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
Steve Solazzo
Jack Tedjai
Dan Thompson
Christopher Vollmar

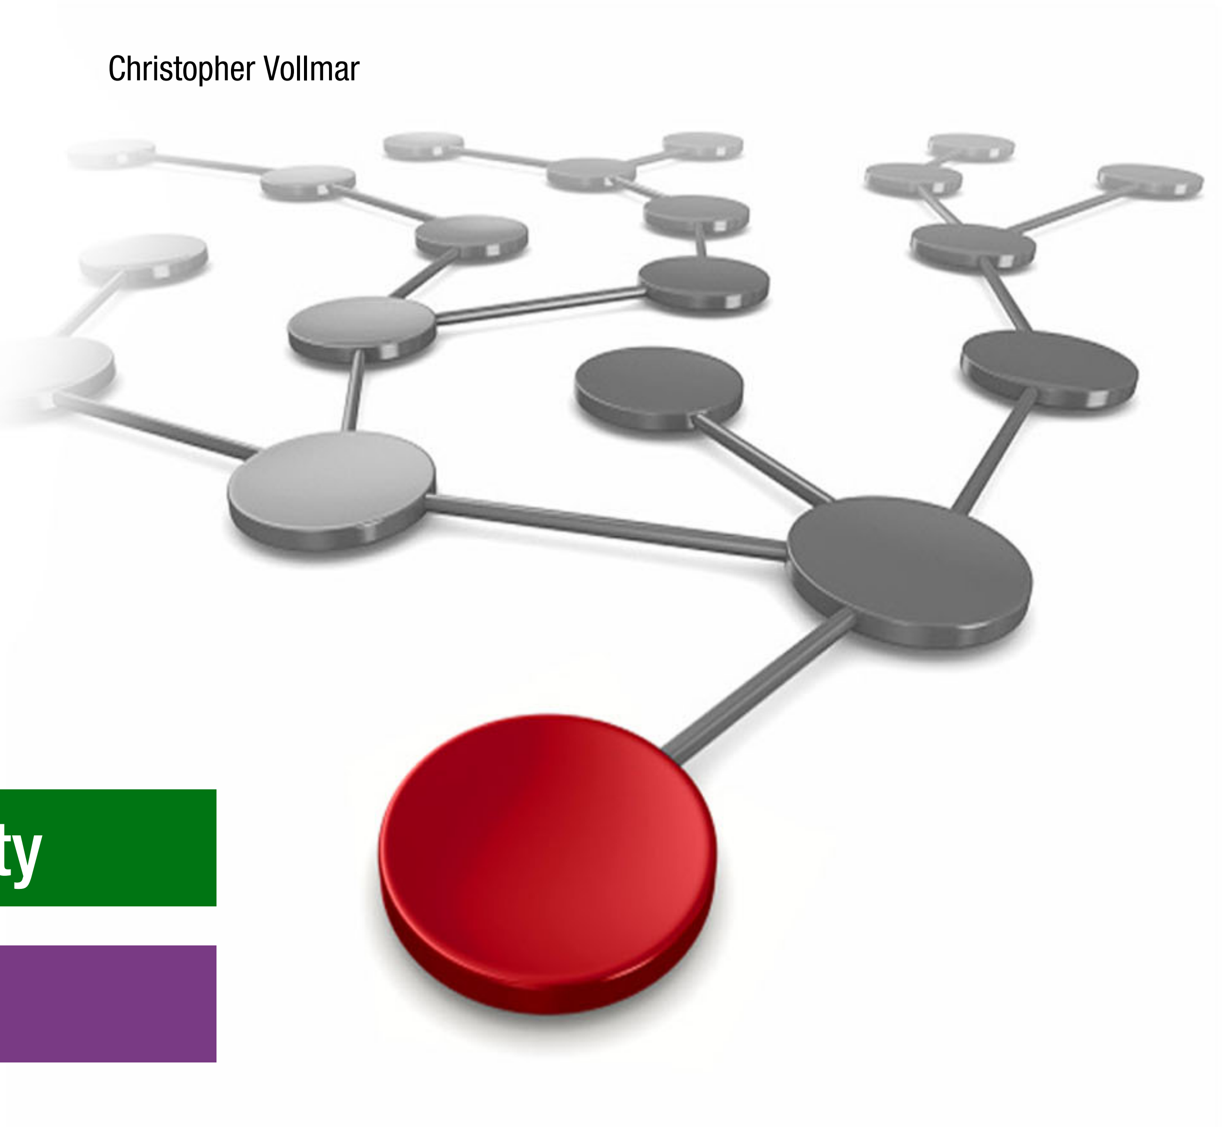

Security

Storage

IBM
Redbooks

IBM Redbooks

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

January 2025

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

First Edition (January 2025)

This edition applies to IBM Storage Defender Data Protect Version 7.1.1 and 7.1.2.

This document was created or updated on January 14, 2025.

Contents

Notices v
Trademarks vi

Preface vii
Authors vii
Comments welcome viii
Stay connected to IBM Redbooks ix

Chapter 1. Introduction to IBM Storage Defender 11
1.1 Overview of IBM Storage Defender 12
1.2 Overview of IBM Defender Data Protect 13
1.3 Overview of IBM Defender Data Management Service 14

Chapter 2. IBM Storage Defender: Data Management Service (DMS) 17
2.1 Data Management Service (DMS) Overview 18
2.2 IBM Storage Defender DMS Dashboard 18
2.2.1 Security 20
2.2.2 System 23
2.2.3 Reporting 25
2.2.4 System View 27

Chapter 3. IBM Storage Defender: Data Protect 29
3.1 Solution Overview 30
3.1.1 Node Based Scale Out Architecture 30
3.2 What is a Cluster and How does It Work? 30
3.3 Physical / Virtual / Robo Implementations 31
3.4 Cluster setup 32
3.4.1 Configuration of a Data Protect Storage Domain 33
3.4.2 Configuration of a Data Protect Protection Group 36
3.4.3 Monitor the Data Protect Protection Group 38
3.5 IBM Storage Ready Nodes 39
3.6 IBM Storage Defender Ready Nodes 40
3.7 Registering with IBM Data Management Service 41
3.8 VMware Workload Use Case 42

Chapter 4. Protecting VMware Workloads With IBM Data Protect 43
4.1 IBM Data Protect Overview 44
4.2 Configuration of Data Protect Sources 44
4.2.1 Configuration of a Data Protect Policy 45
4.3 Recovering Data with IBM Data Protect 47
4.3.1 Recover from a Data Protect Cluster 47
4.3.2 Instant Volume Mount recovery from Data Protect Cluster 50
4.4 Replicating Workloads with IBM Data Protect clusters 53
4.4.1 VMware replication workflow one-to-one model 53
4.5 Monitoring the Protection job list 56
4.6 Recover from Replicated Copies 57

4.6.1 Preparing a failover back to the original source cluster. 59
4.7 CloudArchive the Workloads with IBM Data Protect Cluster. 61
4.7.1 VMware CloudArchive workflow S3 to IBM Cloud Object Storage (ICOS) 61
4.7.2 VMware CloudArchive restore IBM S3 IBM Cloud Object Storage. 65

Chapter 5. Rapid Restore of Virtual Machines 69
5.1 VM Recovery. 70
5.1.1 Recovering VMs 70
5.1.2 Instant Mass Restore 76
5.2 File or Folder Recovery. 77
5.2.1 Indexing and Recovery 78
5.2.2 Recovering Files or Folders 78
5.3 Virtual Disk Recovery 85
5.3.1 Recovering Virtual Disks. 86
5.4 Instant Volume Mount 90
5.4.1 Instant Volume Mounting a VMware virtual disk 91
5.4.2 Teardown of a Mounted Volume. 93

Chapter 6. Integrating IBM Defender Data Protection with IBM Storage Protect 95
6.1 Data Protection Overview 96
6.2 Deployment Exercise Overview 97
6.3 Create an S3 Bucket on the IBM Storage Protect Server 97
6.4 Capture Access Credentials on the SP server 101
6.5 Register the S3 Bucket as External Target on Data Protect Cluster. 101
6.6 Create Polices to Archive Data to The S3 Bucket. 104
6.6.1 Create a Protection Policy for Archiving to the S3 Bucket 104
6.6.2 Configure the Data Protection Group to use the Data Protection Policy. 106
6.7 Monitoring the Protection status 109
6.8 Retrieve from the S3 External Storage 111
6.8.1 Retrieve using Global Search 111
6.8.2 Retrieve using Recoveries Panel 113
6.9 Additional Process Monitoring Examples 115

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.

Trademarks

Preface

This IBM Redbooks publication provides a look at the IBM Storage Defender Data Management Service (DMS) and Data Protect. This includes information to better understand how IBM Storage Defender is used to protect a number of different workloads as well as integrate 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 DMS and Defender Data Protect.

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 based in New Jersey. As a member of the Worldwide Storage Technical Sales Team at IBM, he works with clients, IBM Business Partners, and IBMers around the globe, designing and implementing solutions that address the rapidly evolving cyber 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 using IBM data protection software including DFSMS, DFDSS and Storage Protect. Paul has been with IBM over 4 years and is currently responsible for working with and guiding Storage Protect customers as an IBM Storage Protect/Defender Technical Advisor.

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

Gary Graham is a Brand Technical Specialist covering IBM Defender Data Resilience 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 and currently has a Distinguished Technical Specialist certification from The Open Group.

Richard Hurst is a cyber-resiliency and storage consultant with IBM’s Expert Labs team. With 20+ years experience stretching across multiple disciplines, Richard has provided installation and configuration services as well as providing support for many different name brand products. As an invaluable member of the IBM’s Expert Labs team Richard has worked with multiple IBM products including IBM Defender Data Protect, IBM Storage Ceph and IBM Flashsystems, as well as consulting for IBM’s Cybervault Workshop. Richard Hurst continues to broaden is skill set through Openshift/IBM Fusion as well as Cybersecurity training.

Juan Carlos Jimenez is IBM’s world-wide Data Resiliency Product Manager. He is focused on defining roadmaps, initiatives, and strategy within the various data resiliency software products that he manages alongside his team. Juan Carlos brings an end-to-end view to cyber resilience leveraging his expertise in both storage and security. Juan Carlos developed our Cyber Resiliency Assessment Tool which has been helping numerous enterprises identify and close gaps in their IT environments.

James Morassutti is a Senior Storage Technical Specialist based out of Toronto, Canada. His career in IT spans over 20 years in key areas such x86, Networking, Cyber Resiliency and Storage solutions. James is the National Data Protection SME for Canada, focused on helping clients design solutions to support their Operational Resiliency and protection their critical data to support their Cyber Resiliency Practices. James is passionate about automotive, technology and educated in Mechanical Engineering.

Jack Tedjai is an IBM Certified Expert IT Specialist and IBM Systems subject matter expert, working in the Northern Europe Infrastructure Lab Expert Services organization. He joined IBM in 1998, and has more than 25 years of experience in the delivery of Storage, Storage Virtualize, Backup and Cyber Resilience services for Open Systems. He is mostly involved in architecture and deployments world-wide 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 currently works in the Advanced Technology Group (ATG), IBM Technology, Americas.

Christopher Vollmar Principal, World Wide 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 Cyber Security practices. He is an author of several IBM Redbooks®, an 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

Introduction to IBM Storage Defender

Just a few decades ago, considerations for data resilience were a much simpler. If a company lost or damaged an important file or folder, they'd simply load up the previous day's backup tape, retrieve a copy of the missing data, and return to operating normally from there.

Those days are long gone. Today, the volume of data and diverse range of workloads have made backup and restore operations much 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 may seem like a dangerous place with new and creative attempts to exploit vulnerabilities, careful planning and execution of appropriate data security and data resilience processes can enable organizations to gracefully recover from otherwise dire situations. This Redbooks publication provides guidance on one of IBM's solutions dedicated to these use cases, enabling customers to recover rapidly, and at scale.

