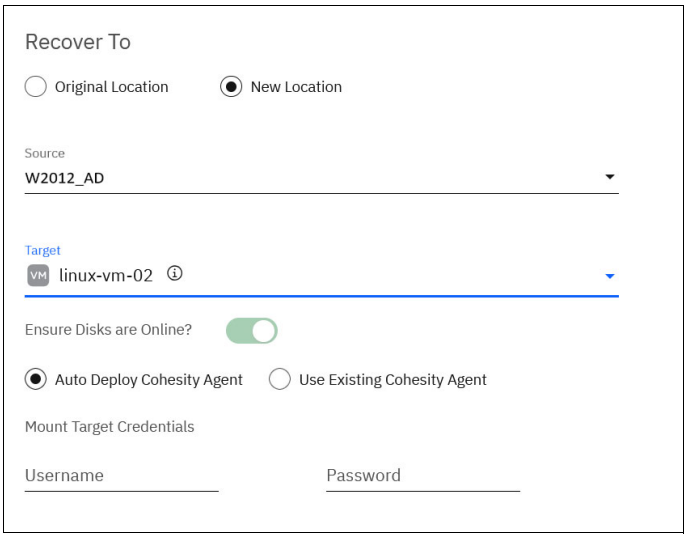


Figure 4-12 Instant Volume Mount recovery to a new VM location

10. Finish the recovery and monitor the workflow (Figure 4-13).



Figure 4-13 Instant Volume Mount recovery from IBM Storage Defender DP: Activity log

The volume is available in the Windows disk manager and can be brought online by using Windows disk management (Figure 4-14).

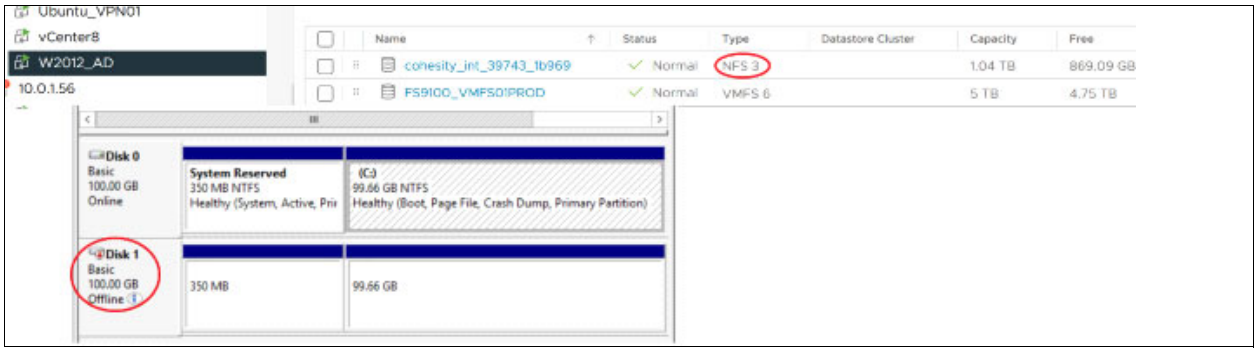


Figure 4-14 Instant Volume Mount recovery from Protect Windows volume status

When access to the volume is no longer needed, use the **Teardown** option on the mounted volume in the IBM Storage Defender DP GUI to remove the volume from the target. During the teardown, the following actions happen:

1. The IBM Storage Defender DP cluster deletes the compute instances of the cloned VM object in the ESXi host.
2. The VM files (such as VMDK and VMX) are deleted on the view that acts as the data store.
3. If no VMs are using the view (data store), the view is unmounted from all ESXi hosts.

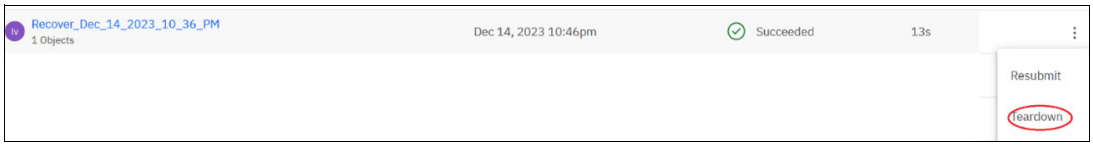


Figure 4-15 Volume recovery: Teardown option

4.4 Replicating workloads with IBM Storage Defender DP clusters

With IBM Storage Defender DP, backups that are ingested in the IBM Storage Defender DP cluster can be replicated to one or more target clusters. IBM Storage Defender DP supports the following replication configurations:

One-to-one model A single production cluster can back up to a disaster recovery site.

One-to-many model A single production cluster can back up to multiple disaster recovery sites.

Many-to-one model Multiple production clusters can back up to a remote cluster.

Many-to-many model Multiple production cluster can back up to multiple remote clusters across several sites.

This technology works with global data deduplication to reduce storage requirements when data is spread across several sites. It helps reduce the network bandwidth that is required for replication of data for DR purposes.

4.4.1 VMware replication workflow one-to-one model

This section describes the replication workflow of VMware and an IBM Storage Defender DP cluster that uses the one-to-one model. The backup data that comes from the source VMware vCenter on the primary host is ingested in the primary IBM Storage Defender DP cluster DC1, which is then replicated to a secondary DR cluster (DC2).

Figure 4-16 on page 45 shows the one-to-one model.

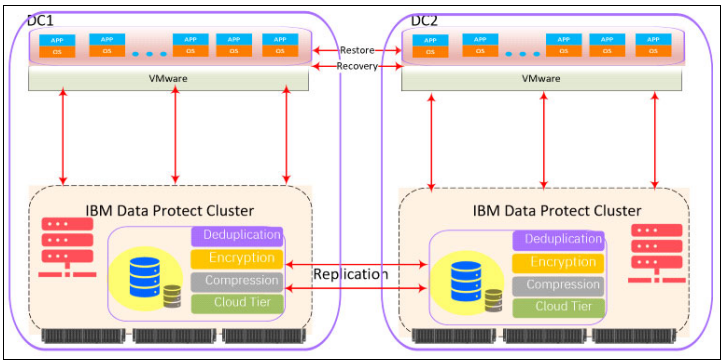


Figure 4-16 Replication workload: One-to-one model

The DR cluster in DC2 can fail over to the VMware vCenter cluster to continue the backup and restore the production workload. After a successful failover, the load can fail back to DC1 and continue using it as the primary cluster.

IP failover and load-balancing

Each node has a virtual IP (VIP) address, and clients can access any node on the cluster by using its VIP. When the node that was directly serving a client goes down, the VIP seamlessly switches over to another available node. The client continues to access the same IP address, and the fully distributed architecture enables the cluster to access client data that is physically present on any node in the cluster. The cluster depends on an external DNS server to do round-robin VIP addresses so that client requests can be load-balanced.

The IBM Storage Defender DP cluster performs source-side data deduplication and sends only changed data in increments to the other site. The cluster also continues to scan and auto-heal on the other site.

Registering a remote cluster

To register a remote cluster, the `ip-address` of the node, `node-cluster`, and the admin account and password are required. In a multiple-node cluster, always use the VIP to register the remote cluster.

Note: A best practice is to use one VIP for each cluster node

After connecting the clusters, you may enable Remote Access [1] and Replication [2], as shown in Figure 4-17. Replication is done by storage domain pairing. Map the storage domains in the source cluster to storage domains in the remote cluster.

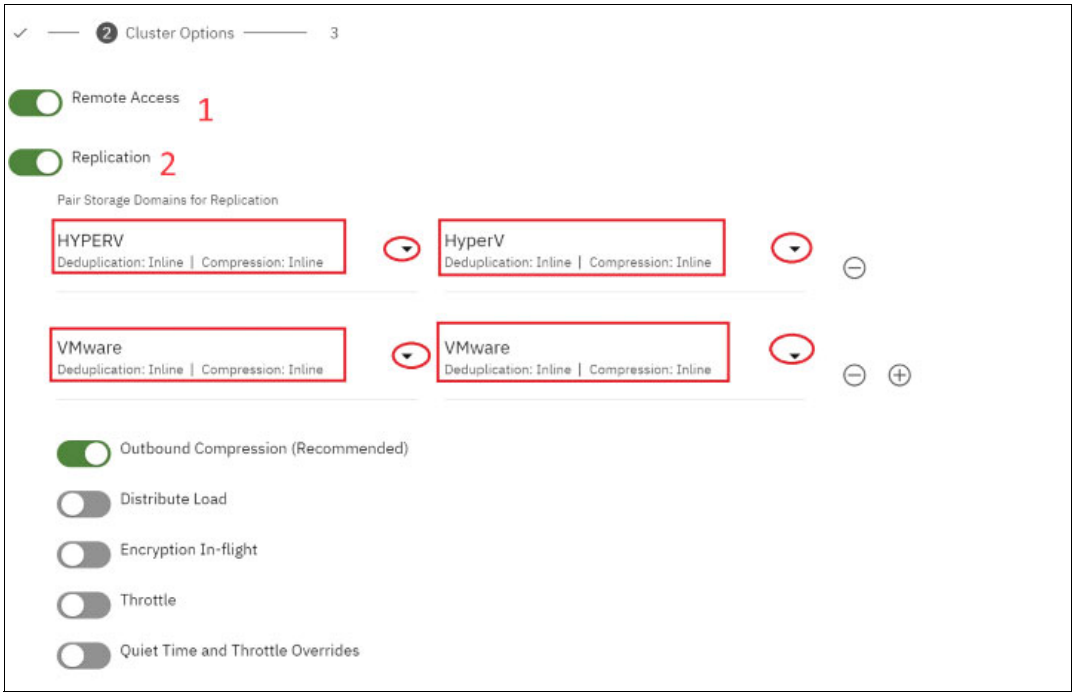


Figure 4-17 Remote cluster setup with paired storage domains for replication

Enable any additional features to apply to the cluster pairing configuration. Once the pairing is set up between the clusters, configure the pairing on the other cluster to point back to this primary cluster.

After the replication setup is finished, update the existing backup policies. From Data Protection, select Policies, and then complete the following steps in the window that is shown in Figure 4-18 on page 47:

1. In the replication field “replicate to”, select **Remote Cluster**.
2. For the Replication Target, select the hostname as the remote replication target.

Note: The following options may be toggled for the cluster:

- Outbound Compression
- Distribute Load
- Encryption In-flight
- Throttle
- Quiet time and Throttle Overrides

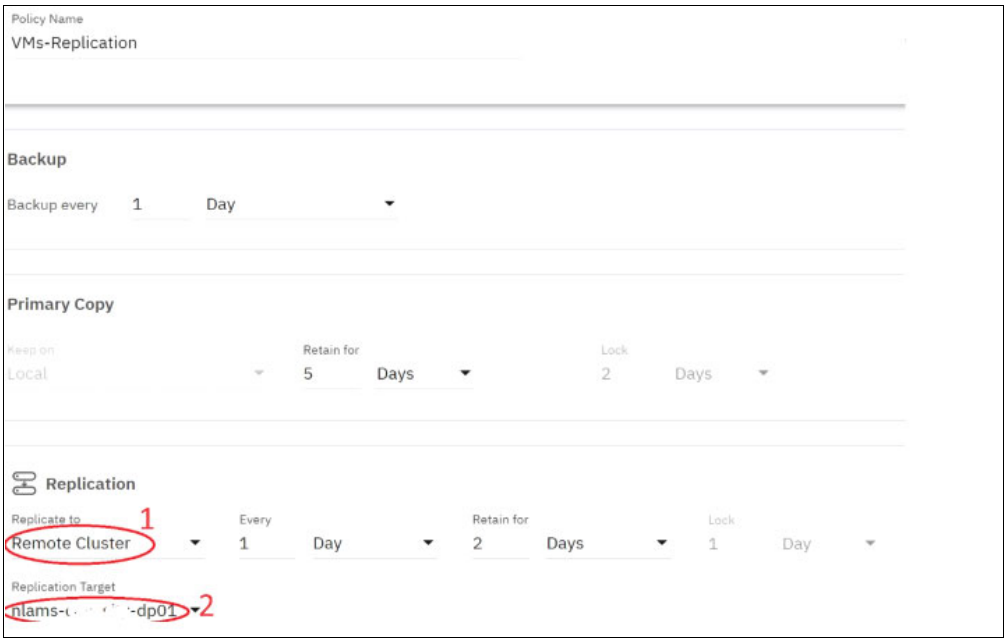


Figure 4-18 Modifying the existing backup policy

After modifying the policies, double-check the Policies Details view to help ensure that the configuration is correct (Figure 4-19).

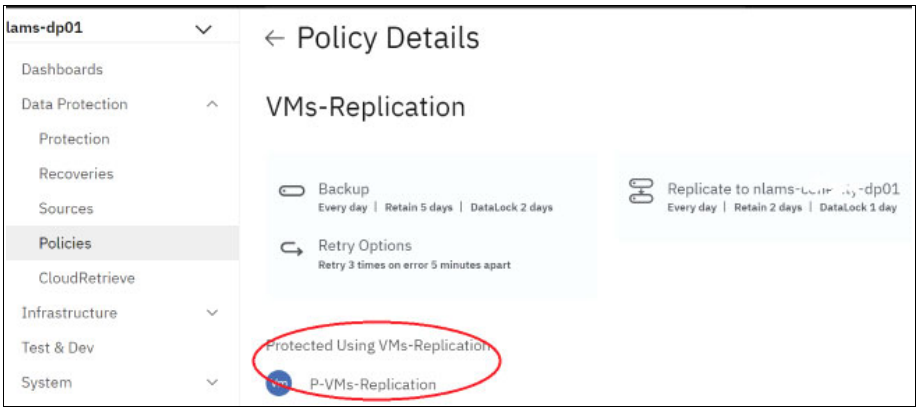


Figure 4-19 Reviewing the backup policy to confirm that the replication feature is enabled

4.5 Monitoring the Protection job list

To monitor Protection jobs, use the menu in the left pane to go to the IBM Storage Defender DP section, and then complete the following steps:

1. Under Data Protection, select **Protection** and click the Protection Group. For this example, the group name is “DP01-P-VMs-Replication” (Figure 4-20).

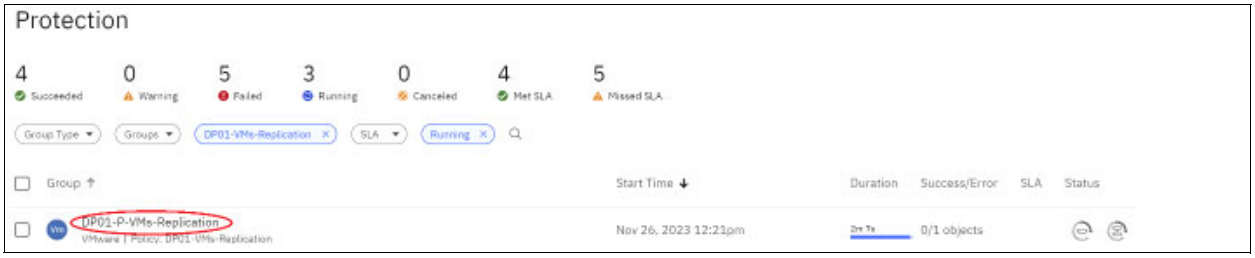


Figure 4-20 Protection group status

After you select the Protection Group, you see a list of the replication tasks that are associated with the group (Figure 4-21).

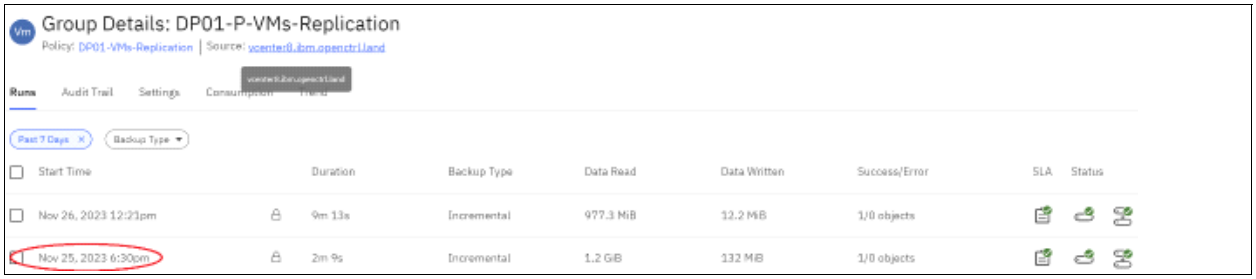


Figure 4-21 Backup and replication task details

2. In the list of jobs for the group, click the job that you want to monitor. In this example, select the job “Nov 25, 2023 6:30 PM”, and select the **Replication** tab (Figure 4-22). You see the details of the jobs replication process. In this example, the replication job finished successfully with 16 GiB logical data backed up and 1 GiB of data transferred to the target replication group.

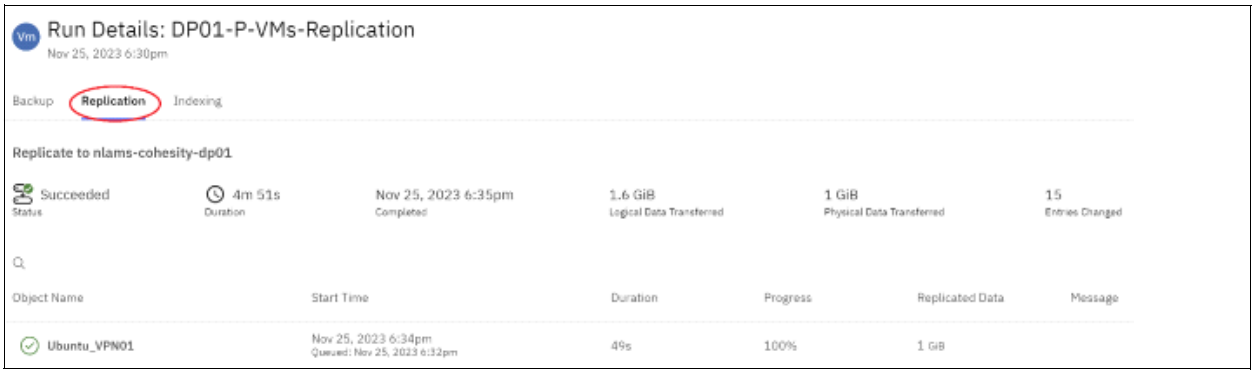


Figure 4-22 Replication task details

4.6 Recovering from replicated copies

This section describes how to recover data from replicated copies if a disaster strikes the source cluster. You can also use this process when there is a planned maintenance outage of the source cluster to maintain access to data and systems. When the IBM Storage Defender DP cluster or the source is down or not available, you can fail over Protection Groups to the replication partner to maintain continuity.

After you create a Protection Group on a cluster with a replication schedule, once the first snapshot is replicated, the Protection Group is marked as Failover Ready and the Protection Group becomes inactive. The objects (VMs) and policy are no longer associated with the inactive Protection Group.

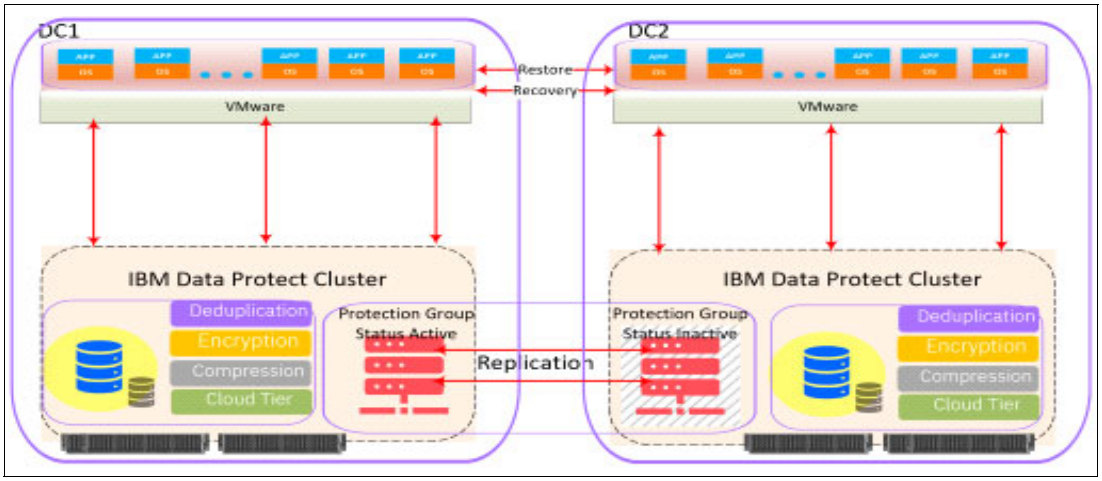


Figure 4-23 Protection Group Active / Inactive status

When an active Protection Group is deactivated, the Protection Group becomes Failover Ready. The selected VMs and policy are still associated with the Protection Group on the source cluster. Snapshots are no longer captured by the original Protection Group on the original capturing cluster. However, the existing snapshots that are captured by the Protection Group are not deleted; they remain stored in the associated storage domain.

In this example, the active Protection Group on the source cluster is DP01-P-VMs-Replication. The following actions can be taken on this Protection Group: Run Now, Pause Future Runs, Deactivate, Edit, or Delete (Figure 4-24).

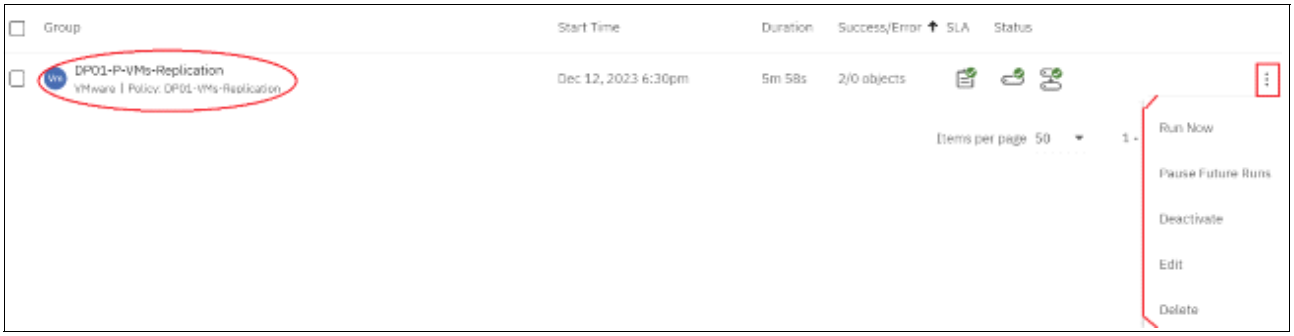


Figure 4-24 Protection Group actions

On the replicated cluster, the Protection Group is marked as inactive and Failover Ready. This Protection Group has only the options Failover and Delete (Figure 4-25).

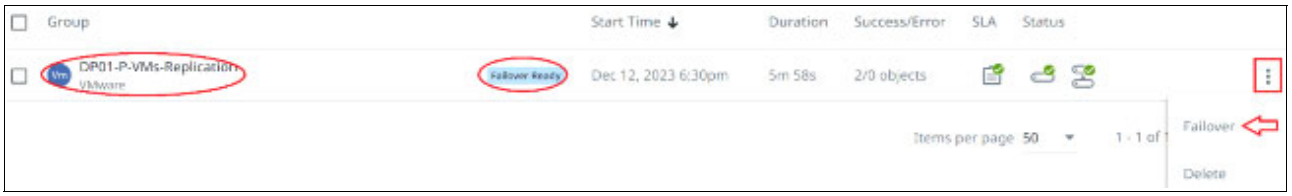


Figure 4-25 Inactive Protection Group actions

When you activate the Failover feature, specify the source vCenter server where the recovered VMs are placed. Figure 4-26 shows that vcenter8 is specified as the new source.

Note the following message: “On Failover this Protection Group will be activated on this Cluster. If this Inactive Protection Group was created by replication, this Failover causes the rejection of the incoming replicated Snapshots created by the associated Protection Group on the Primary Cluster. As part of the Failover, you can recover VMs from Snapshots located on this Cluster. Specify a Source to place the recovered VMs.”



Figure 4-26 Protection Group that is configured for failover for continuous protection

After the Failover feature is activated, the Protection Group DP01-P-VMs-Replication now has the following options that are available: Run Now, Pause Future Runs, Deactivate, Edit, and Delete (Figure 4-27).



Figure 4-27 Protection Group available actions after failover

The Protection Group is now ready to continue protecting the VMs on the replicated cluster. You now may edit the objects that require protection, as shown in Figure 4-28.