In this chapter:

- 1.1, "Overview of IBM Storage Defender" on page 12
- 1.2, "Overview of IBM Defender Data Protect" on page 13
- 1.3, "Overview of IBM Defender Data Management Service" on page 14

1.1 Overview of IBM Storage Defender

Organizations today need a data resilience strategy that encompasses every aspect of their on-premises and cloud environments. One which supports 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 includes virtual machines (VMs), databases, applications, file systems, SaaS workloads, and containers.

IBM Storage Defender features a combination of exceptional scalability, multiple layers of cyber resilience, broad application support, and cost-saving data reduction technologies. By using 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 by support for logical air gap to object storage (WORM technology) as well as the ability to physically air gap data to tape.

Key capabilities of IBM Storage Defender include:

Data Resilience and Compliance

- Set policies and standards to ensure resilience compliance across the data estate. 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 leveraging AI infused technology. It includes a data resiliency service that enhances existing security systems by including storage-specific malware and anomaly detection, as well as providing a trust index to help IT leaders decide where to prioritize the allocation of resources.

Safe and fast recovery, at scale

- Can enable organizations to validate, recover, and restore data more quickly. Completely restore data at scale from an immutable backup or snapshot for each workload, very quickly (eg: 1000s of VMs in minutes).

Flexible licensing

- Licensing is based on resource units (RUs), 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®, FlashSystem, IBM Storage Scale, IBM Storage Ceph, and IBM Storage 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.

Defender is comprised of various components designed to meet customers resilience needs. In this Redbook, we will take a deep dive into two of those components, Defender Data Protect and the Defender Data Management Service.

1.2 Overview of IBM Defender Data Protect

Defender Data Protect (DDP) is one of the many components of IBM Storage Defender. This component offers data management and resiliency for the broadest workload support in the industry.

Defender Data Protect supports the following workloads:

- Hypervisors: VMware, M. Hyper V, Nutanix AHV, and Oracle VM (OVM)
- Databases: Oracle, Oracle (OVM), Oracle RAC, SAP HANA, SAP Oracle, SAP DB2®, SAP MS SQL, SAP Sybase ASE, Sybase IQ & ASE, IBM DB2, MS SQL, Hadoop, IRIS & Cache (EPIC)
- Modern Databases: Cassandra, CouchbaseDB, MySQL, Hbase, MongoDB, PostgreSQL, and Hive
- Cloud-Native Databases: CockroachDB, AWS RDS, and AWS Aurora
- Cloud Applications: AWS VM (EC2), M365, Exchange Online, Azure VM, and Google Compute
- Physical: Windows, AIX®, Linux, and Solaris
- Containers: Kubernetes, and Tanzu
- File Systems: NetApp, IBM Storage Scale, Google EFS, Elastifile, and Pure flash arrays
- Applications: Exchange, Microsoft Active Directory, and Microsoft SharePoint

In this Redbook we will deep dive into how this solution protects the most critical workloads for modern enterprises.

IBM Storage Defender Data Protect boasts a scale-out architecture comprised of clusters. These clusters can be deployed virtually, in the cloud, or on promise through physical nodes. These physical nodes include CPU, Memory, Storage, Network, Operating System, File System, and the Backup Software. An example of these nodes is the IBM Defender Ready Node. By leveraging this cluster and node architecture, Defender Data Protect can execute data management operations like backups, cloning, and restores rapidly, at scale. This is possible by equally spreading the workload or action among all nodes in a cluster. Lastly, upgrades, and expansions can be done easily and non-disruptively by simply adding more nodes to a cluster.

Some of the key capabilities of IBM Storage Defender Data Protect that will be covered in this document are:

Integrated Cybersecurity:

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

Instant Mass Restore:

- Enables users to restore a high number of VMs instantly. For example, 200 VMs in around 10 minutes and 2000 VMs in about 40 minutes. This drastically reduced downtown after an incident like a large-scale malware attack or ransomware attack.

Global Actionable Search:

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

Fast Cloning:

- Extremely fast cloning of large databases for devOps, testing, and other development use cases. For example cloning 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. This reduces the capacity requirements and lowers licensing costs.

Primary Storage Integration:

- Defender Data Protect integrates with the IBM Storage FlashSystem family to backup volumes to Defender Data Protect via volume snapshots, recover from either on-array snaps or offloaded volume backups, and coordinates HW snaps for VM backup to minimize VM stuns.

1.3 Overview of IBM Defender Data Management Service

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

Users can connect Defender Data Protect Clusters, IBM Storage Protect Servers, IBM FlashSystems, and other assets into the service to drive end-to-end data resiliency operations from a single pane of glass interface.

Other noteworthy DMS capabilities include:

Quorum:

- This function limits certain actions, including destructive actions from being carried out by single users. This is achieved by its permission-based nature where user 1 requests an action and a second and/or third user needs to approve the request before it is executed. This could also be considered a form of Two Person Integrity (TPI) checking. This prevents some destructive attacks and reduces 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 the best practice and business needs.

Simulations:

- This functionality 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. This empowers users to proactively plan for various situations, such as acquiring new nodes, integrating new workloads, optimizing current workloads, and more. Simulations can be created with scenarios using specific clusters and time periods to help better understand and plan environment changes.

Reports:

- This function allows users to create and view an overall summary of the data protection jobs and storage systems. Additionally, users can analyze data at the granular level using powerful filtering options. Filter, schedule, email, and download reports to ensure users who needs detailed information on the environment and its status get what they need when they need it.

2

IBM Storage Defender: Data Management Service (DMS)

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

This Chapter includes the following topics:

- 2.1, “Data Management Service (DMS) Overview” on page 18
- 2.2, “IBM Storage Defender DMS Dashboard” on page 18

2.1 Data Management Service (DMS) Overview

IBM Storage Defender Data Management Service (DMS) is a SaaS-based management platform that provides the ability to have a single view and global management interface for all IBM Storage Defender Data Protect clusters in your environment. Whether the cluster is on-premises, cloud or Virtual Edition and regardless of cluster size, the 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 take action directly from the search results page. For example, search for all unprotected VMs and create jobs to protect them.
- **SmartAssist:** Provides Defender DMS recommendations based on capacity forecasting and disk failure prediction.
- **Security tools:** Detect threats and other anomalies across clusters with 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 fifteen minutes.

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

2.2 IBM Storage Defender DMS Dashboard

The DMS dashboard summary provides a comprehensive view of the various aspects related to managed Data Protect clusters, this includes cluster health, protection status and reporting.

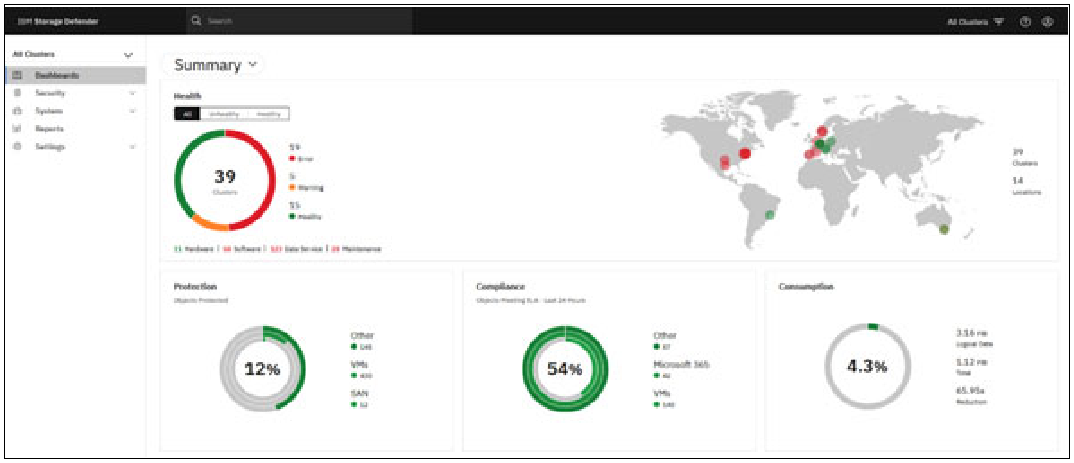

Figure 2-1 IBM Storage Defender DMS summary panel

In the summary dashboard users can easily view the following key functions of DMS:

- The total number of healthy and unhealthy clusters, alert summary, and location of the Data Protect Clusters

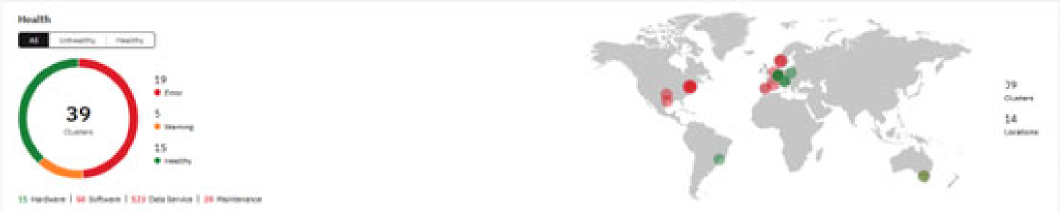

Figure 2-2 DMS cluster health view

- Summary of total objects that have been associated with a backup job

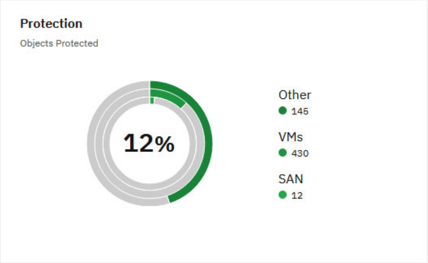

Figure 2-3 DMS backup job protection summary

- Objects that have met their SLA's in the last 24 hours

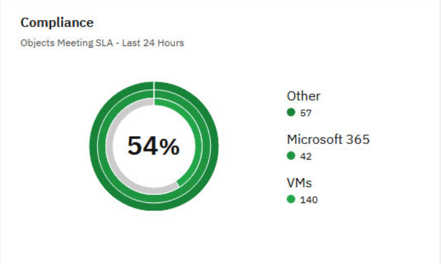

Figure 2-4 DMS backup job status summary

- Storage consumption and data reduction ratios

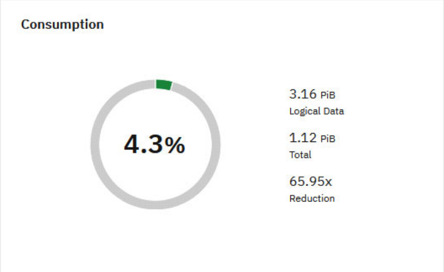

Figure 2-5 DMS storage consumption summary

The Data Protection view provides a topology of the managed data protect clusters and their archive and cloud tier targets as well as a summary of the recoveries made in the last 30 days.

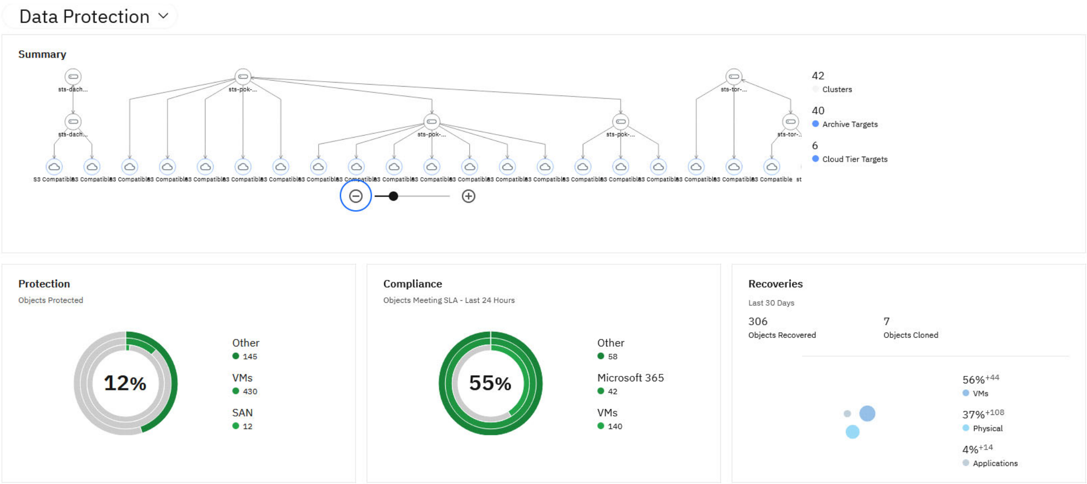

Figure 2-6 DMS data protection topology overview panel

The cloud view provides a summary of cloud activity including data archived\restored as well as consumption by provider.

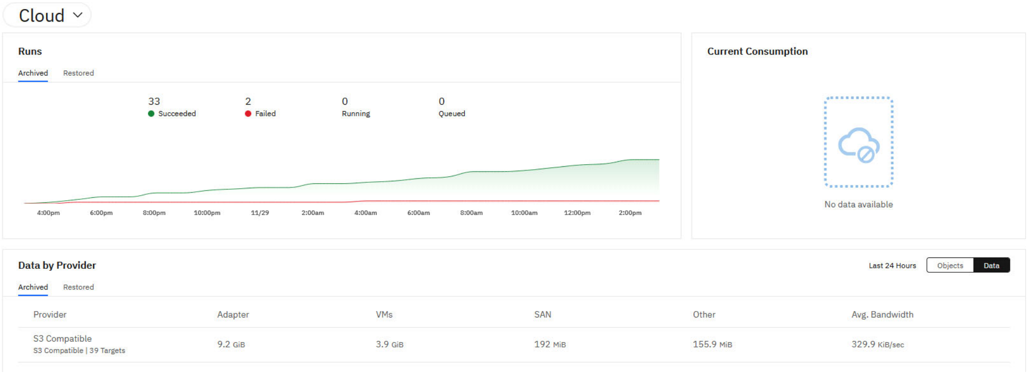

Provider	Adapter	VMs	SAN	Other	Avg. Bandwidth
S3 Compatible S3 Compatible | 39 Targets	9.2 GiB	3.9 GiB	192 MiB	155.9 MiB	329.9 KiB/sec

Figure 2-7 DMS storage provider overview panel

2.2.1 Security

The Security dashboard provides tools which assist administrators to quickly understand their current security stance and provides actionable insights with recommendations related to any potential vulnerabilities and 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 then provides recommendations to improve scores based on a number of key attributes. The scores, icons and other information provided in this feature indicate how the organization’s deployment and configuration compares with IBM’s minimum recommended practices. These practices are designed to supplement (but are not a substitute for) a robust and comprehensive information security program managed by an organization’s designated experts.

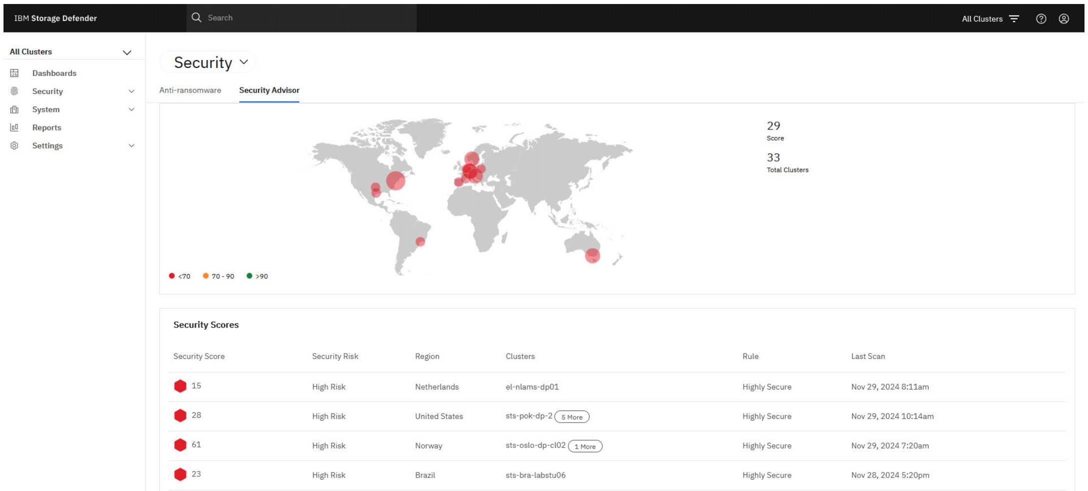

29
Score

33
Total Clusters

Security Scores

Security Score	Security Risk	Region	Clusters	Rule	Last Scan
15	High Risk	Netherlands	el-nlams-dp01	Highly Secure	Nov 29, 2024 8:11am
28	High Risk	United States	sts-pok-dp-2 5 More	Highly Secure	Nov 29, 2024 10:14am
61	High Risk	Norway	sts-oslo-dp-cl02 1 More	Highly Secure	Nov 29, 2024 7:20am
23	High Risk	Brazil	sts-bra-labstu06	Highly Secure	Nov 28, 2024 5:20pm

Figure 2-8 DMS security overview panel

The Security view also summarizes threats detected across backed up objects and their attributes. This is available under the Anit-ransomware header.

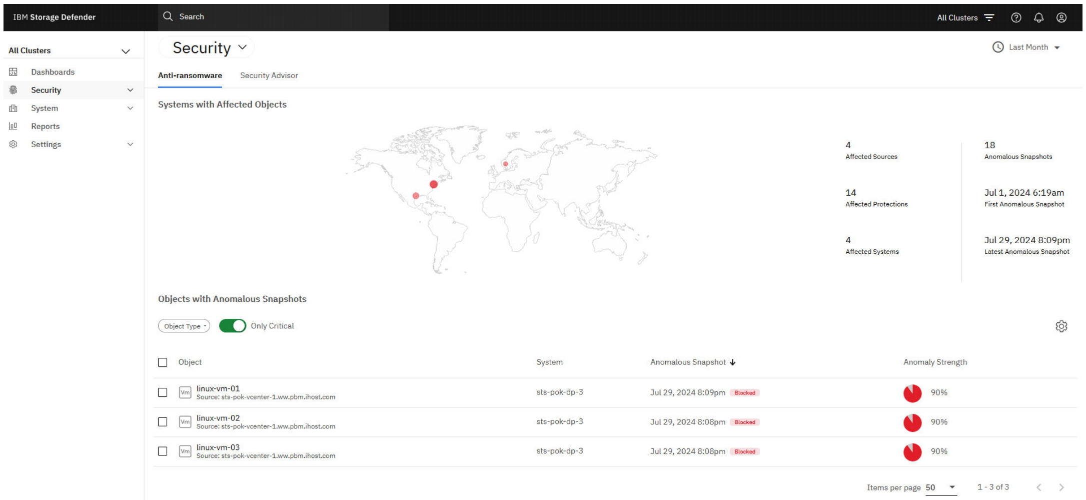

Systems with Affected Objects

4
Affected Sources

14
Affected Protections

4
Affected Systems

18
Anomalous Snapshots

Jul 1, 2024 6:19am
First Anomalous Snapshot

Jul 29, 2024 8:09pm
Latest Anomalous Snapshot

Objects with Anomalous Snapshots

Object	System	Anomalous Snapshot	Anomaly Strength
linux-vm-01 Source: sts-pok-vcenter-1.ww.pbm.ihost.com	sts-pok-dp-3	Jul 29, 2024 8:09pm Blocked	90%
linux-vm-02 Source: sts-pok-vcenter-1.ww.pbm.ihost.com	sts-pok-dp-3	Jul 29, 2024 8:08pm Blocked	90%
linux-vm-03 Source: sts-pok-vcenter-1.ww.pbm.ihost.com	sts-pok-dp-3	Jul 29, 2024 8:08pm Blocked	90%

Items per page 50 1 - 3 of 3

Figure 2-9 DMS security threat overview panel

Users have the ability to customize the action thresholds, manage notifications, and block recoveries that have been tagged as a security risk.

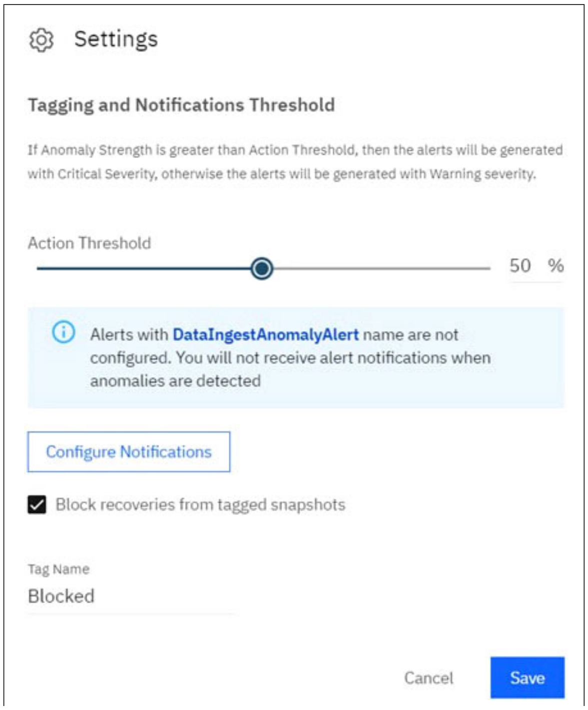

Settings

Tagging and Notifications Threshold

If Anomaly Strength is greater than Action Threshold, then the alerts will be generated with Critical Severity, otherwise the alerts will be generated with Warning severity.

Action Threshold

50 %

Alerts with **DataIngestAnomalyAlert** name are not configured. You will not receive alert notifications when anomalies are detected

☑ Block recoveries from tagged snapshots

Tag Name

Blocked

Figure 2-10 DMS notification settings panel

Quorum

As a part of the Security section from DMS in the left-hand navigation tab you can explore as well as the Quorum function. In the nature of Two-Person Integrity (TPI), Quorum approvals is a feature of DMS than ensures a pre-defined approver, or group of approvers must approve of actions requested by users prior to those actions taking place. Quorum helps eliminate risk of destructive operations being performed by administrators due to malicious or accidental actions.

The quorum dashboard allows users to view all quorum requests that are pending approval as well as the current user’s requests. From here, 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.

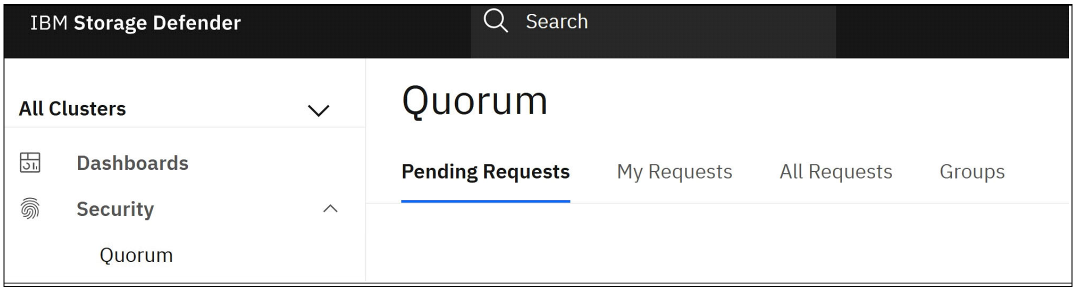

Figure 2-11 Quorum dashboard

Creating a Quorum Group

Leveraging the Quorum function, users are able to create Quorum Group, that allows the restriction of a variety of pre-designated functions in a selectable fashion referred to as “Operations” which will be governed by the group. From there administrators are able to select the users that need to approve the actions, how many of the group that need to apply, for example 3 out of 5 or 2 of 4 etc. The timeline for approval is also a selectable option, allowing the request a certain amount of to be approved or else it will decline automatically. Members of the Quorum Group will be notified by email of a pending approval.