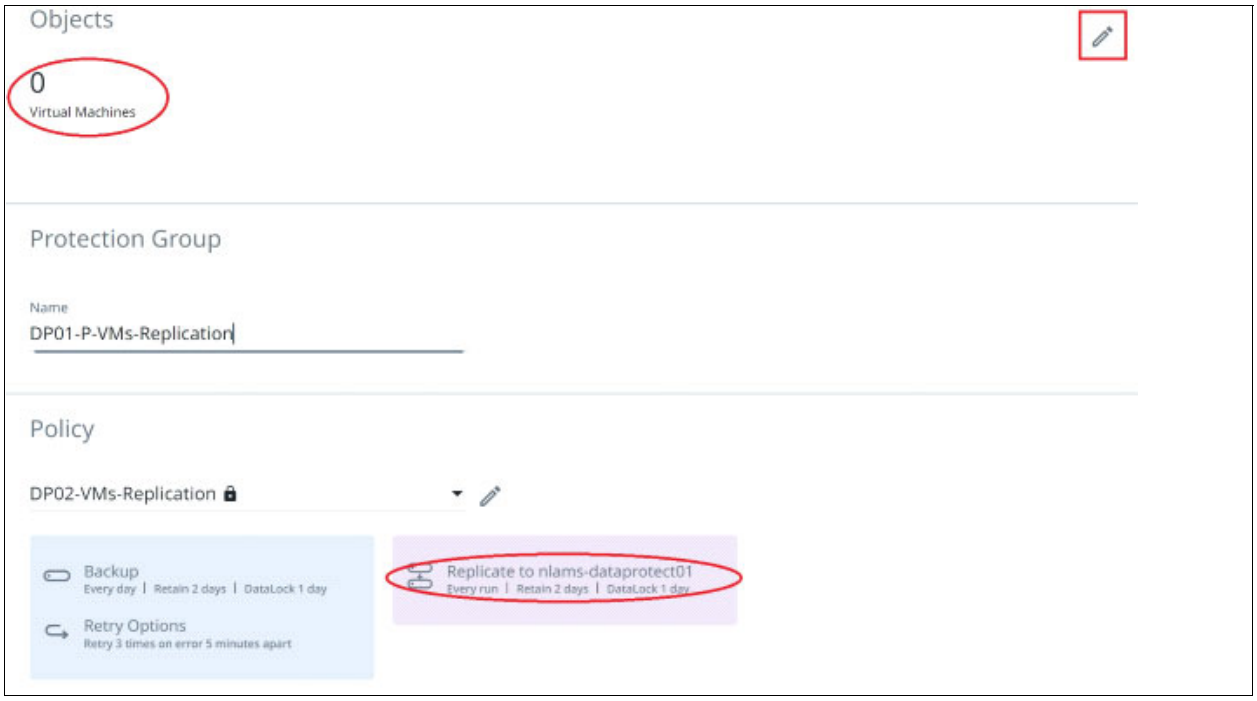


Figure 4-28 Failover Protection Group for continuous protection: Editing a workload

Note: If a failback to the original cluster is required, select **Deactivate** on the Protection Group for the source cluster.

4.6.1 Preparing a failover back to the original source cluster

When the original source cluster is returned to an available and accessible status, the system can be reverted to running on the primary cluster by clicking **Deactivate** on the backup cluster. If the backup Protection Group is not deactivated, the local backup job continues to run successfully, but the replication tasks end in an error state (Figure 4-29).



Figure 4-29 Fallback Protection Group for continuous protection: Replication error

In this example, the backup finished successfully with 287.2 MB of incremental data written (Figure 4-30).



Figure 4-30 Fallback Protection Group for continuous protection: Local backup status

If the original source cluster is running, then you can click **Deactivate Only** to make the Protection Group Failover Ready (Figure 4-31).



Figure 4-31 Fallback Protection Group deactivate option

Once you click Deactivate, you see the following options for the Fallback cluster (Figure 4-32).

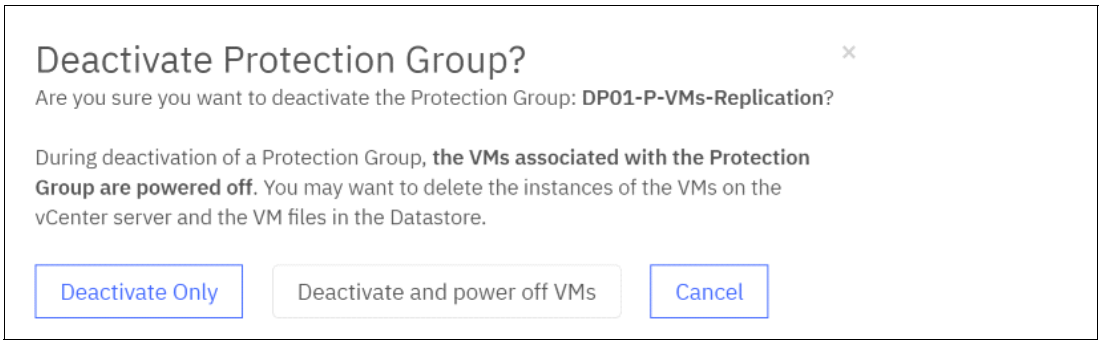


Figure 4-32 Fallback Protection Group for continuous protection: Deactivate Only confirmation

Note: The **Deactivate and power off VMs** option is valid only if there is an Active/Passive data center or a vCenter on each data center.

After you click **Deactivate Only**, resynchronization begins, as shown in Figure 4-33 on page 53.



Figure 4-33 Fallback Protection Group for continuous protection: Re-sync after Deactivate Only

In this example, once the replication task finishes, the resynchronization is complete. The details show that the replication task completed within 4 m 57 s and replicated 1.7 GB of data (Figure 4-34).

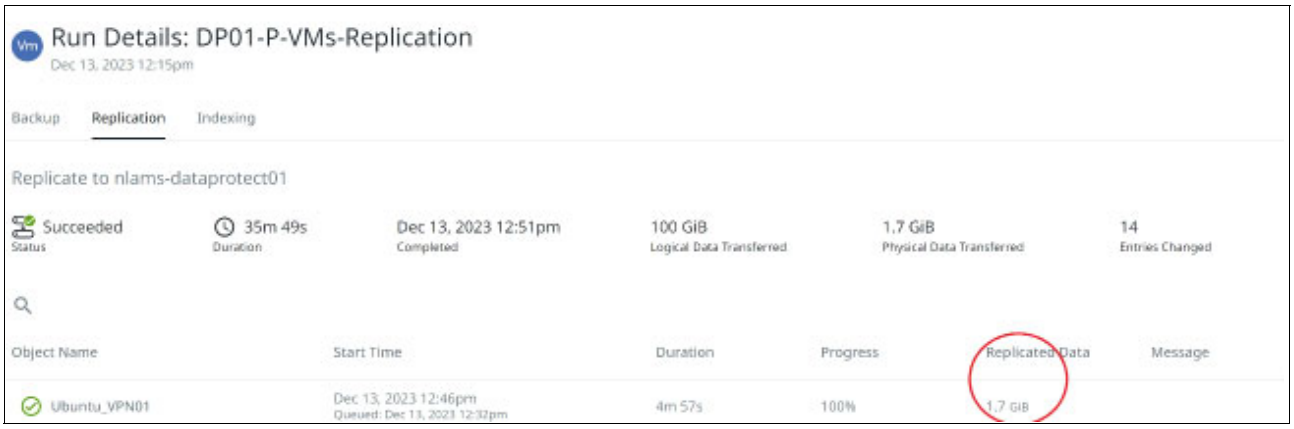


Figure 4-34 Fallback Protection Group for continuous protection: Re-Sync incremental data written

Now that replication is complete, you see 1.7 GiB as replicated data.

4.7 Archiving workloads with CloudArchive and IBM Storage Defender DP cluster

In today's technology-driven world, it is a major challenge having to store, protect, and maintain the availability of the ever-growing amount of data that is created, such as sensitive HR data and healthcare records to large media and entertainment files. As many organizations store their growing data collections on network-attached storage ((NAS) or GPFS) and run most of their workloads on a VMware infrastructure, they should focus on NAS and VMware data protection, that is, how to store, protect, and retain all that data.

CloudArchive Direct is a policy-driven archival feature in IBM Storage Defender DP that addresses these challenges by streaming data directly to lower-cost storage on an external target without storing local backups. Although IBM Storage Defender DP does not store the data, it indexes it and stores the metadata locally for fast search and recovery while offering options to compress and encrypt the data.

IBM Storage Defender DP CloudArchive Direct can help you reliably archive data set securely to a supported S3 object store, which includes on-premises IBM Cloud Object Storage.

4.7.1 VMware CloudArchive workflow S3 to IBM Cloud Object Storage

This section describes the CloudArchive workflow of VMware and provides an example of a one-to-one configuration by using IBM Storage Defender DP CloudArchive Direct. The backup data originates from the source VMware vCenter on the primary host, and then it is ingested by the primary IBM Storage Defender DP cluster DC1. This data is replicated to a secondary (DR) cluster DC2. In this configuration, there is also an optional step of using the on-premises IBM Cloud Object Storage erasure coding (EC) process solution for long-term data retention.

Figure 4-35 shows the example one-to-one CloudArchive workflow.

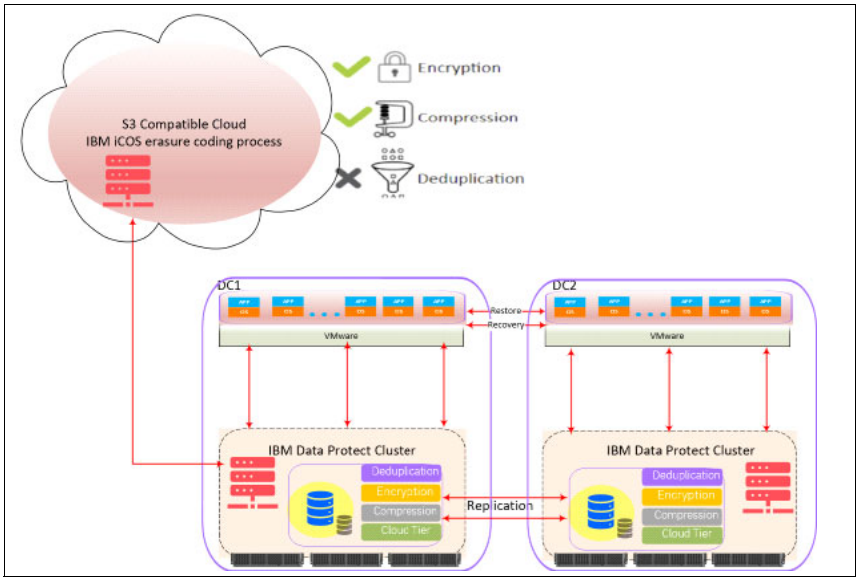


Figure 4-35 One-to-one CloudArchive workflow

Registering external targets as S3 compatible

To register an external target, you need the `ip-address` of the endpoint (IBM Cloud Object Storage Accessor), the Bucket name, the administrator Access Key ID, and the Secret Access Key.

From the **Infrastructure** tab, select **External Targets**, and then select **Add external target**. For this example, we completed the External Target fields as shown in Figure 4-36 on page 55:

- Purpose: Archival
- Storage Type: S3Compatible
- Storage Class: Regular
- Bucket Name: COSbucket

- Access Key ID: COSkeyID
- Secret Access Key: xxxxxxxxxxxxx
- Endpoint: 10.0.2.x
- AWS Signature Version: Ver 4
- External Target Name: AnyName
- Archival Format: Incremental with Periodic Full

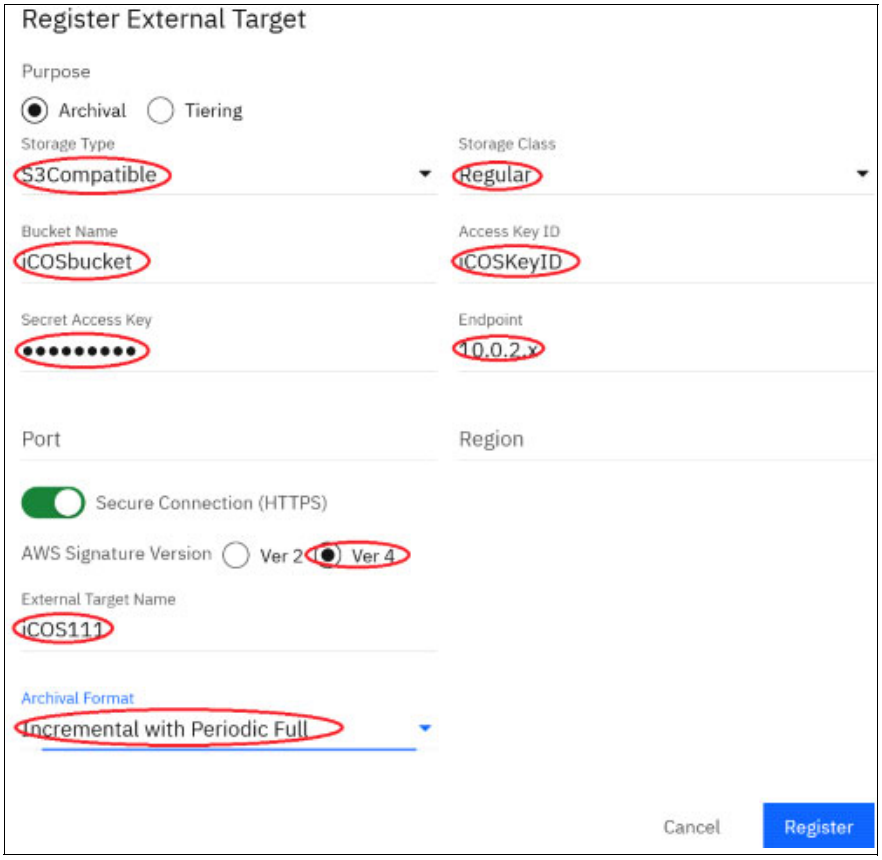


Figure 4-36 Registering an external target: IBM Cloud Object Storage configuration example

Select **Register** to complete the registration process and finish creating the external target.

Incremental with Periodic Full Archive

The first archive of a Protection Group is always a Reference Full Archive, which forms the basis for Incremental Archives. With Incremental Archives, the next archive after a reference full is an Incremental Archive that designates the Reference Full Archive as its reference.

An Incremental Archive contains only data that is different between its reference's snapshot and the snapshot that is being archived. Once initiated, the changed blocks are identified and any changed data blocks are compared to the reference snapshot. Only the necessary data is sent to the external target. If deduplication is enabled and data blocks have not changed, only metadata is transferred.

To reconstruct this snapshot of an Incremental Archive, both the incremental and the reference full must be processed. Then, the next Incremental Archive is an increment on top of the previous Incremental Archive, that is, it designates the previous Incremental Archive as its reference. Incremental Archives form an “archival chain” that starts at the selected archive and ends with a Reference Full Archive.

To limit the number of archives that must be processed to restore a snapshot, the chain is “broken” every 90 days. When the threshold is reached, the next archive will not be an Incremental Archive, but rather a periodic full archive. A regular, non-Reference Full Archive is full in the sense that it covers all previous increments since the Reference Full Archive. However, the periodic full archive still refers to the Reference Full Archive as its base, similar to the initial Incremental Archive.

The next Incremental Archive after a full archive uses the full archive as its reference (Figure 4-37).

The threshold is not configurable through the web interface, and it is always 90 days regardless of how many archives are performed.

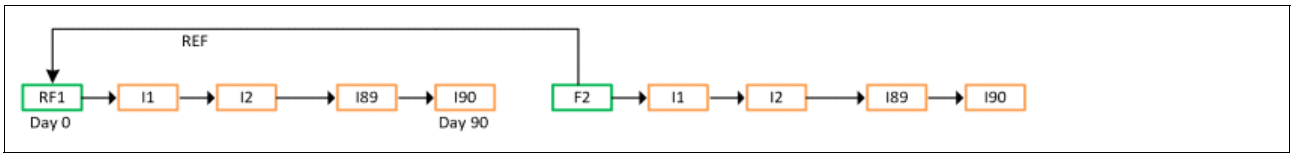


Figure 4-37 Incremental Archive reference workflow: Using Reference Full Archive

During Incremental Archives, the current Reference Full Archive might become too inefficient, that is, its utilization as a reference drops below 50%. If this situation occurs, the Reference Full Archive is retired. Then, a new Reference Full Archive is created (Figure 4-38).

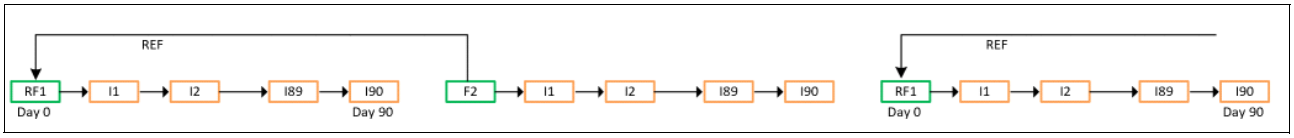


Figure 4-38 Reference Full Archive workflow: New base Reference Full Archive

Retiring a Reference Full Archive means that the next time a Full Archive must be created, it is created as a Reference Full Archive rather than a regular non-Reference Full Archive. The next Incremental Archive uses the new Reference Full Archive as its reference base.

CloudArchive Direct Incremental Forever

Enabling this option performs a Full Archive and then continues with Incremental Forever archival of data from the IBM Storage Defender DP cluster to an external target. Incremental Forever archival eliminates the need for periodic Full Archive and supports global deduplication, which improves utilization of the target storage location.

This option is enabled by default when registering a new external target. To enable this option for an existing registered external target, edit the registered target to enable this option. The migration from Incremental Archival to Incremental Forever Archival takes place during the next reference archive creation.

Note: As a best practice, use Incremental Forever and turn on Source-Side Deduplication. This approach increases cloud performance and reduces the amount of data that the IBM Storage Defender DP cluster transfers to the external target.

Figure 4-39 shows the Default settings when registering an external target when you use CloudArchive Direct Incremental Forever.

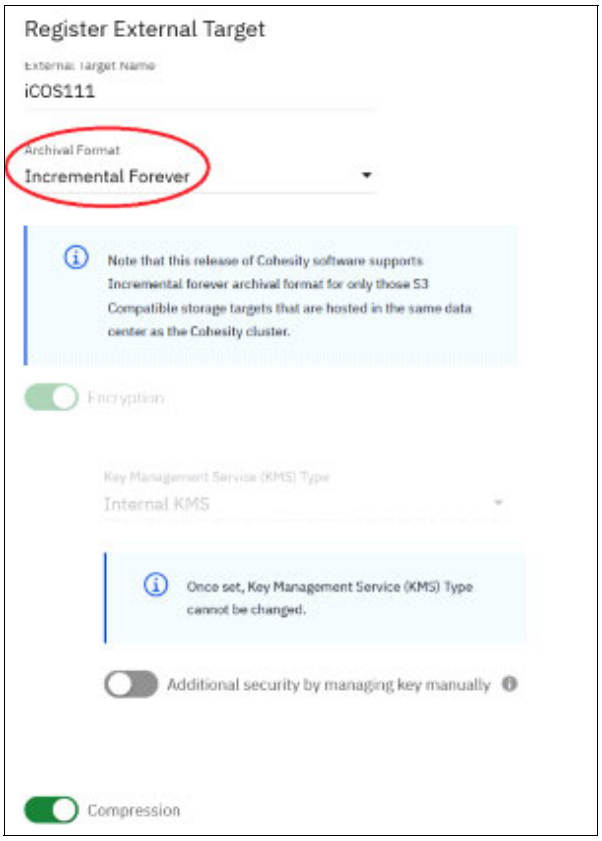


Figure 4-39 Default settings when registering an external target

4.7.2 VMware CloudArchive restore from IBM S3 IBM Cloud Object Storage

The following section provides an example of selecting and restoring an archive from IBM S3 IBM Cloud Object Storage while using VMware CloudArchive. To perform the restore, select the **Data Protection** tab, select **Recoveries**. and then select **Recover**. From here, filter the selections to find the VM archives.

Select the VM from the list to restore, then click the pencil icon to specify the location of the cloud and the recovery point, as shown in Figure 4-40.

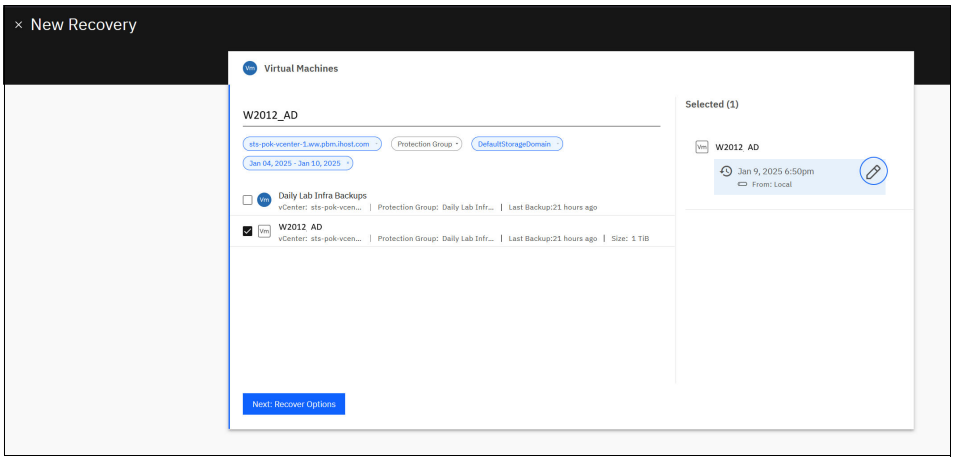


Figure 4-40 Virtual Machine restore settings pencil icon

Figure 4-41 shows the point in time recovery from the cloud location.

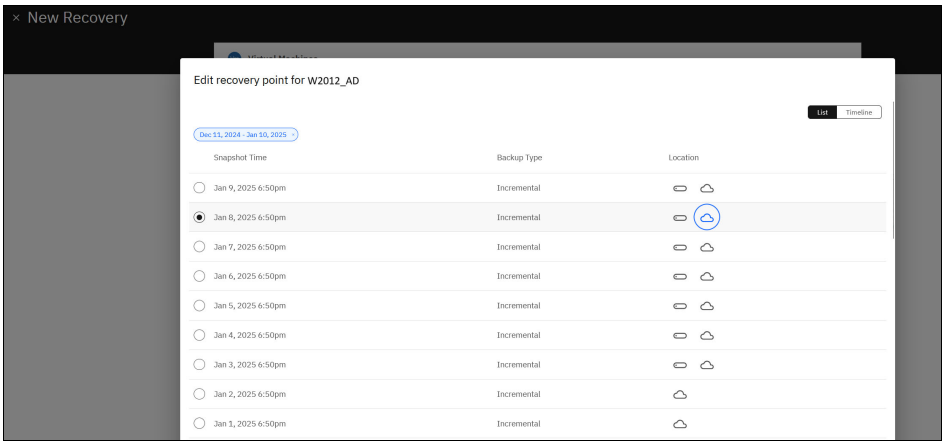


Figure 4-41 Point in time recovery from cloud location

Once the archive is chosen and selected, you can use the Recovery Method prompt to choose your recovery methods and options. For this example, we used **Instance Recovery** and **Keep Existing VM**. This option powers off and renames the existing VM as `_obsolete_xxx` (Figure 4-42 on page 59).

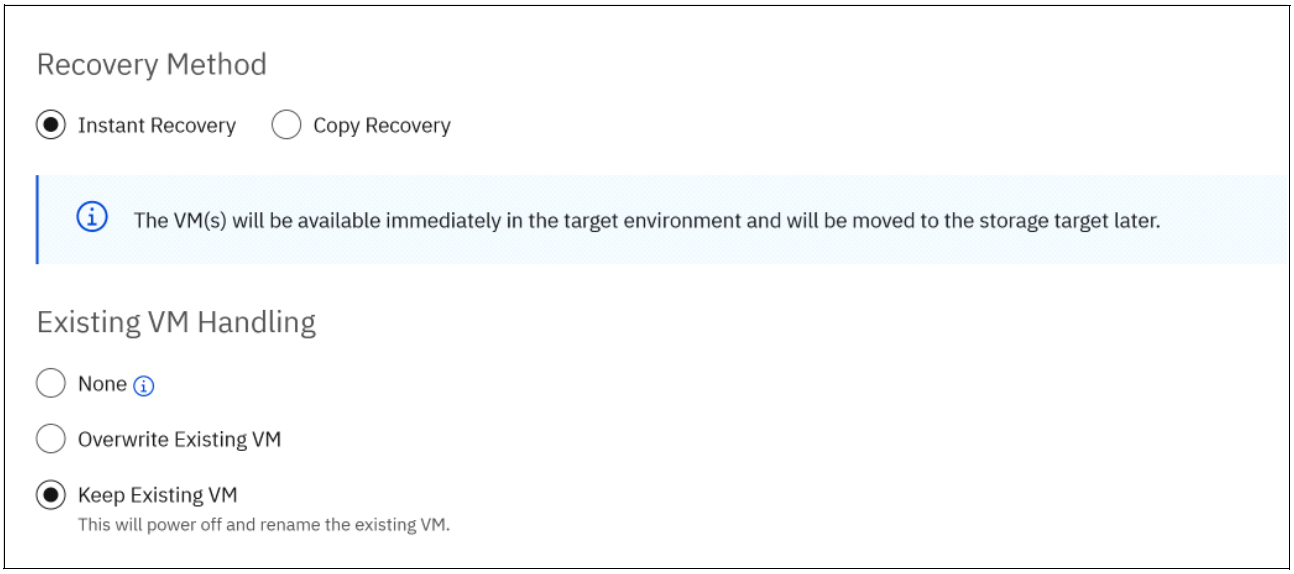


Figure 4-42 Recovery Method in CloudArchive

In the activity log for the recovery process, you can see that the recovery is from the IBM Cloud Object Storage location, which is indicated by the kListVaultFSfiles task (Figure 4-43).