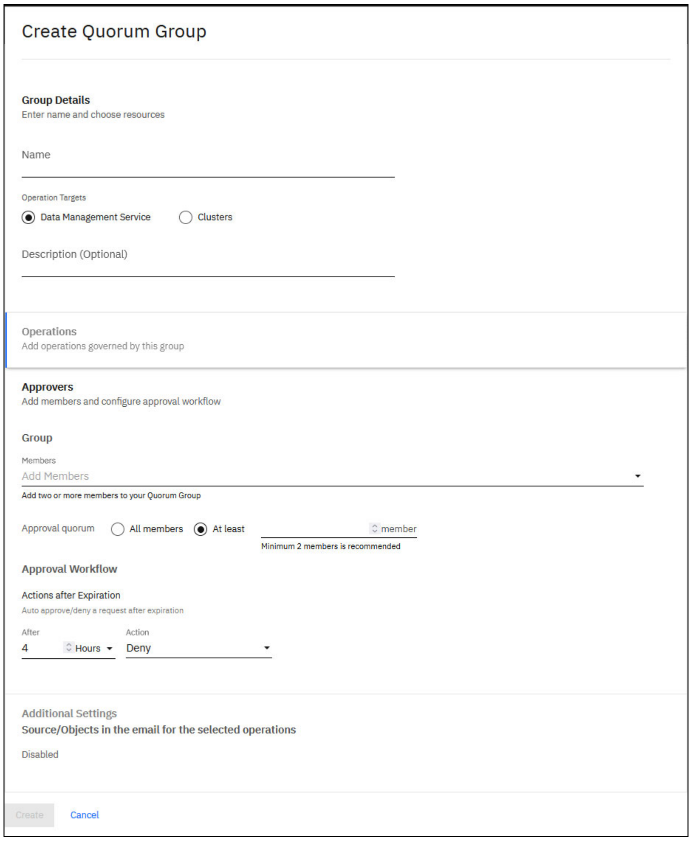

Create Quorum Group

Group Details
Enter name and choose resources

Name

Operation Targets
● Data Management Service ○ Clusters

Description (Optional)

Operations
Add operations governed by this group

Approvers
Add members and configure approval workflow

Group
Members
Add Members
Add two or more members to your Quorum Group

Approval quorum ○ All members ● At least member
Minimum 2 members is recommended

Approval Workflow

Actions after Expiration
Auto approve/deny a request after expiration

After: 4 Hours
Action: Deny

Additional Settings
Source/Objects in the email for the selected operations

Disabled

Figure 2-12 Create Quorum Group panel

2.2.2 System

The System tab provides access to the various tools for users to manage the Data Protect Deployment. By selecting the system tab users are presented with the following options:

Health Panel

The health panel (Figure 2-13 on page 24) provides administrators with a consolidated list of alerts that have been generated by the various clusters in the Data Protect deployment.

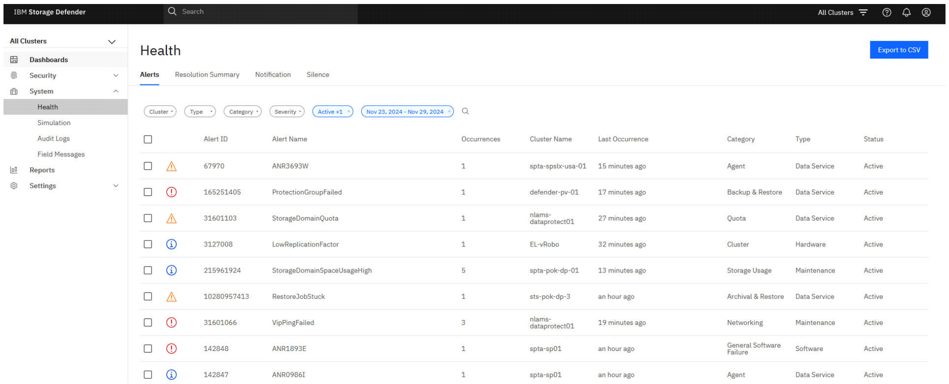

Alert ID	Alert Name	Occurrences	Cluster Name	Last Occurrence	Category	Type	Status
67970	ANR3693W	1	spta-spslx-usa-01	15 minutes ago	Agent	Data Service	Active
165251405	ProtectionGroupFailed	1	defender-pv-01	17 minutes ago	Backup & Restore	Data Service	Active
31601103	StorageDomainQuota	1	nlams-dataprotect01	27 minutes ago	Quota	Data Service	Active
3127008	LowReplicationFactor	1	EL-vRobo	32 minutes ago	Cluster	Hardware	Active
215961924	StorageDomainSpaceUsageHigh	5	spta-pok-dp-01	13 minutes ago	Storage Usage	Maintenance	Active
10280957413	RestoreJobStuck	1	sts-pok-dp-3	an hour ago	Archival & Restore	Data Service	Active
31601066	VipPingFailed	3	nlams-dataprotect01	19 minutes ago	Networking	Maintenance	Active
142848	ANR1893E	1	spta-sp01	an hour ago	General Software Failure	Software	Active
142847	ANR0986I	1	spta-sp01	an hour ago	Agent	Data Service	Active

Figure 2-13 IBM Storage Defender Health panel showing a list of Alerts

Alerts

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

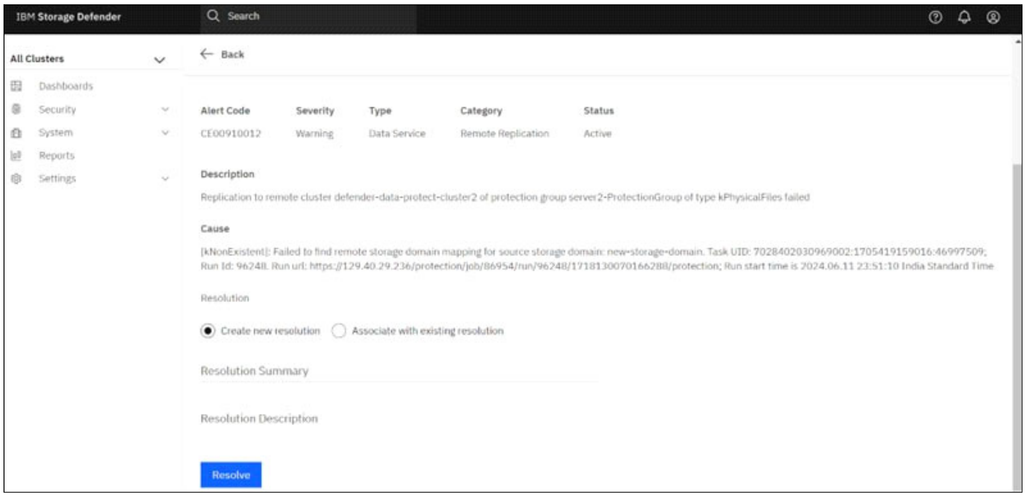

Alert Code	Severity	Type	Category	Status
CE00910012	Warning	Data Service	Remote Replication	Active

Description

Replication to remote cluster defender-data-protect-cluster2 of protection group server2-ProtectionGroup of type kPhysicalFiles failed

Cause

[kNonExistent]: Failed to find remote storage domain mapping for source storage domain: new-storage-domain. Task UID: 7028402030969002:1705419159016:46997509; Run Id: 96248. Run url: https://129.40.29.236/protection/job/86954/run/96248/1718130070166288/protection; Run start time is 2024.06.11 23:51:10 India Standard Time

Resolution

● Create new resolution ○ Associate with existing resolution

Resolution Summary

Resolution Description

Figure 2-14 Alert details panel

- The resolution summary tab provides administrators with a consolidated view of all the previously generated resolutions
- The silence tab allows users to create rules to suppress certain notifications based on the cluster, severity of the alert, category, type and names
- The notification tab allows users to customize notification delivery via email or webhooks, which allows customization on the alerting and the type of alerts as depicted in Figure 2-15 on page 25 (below)
- Users are able to create alerts specific to clusters, alert levels, preset alert types, alert categories and alert names with deep granularity

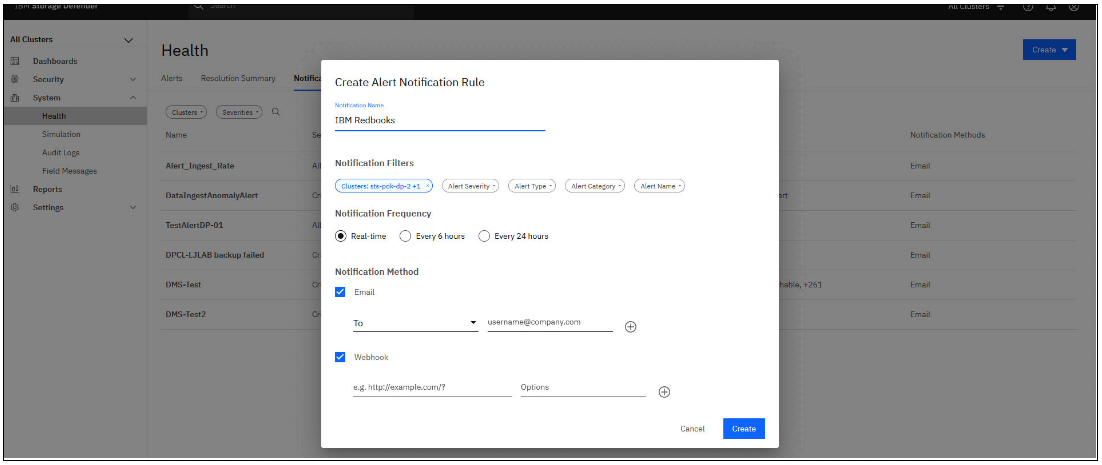

Name	Se		Notification Methods
Alert_Ingest_Rate	All		Email
DataIngestAnomalyAlert	Cr	ert	Email
TestAlertDP-01	All		Email
DPCL-LJLAB backup failed	Cr		Email
DMS-Test	Cr	hable, +261	Email
DMS-Test2	Cr		Email

Create Alert Notification Rule

Notification Name
IBM Redbooks

Notification Filters
Clusters: sts-pok-dp-2 +1 × | Alert Severity | Alert Type | Alert Category | Alert Name

Notification Frequency
● Real-time ○ Every 6 hours ○ Every 24 hours

Notification Method
☑ Email
To username@company.com
☑ Webhook
e.g. http://example.com/? Options

Cancel Create

Figure 2-15 Create notification rule dialog panel

Simulation

The Simulation tab (Figure 2-16 on page 25) provides administrators with the ability to simulate new workloads against the existing deployment to help predict storage consumption over time. This is a powerful tool to help administrator proactively plan, if and when additional storage capacity may be required to support environment changes.

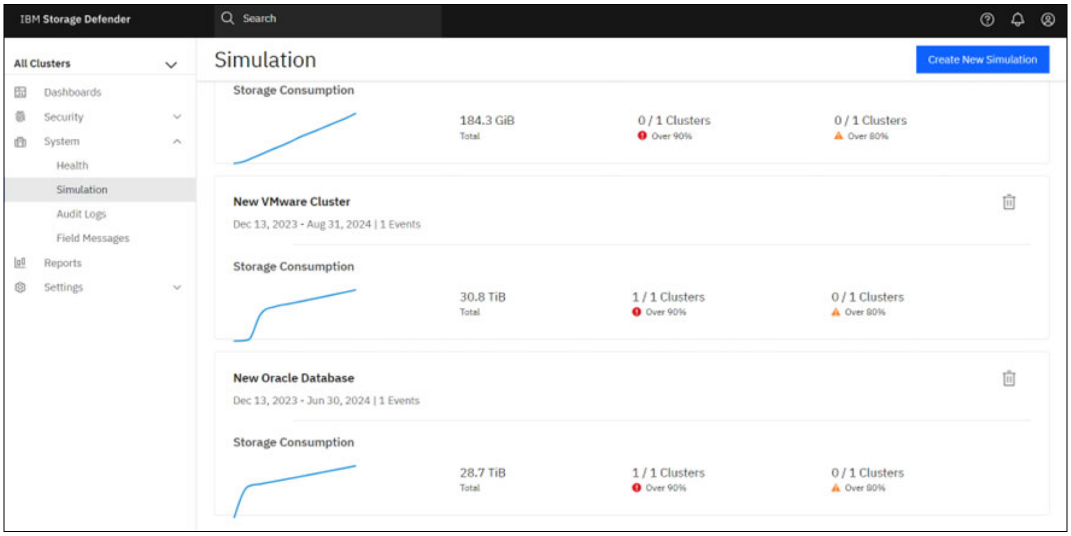

Storage Consumption
184.3 GiB Total | 0 / 1 Clusters Over 90% | 0 / 1 Clusters Over 80%

New VMware Cluster
Dec 13, 2023 - Aug 31, 2024 | 1 Events
Storage Consumption
30.8 TiB Total | 1 / 1 Clusters Over 90% | 0 / 1 Clusters Over 80%

New Oracle Database
Dec 13, 2023 - Jun 30, 2024 | 1 Events
Storage Consumption
28.7 TiB Total | 1 / 1 Clusters Over 90% | 0 / 1 Clusters Over 80%

Figure 2-16 Simulation details panel

To create a new simulation, select the blue toggle on the upper right corner of the screen. From here users can give the simulation a name as well as an end date for the simulation. It will create the ability to model the effect on the IBM Defender Data Protect Cluster of a variety of factors that include adding resources, adding new workloads (protection groups) or expanding the protection groups or Storage domains.