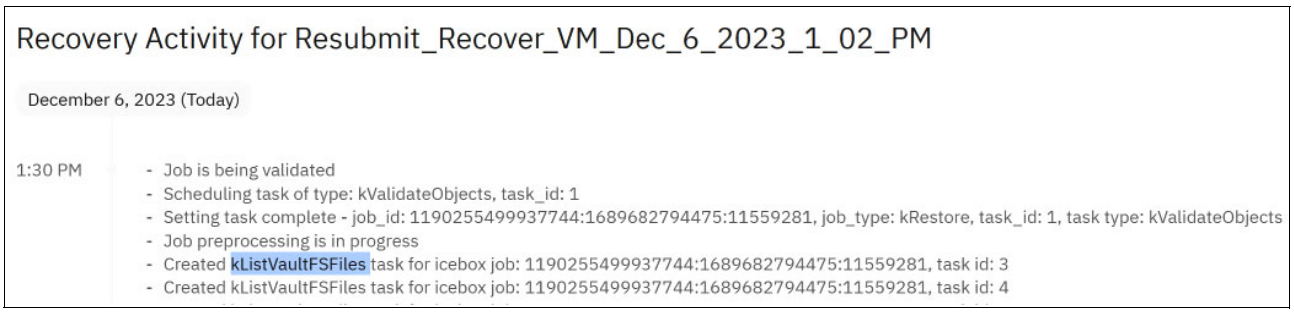


Figure 4-43 Recovery Activity with CloudArchive task creation details

Figure 4-44 shows a summary of the Recovery Activity.

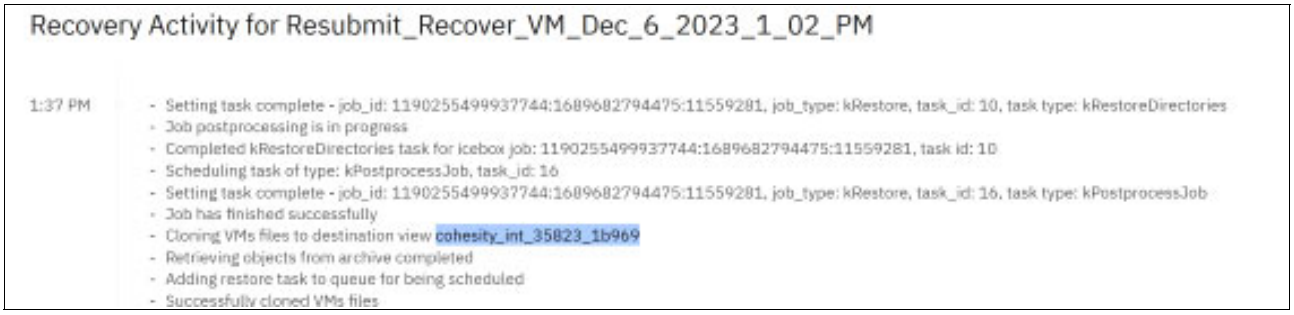


Figure 4-44 Recovery Activity with CloudArchive full action summary

The recovery snapshot is exported, and you can see it in the NFS view for the vCenter data center. The VM is then relocated to the original data store. After the VM is successfully cloned, the NFS view is removed (Figure 4-44 on page 59).

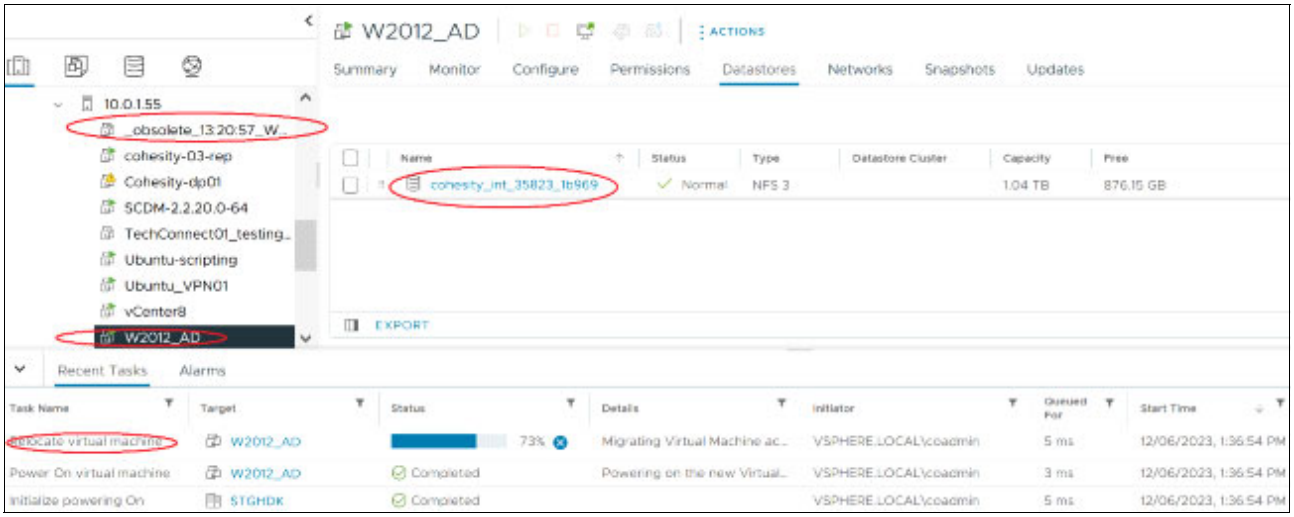


Figure 4-45 Instant recovery CloudArchive to vCenter

From the vCenter tasks, you can monitor the VM creation and VMotion processes for the recovered VM. Once the tasks are complete, the recovery process is complete.

5

Rapid restore of virtual machines

In Chapter 4, “Protecting VMware workloads with IBM Data Protect” on page 33, you learned about the various features and capabilities that IBM Storage Defender Data Protect (DP) offers for the protection of VMware virtual machine (VM) workloads. This chapter explores the different VMware recovery options that are available with IBM Storage Defender DP.

This chapter describes the following topics:

- 5.1, “Virtual machine recovery” on page 62
- 5.2, “File or folder recovery” on page 69
- 5.3, “Virtual disk recovery” on page 77
- 5.4, “Instant Volume Mount” on page 82

© Copyright IBM Corp. 2025.

5.1 Virtual machine recovery

By using IBM Storage Defender DP, you can recover one or more protected VMs from snapshots that you created previously by using an IBM Storage Defender DP Protection Group. (For more information about protecting VMware workloads, see Chapter 4, “Protecting VMware workloads with IBM Data Protect” on page 33). You can directly recover VMs from an IBM Storage Defender DP cluster or from a currently registered archival external target, and the destination can be the original location or a new location.

5.1.1 Recovering VMs

To recover a VMware VM from a IBM Storage Defender DP backup, complete the following steps:

1. In the left pane of the DMS GUI, select **Data Protection** → **Recoveries**, as shown in Figure 5-1.

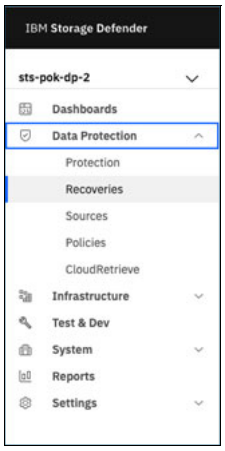


Figure 5-1 IBM Storage Defender Data Protection menu

2. Use the **Recover** menu at the upper right of the **Recoveries** window and select **Recover** → **Virtual Machines** → **VMs**, as shown in Figure 5-2.

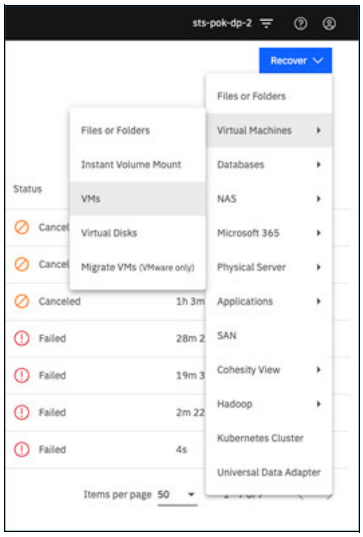


Figure 5-2 Virtual Machines Recovery window

3. In the New Recovery window, enter a wildcard search term to search for individual VMs or protection groups, as shown in Figure 5-3.

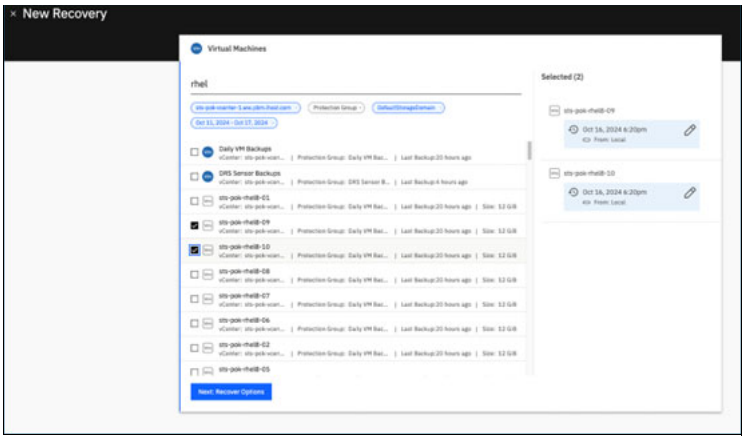


Figure 5-3 New recovery VM search window

Note: For each VM selected, IBM Storage Defender DP automatically selects the most recent snapshot. To select an earlier snapshot for any of the selected VMs, click the pencil icon for that VM in the right column, and then use the dialog box to select the snapshot, as shown in Figure 5-4.

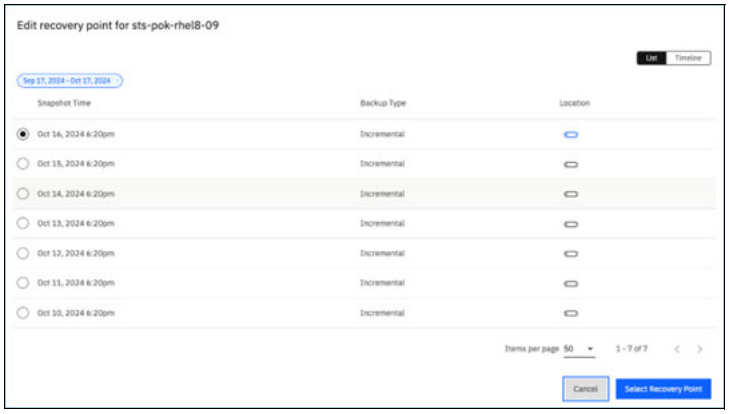


Figure 5-4 Alternative restore point selection window

Note:

- Selecting a Protection Group (In our example, “Daily VM Backups” or “DRS Sensor Backups”) restores all the VMs that are present in that Protection Group at the time of the selected backup.
- In a single restore job, only VMs or Protection Groups with snapshots in the same storage domain and data protection source can be selected.
- In this example, there are only local (on-cluster) snapshots that are available for recovery. If the example IBM Storage Defender DP cluster created archive copies on external targets like cloud storage or IBM Spectrum Protect, those snapshots are indicated by cloud icons under the Location column. Click the cloud icon to recover from an archival snapshot.

4. Once you have selected the snapshots, click **Next: Recover Options** to proceed.
5. In the next window, note the summary of the selected VMs at the top of the page, as shown in Figure 5-5. If necessary, click the pencil icon to edit your selection.

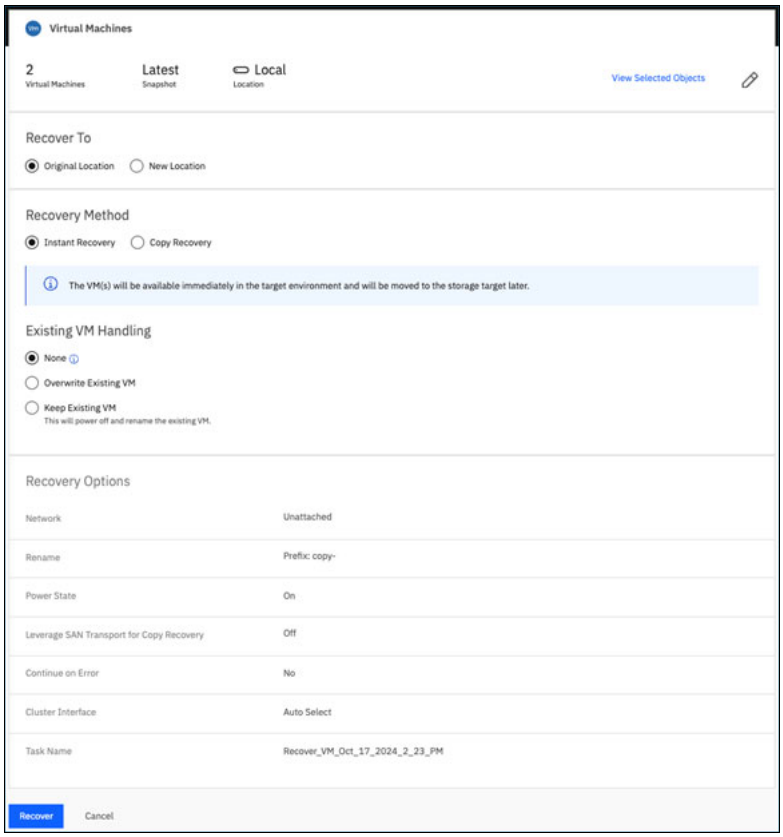


Figure 5-5 Virtual machine recovery summary window

The Recover To value is set to **Original Location**. In our example, the two selected VMs are restored to their original location. To recover the selected VMs to an alternative location, select **New Location**, then use the **Registered Source** drop-down menu to either select from an existing registered ESX host or vCenter, or register a new one, as shown in Figure 5-6.

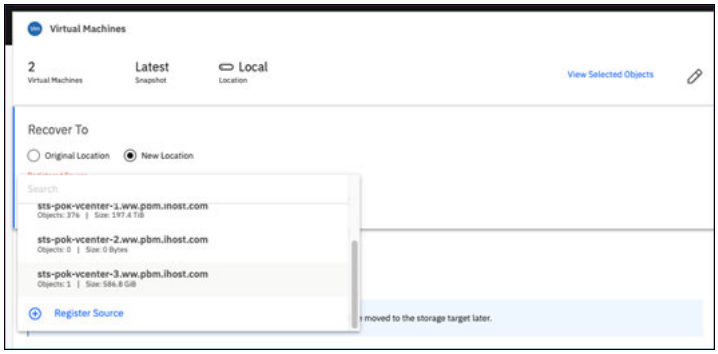


Figure 5-6 Virtual Machine recovery setting window

The Recovery Method is set to **Instant Recovery**.

Note: When recovering VMs, two recovery methods are supported:

- Instant Recovery: The fastest method to get the VMs running. In Instant Recovery, the IBM Storage Defender DP cluster creates a writeable snapshot of the immutable snapshots that were selected in step 4 on page 64, and mounts it as temporary NFS data stores in the target VMware environment. The VMs that are present in this data store are then registered to the target vCenter and powered on. After the last VM in the recovery job is powered on, then IBM Storage Defender DP orchestrates the storage vMotion of the VMs from the NFS data store to the target data store. Once the vMotion tasks complete, the IBM Storage Defender DP cluster unmounts the temporary data store. The benefit of Instant Recovery is that the VMs that are recovered can be powered on before any data is moved from the backup to the target data store. Instant Recovery provides the foundation for Instant Mass Restore (IMR), which is described in 5.1.2, “Instant Mass Restore” on page 68.
- Copy Recovery: The fastest method to get the VM files out of the backup system and back to the VMware datastore. It copies the VM files from the IBM Storage Defender DP cluster to the VMware target data store, and only once the VM files are fully restored to the data store can the VMs be powered on. To reiterate, Copy Recovery provides a faster method to move the files back to the data store than Instant Recovery, but the VMs cannot be powered on until the VM files are fully restored.

When you perform a VM restore and select **Original Location**, IBM Storage Defender DP offers multiple options for handling existing VMs in the original location:

- **None**: VM recovery fails if there is an existing VM in the recovery location.
- **Overwrite Existing VM**:
  - For an Instant Recovery, the existing VM is deleted from the target system before the restoration.
  - For a Copy Recovery, IBM Storage Defender DP attempts a differential restore, where only the blocks that differ between the existing VM and the backup copy are restored from the backup. This approach can result in a faster overall restore time. An example use case for differential restore is recovering from a large-scale operating system patching error.
- **Keep Existing VM**: The existing VM is powered off and renamed before the restoration.

Note the available Recovery Options that are shown in Figure 5-7.

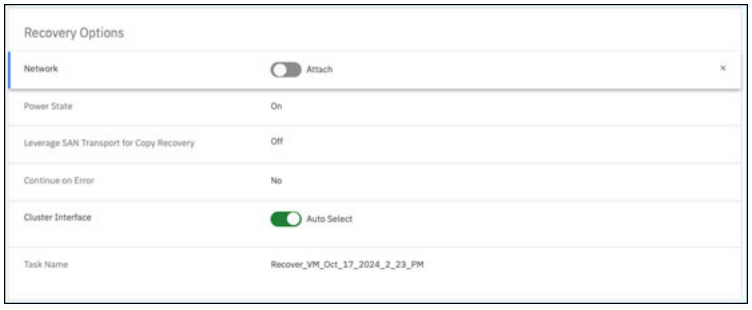


Figure 5-7 VM Recovery Options window

– Network:

  • For recovery to the Original Location, the default value is **Detached**, where all VMs power on detached from any network. If you select **Attached**, all VMs are powered on attached to their original virtual network. Also, selecting **Attached** unlocks the **Start Connected** option. If you select **Start Connected**, the VMs connect to the original networks when the VMs restart. If you clear **Start Connected**, the VMs do not connect to any virtual networks on restart.

Note: When recovering to the Original Location, IBM Storage Defender DP recovers the original MAC address for VMs that had manual MAC addresses at the time of the backup. If the VMs were set to have automatically generated MAC addresses at the time of the backup, the recovered VMs have the same setting.

  • For recovery to a New Location, the default value is **Detached**, where all VMs power on detached from any network. If you select **Attached**, a drop-down menu opens, where you can select the destination virtual network to attach the VMs. This network will be attached to all the VMs in the recovery job. To override this setting, click the **Add Network Override** link and select separate network settings for each VM in the recovery job.
  • Also, selecting **Attached** unlocks the **Start Connected** and **Preserve MAC Address** options. If you select **Start Connected**, the VMs connect to the selected network when the VMs restart. If the option is not selected, the VMs are not connected to any virtual networks after a restart. If you select **Preserve MAC Address**, VMs that were set to have manual MAC addresses are restored with their original MAC addresses.

Note: When recovering to a New Location, VMs that were set to have automatically generated MAC addresses at the time of the backup now have manual MAC addresses.

– Rename: You may enter a prefix or suffix to the names of the recovered VMs.

– Power State: The default value is **Power On**. VMs are powered on after restoration.

– Leverage SAN Transport for Copy Recovery: Instructs IBM Storage Defender DP to use SAN transport for data stores that are on Fibre Channel (FC) or iSCSI-based storage arrays. Enable this option to enable SAN transport mode over FC or iSCSI for the recovery of VMware VMs.

Note: SAN transport over FC requires that FC is configured and zoned between the IBM Storage Defender DP cluster and the target ESX hosts.

– Continue on Error: The default value is **No**. The recovery job fails at the first error when recovering a VM. If you select **Yes**, the protection job continues to run even if errors are encountered.

– Cluster Interface: Only available for Instant Recoveries. The default value is **Auto Select**. The IBM Storage Defender DP cluster automatically selects the correct interface group to use for the recovery. To manually select the Interface Group, click the **Cluster Interface** row, disable the **Auto Select** slider, and use the **Interface Group** drop-down menu to select the Interface Group to use for the recovery job.

– Task Name: Optionally, change the default name of the recovery job.

6. Click **Recover** to initiate the restore job. You see the **recovery summary** window, as shown in Figure 5-8.



Figure 5-8 Recovery Summary window

7. If the details look correct, click **Start Recovery** to initiate the restore. You are redirected to the Recoveries window, where you see the restore job running, as shown in Figure 5-9.

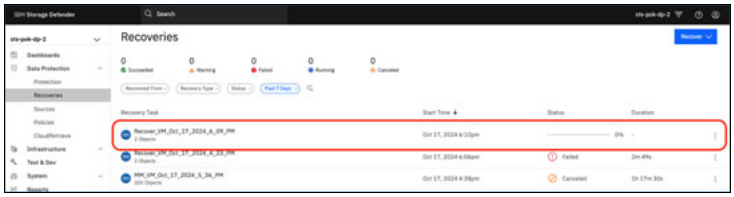


Figure 5-9 Recoveries details window

8. Click the recovery job to load its details, as shown in Figure 5-10.

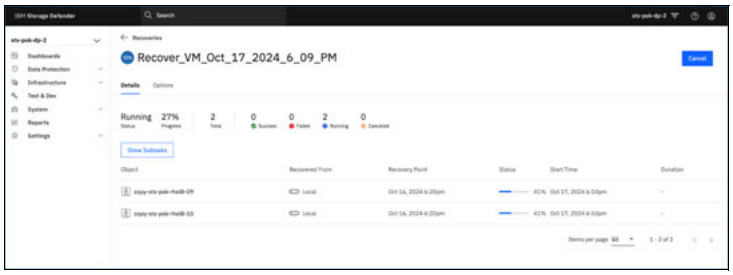


Figure 5-10 Recovery job details window

9. Click **Show Subtasks** to view the log of the recovery job. An example of the log is shown in Figure 5-11.

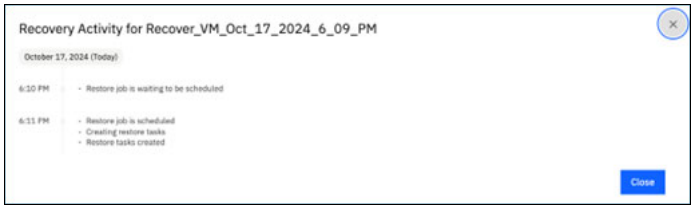


Figure 5-11 Recovery job log details

Figure 5-12 shows the successful completion status of the recovery job.

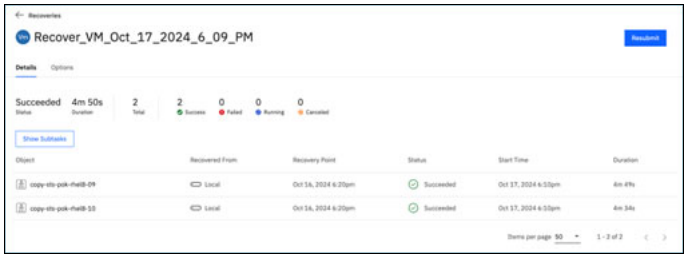


Figure 5-12 Recovery job completion status details

Figure 5-13 confirms in the vCenter that your VMs were successfully restored.

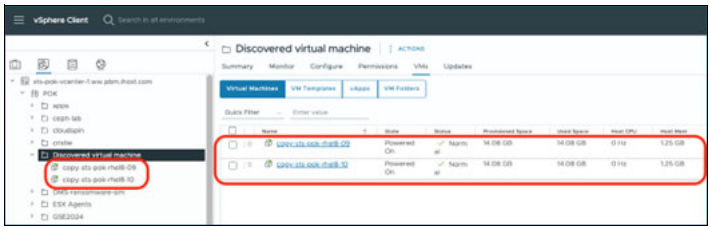


Figure 5-13 Restored VMs displayed in vCenter

Now that you explored the fundamental capabilities of VMware VM restoration, you can examine how IBM Storage Defender DP uses these capabilities along with its scale-out architecture to support large-scale, raid restore of VMs.

5.1.2 Instant Mass Restore

In this section, you select the VMs to recover in the recover task. In a recover task, you can select multiple VMs to restore. The ability to select more than one VM in a recovery task is called IMR. With IMR, you can use the Cohesity NFS views as the data store, allowing the recovery process to happen faster or instantly and circumventing the requirement of restoring the VMs to the original data store before starting.

Before you can recover Protected Objects (VMs), you must have a snapshot.

The key to recovery from a large-scale data corruption incident or other major intrusion is to bring the business services back online quickly. The ability to select multiple VMs in a recovery task and to recover them by using the Instant Recovery method is called IMR. In IMR, the IBM Storage Defender DP cluster performing the recovery follows the same process described in 5.1.1, “Recovering VMs” on page 62, but enables Instant Recovery at scale. IMR facilitates the rapid recovery of hundreds or even thousands of VMs.

Figure 5-14 on page 69 shows the use cases that we tested during a IBM Storage Defender DP proof of concept. It illustrates various different recovery scenarios, including multiple types of large-scale restoration in a VMware environment. This set of use cases highlights the power of IBM Storage Defender DP to drive workload restoration at scale and speed.