2.2.3 Reporting

The built in reporting tool allows users to schedule or create ad-hoc reports organization wide or narrowed down to specific clusters or workloads. Reports can be filtered, scheduled, emailed or downloaded as desired.

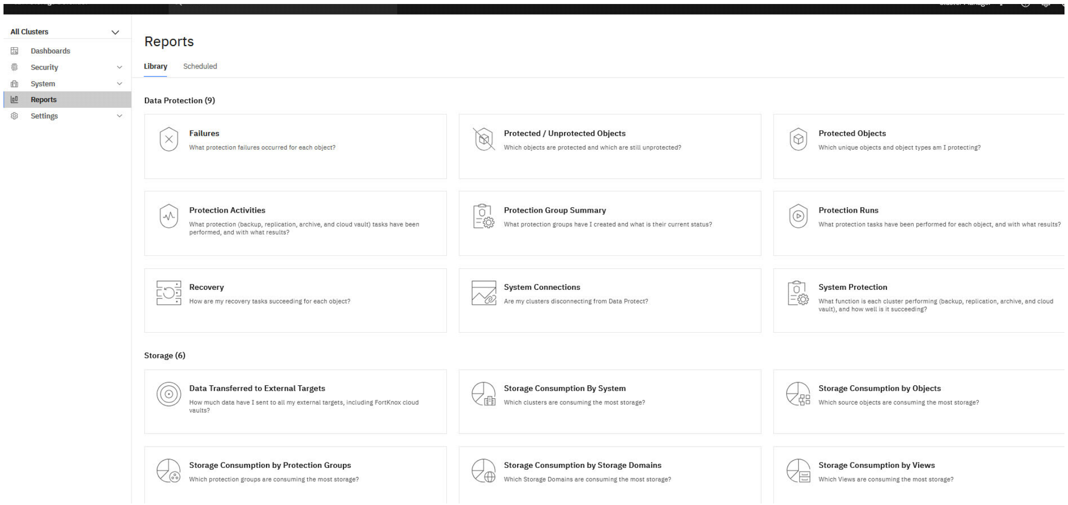

Data Protection (9)

Failures
What protection failures occurred for each object?

Protected / Unprotected Objects
Which objects are protected and which are still unprotected?

Protected Objects
Which unique objects and object types am I protecting?

Protection Activities
What protection (backup, replication, archive, and cloud vault) tasks have been performed, and with what results?

Protection Group Summary
What protection groups have I created and what is their current status?

Protection Runs
What protection tasks have been performed for each object, and with what results?

Recovery
How are my recovery tasks succeeding for each object?

System Connections
Are my clusters disconnecting from Data Protect?

System Protection
What function is each cluster performing (backup, replication, archive, and cloud vault), and how well is it succeeding?

Storage (6)

Data Transferred to External Targets
How much data have I sent to all my external targets, including FortKnox cloud vaults?

Storage Consumption By System
Which clusters are consuming the most storage?

Storage Consumption by Objects
Which source objects are consuming the most storage?

Storage Consumption by Protection Groups
Which protection groups are consuming the most storage?

Storage Consumption by Storage Domains
Which Storage Domains are consuming the most storage?

Storage Consumption by Views
Which Views are consuming the most storage?

Figure 2-17 DMS Reports panel

DMS provides a comprehensive list of built-in reports to help users gather detailed information about the environment. At the time of publication, 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. Users can further customize the results be selecting from a number of filtered options available within each report.

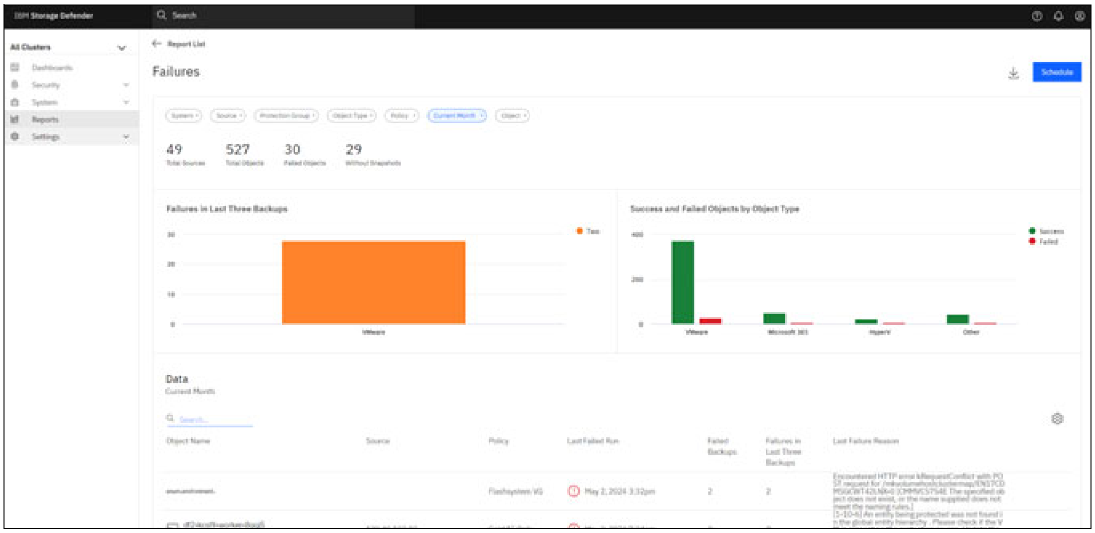

Figure 2-18 DMS report example

Reports can be downloaded locally by selecting the download icon and choosing the desired file format or scheduled for delivery in a variety of formats.

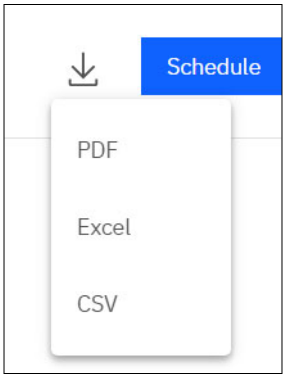

Figure 2-19 Download report formats

2.2.4 System View

The DMS consolidates system information of all the Data Protect clusters under management to provide a single pane of glass view of the health organization wide. The DMS Dashboard provides the following information:

- Health information of all clusters. This includes healthy clusters as well as clusters that have warnings or errors.
- A protection view. Which reports on the total number of objects currently under protection.
- A compliance view. Which reports on the total number of objects that have met their SLA’s over the last 24 hours.
- A Consumption view. Which consolidates the total storage consumed as well as available capacity and overall data reduction ratios.

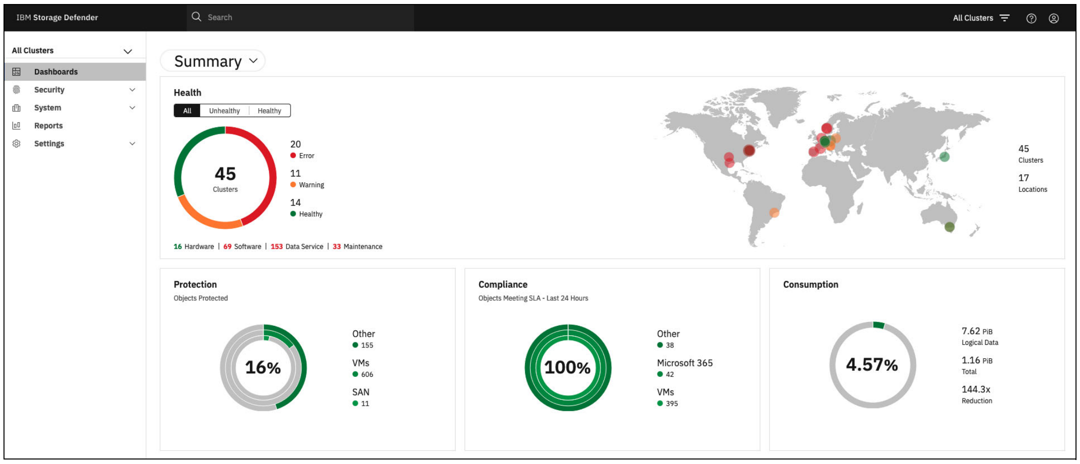

Figure 2-20 IBM Storage Defender Data Management Service Dashboard "System Board"

3

IBM Storage Defender: Data Protect

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

The sections below, we will introduce the new architecture and key components, then explore several deployment options to build an IBM Defender Data Protect and Replica solution to service Instant Mass Restore capabilities for hundreds of VMs at once.

This chapter provides, describes, discusses, or contains the following:

- 3.1, “Solution Overview” on page 30
- 3.2, “What is a Cluster and How does It Work?” on page 30
- 3.3, “Physical / Virtual / Robo Implementations” on page 31
- 3.4, “Cluster setup” on page 32
- 3.5, “IBM Storage Ready Nodes” on page 39
- 3.6, “IBM Storage Defender Ready Nodes” on page 40
- 3.7, “Registering with IBM Data Management Service” on page 41
- 3.8, “VMware Workload Use Case” on page 42

3.1 Solution Overview

The diagram below shows an IBM Defender Data Protect environment with the management plane called the Data Management service (DMS for short). The DMS is cloud based and, is licensed using a SaaS (Software as a Service) subscription. The Data Protection Resources below the management plane can be scaled out to fulfill the services requested by the DMS.

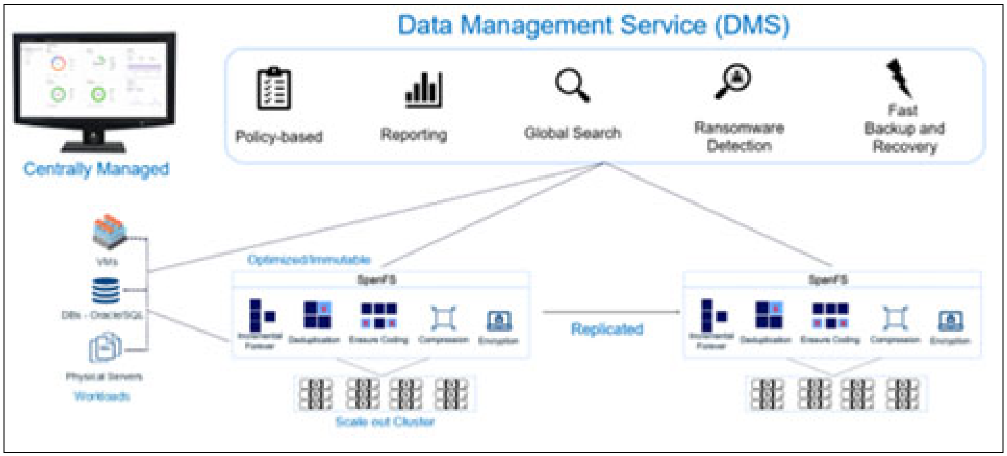

Figure 3-1 IBM Storage Defender Data Protect Architecture Overview

3.1.1 Node Based Scale Out Architecture

IBM Defender Data Protect 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. A cluster can be created on a physical node or a virtual appliance that runs the IBM Defender Data Protect software. Each node has CPU, memory, network and storage capacity resources that are managed by the Data Protect software. Nodes can operate as a single node cluster (ROBO node, Virtual Edition node) or in groups of nodes called clusters. A minimum of four nodes are required to form an IBM Data Protect cluster. Three node clusters are possible, but the recommendation for n+2 node redundancy would require four nodes to allow the cluster to run with up to two nodes being unavailable.

3.2 What is a Cluster and How does It Work?

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

SpanFS is an integrated part of Defender Data Protect and provides file services for the data protection cluster, spreading the IO 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. This reduces the capacity requirements and lowers software licensing costs. IBM Defender Data Protect provides immutable storage, data encryption and multi-site replication to increase data resiliency.

Immutability ensures 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 Data Protect cluster can write to the file system through trusted APIs.

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

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

The in line 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 is also default always on capability but can be turned off in selected Storage Domains. IBM Defender Data Protect leverages in line ZSTD compression. Replication - Organizations can achieve enterprise-level resiliency with site-to-site replication between IBM Defender Data Protect clusters. All data on a single IBM Storage Defender Data Protect cluster can be replicated to one or more clusters through the use of Protection Groups. The Protection Group specifies the objects (like databases, physical servers, VMs, Views etc.) to be backed up and replicated. Full and partial fail-over functions are fully automated IBM Defender Data Protect at ether the cluster level or the DMS level for simple, seamless recovery.

3.3 Physical / Virtual / Robo Implementations

A cluster can be built to run in traditional datacenters (like Production and D/R), or on the edge of the enterprise at a remote office or branch office using physical or virtual nodes, or out in the public cloud using virtual edition.

The table below shows multiple deployment options for Virtual and Physical nodes as well as some general suggestions and deployment examples. In this table FETB = Front End TB or the amount of data in a workload to be protected.

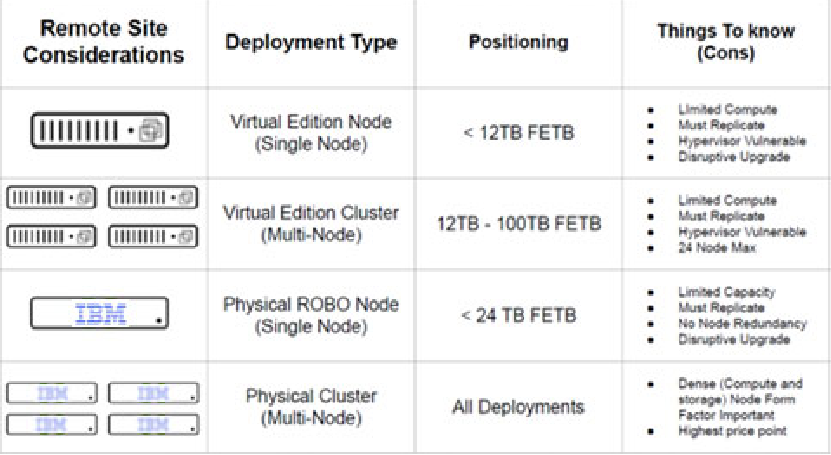

Remote Site Considerations	Deployment Type	Positioning	Things To know (Cons)
	Virtual Edition Node (Single Node)	< 12TB FETB	• Limited Compute • Must Replicate • Hypervisor Vulnerable • Disruptive Upgrade
	Virtual Edition Cluster (Multi-Node)	12TB - 100TB FETB	• Limited Compute • Must Replicate • Hypervisor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

Figure 3-2 Shows a comparison of Virtual and Physical deployment options

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

Physical deployment options include the ROBO (Remote Office / Branch Office) single node that offers more compute power than the single virtual node, and limited capacity. Since this is a single node, there is no node redundancy and there is limited disk capacity. To protect the ROBO node, set up replication to a Data Protect cluster to protect the backup data and provide an additional copy of the backup data at a centrally managed site.

While IBM Defender Data Protect is designed with IBM Storage Ready Nodes in mind it does not require the use of a particular vendor's hardware offerings. Several different vendor platforms have been tested and approved for use including Dell, Cisco, HP and others. For more information on vendor platforms see the following link (requires login with IBM ID):
https://www.ibm.com/support/pages/node/6985577

Physical clusters are the best options for scaling performance and capacity to meet the needs of Production backup windows and replication to the Disaster Recovery site. Also note the cluster resources are available for Cyber Resiliency design considerations like Instant Mass Restore capability and recovery assurance testing.

3.4 Cluster setup

Details for setting up the cluster will vary depending on the choice of hardware or virtual servers being used for the cluster nodes. See the following link for details (requires login with IBM ID): https://ibm.biz/BdGrAv

For the exercise in this Redbook, we will be summarizing the cluster setup process using IBM Storage Ready Nodes

A pre-requisite to having the services team onsite would be to complete the following tasks:

- Racking the cluster
- Cabling the cluster
- Verifying prerequisites for cluster setup
- Configuring the BIOS on the IBM Storage Ready Nodes

3.4.1 Configuration of a Data Protect Storage Domain

Use the following steps to configure a Defender Data Protect Storage Domain

1. Expand settings
2. Select summary
3. Click storage domains
4. Click on Create Storage Domain

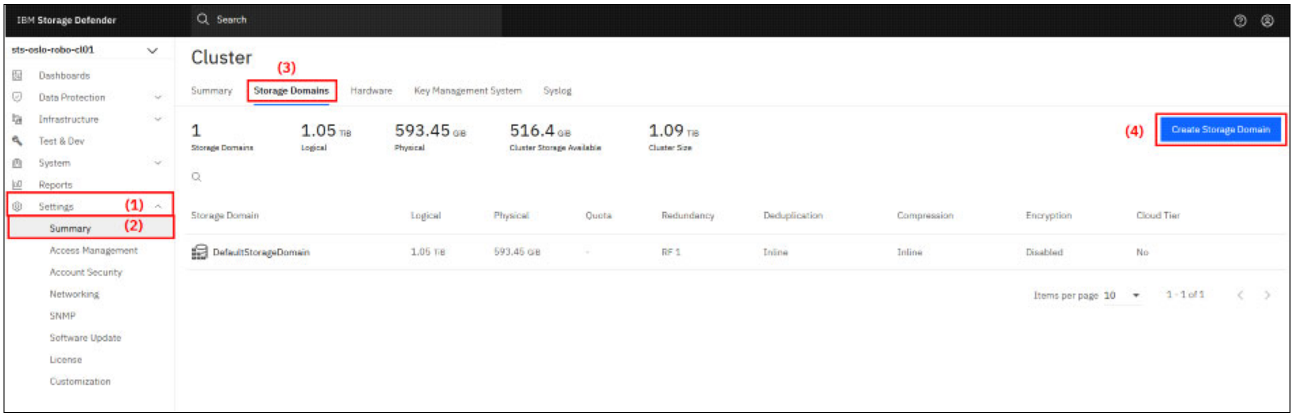

1 Storage Domains
1.05 TiB Logical
593.45 GiB Physical
516.4 GiB Cluster Storage Available
1.09 TiB Cluster Size

Storage Domain	Logical	Physical	Quota	Redundancy	Deduplication	Compression	Encryption	Cloud Tier
DefaultStorageDomain	1.05 TiB	593.45 GiB	-	RF 1	Inline	Inline	Disabled	No

Items per page 10 1 - 1 of 1

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 deduplication and compressions if requirements differ from default

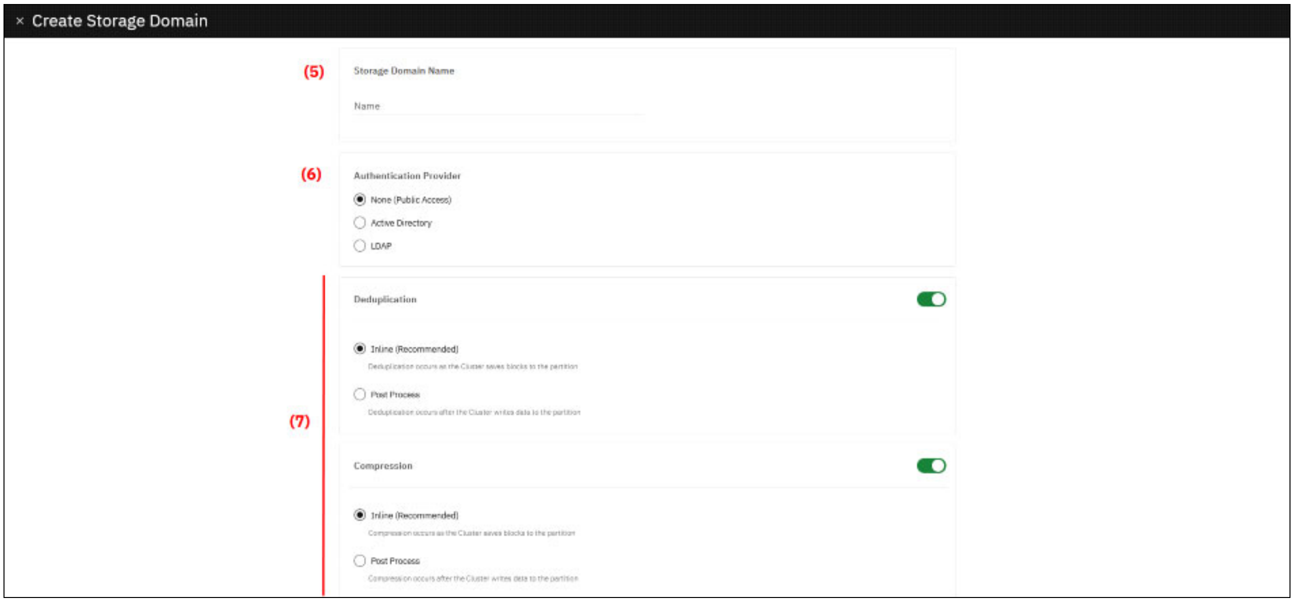

Storage Domain Name

Name

Authentication Provider

● None (Public Access)

○ Active Directory

○ LDAP

Deduplication

● Inline (Recommended)

Deduplication occurs as the Cluster saves blocks to the partition

○ Post Process

Deduplication occurs after the Cluster writes data to the partition

Compression

● Inline (Recommended)

Compression occurs as the Cluster saves blocks to the partition

○ Post Process

Compression occurs after the Cluster writes data to the partition

Figure 3-4 Storage Domain configuration steps 5-7

> **Note:** Encryption is either on or off. It cannot be modified after creation. IF encryption needs to be disabled, you will have to delete the Storage Domain and recreate as needed or create another storage domain with the same setting including encryption.