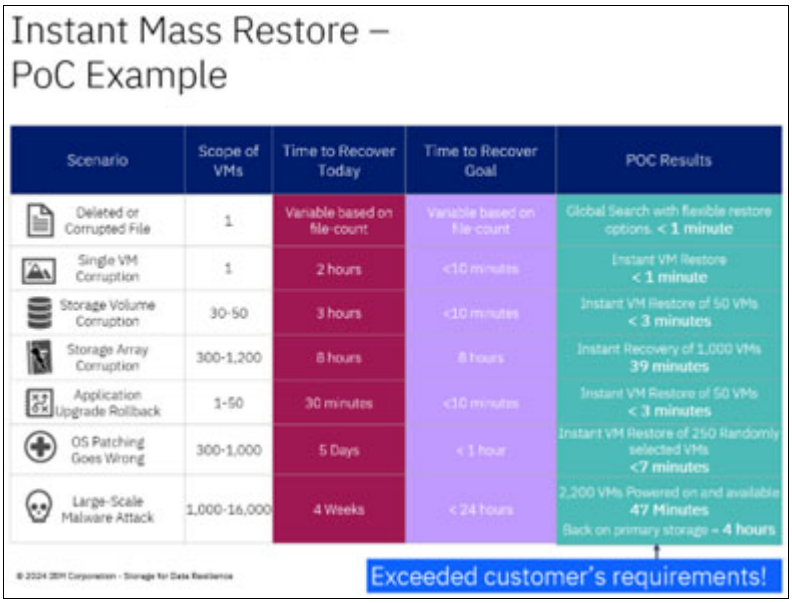


Figure 5-14 Instant Mass Restore proof of concept: Details

Figure 5-15 shows the recovery of 200 VMs.

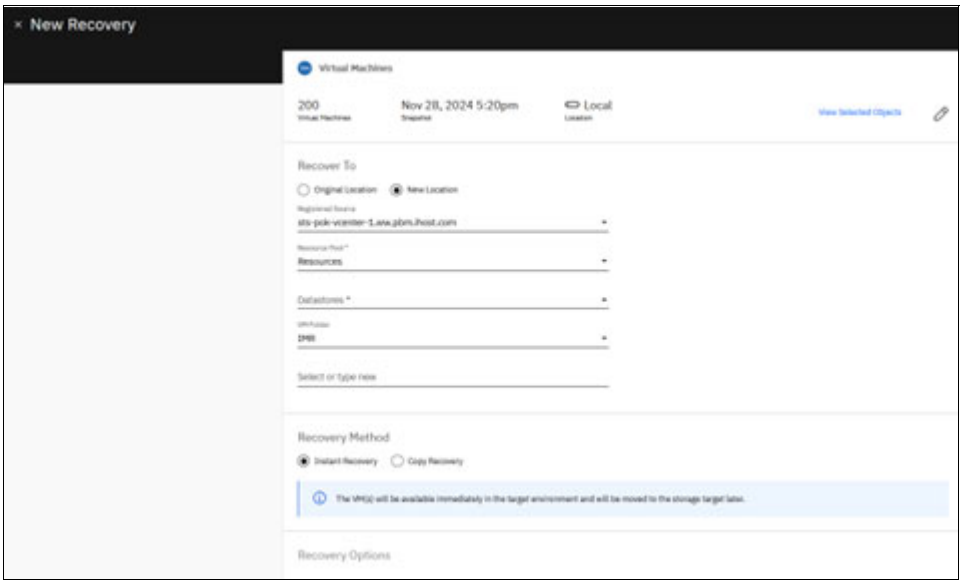


Figure 5-15 Instant Mass Restore finished restore example

5.2 File or folder recovery

In addition to full VM restoration, IBM Storage Defender DP can recover files and folders from a VM backup to either their original location or a new location. Alternatively, you can download the files and or folders that are in a IBM Storage Defender DP backup directly through the IBM Storage Defender DP GUI (DMS). Optionally, you can instruct IBM Storage Defender DP to preserve the file and folder permissions and attributes at the time of the backup.

5.2.1 Indexing and recovery

Fundamental to the ability of IBM Storage Defender DP to offer granular file and folder level restore capabilities from VMware VM backups is file system indexing. IBM Storage Defender DP by default attempts to index the file systems of the Windows and Linux VMs that it backs up. Although indexing is an optional parameter for VM backups, most users find that it is beneficial to leave it enabled because IBM Storage Defender DP can include this index data in a global catalog, which enables IBM Storage Defender users to use a single search interface for recovering all the file data across their infrastructure.

Note: IBM Storage Defender DP supports recovering files or folders from Windows VMs to Windows VMs, and from Linux VMs to Linux VMs only.

File recovery and Instant Volume Mount (IVM) from a IBM Storage Defender DP backup of a Windows VM with Windows deduplication that are enabled for one or more volumes is supported only when the target Windows machines have deduplication that is installed. (Deduplication does not need to be enabled on the target machine volumes, but only installed.)

5.2.2 Recovering files or folders

To recover files or folders by using IBM Storage Defender DP, complete the following steps:

1. Select **Data Protection** → **Recoveries**, as shown in Figure 5-16.

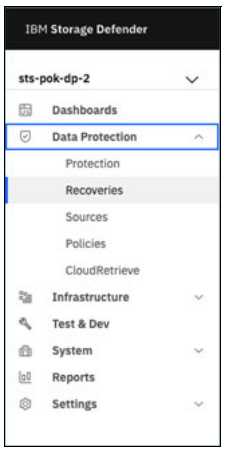


Figure 5-16 IBM Storage Defender Data Protection menu

2. Select **Recover** → **Virtual Machines** → **Files or Folders**, as shown in Figure 5-17 on page 71.

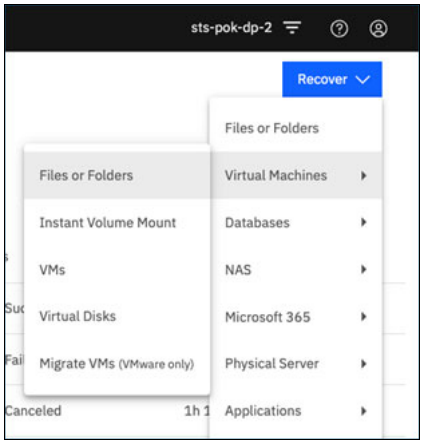


Figure 5-17 VM Recovery drop-down menu

3. In the New Recovery window, enter a wildcard search term to search for individual files or folders, as shown in Figure 5-18.

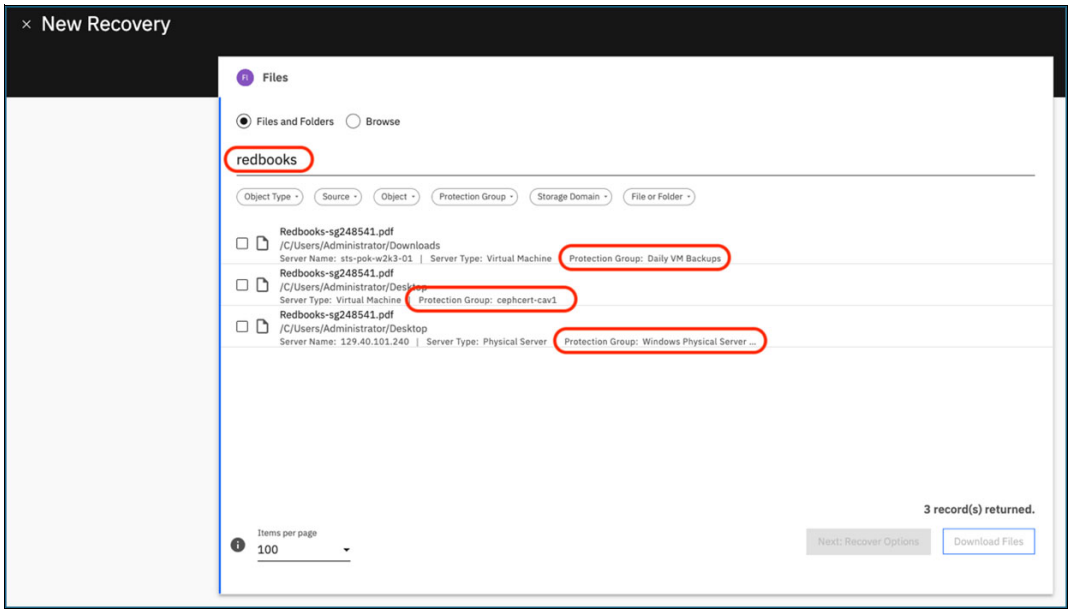


Figure 5-18 VM New Recovery file search window

Note: You can use the Files and Folders search feature to search across all Protection Groups within a DMS instance. In our example, the search term “Redbooks” returns three file results from three separate Protection Groups.

A single File and Folder Recovery job can restore multiple files or folders that were backed up as part of the same Protection Group. Use separate File/Folder Recovery jobs to recover files or folders that were backed up in separate Protection Groups.

4. Select the checkbox for your files or folders to add them to the Recovery Cart. When you select an item, the search results immediately filter the search results to show only items from the same Protection Group and Recovery Point, as shown in Figure 5-19.

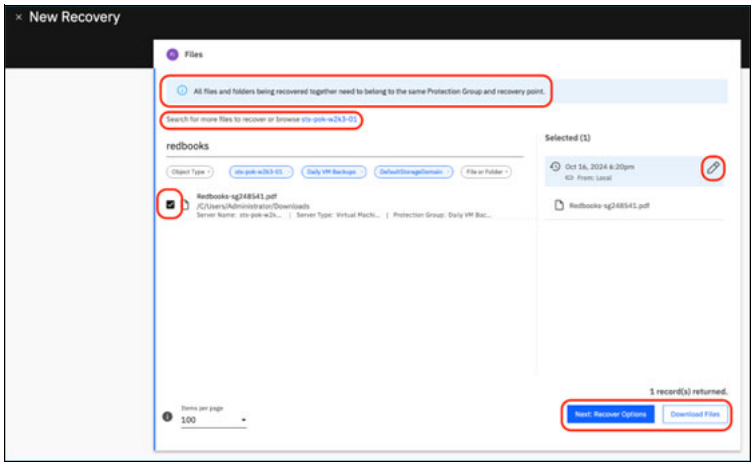


Figure 5-19 Files selection for recovery example

5. Optionally, click the pencil icon to select an earlier version. (Again, IBM Storage Defender DP defaults to the latest backup copy.)

Note: The option to select an earlier recovery point is available only when a single file is selected for recovery.

6. Optionally, click the link to the source VM, which in our example is **sts-pok-w2k3-01**. Figure 5-20 shows how to browse the index of backed up files from this VM and select more files or folders for recovery.

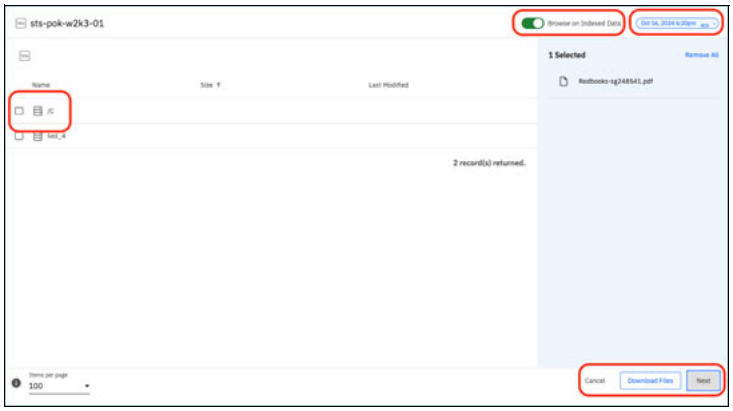


Figure 5-20 Browsing indexed data for a specific VM

Note the Browse on Indexed Data slider at the top of the page. By default, IBM Storage Defender DP presents an indexed (cached) representation of the file system of this VM at the date and time of the selected Recovery Point shown at the upper right. Browsing indexed data enables better search performance and displays only files and folders that are indexed. Hidden paths and manually excluded paths are not shown in this mode.

Turning off the slider triggers IBM Storage Defender DP to empty the Recovery Cart, mount the selected backup internally, and refresh the file browser to show all the files that are available in the selected Recovery Point.

If you click the Recovery Point timestamp, you can select an alternative Recovery Point, as shown in Figure 5-21.

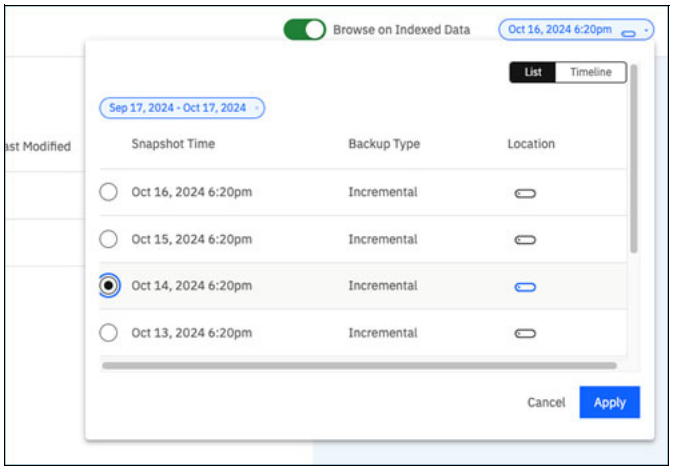


Figure 5-21 Alternative Recovery Point selection

In our example, we click **Cancel** on the Recovery Point selection window, and then click **Cancel** again to exit the file browser and return to the original Recovery Cart, as shown in Figure 5-22.

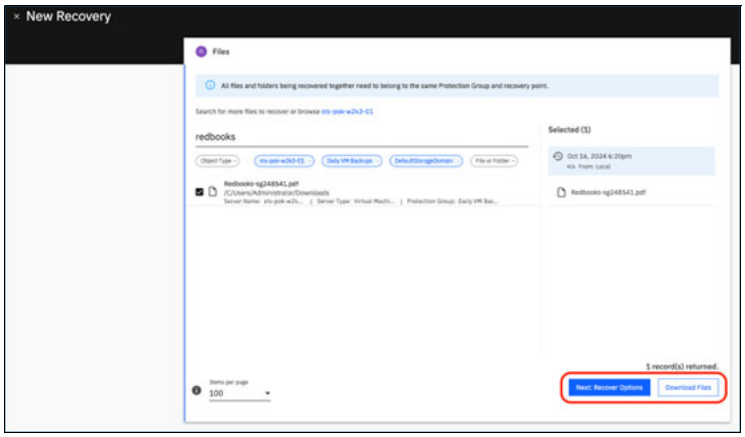


Figure 5-22 Return to recovery cart

7. From the original **Recovery Cart**, you can select either of the following options:
   - **Next: Recover Options** to proceed with recovering the selected items.
   - **Download Files** to download a compressed archive of the selected items directly through the browser.

In our example, we choose **Next: Recover Options**. The New Recovery page displays a summary of the selected items at the top of the window (Again, click the pencil icon to edit the items selected for recovery) followed by more recovery parameters and options, as shown in Figure 5-23.

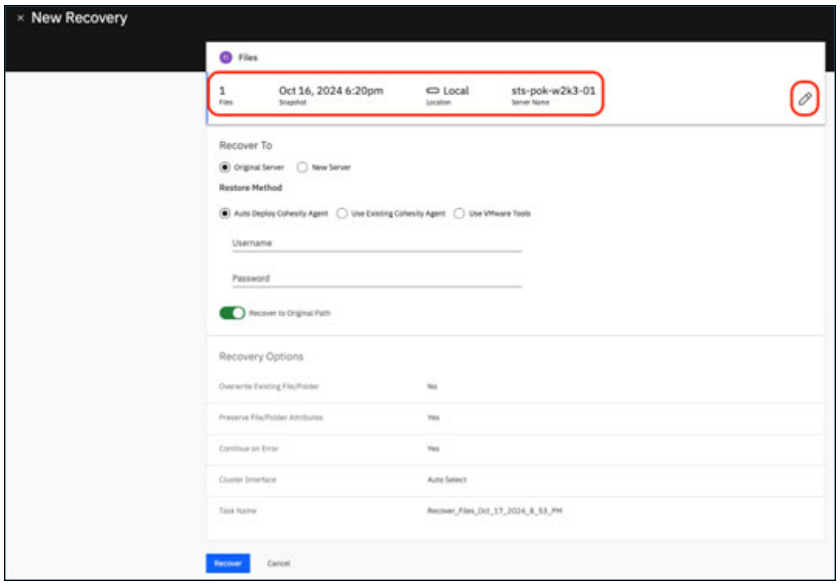


Figure 5-23 New recovery summary window

In the Recover To field, select either the **Original Server** or a **New Server** as the target for the file or folder restore. When the default selection, **Original Server**, is selected, the Recover to Original Path slider is set to **On**, as shown in Figure 5-24.

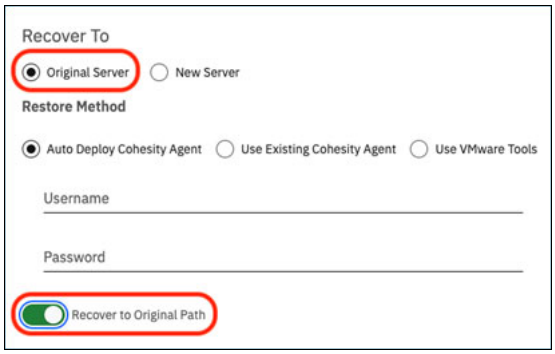


Figure 5-24 Recovery to original location settings dialog box

- To recover the items to an alternative path on the Original Server, turn off the slider and enter the path for restoration.
- To recover the items to a New Server, select **New Server**, and then use the **Source** and **Target** drop-down menus to select the destination ESX host or vCenter server and Target VM. Optionally, you can use the **Source** drop-down menu to register a new Source server.

In our example, we select the already registered vCenter server **sts-pok-vcenter-1.ww.pbm.ihost.com** as our Source, and then select the VM **sts-pok-rhel8-10-8** as the Target for restore, as shown in Figure 5-25 on page 75.

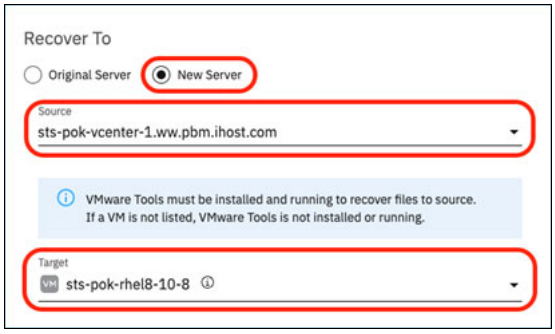


Figure 5-25 Recovery to new server settings window

Note: The term *Server* in this dialog box, whether Original or New, is a bit misleading. Here *Server* denotes the *Virtual Server* (VM) to which you want to recover the items.

Also, *Source* is used here in the context of registered IBM Storage Defender DP protection sources. When you select **New Server,** the **Source** drop-down menu loads the registered VMware data protection sources in this DMS instance (vCenter Servers or ESX hosts) and the **Target** drop-down menu dynamically loads based on the selected Source, which enables you to select the target VM to which you want to recover the items.

8. Regardless of whether the items are recovered to the Original Server or a New Server, you must select the Restore Method. Here are the available methods:
   - **Auto Deploy Cohesity Agent**: IBM Storage Defender DP deploys the Cohesity agent by using the supplied credentials, and then recovers the items to the Target VM by using the agent.
   - **Use Existing Cohesity Agent**: IBM Storage Defender DP uses the existing Cohesity agent on the Target VM to restore the items.
   - **Use VMware Tools**: IBM Storage Defender DP uses VMware Tools to perform the recovery of the items. VMware Tools must be installed and running on the Target VM.

   In our example, we know that VMware Tools are already installed on the original source VM, so we select **Original Server**, as shown in Figure 5-26. Select **Use VMware Tools**, enter the user credentials to access the Target VM, and select **Recover to Original Path**.

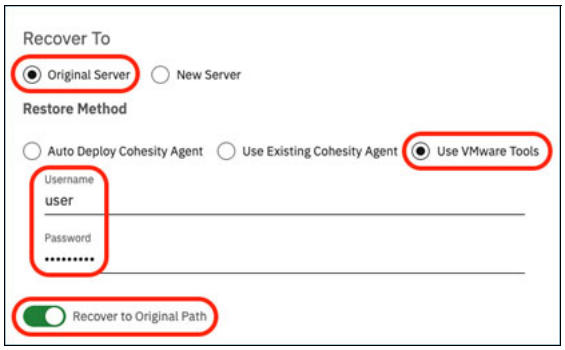


Figure 5-26 Recovery to original location settings window

IBM Storage Defender DP offers the following additional Recovery Options for VMware file and folder recovery:

- **Overwrite Existing File/Folder**: The default value is **No**; If you select **Yes**, existing files might be overwritten as part of the recovery job.
- **Preserve File/Folder Attributes**: The default value is **Yes**. By default, IBM Storage Defender DP preserves the ACLs, permissions, and timestamps for all recovered items. If you select **No**, the ACLs and permissions are not preserved.

Note: When recovering both folders and files, folders receive new timestamps and files are restored with their timestamps from the time of the backup. For file-only recovery, the files receive new timestamps.

- **Continue on Error**: The default value is **Yes**. The recovery job will on the first error encountered. If set to **No**, the recovery continues when errors are encountered.
- **Continue on Error**: The default value is **YES**. This option enables the recovery task to continue even if one of the objects encounters an error. If set to **NO**, the recovery operation fails if any of the objects that are being recovered encounters an error.
- **Cluster Interface**: The default value is **Auto Select**. The IBM Storage Defender DP cluster automatically selects the correct interface group to use for the recovery. To manually select the Interface Group, click the **Cluster Interface** row, disable the **Auto Select** slider, and use the **Interface Group** drop-down menu to select the Interface Group to use for the recovery job.
- **Task Name**: Optionally, change the default name of the recovery job.

9. To initiate the recovery job, click **Recover**. You are redirected to the Recoveries window, where you see the restore job running, as shown in Figure 5-27.

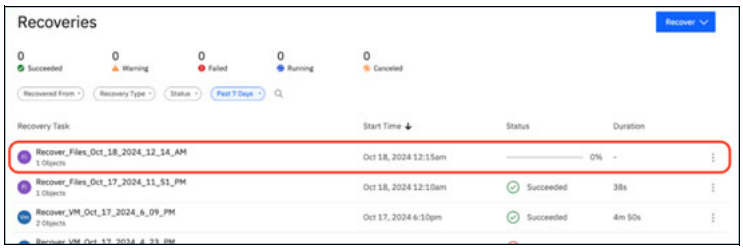


Figure 5-27 Recoveries status dashboard

10. Click the recovery job to load its details, as shown in Figure 5-28.

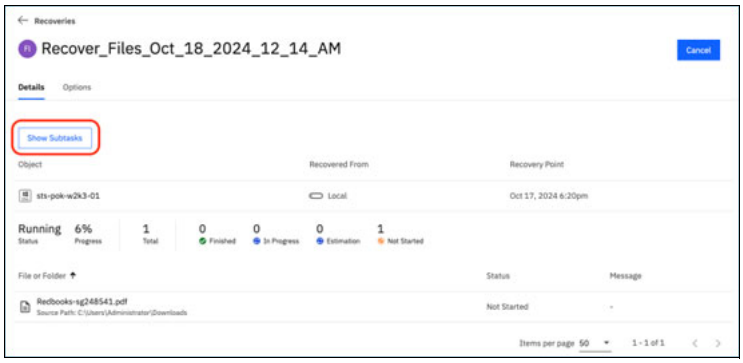


Figure 5-28 File recovery job details window

11. Click **Show Subtasks** to view the log of the recovery job, as shown in Figure 5-29.



Figure 5-29 File recovery log details

12. While the recovery process is running, you can monitor the Recoveries window to see the status of the job. Figure 5-30 shows the successful completion status of the recovery job.

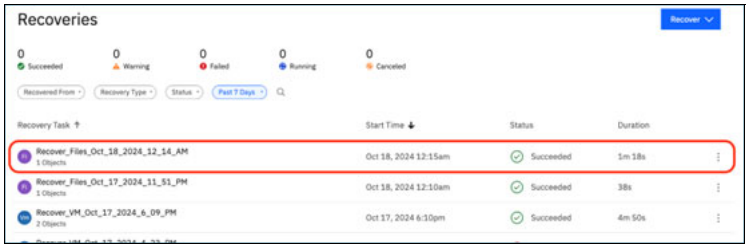


Figure 5-30 Recovery job status window

5.3 Virtual disk recovery

IBM Storage Defender DP supports the recovery of VMware VM virtual disks from snapshot backups that it creates. The virtual disks can be recovered to the original VM or to a different VM that is registered within the same vCenter.

The process of virtual disk recovery consists of the following steps:

1. IBM Storage Defender DP creates a writable clone of the selected Virtual Machine Disk (VMDK) file from the selected Recovery Point and mounts it to the target ESX hosts as a temporary NFS data store from the IBM Storage Defender DP cluster.
2. IBM Storage Defender DP attaches the new VMDK virtual disks to the target VM and initiates a Storage vMotion of the new VMDK from the temporary data store to the target data store.
3. The temporary data store is automatically removed on vMotion completion.

Note: After virtual disk recovery completes, you might need to perform other operating system-level actions to make the recovered disks online and available.