8. Enable encryption if required
9. Expand more option
10. Enable quota and alert options if required
11. Set the quota limit for each as needed

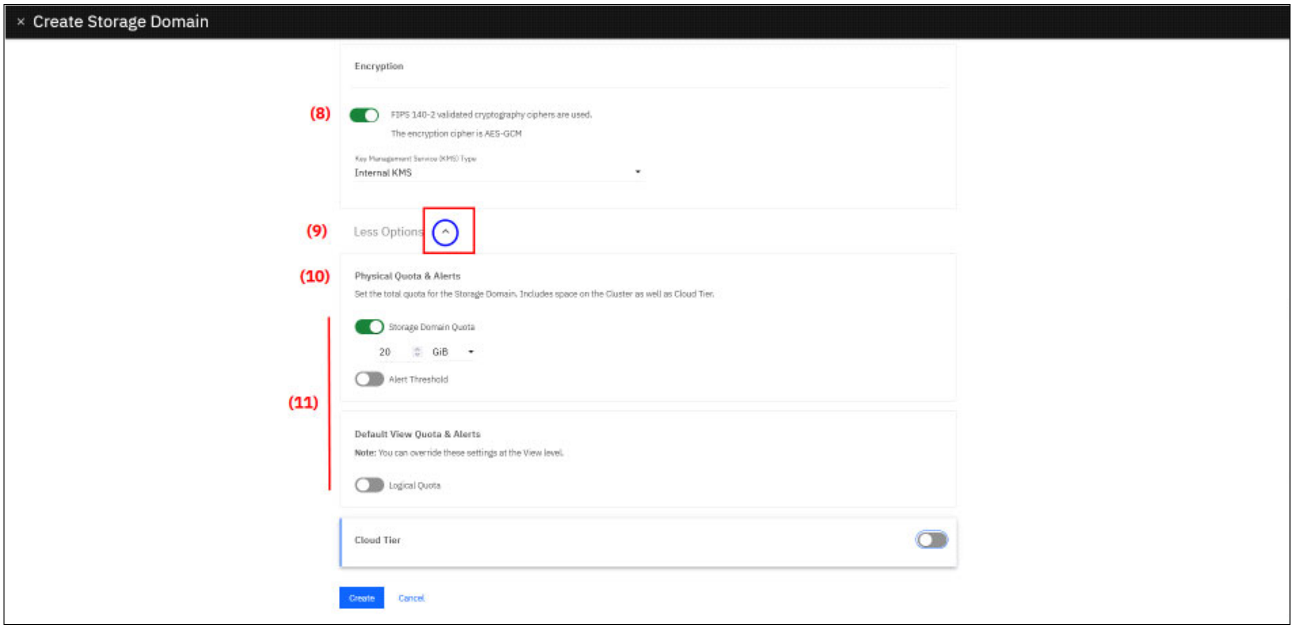

Encryption

FIPS 140-2 validated cryptography ciphers are used.

The encryption cipher is AES-GCM

Key Management Service (KMS) Type

Internal KMS

Less Options

Physical Quota & Alerts

Set the total quota for the Storage Domain. Includes space on the Cluster as well as Cloud Tier.

Storage Domain Quota

20 GiB

Alert Threshold

Default View Quota & Alerts

Note: You can override these settings at the View level.

Logical Quota

Cloud Tier

Figure 3-5 Storage Domain configuration steps 8-11

12. Enable Cloud tier if designed for external S3 target
13. Select the external target
14. Specify the Storage Type, Storage class and other options, then Register

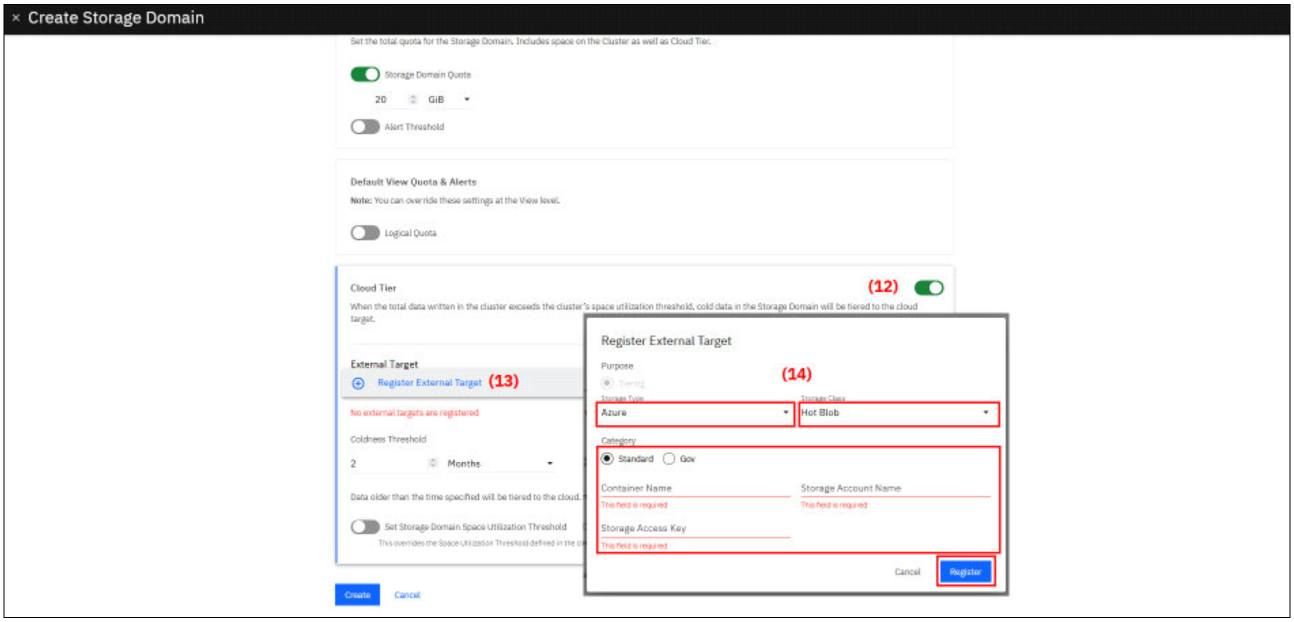

× Create Storage Domain

Set the total quota for the Storage Domain. Includes space on the Cluster as well as Cloud Tier.

Storage Domain Quota: on

20 GiB

Alert Threshold: off

Default View Quota & Alerts

Note: You can override these settings at the View level.

Logical Quota: off

Cloud Tier: on

When the total data written in the cluster exceeds the cluster's space utilization threshold, cold data in the Storage Domain will be tiered to the cloud target.

External Target

Register External Target

No external targets are registered

Coldness Threshold

2 Months

Data older than the time specified will be tiered to the cloud.

Set Storage Domain Space Utilization Threshold: off

This overrides the Space Utilization Threshold defined in the cl

Register External Target

Purpose

● Tiering

Storage Type: Azure

Storage Class: Hot Blob

Category

● Standard ○ Gov

Container Name

This field is required

Storage Account Name

This field is required

Storage Access Key

This field is required

Figure 3-6 Storage Domain configuration steps 12-14

15. Configure Threshold for moving data to Cloud tier
16. Select the required fault tolerance and redundancy requirement
17. If required enable Erasure Coding
18. Select Erasure Coding (EC) configuration from drop down
19. Click create

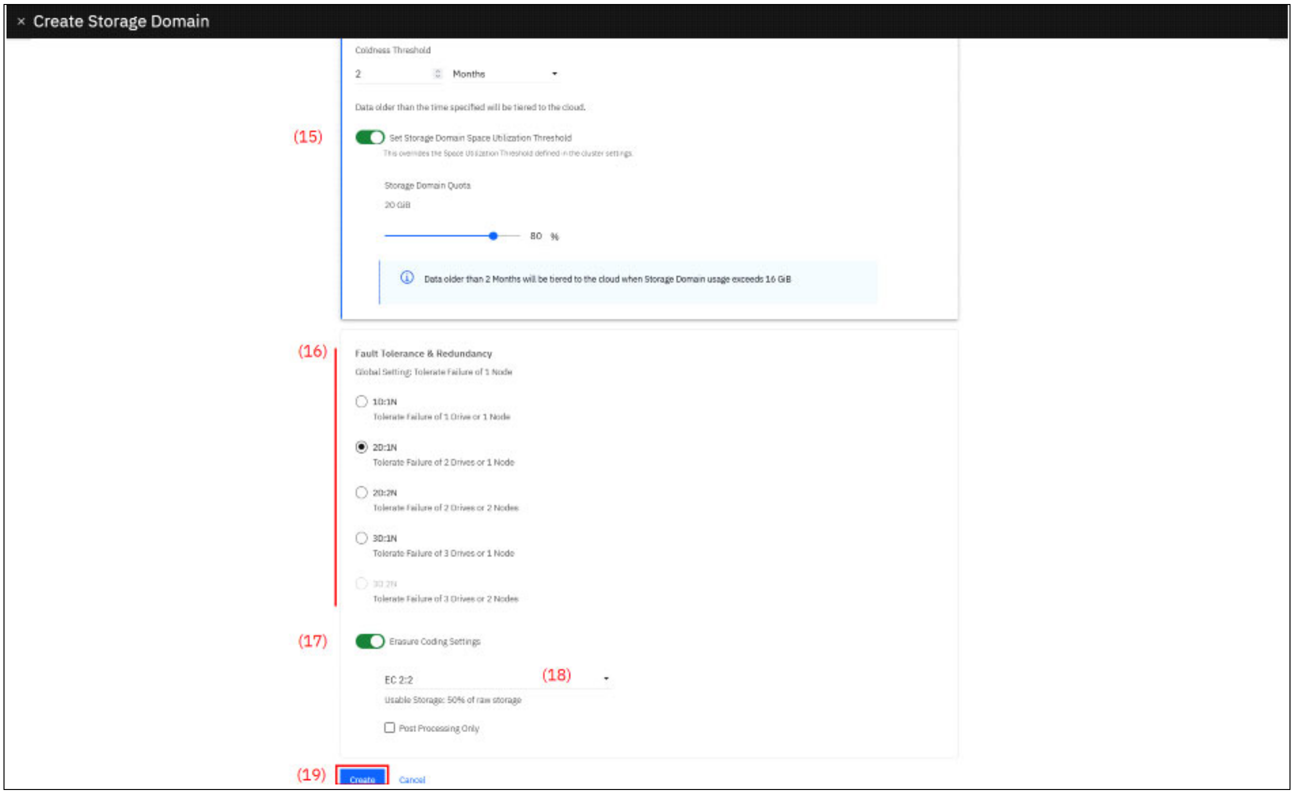

× Create Storage Domain

Coldness Threshold

2 Months

Data older than the time specified will be tiered to the cloud.

(15) Set Storage Domain Space Utilization Threshold

This overrides the Space Utilization Threshold defined in the cluster settings.

Storage Domain Quota

20 GiB

80 %

Data older than 2 Months will be tiered to the cloud when Storage Domain usage exceeds 16 GiB

(16) Fault Tolerance & Redundancy

Global Setting: Tolerate Failure of 1 Node

○ 10:1N
Tolerate Failure of 1 Drive or 1 Node

● 20:1N
Tolerate Failure of 2 Drives or 1 Node

○ 20:2N
Tolerate Failure of 2 Drives or 2 Nodes

○ 30:1N
Tolerate Failure of 3 Drives or 1 Node

○ 30:2N
Tolerate Failure of 3 Drives or 2 Nodes

(17) Erasure Coding Settings

EC 2:2 (18)

Usable Storage: 50% of raw storage

□ Post Processing Only

Figure 3-7 Storage Domain configuration steps 15-19

A note on Fault Tolerance and Redundancy:

Data Protect supports the configuration of fault tolerance at the cluster level, as well as at the storage domain level. When deploying a cluster, a cluster-level fault tolerance is selected. This fault tolerance setting becomes the default value for any new storage domains created on the cluster. Optionally, a different fault-tolerance setting can be selected when creating a new storage domain or editing an existing domain. Modification of storage domain fault tolerance is a non-disruptive action and can be performed without taking an outage.

This feature proves invaluable in environments that grow over time. As the amount of physical hardware within a Data Protect cluster grows to meet the increasing storage demands, or changes are made to the protected workloads, the likelihood of a hardware change increases too. Data Protect can adapt to provide increased resilience in such scenarios without requiring end-user disruption.

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

It’s also important to note that the available Erasure Coding options in step 18 in Figure 3-7 depend directly on the Fault Tolerance value selected in step 16.

3.4.2 Configuration of a Data Protect Protection Group

This section demonstrates how to create a Protection Group, combining all of the previously configured base components together to protect a VMware cluster or ESXi host.