5.3.1 Recovering virtual disks

To recover a VMware virtual disk by using IBM Storage Defender DP, complete the following steps:

1. Select **Data Protection** → **Recoveries**, as shown in Figure 5-31.

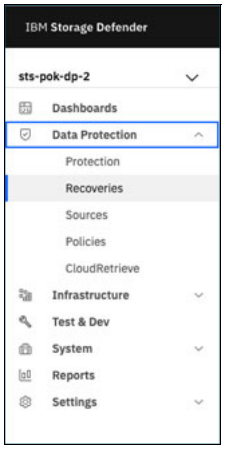


Figure 5-31 Storage Defender Data Protection menu

2. At the upper right of the Recoveries window, select **Recover** → **Virtual Machines** → **Virtual Disks**, as shown in Figure 5-32.

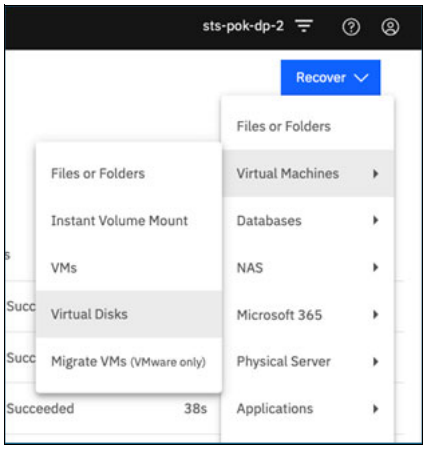


Figure 5-32 Virtual disks recovery option

3. In the New Recovery window, enter a wildcard search term to search for the VM whose disks that you want to recover, and then select the VM, as shown in Figure 5-33 on page 79. IBM Storage Defender DP defaults to the most recent Recovery Point. You can use the pencil icon to select an earlier Recovery Point if you want. In our example, we leave the default Recovery Point selected and click **Next: Recover Options**.

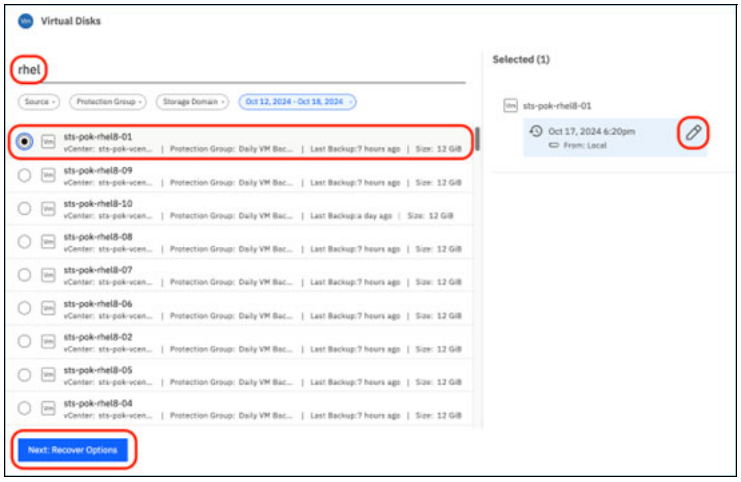


Figure 5-33 Virtual disk recovery search window

The New Recovery window (Figure 5-34) shows a summary of the selected source VM and associated Recovery Point. The Recover To option defaults to Original Location. Select **New Location** to restore the virtual disks to a different VM.

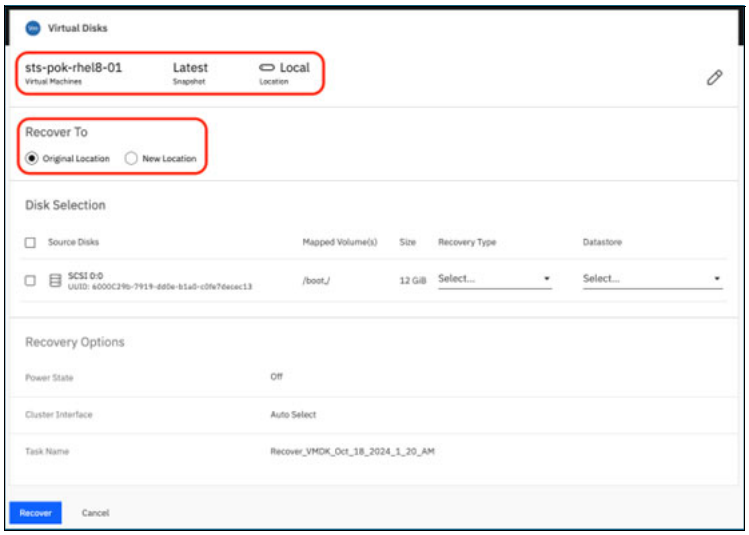


Figure 5-34 Virtual disks recovery options window

4. Next, select the boxes for the disks that you want to recover (Figure 5-35), and for each disk, select the recovery type.

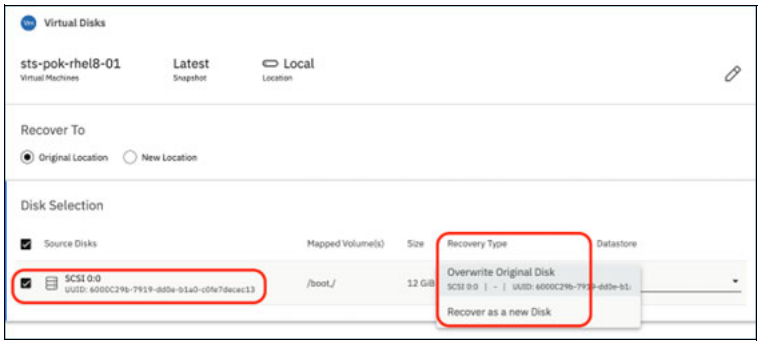


Figure 5-35 Virtual disks recovery window disk selection

Here are the available Recovery Types:

- **Overwrite Original Disk**: The original disk is overwritten in the original data store with the recovered virtual disk. As a best practice, quiesce the VM disk I/O before you perform the disk recovery.
- **Recover as a new Disk**: The recovered virtual disk is attached to the target VM as a new virtual disk. Choose this option if the virtual disk was deleted from the VM since the time of the backup.

If **Recover as a new Disk** is selected, you must select a target data store for the new disk with the **Datastore** drop-down menu, as shown in Figure 5-36.

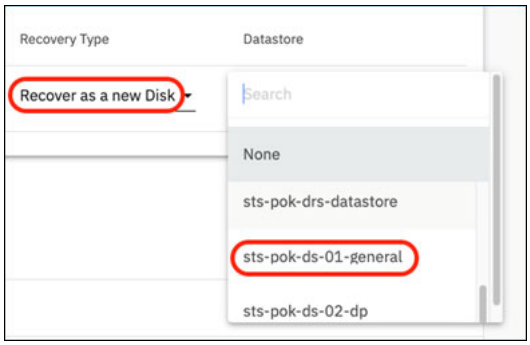


Figure 5-36 Virtual disk recovery window "Recovery as a new disk option"

Note: If recovery fails with **Overwrite Original Disk** selected, the original disk will not be available. If this situation occurs, retry the recovery to restore the disks.

For recoveries where one or more disks are overwritten, IBM Storage Defender DP automatically powers off the target VM before the recovery. Optionally, the target VM can be powered on automatically after successful recovery.

5. In our example (Figure 5-37 on page 81), we select **Original Location** and choose **Recover as a new Disk**, and then select **sts-pok-ds-01-general** as the target data store.

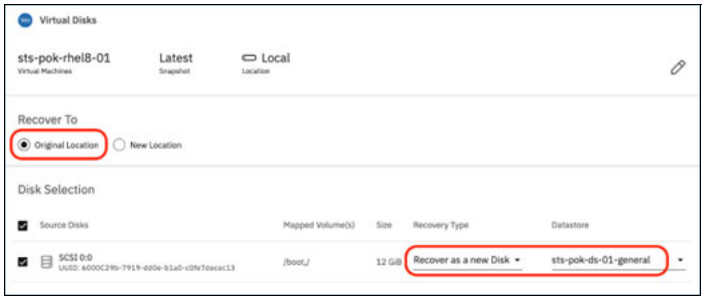


Figure 5-37 Virtual disk recovery options window

Here are the Virtual Disk Recovery Options:

- **Power State**: The default value is **Off**. The VMs are not powered off before the restoration.
- **Cluster Interface**: The default value is **Auto Select**. The IBM Storage Defender DP cluster automatically selects the correct interface group to use for the recovery. To manually select the Interface Group, click the **Cluster Interface** row, disable the **Auto Select** slider, and use the **Interface Group** drop-down menu to select the Interface Group to use for the recovery job.
- **Task Name**: Optionally, change the default name of the recovery job.

6. Click the **Power State** row under Recovery Options (Figure 5-38). The value changes from **Off** to a slider that is labeled **Power off VMs before Restore**.



Figure 5-38 Recovery options window power options

7. Turn off the **Power off VMs before Restore** slider (Figure 5-39). There is now an option that is named **Power on VMs after Restore**.



Figure 5-39 Recovery options window power options

Note: The Power State options can be helpful when recovering virtual disks to VMs that do not support “Hot Add” disk operations.

8. In our example (Figure 5-40), we revert to the default Power State options and click **Recover**. The browser redirects to the Recoveries window, and we can see our virtual disk restore job running.



Figure 5-40 Recoveries status window

9. Click **Show Subtasks** (Figure 5-41) to view the recovery job logs (Figure 5-42).

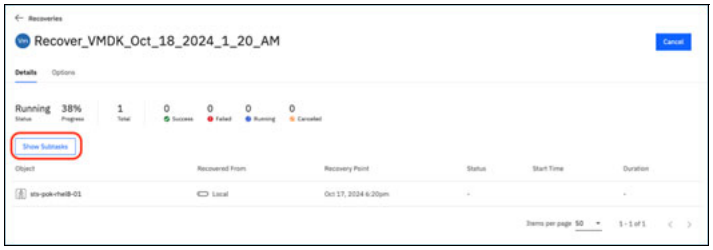


Figure 5-41 Recovery details window

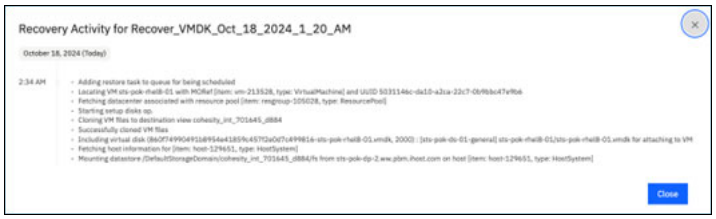


Figure 5-42 Recovery job log details

Figure 5-43 shows that the virtual disk recovery job completed successfully.

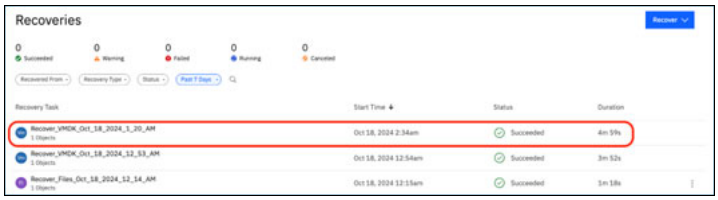


Figure 5-43 Recoveries job summary window

5.4 Instant Volume Mount

Use IVM to make a selected backup volume available to a target server target without needing any automated movement of the backup data from the IBM Storage Defender DP system. Common use cases include granular recovery of Microsoft SQL, Microsoft Exchange, and Microsoft Sharepoint databases, where the database or application administrator can perform granular recovery operations manually or by using scripts. IVM supports mounting to the original VM or a new VM.

Note: IVM is available only for backup volumes that are stored locally on the IBM Storage Defender DP cluster. IVM is not available from cloud or archive backup copies.

The default path of the recovered file or instantly mounted volume is `<Cluster_IP_address>/<Cloned view name>\<PATH>`. To access this location, both the target for the IVM and the IBM Storage Defender DP cluster must be in the same Active Directory domain.

5.4.1 Using Instant Volume Mount to mount a VMware virtual disk

To recover VMware virtual disk by using IBM Storage Defender DP, complete the following steps:

1. Select **Data Protection** → **Recoveries**, as shown in Figure 5-44.

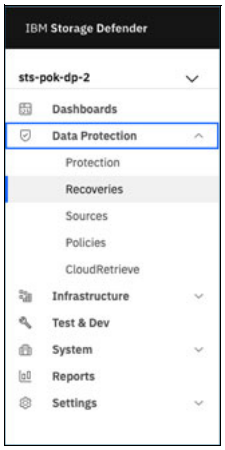


Figure 5-44 Storage Defender Data Protection menu

2. At the upper right of the **Recoveries** window, select **Recover** → **Virtual Machines** → **Instant Volume Mount**, as shown in Figure 5-45.

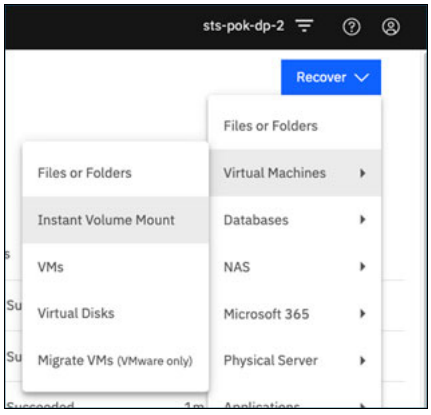


Figure 5-45 Instant Volume Mount option

3. In the New Recovery window, which is shown in Figure 5-46, enter a wildcard search term to search for the VM whose disks that you want to instantly mount, and then select the VM. IBM Storage Defender DP defaults to the most recent Recovery Point. You may use the pencil icon to select an earlier Recovery Point. In our example, we leave the default Recovery Point selected and click **Next: Recover Options**.

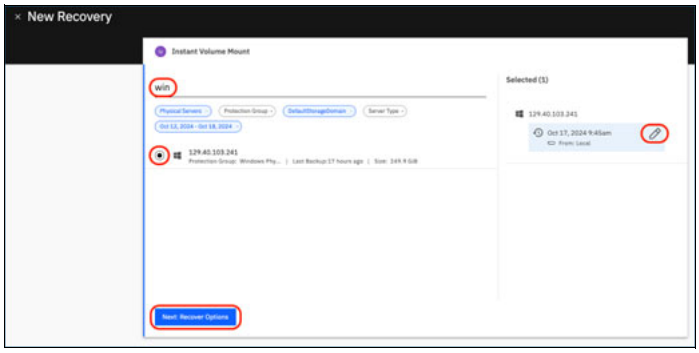


Figure 5-46 Recovery point options window

4. The New Recovery window shows a summary of the selected source VM and its associated Recovery Point, as shown in Figure 5-47. The Select Volumes slider defaults to **Off**, so you can select one or more specific volumes for instant mounting in the drop-down menu. Optionally, you can turn off the slider to select all volumes that were backed up with this VM at the time of the selected Recovery Point.

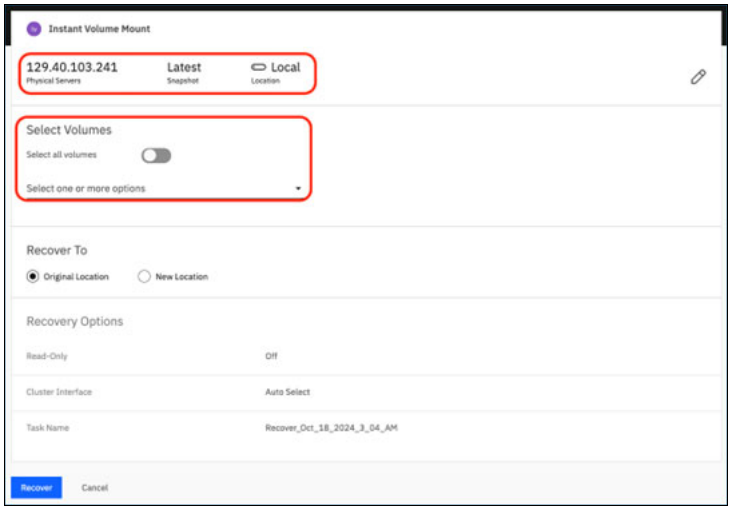


Figure 5-47 Instant Volume mount options window

5. The Recover To option defaults to **Original Location**. Select **New Location** to instantly mount the selected virtual disks to a different VM.

   Here are more Recovery Options for VMware IVM:

   – **Read-Only**: The default value is **Off**. The selected volumes are mounted to the target VM in read-only mode.

   – **Cluster Interface**: The default value is **Auto Select**. The IBM Storage Defender DP cluster automatically selects the correct interface group to use for the recovery. To manually select the Interface Group, click the **Cluster Interface** row, disable the **Auto Select** slider, and use the **Interface Group** drop-down menu to select the Interface Group to use for the recovery job.

– **Task Name**: Optionally, change the default name of the recovery job.

6. In our example (Figure 5-48), we choose to recover the F drive to a New Location.

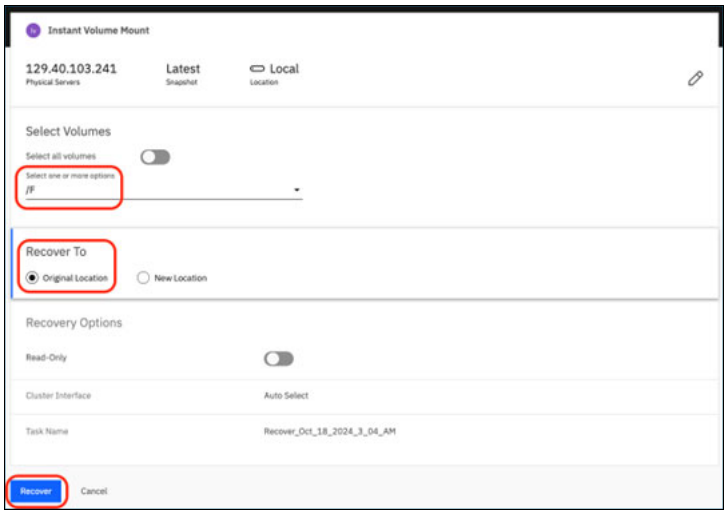


Figure 5-48 Instant volume mount recover to options

Figure 5-49 shows that the IVM job was successfully initiated.



Figure 5-49 Recoveries job log window

Once the mount job completes the volume will be mounted and accessible on the target VM.

5.4.2 Tearing down a mounted volume

Once the mounted volume is no longer needed, you can unmount it from the target host by selecting **Teardown** from the upper right menu, as shown in Figure 5-50.



Figure 5-50 Recovery window: Teardown option

To confirm that you want to unmount the volume, click Teardown, as shown in Figure 5-51.

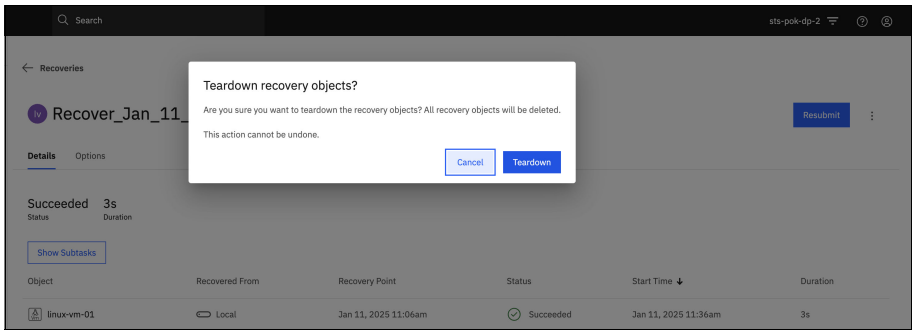


Figure 5-51 Teardown recovery object confirmation dialog

6

Integrating IBM Storage Defender DP with IBM Storage Protect

This chapter demonstrates how to configure IBM Storage Defender Data Protect (DP) to send archives to an external S3 target by using the S3 agent. The target in this example is an IBM Storage Protect server. This type of configuration provides a Glacier-like S3 to tape archive for longer-term retention use cases. We show use cases, examples about how to set up an IBM Storage Defender DP policy and a Protection Group, along with an example data retrieval.

This chapter describes the following topics:

- 6.1, “IBM Storage Defender DP overview” on page 88
- 6.2, “Deployment exercise overview” on page 89
- 6.3, “Creating an S3 bucket on an IBM Storage Protect server” on page 90
- 6.4, “Capturing access credentials on an IBM Storage Protect server” on page 93
- 6.5, “Registering the S3 bucket as an external target on a DP cluster” on page 93
- 6.6, “Creating policies to archive data to the S3 bucket” on page 96
- 6.7, “Monitoring the Data Protection status” on page 102
- 6.8, “Retrieving data from S3 external storage” on page 104
- 6.9, “Additional process monitoring examples” on page 108

© Copyright IBM Corp. 2025.

6.1 IBM Storage Defender DP overview

IBM Storage Defender DP supports sending data to external S3 targets. These S3 external target resources may be in a public cloud, private cloud, or on-premises S3 capable storage devices. For example, you might use an S3 compatible object storage on the cloud or S3 object storage on-premises (IBM Cloud Object Storage) for configurations that fit a wide range of customer needs and requirements.

Figure 6-1 shows an example of workload protection with IBM Storage Defender DP archive to S3 Storage.

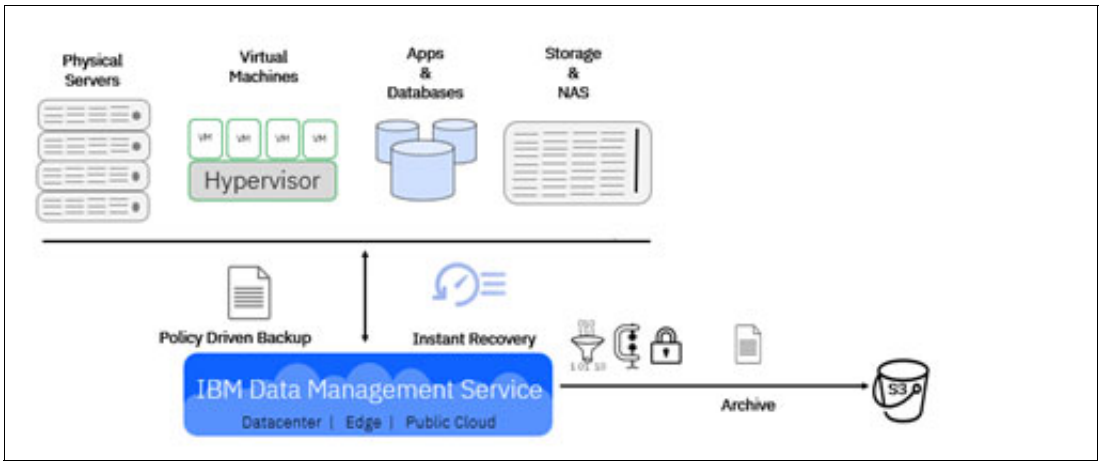


Figure 6-1 Workload protection with IBM Storage Defender DP archive to S3 Storage example

Note: IBM Storage Defender DP archive to IBM Storage Protect tape is intended for weekly or monthly archives. Each archive is a full copy of the Protection Group from IBM Storage Defender DP to IBM Storage Protect. Incremental functions for archives are not available at the time of writing.

S3 offload use cases and scenarios

There are several S3 offload use case and scenarios:

- The IBM Storage Protect Server S3 agent can be set up by other components in the IBM Storage Defender stack. For example, it is possible to use the IBM Storage Protect S3 Agent to provide an object storage bucket for IBM Storage Defender DP and IBM Fusion.
- You must have a weekly copy that goes to tape, following the 3-2-1 rule. The S3 Archive to tape function protects a FULL copy of the specified data by creating a physical air gap once the data is sent to tape. The archive copies can also be on a separate device and can be in a separate location on separate media for added protection.

- Data is protected by the IBM Storage Defender DP cluster with options to replicate the data so there are two copies for disaster recovery (DR) and resiliency. A weekly third archive copy to tape can be added for air-gapping and longer-term retention, often needed for legal requirements.
- Here is a great use for existing tape libraries in the environment. When an archive copy is sent to the IBM Storage Protect S3 bucket, the data is cached on the IBM Storage Protect S3 cold data cache pool. Then, it is migrated to physical tape volumes that can be managed on-site or off-site for longer-term retention needs. This cold data cache pool is helpful in reducing the number of tape drives that is required to meet the data throughput. Also, this disk cache is used for restores.

Compatibility details

A complete list of S3 storage compatibility for IBM Storage Defender DP is available at the IBM Storage Defender DP Support pages. At the time of writing, the current release of IBM Storage Defender DP is Version 7.1, and the current release of IBM Storage Protect is Version 8.1.20. For more information about IBM Storage Defender DP, see IBM Storage Defender Data Protect.

6.2 Deployment exercise overview

This section take a closer look at how the IBM Storage Protect S3 Agent integration is set up for use with IBM Storage Defender DP.

Figure 6-2 shows the S3 Storage options at the upper left, with the data management service at the bottom. In the middle is a time line that shows the steps that are required to complete this integration.

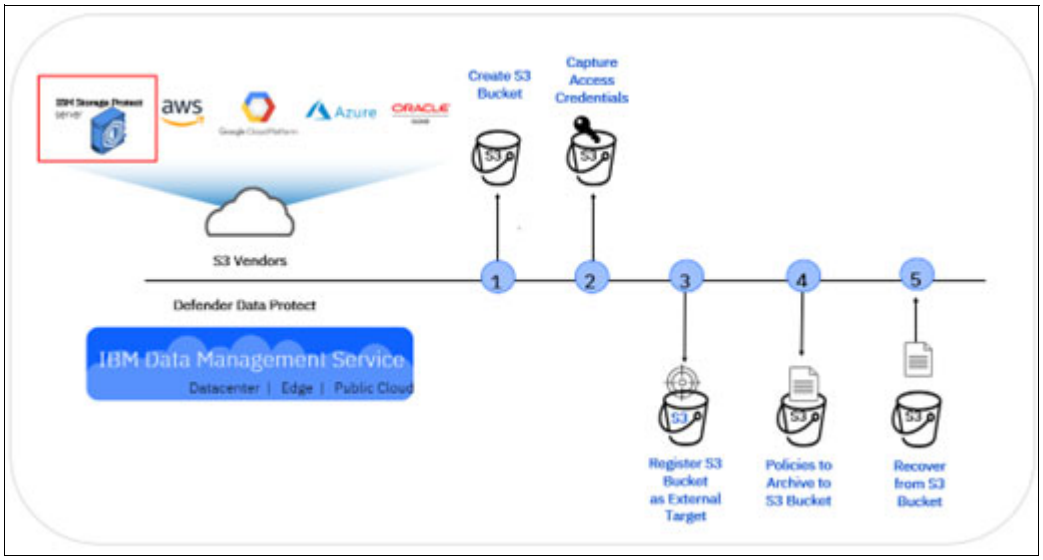


Figure 6-2 Using S3 compatible storage with IBM Storage Defender DP

Here are the high-level steps to complete:

1. Create an S3 bucket on the IBM Storage Protect server.
2. Capture the access credentials on the IBM Storage Protect server.
3. Register the S3 bucket as an external target on the IBM Storage Defender DP cluster.

4. Archive to the S3 bucket.
5. Recover from the S3 bucket.

6.3 Creating an S3 bucket on an IBM Storage Protect server

These steps can be performed on a server with either IBM Storage Protect or IBM Storage Defender licensing. IBM Storage Protect servers must be Version 8.1.20 or later.

6.3.1 How to set up an IBM Storage Protect server to receive IBM Storage Defender DP archives

The command in Example 6-1 creates the `coldcache` storage pool (a FILE type device class) for the `cold-data-cache` storage pool. The `nextstgpool` option is used to define the tape storage pool where archives are stored. In the command, the tape pool is defined as `coldtapestg`.

Note: When creating a cache, ensure that it is large enough for the largest retrieval that might need to be done, and add additional space for any additional incoming archives that might need to be processed.

This storage pool is used as a temporary cache location for the IBM Storage Defender DP archives before migrating them to the tape pool.

For recovery, the IBM Storage Defender DP Protection Group is brought back from this tape storage and placed in the `cold-data-cache` storage pool where it is held for 7 days. Therefore, the size of the IBM Storage Defender DP groups and the size of the virtual machines (VMs) that are protected must be accounted for when determining the amount of storage pool space that is required for successful archive and retrieve operations.

Example 6-1 Defining a coldcache storage pool on the IBM Storage Protect server

```
TSMSRC> define stgpool coldcache stgt=colddatacache nextstgpool=coldtapestg
directory=/sp-src/coldcache migpro=3
Migpro=# is the number of tape drives used

ANR2200I Storage pool COLDCACHE defined (device class COLDCACHEDEVCLASS)
```

Now, complete the following steps:

1. Create the S3 Storage Agent.

   To create the S3 agent, use the `define server` command. In Example 6-2, the server is defined as `s3-dp`.

Example 6-2 Defining the Storage Agent that connects to the S3 enabled storage

```
TSMSRC> def server s3-dp hla=<ip-address-of-IBM Storage Protect-server> lla=9000
objecta=yes

ANR4601I A configuration file for the object agent was created in the instance directory
and is ready to set up. To configure and start a service for object agent S3-DP, run the
following command: "/opt/tivoli/tsm/server/bin/spObjectAgent" service install
"/sp-src/sp_home/S3-DP/spObjectAgent_S3-DP_1500.config"
ANR1660I Server S3-DP defined successfully.
```

After you run the `define server` command, run the command in the resulting ANR4601I message as root from the host CLI, as shown in Example 6-3.

Example 6-3 Service install command execution

```
"/opt/tivoli/tsm/server/bin/spObjectAgent" service install
"/sp-src/sp_home/S3-DP/spObjectAgent_S3-DP_1500.config"

2023-11-20 14:14:37.558414 I | Installed and started system service as
spoa9000S3-DP.
```

From the IBM Storage Protect Operations Center, you can see the results of this command and view or copy the certificate information from the **Servers/Object Agent** menu (Figure 6-3).

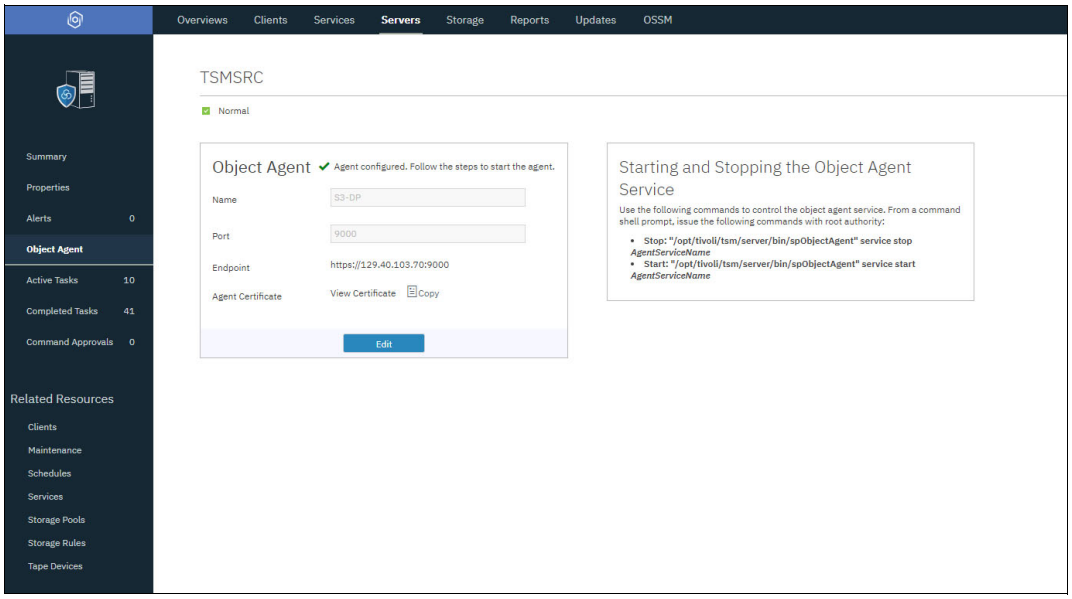


Figure 6-3 Operations Center GUI showing the Object Agent Configuration details

2. From the IBM Storage Protect CLI, define the object domain.

   In Example 6-4, the standard pool `s3-fusion` is a container storage pool that is used for IBM Fusion backups. The `coldcache` pool is the name of the pool that we created in Example 6-1 on page 90. The archives from IBM Storage Defender DP first land in this `coldcache` pool, and then they auto migrate to the `nextstgpool`, which is set to `coldtapestg`.

Example 6-4 The define objectdomain command

```
TSMSRC> define objectdomain S3-DP-fusion standardpool=s3-fusion coldpool=coldcache

ANR1500I Policy domain S3-DP-FUSION defined.

ANR1510I Policy set STANDARD defined in policy domain S3-DP-FUSION.

ANR1520I Management class STANDARD defined in policy domain S3-DP-FUSION, set
STANDARD.

ANR1520I Management class COLD defined in policy domain S3-DP-FUSION, set
STANDARD.
```

```
ANR1530I Backup copy group STANDARD defined in policy domain S3-DP-FUSION, set
STANDARD, management class STANDARD.

ANR1530I Backup copy group STANDARD defined in policy domain S3-DP-FUSION, set
STANDARD, management class COLD.

ANR1538I Default management class set to STANDARD for policy domain S3-DP-FUSION,
set STANDARD.

ANR1554W DEFAULT Management class STANDARD in policy set S3-DP-FUSION STANDARD
does not have an ARCHIVE copygroup: files will not be archived by default if this
set is activated.

ANR1514I Policy set STANDARD activated in policy domain S3-DP-FUSION.
```

3. Register the object client node on your IBM Storage Protect server.

   In Example 6-5, the `register node` command registers a node on the IBM Storage Protect server to store the archive data. This command uses the same domain name that was used in Example 6-4 on page 91.

Example 6-5 Registering the node command

```
TSMSRC> reg node spta-pok-dp-03 type=objectclient domain=s3-dp-fusion

ANR2470I The new authentication credentials for object client node SPTA-POK-DP-03
are: Access Key ID: ********************, Secret Access Key:
****************************************.

ANR2060I Node SPTA-POK-DP-03 registered in policy domain S3-DP-FUSION.

ANR4200I The password for node SPTA-POK-DP-03 is now case-sensitive.
```

4. Create the S3 bucket that is used by IBM Storage Defender DP to send archives to IBM Storage Protect.

   On the IBM Storage Protect server, download and install the MinIO client utility from here. The rpm name that is used in this example is `mcli-20231115224558.0.0.x86_64.rpm`. This version is the minimum one of the RPM that should be used, later versions that are available can be found from the same link.

   As root, from the directory where the `.rpm` file is, run the `mcli alias set alias-name end-point access-key secret-key --insecure` command, as shown in Example 6-6.

Example 6-6 Setting the mcli alias

```
mcli alias set s3-dp {"C:\\Users\\014602649\\Downloads\\h"} --insecure

mcli: Configuration written to `/root/.mcli/config.json`. Update your access
credentials.
mcli: Successfully created `/root/.mcli/share`.
mcli: Initialized share uploads `/root/.mcli/share/uploads.json` file.
mcli: Initialized share downloads `/root/.mcli/share/downloads.json` file.

Added `s3-dp` successfully.
```

5. Create the bucket.

In Example 6-7, `s3-dp` is the alias name from Example 6-6 on page 92 and `s3-dp-minio` is the bucket name that was used.

As root, run the `mcli mb alias-name/bucket-name` command.

Example 6-7 S3 bucket creation

```
mcli mb s3-dp/s3-dp-minio –insecure

Bucket created successfully `s3-dp/s3-dp-minio`
```

6.4 Capturing access credentials on an IBM Storage Protect server

The information in this section is needed to register the IBM Storage Protect bucket that was created in a new IBM Storage Defender DP external target.

To continue with the next configuration steps, you need the following items:

- Bucket Name.
- Access Key ID.
- Secret Access Key.
- Endpoint: The IP address of your IBM Storage Protect server.
- Port number: The LLA that is used on the defined server in Example 6-2 on page 90.

6.5 Registering the S3 bucket as an external target on a DP cluster

Note: IBM Storage Defender DP cluster must be Version 7.1 or later.

In IBM Storage Defender DP, select **AWS Signature version v4**. For the External Target Name, use the IBM Storage Protect Node that was registered in Example 6-5 on page 92. This node is the IBM Storage Protect node under which the archives are stored.

To register an external target with your cluster, log in to the IBM Storage Defender dashboard and complete the following steps:

1. Select **Data Management** to start the Data Management Service (Figure 6-4).

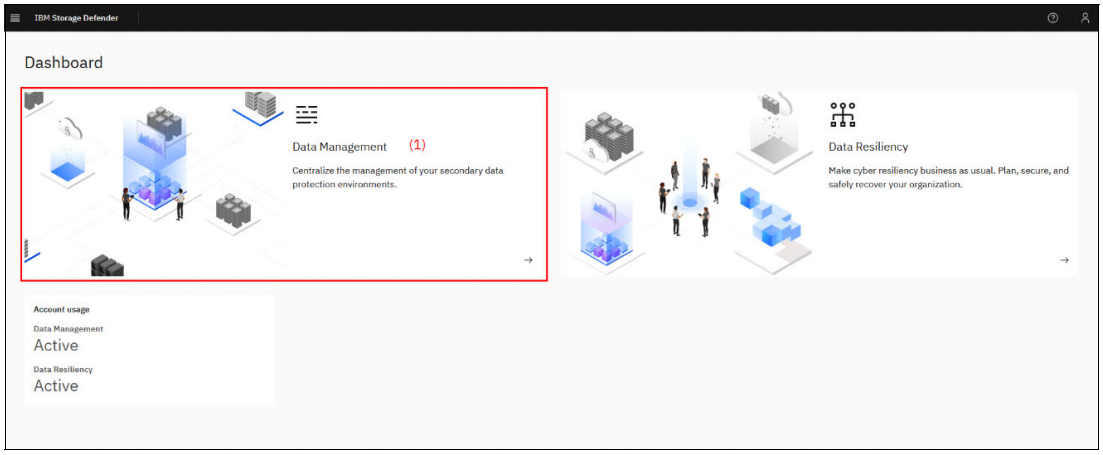


Figure 6-4 Logging in to the IBM Storage Defender dashboard: Step 1

2. Select the **Administrator** level access to do configuration tasks on the cluster (Figure 6-5).

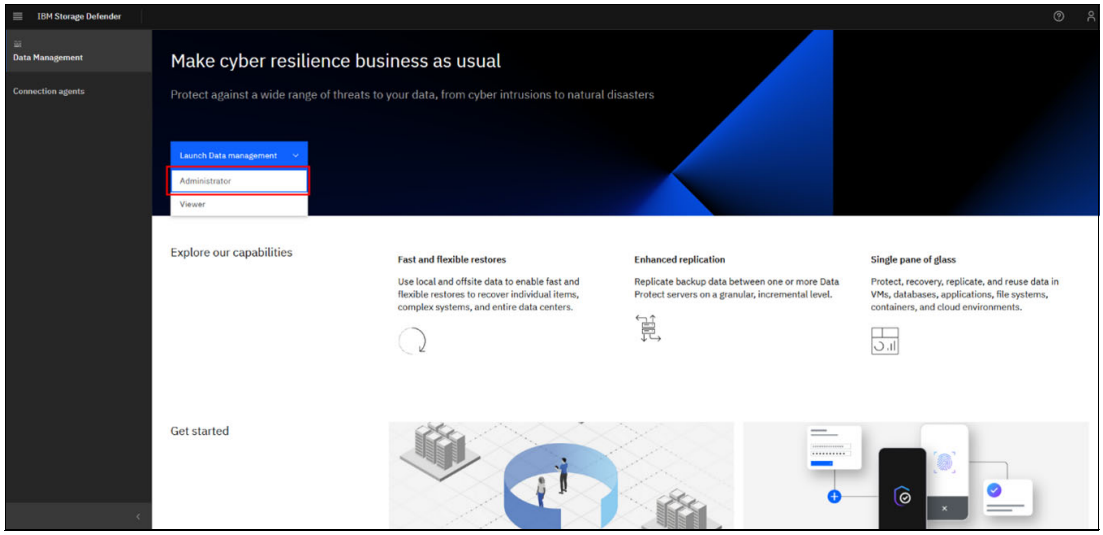


Figure 6-5 Logging in to IBM Storage Data Management Service: Step 2

3. Select the cluster resource by clicking the drop-down menu (Figure 6-6 on page 95) and making the following selections:
   a. Expand **Infrastructure**.
   b. Select **External Targets**.
   c. Click **Add External Target**.

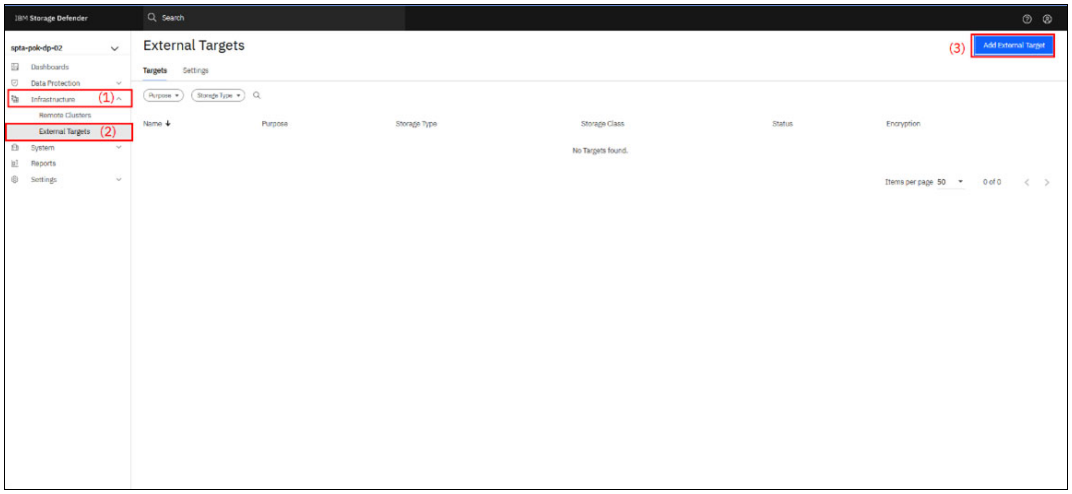


Figure 6-6 Registering an S3 external target in IBM Storage Defender DP: Steps 3a - c

4. Select Purpose (**Archival**).
5. Select Type (**S3 Compatible**).
6. Select Storage Class (**Tape Based for IBM Storage Protect S3 Agent to tape**).
7. Enter the Bucket Name.
8. Enter the Access Key ID.
9. Enter the Secret Access Key that you captured when you created your bucket.
10. Enter the Endpoint IP address or fully qualified domain name (FQDN).
11. Enter the port number, and then scroll down for additional options to continue the External Target Registration. Do not click register now.

    Figure 6-7 shows steps 4 - 11.

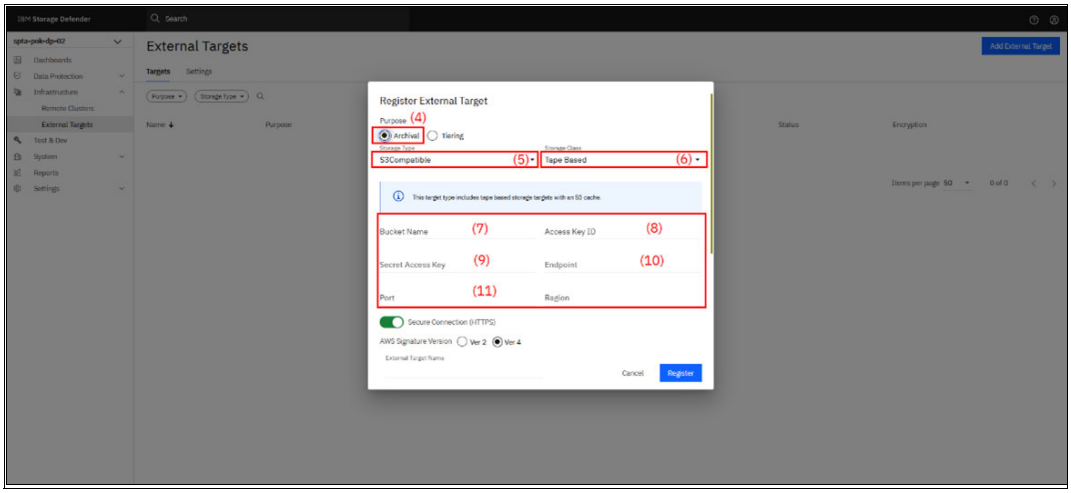


Figure 6-7 Registering an S3 external target in IBM Storage Defender DP: Steps 4 - 11

Note: Do not include `http://` or `https://` in the mount path.

Secure Connection (HTTPS) Is enabled by default. If your S3 bucket is accessible through HTTP and not HTTPS (that is, without SSL security), disable this option.

12. Set the AWS Signature version to **Ver. 4**.
13. Set the External Target Name as the IBM Storage Protect node name.
14. Set additional options for Encryption. (See the note about the Key Management Service type in the GUI window) If you want to enable manual key management for extra security, enable it here.

Important: With this option enabled, a cluster must have the correct key to access data from the archive. You can download the key file (only once) after you register your bucket. This key is required when you use CloudRetrieve. If you do not have it, you may recover data to its original cluster, but you cannot retrieve it onto a new cluster (for example, in a DR scenario).

15. Edit the setting for **Compression** as needed.
16. Enable and configure the **Bandwidth Throttling** settings as needed. The upload and download speeds can be throttled separately. Throttling can be enabled all the time or only during specific days and times.
17. Click **Register** to complete form submission.

    Figure 6-8 shows steps 12 - 17.

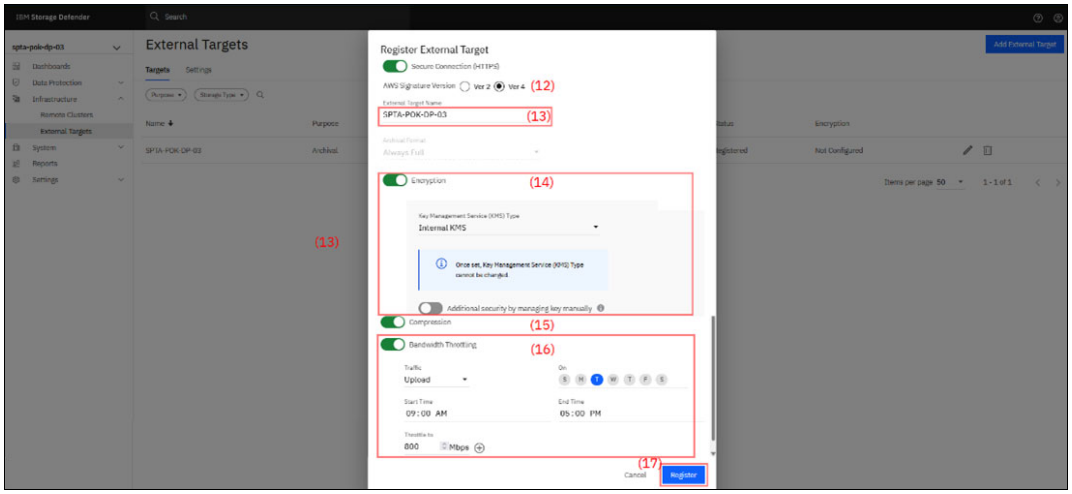


Figure 6-8 Register an S3 external target in IBM Storage Defender DP: Steps 12 - 17

The S3 bucket is now registered and available as an external target in IBM Storage Defender DP. This target can be selected when you create a Data Protection policy for a Data Protection Job.

6.6 Creating policies to archive data to the S3 bucket

This section is split into two subsections because there are multiple policies that must be defined in IBM Storage Defender DP to correctly configure the system for this use case.

6.6.1 Creating a Data Protection policy for archiving to the S3 bucket

In IBM Storage Defender DP, associate the external S3 target name from step 3 on page 94 to a Data Protection policy in IBM Storage Defender DP.

To create a Data Protection policy, complete the following steps:

1. Expand **Data Protection**.
2. Select **Policies**.
3. Click Create Policy to begin creating the policy.

   Figure 6-9 shows steps 1 - 3.

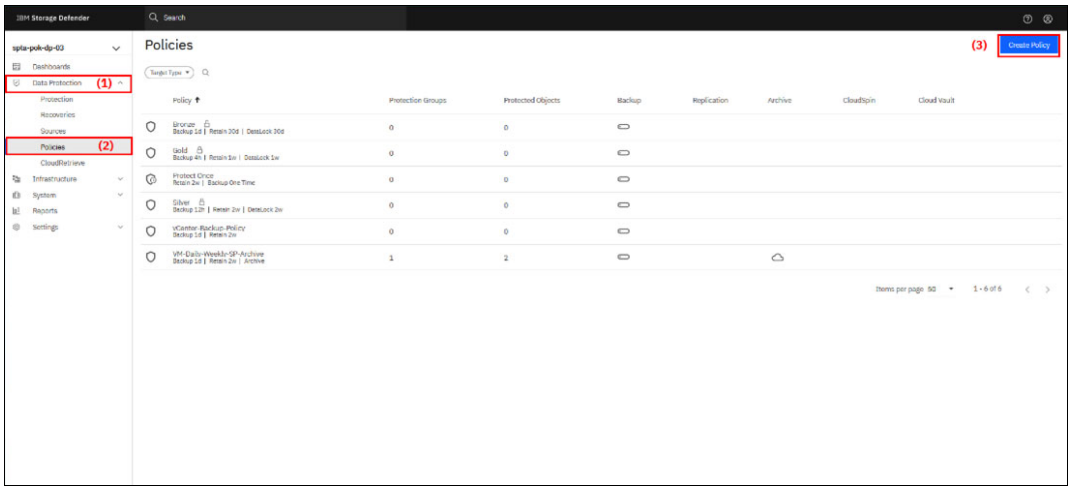


Figure 6-9 Creating a Data Protection policy for archive to the S3 External Target: Steps 1 - 3

4. In the dialog box that opens, select **More Options** (Figure 6-10).

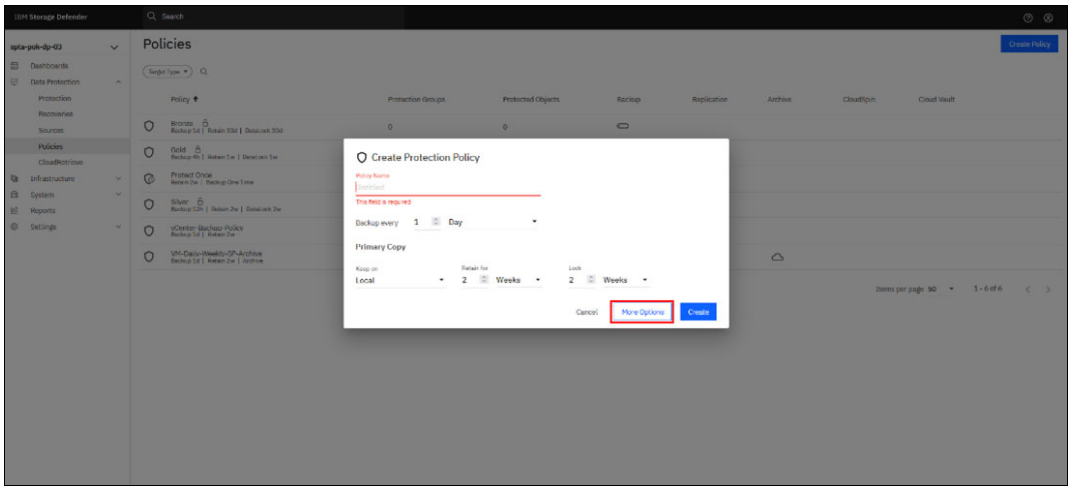


Figure 6-10 Creating a Data Protection policy for archive to the S3 External Target: Step 4

5. Name the new policy.
6. Disable the **Data Lock** option.
7. Click **Add Archive** to show additional options.

Figure 6-11 shows steps 5 on page 97 - 7 on page 97.

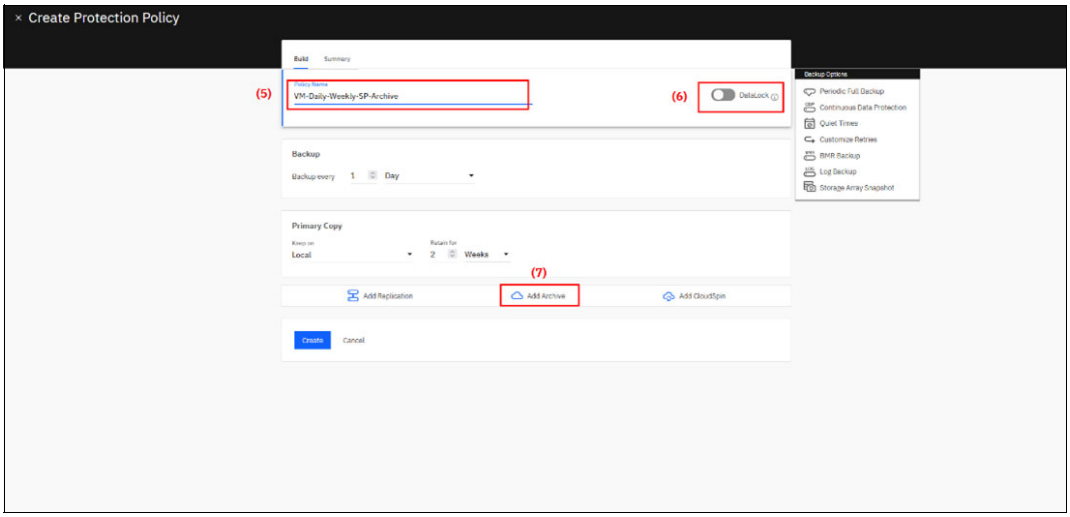


Figure 6-11 Data Protection policy for archive to the S3 External Target: Steps 5 - 7

8. In the Archive section, set the External Target, Frequency, and Retention settings.

> **Note:** The recommended frequency is weekly or monthly because this copy will be a full copy.

9. Click **Create** to create the policy.

Figure 6-12 shows steps 8 - 9.

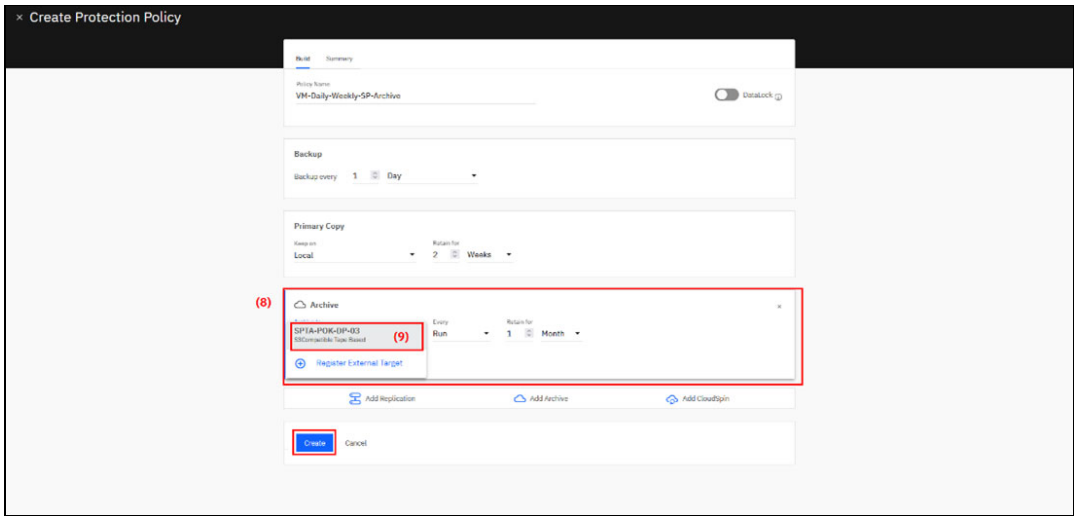


Figure 6-12 Creating a Data Protection policy for archive to the S3 External Target: Steps 8 and 9

6.6.2 Configuring the Data Protection Group to use the Data Protection policy

The Data Protection Group includes the definition of the workloads that are protected (VMs in this example) and the Data Protection Policy that is applied to that group.

To configure the Data Protection Group, complete the following steps:

1. Expand **Data Protection**.
2. Select **Protection**.
3. Select the **Protect** drop-down menu.
4. Select the workload type that you want to protect (in this example, **Virtual Machines**). Figure 6-13 shows steps 1 - 4.

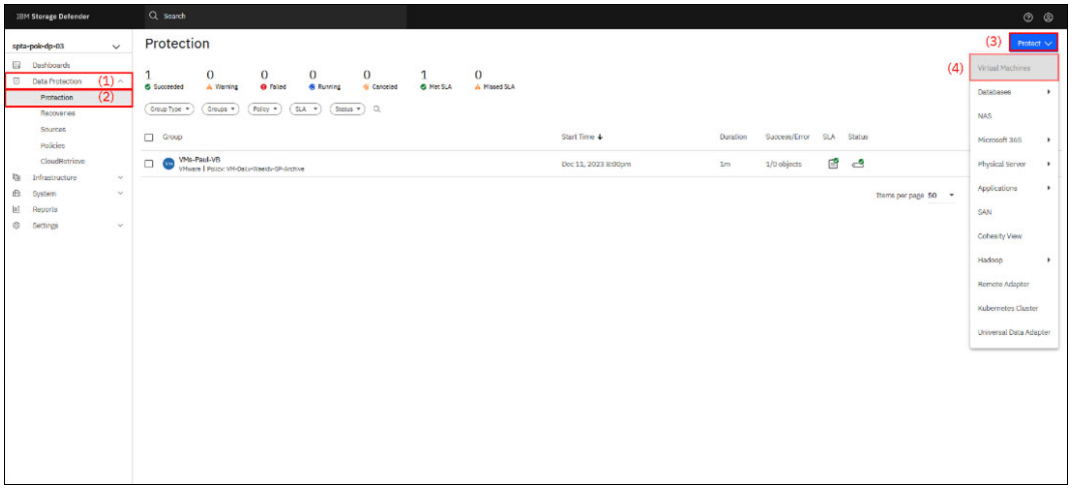


Figure 6-13 Creating a Data Protection policy by using the S3 External Target: Steps 1 - 4

5. In the New Protection dialog box, select **Add Objects** (Figure 6-14).

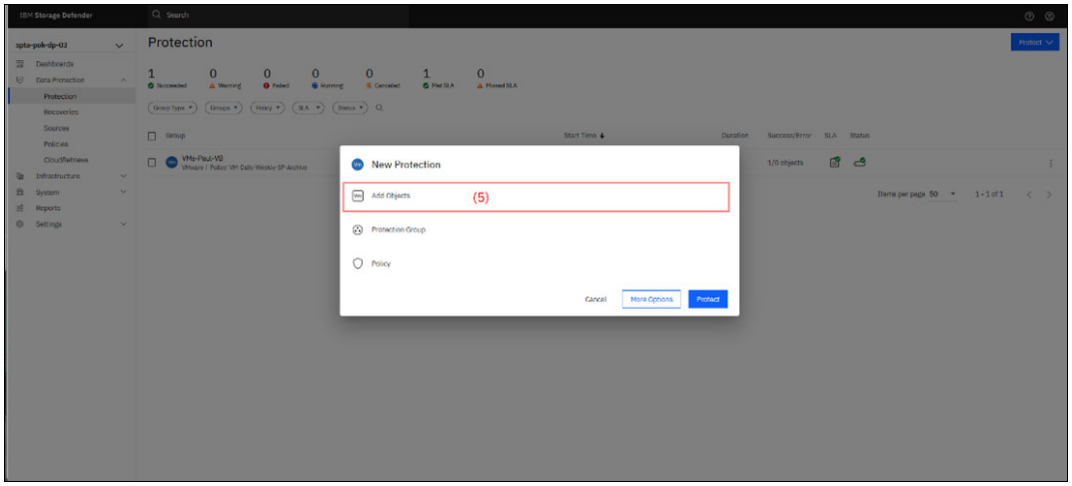


Figure 6-14 Virtual Machine New Protection options: Step 5

6. Select the IP address of the vCenter server.
7. Search for VMs by using a name or pattern.
8. Select any ESX Hosts or individual VMs by name to protect.
9. Click **Continue**.

Note: Be aware of the number and size of VMs in this Protection Group. The selected VMs are managed in the Protection Group for backup and recovery operations. The sizing for the IBM Storage Protect cold cache pool is based on ingest and the restore pattern.

Figure 6-15 shows steps 6 on page 99 - 9 on page 99.

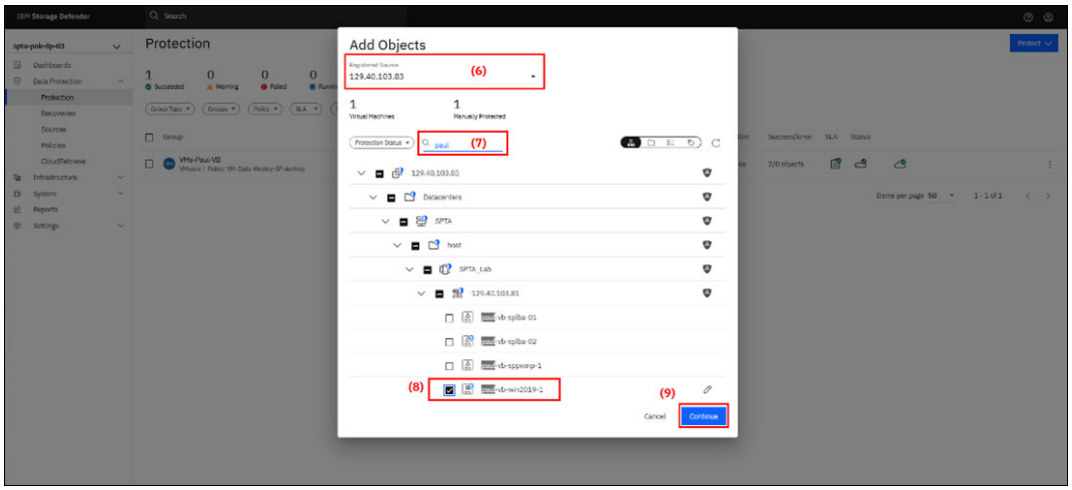


Figure 6-15 Add Objects window: Steps 6 - 9

10. Enter a name for the new Protection Group.
11. Select the policy that you created in 6.6.1, “Creating a Data Protection policy for archiving to the S3 bucket” on page 96.

    Figure 6-16 shows steps 10 and 11.

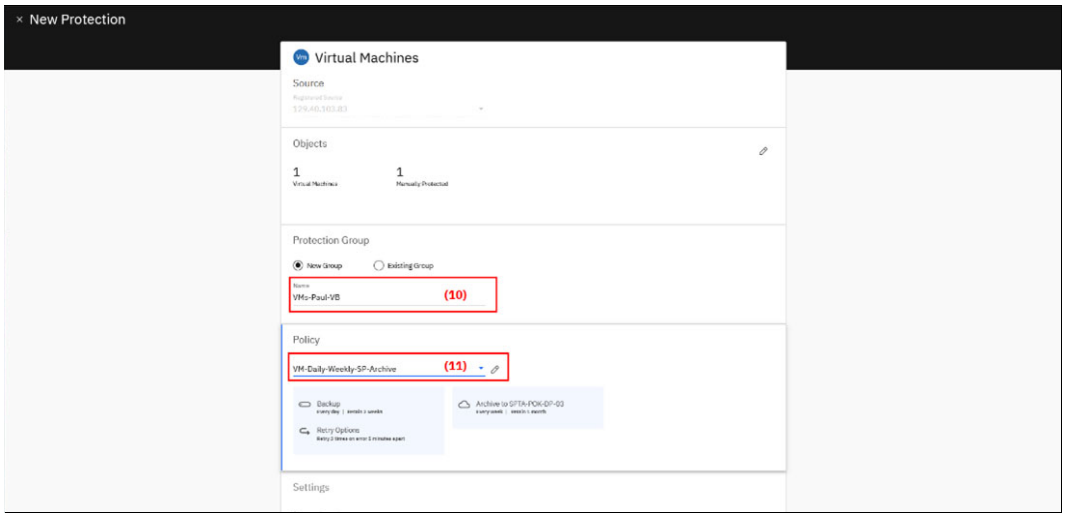


Figure 6-16 Virtual Machines window: steps 10 and 11

Scroll down to see more options and additional settings.

12. Review the settings for the Storage Domain.
13. Review the additional settings if needed.
14. Click **Protect** to complete the configuration.

Figure 6-17 shows steps 12 on page 100 - 14 on page 100.

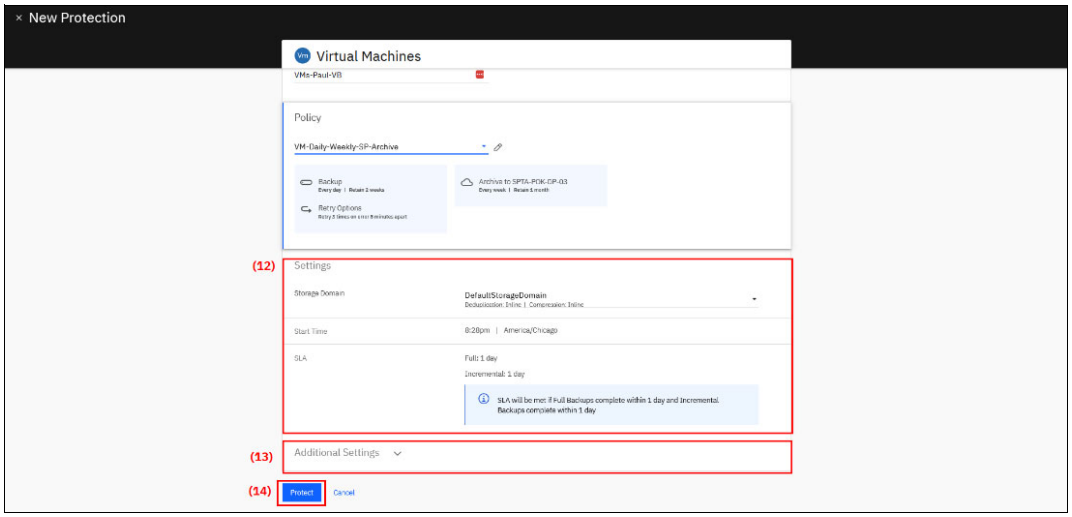


Figure 6-17 Virtual Machines window: Steps 12 - 14

The Data Protection Group was successfully associated with the Data Protection policy to protect the VMs in the group.

Note: The archive to S3 is processed like a full backup job, not an incremental job. As a best practice, limit the frequency to a weekly or monthly archive to meet your retention needs.

6.7 Monitoring the Data Protection status

The statuses of Data Protection jobs are shown in the Protection window. You can review the overall status for multiple protection jobs at a glance in this window, and click the job name to get additional details about the job.

The Data Protection Group Detail window (Figure 6-18) shows the daily VM backup statistics and the following items:

1. The SLA policy
2. The job status
3. A weekly archive to cloud, which in this case in our S3 external target

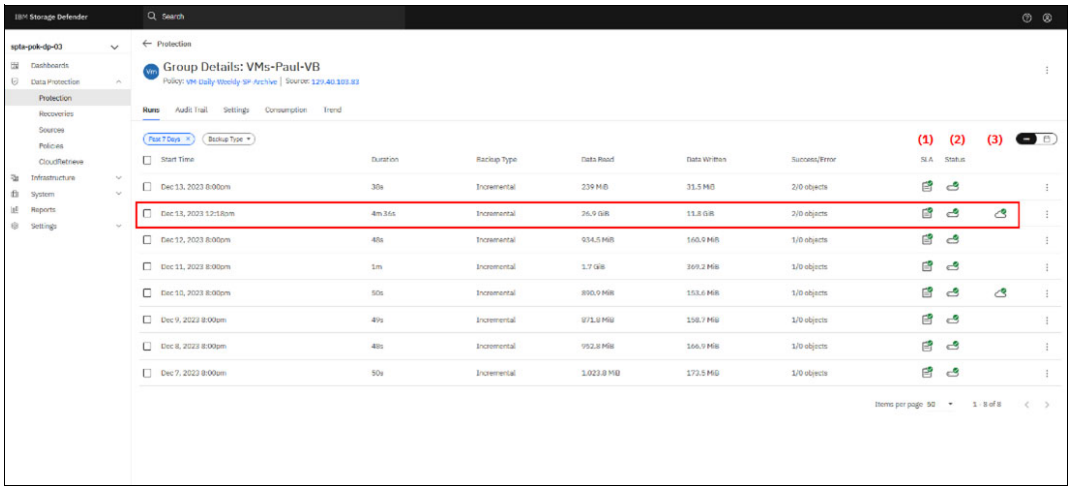


Figure 6-18 Reviewing the Protection Group details and status

In this example configuration, a daily backup with a weekly archive to S3 has been running for a couple of weeks. In Figure 6-19, a new VM was added to the Protection Group for both the daily backup and weekly archive job on 12/13/23. The full VM backup was taken, and we allowed the archive to S3 to be pushed after adding a VM to the Protection Group.

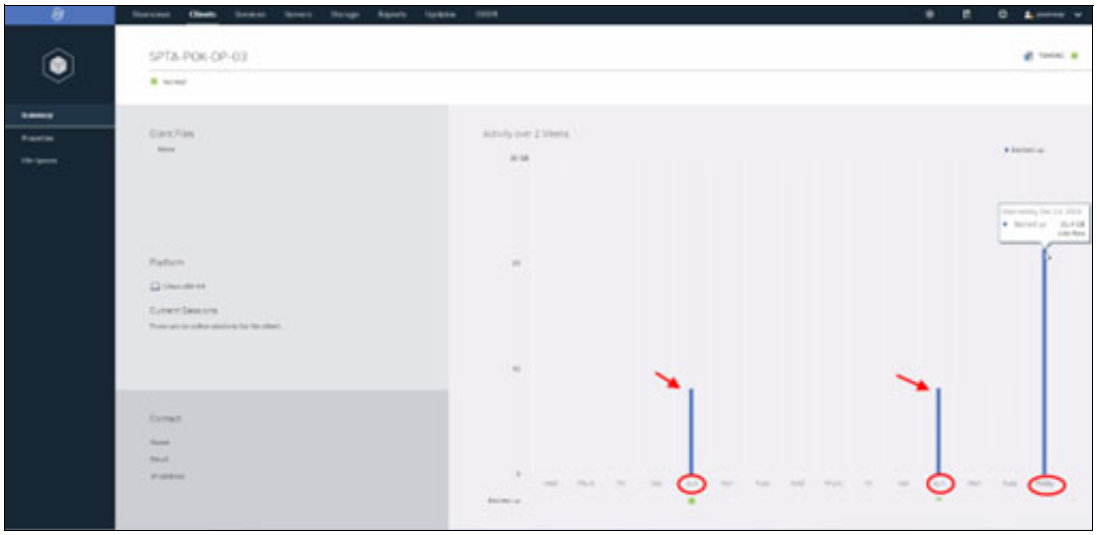


Figure 6-19 Reviewing the activity of the S3 Agent node on the IBM Storage Protect server

Using the Operations Center GUI, you can monitor the weekly activity of an S3 Agent node on an IBM Storage Protect server. Figure 6-19 on page 102 shows that there was an increase in amount of data was being stored after a new VM was added to the Data Protection group, which uses the S3 archive `coldcache` pool. For more information about `coldcache` migration and volume status, use the operations center to run queries on the IBM Storage Protect server, as shown in Figure 6-20 and Figure 6-21.

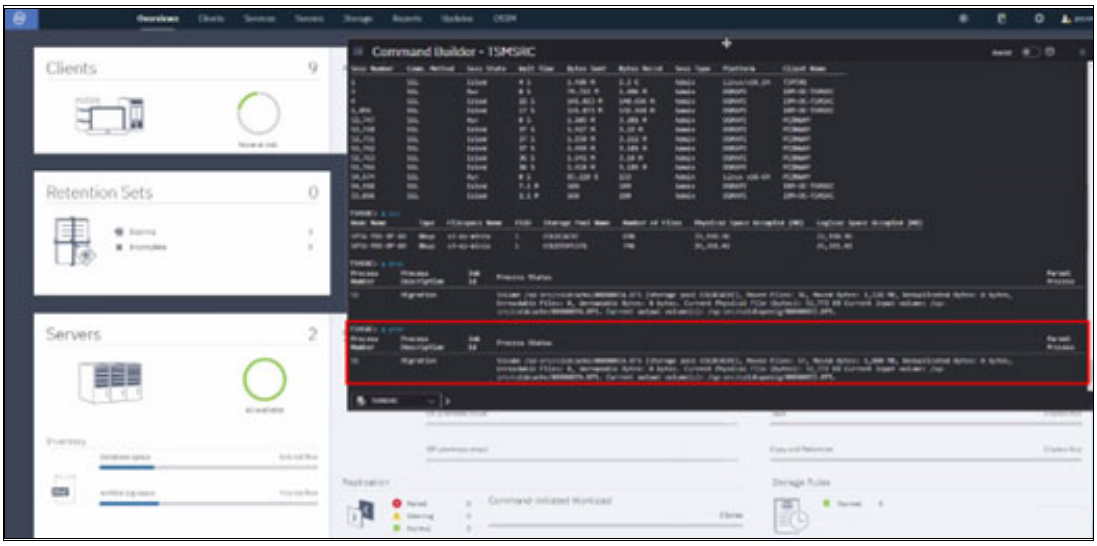


Figure 6-20 IBM Storage Protect server showing migration from a cold cache to cold tape

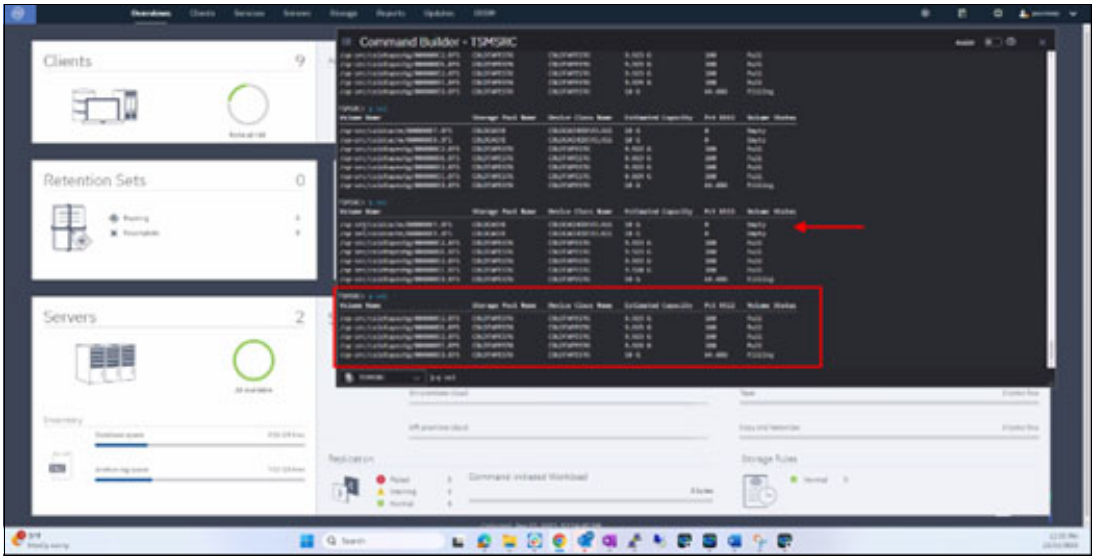


Figure 6-21 IBM Storage Protect server showing the volume status after migration

6.8 Retrieving data from S3 external storage

With the flexible search tools that are available in IBM Storage Defender DP, there are multiple ways to find your data and create restore jobs for protected workloads.

6.8.1 Retrieving data by using Global Search

In the following example, use the Global Search bar to retrieve a VM from the S3 Archive by completing the following steps:

1. Use the search bar and type the name of the VM.
2. With the VM selected, you see the following panes:
   - Search Results
   - Virtual Machine information
   - Protection Group information

   Figure 6-22 shows steps 1 and 2.

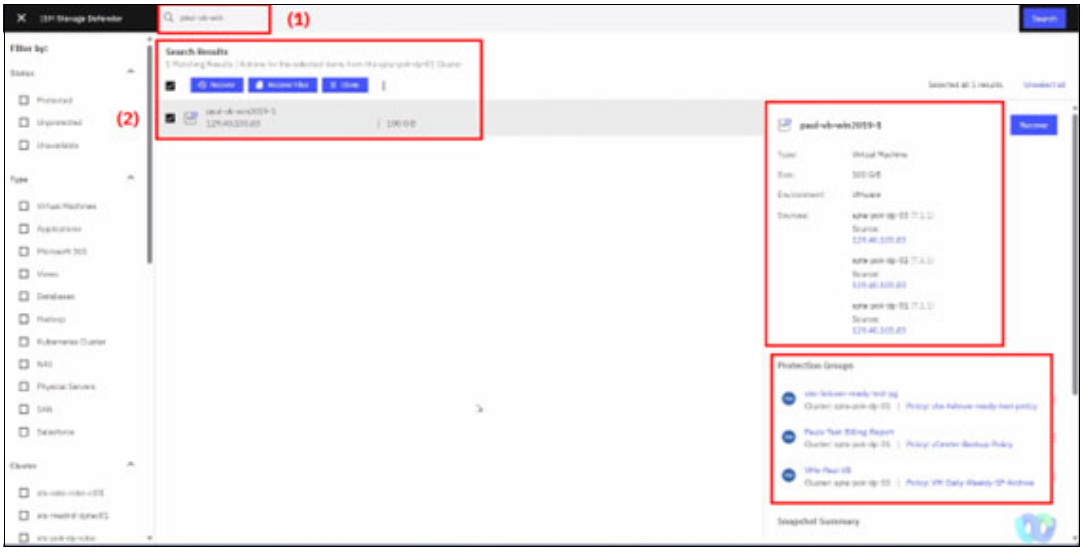


Figure 6-22 Recovery job: Steps 1 and 2

3. Click the pencil icon to open the Edit Recovery Point window for the VM. The Edit Recovery Points table shows the snapshots that are available by timestamp and the location of the snapshot images.
4. Select the snapshot and click the cloud icon to select the S3 location.
5. Click **Select Recovery Point**.

   Figure 6-23 on page 105 shows steps 3 - 5.

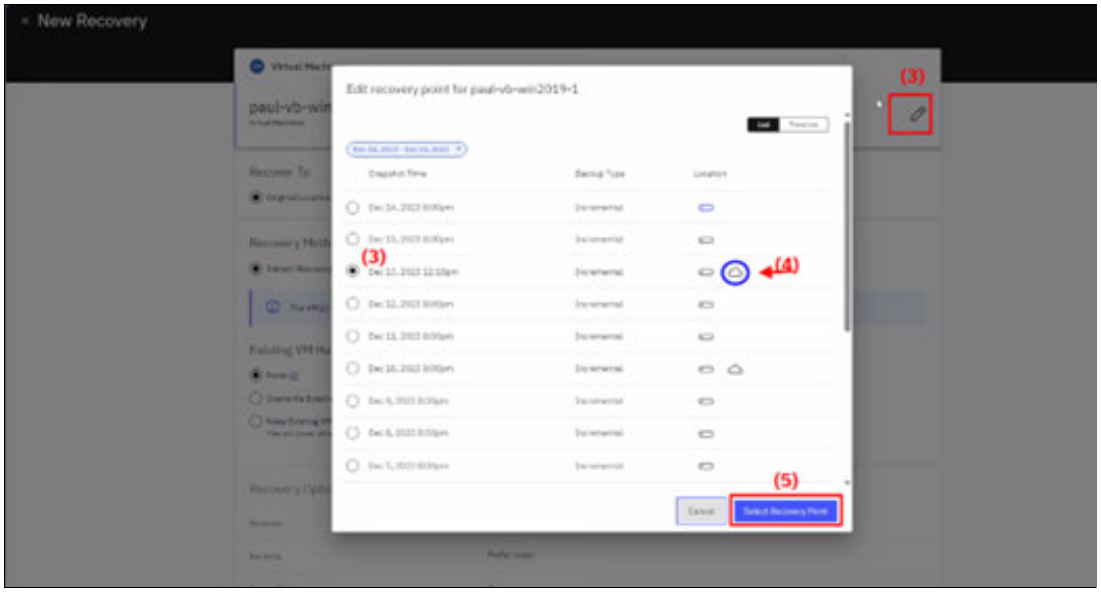


Figure 6-23 VM Recovery job: Steps 3 - 5

6. In the Recovery Method window, select **Copy Recovery**.
7. Scroll down and click **Recover** to start the recovery job.

   Figure 6-24 shows steps 6 and 7.

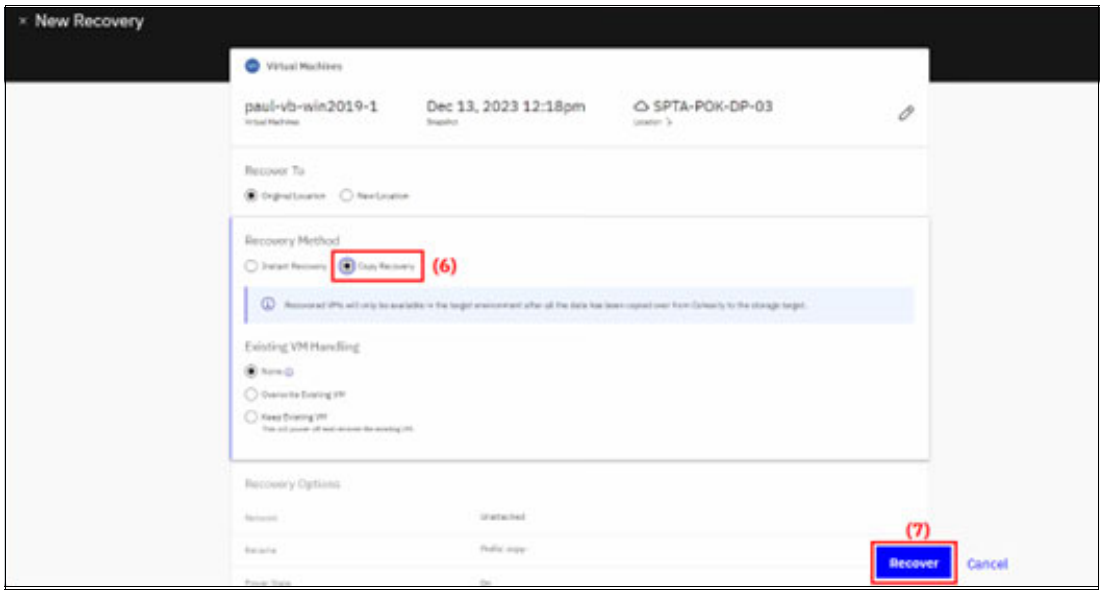


Figure 6-24 VM Recovery job: Steps 6 and 7

Now, the recovery of the selected backup data proceeds, and the VM data is retrieved from the S3 target to the cache storage pool on the IBM Storage Protect server. Then, the backup data is sent to the IBM Storage Defender DP cluster and recovered to the vCenter.

6.8.2 Retrieving data by using the Recoveries window

Another way to retrieve archived data is to use the **Recoveries** option under Data Protection in the IBM Storage Defender DP GUI.

To retrieve data by using the Recoveries window, complete the following steps:

1. In the left pane of the GUI, click **Data Protection**.
2. Click **Recoveries**.
3. Click **Recover** to expand the menu details.
4. Click **Virtual Machines**.
5. Click **VMs**.

   Figure 6-25 shows steps 1 - 5.

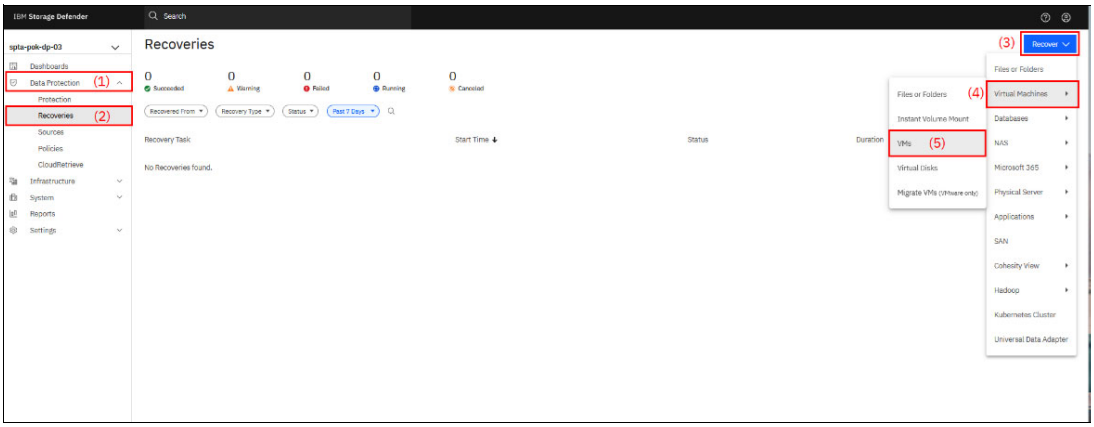


Figure 6-25 Performing a VM Recovery job by using the Recoveries GUI: Steps 1 - 5

6. Use the search bar in the window to enter the name of the VM.
7. Select the VM by selecting the checkbox next to its name.
8. Click the pencil icon to select the Recovery Point and Location.

   Figure 6-26 shows steps 6 - 8.

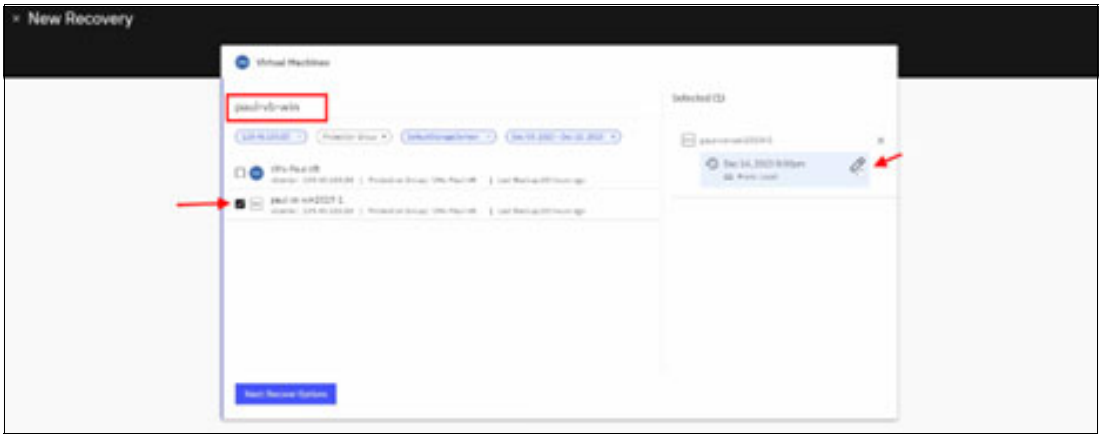


Figure 6-26 Performing a VM Recovery job by using the Recoveries GUI: Steps 6 - 8

9. For the Recovery Point and Location, click the cloud icon to select the S3 External Target.

10. Click **Select Recovery Point**.

Figure 6-27 shows steps 9 and 10.

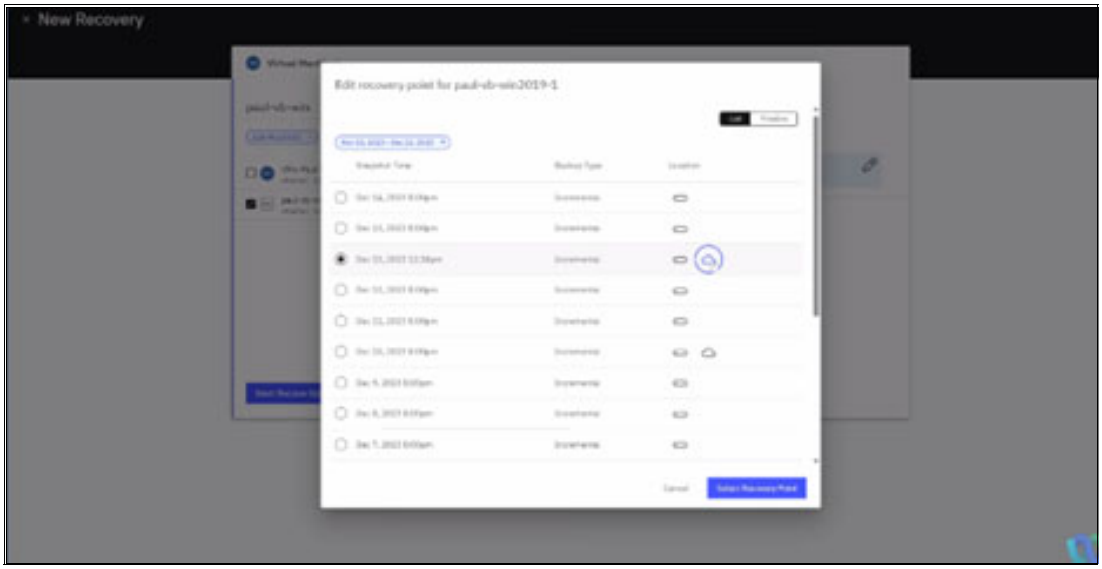


Figure 6-27 Performing a VM Recovery job by using the Recoveries GUI: Steps 9 and 10

11. In the Virtual Machine window, select the **Recovery** method for Copy Recovery.
12. Scroll down and click **Recover**.

Figure 6-28 shows steps 11 and 12.

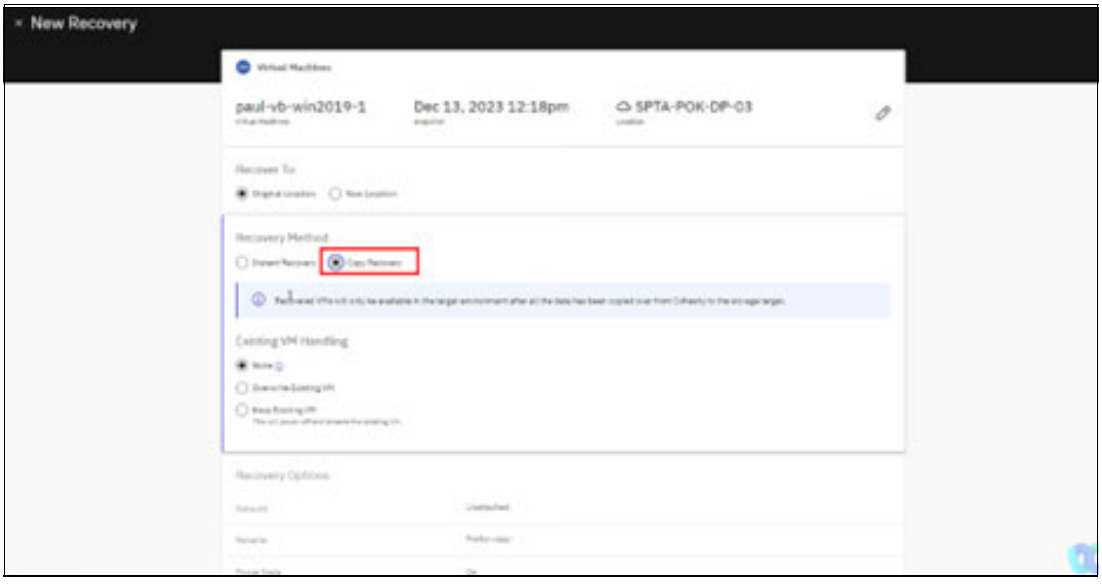


Figure 6-28 VM Recovery job: Steps 11 and 12

The recovery retrieves the archive data.

6.9 Additional process monitoring examples

The following figures show examples that are related to checking the status of archive and retrieval operations. By using these windows, you can conveniently monitor processes across the environment, from the ESX server to the IBM Storage Defender DP cluster and the IBM Storage Protect S3 Agent.

Figure 6-29 shows a virtual machine recovery job in progress.

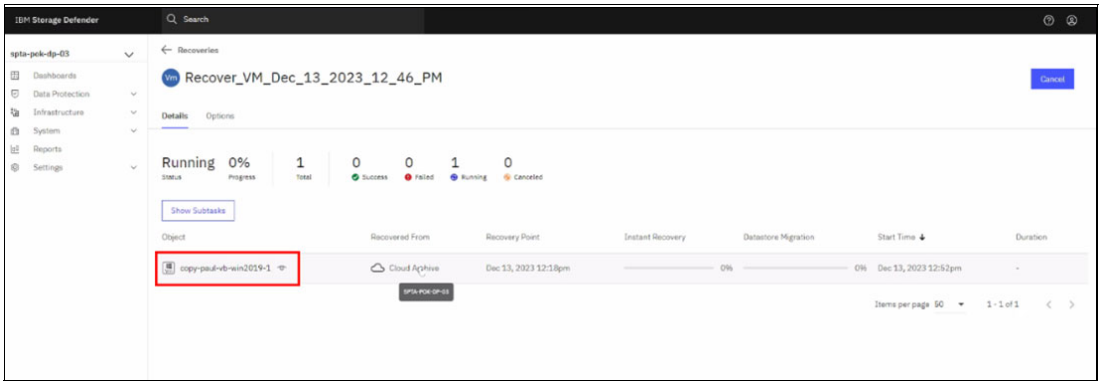


Figure 6-29 Virtual machine recovery job in progress

On the IBM Storage Protect server, you can use the Q Process command to monitor a retrieval from the cold tape pool to the cold cache pool (Figure 6-30).

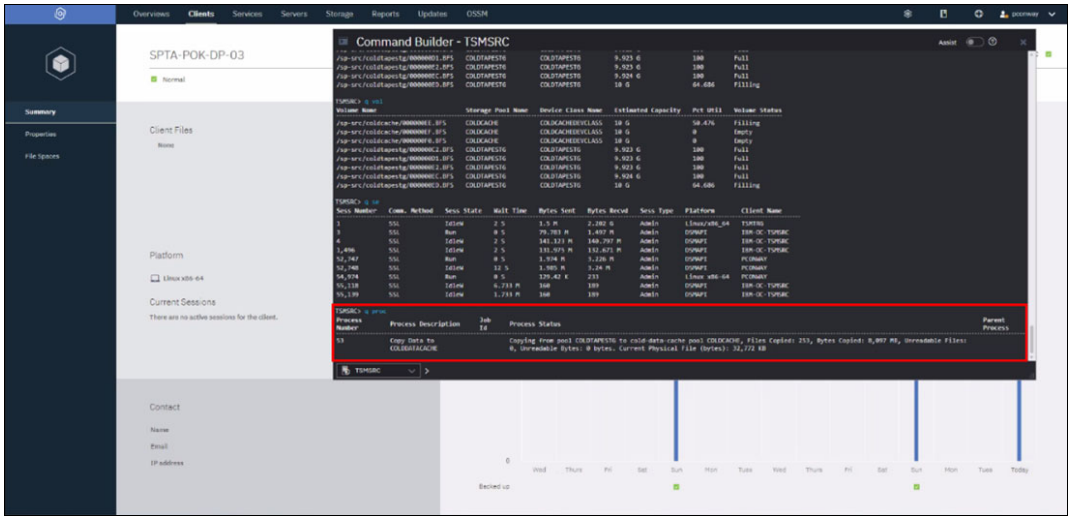


Figure 6-30 Query process output from the IBM Storage Protect server showing a data copy to the cold cache pool

The `query volume` command shows the archived snapshot as restored to the S3 cold cache pool. The **`query session`** command shows the data that was sent from the IBM Storage Protect server to the IBM Storage Defender DP cluster (Figure 6-31 on page 109).

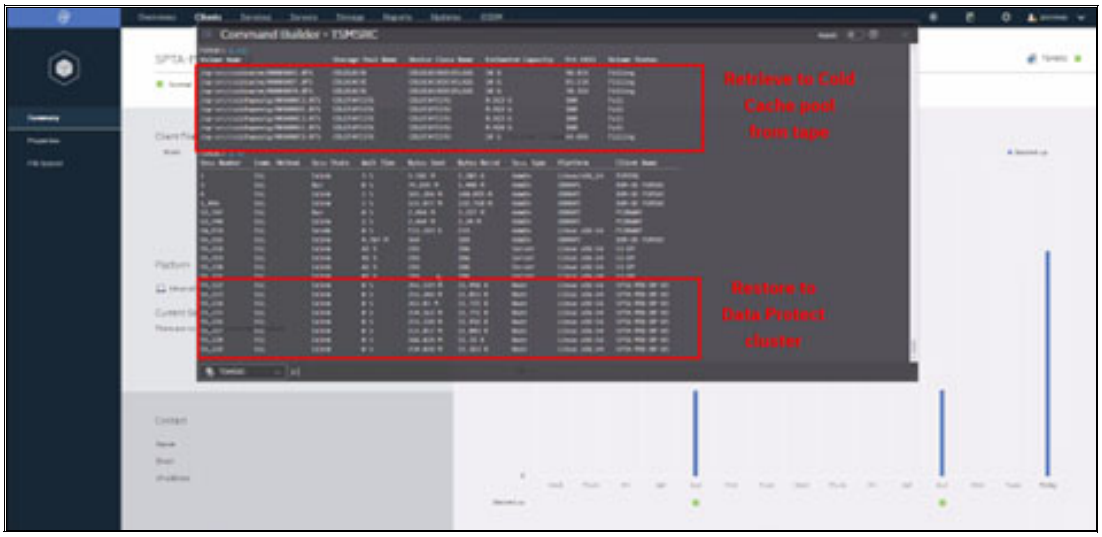


Figure 6-31 IBM Storage Protect query volume command and query session command outputs for recovery related cold cache data movement

In IBM Storage Defender DP, click the Recovery job name to get the status of the job with Recovery Activity details and completion status (Figure 6-32).



Figure 6-32 Job status

To see that the VM was copied and powered on, use the vCenter Client (Figure 6-33)

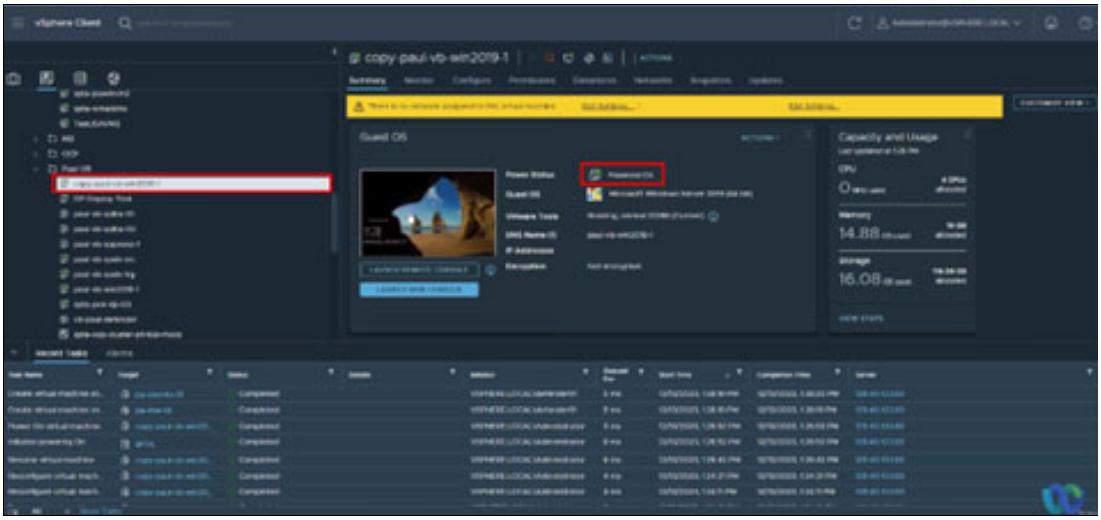


Figure 6-33 The VM was copied and powered on

Abbreviations and acronyms

ATG	Advanced Technology Group
BMR	Bare Machine Recovery
DMS	Data Management Service
DP	Data Protect
EC	Erasure Coding
FC	Fibre Channel
FETB	Front End TB
FQDN	fully qualified domain name
IMR	Instant Mass Restore
IVM	Instant Volume Mount
OVM	Oracle VM
QoS	Quality of Services
ROBO	remote and branch office
SaaS	Software as a Service
TPI	Two-Person Integrity
VIP	Virtual IP address
VM	virtual machine
VMDK	Virtual Machine Disk

© Copyright IBM Corp. 2025.

Redbooks

IBM Storage Defender: IBM Data Management Service and IBM Data Protect

(0.2”spine)
0.17”<->0.473”
90<->249 pages



SG24-8554-00

ISBN 0738462047

Printed in U.S.A.

Get connected