1. Expand Data Protection

2. Select Protection
3. Click on Protect
4. Choose Virtual machine

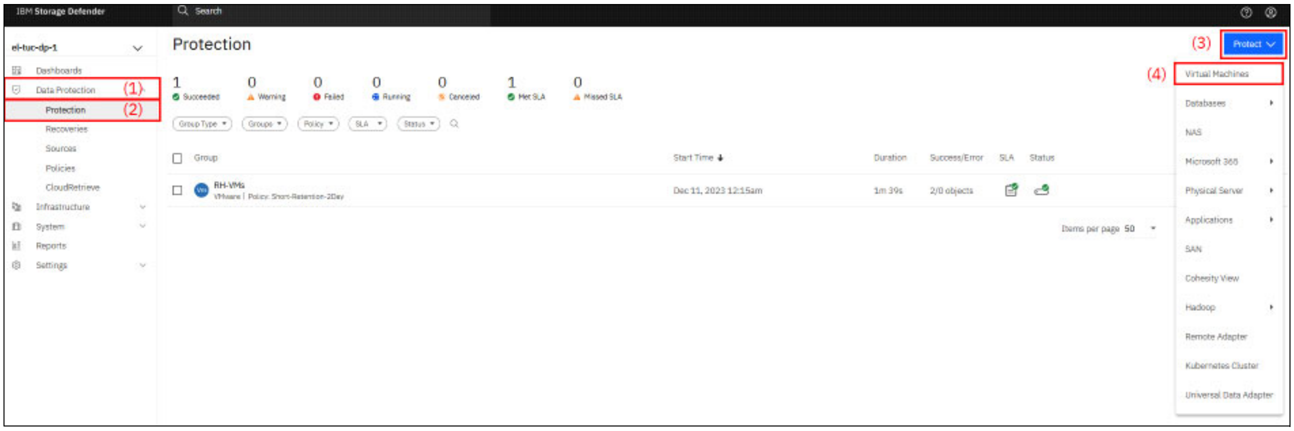

Figure 3-8 Data Protection Policy for virtual machines configuration steps 1-4

A new Protection Wizard window will pop up to complete the configuration:

5. Click Add Objects
6. Select your registered VMWare source
7. Search for and select the desired VMs/hosts for protection
8. Click continue

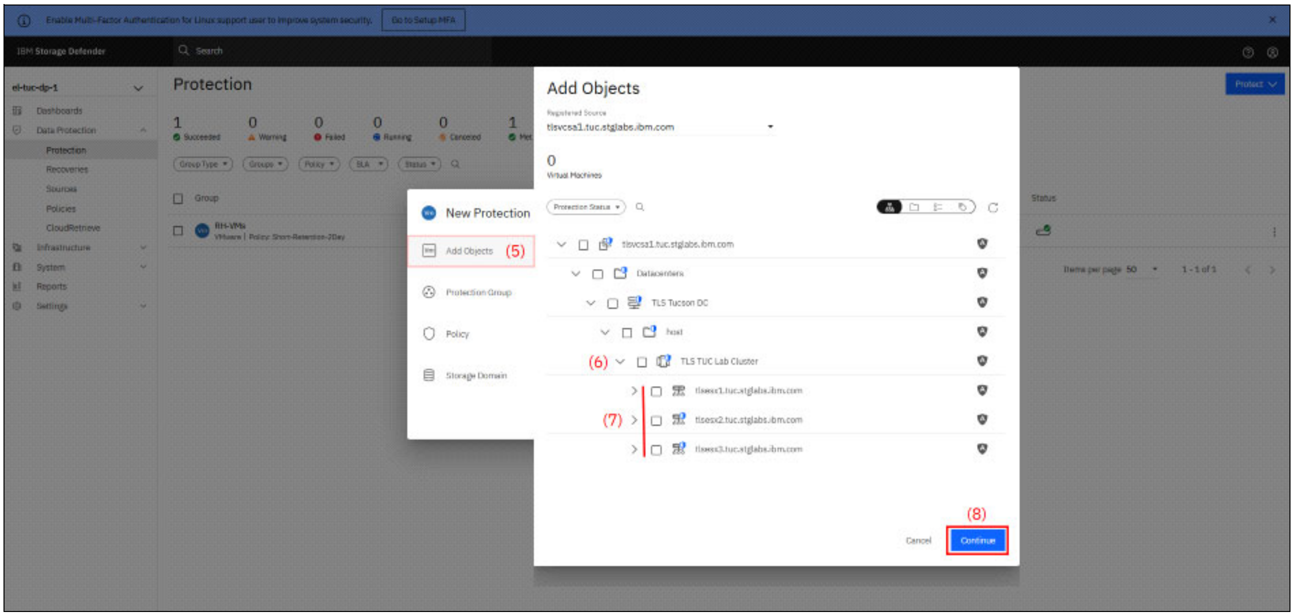

Figure 3-9 Virtual Machine Data Protection Policy configuration steps 5-8

> **Note:** It is possible to view and select VM's and objects by the folder, list of tag views however it is suggested to always use VMware vSphere tags support when selecting objects to ensure consistent data management practices. To use tags select the Tags icon (1) and expand the tags list (2) then select the desired VMware tag (3).

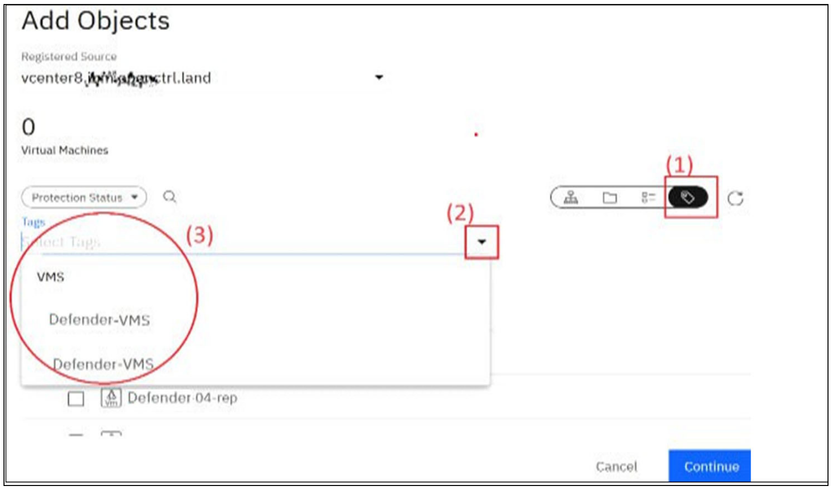

Figure 3-10 Virtual Machine Data Protection Policy configuration, vSphere tags

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

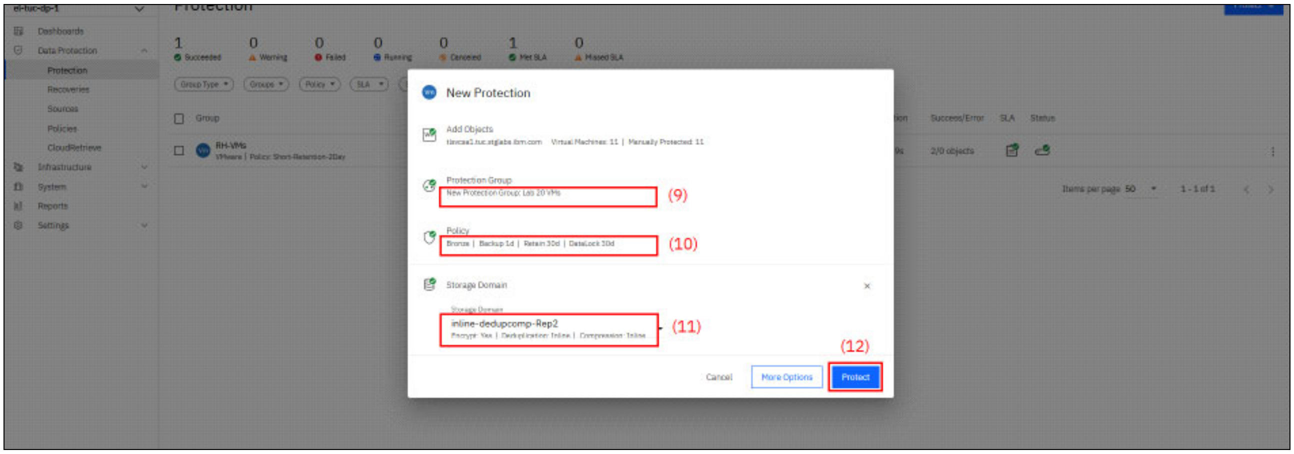

Add Objects: tlevcsa1.tuc.stglabs.ibm.com Virtual Machines: 11 | Manually Protected: 11

Protection Group: New Protection Group: Lab 20 VMs

Policy: Bronze | Backup 1d | Retain 30d | DataLock 30d

Storage Domain: inline-dedupcomp-Rep2
Encrypt: Yes | Deduplication: Inline | Compression: Inline

Figure 3-11 Protection Policy for virtual machines configuration steps 8-12

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

3.4.3 Monitor the Data Protect Protection Group

This section demonstrates how to monitor a Protection Group task:

1. Expand Data Protection
2. Select Protection
3. Usage filter option and typing the VMware Tags “VMtags”

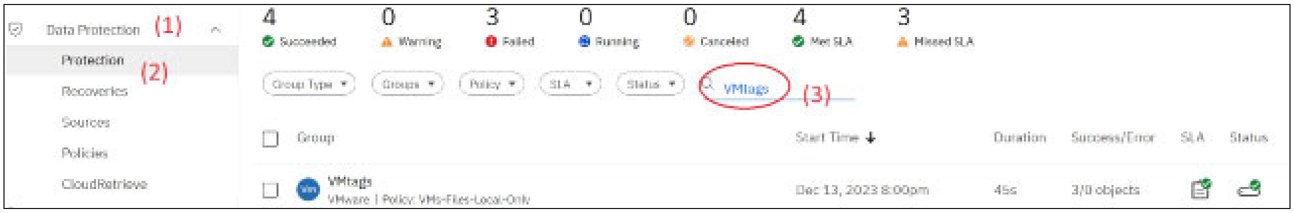

☐	Group	Start Time	Duration	Success/Error	SLA	Status
☐	VMtags VMware | Policy: VMs-Files-Local-Only	Dec 13, 2023 8:00pm	45s	3/0 objects		

Figure 3-12 Monitoring a Protection Group, steps 1-3

4. click on the Protection Group VMtags and select date and time to view which VMs are in the Protection Group backup set.

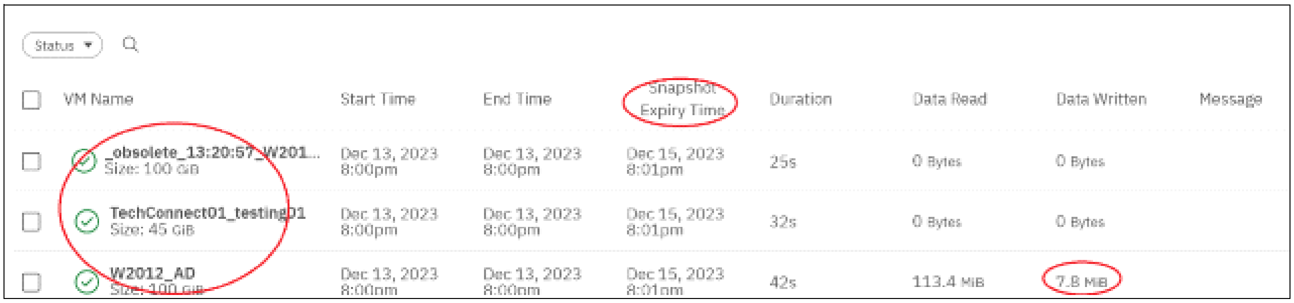

☐	VM Name	Start Time	End Time	Snapshot Expiry Time	Duration	Data Read	Data Written	Message
☐	_obsolete_13:20:57_W201... Size: 100 GiB	Dec 13, 2023 8:00pm	Dec 13, 2023 8:00pm	Dec 15, 2023 8:01pm	25s	0 Bytes	0 Bytes	
☐	TechConnect01_testing01 Size: 45 GiB	Dec 13, 2023 8:00pm	Dec 13, 2023 8:00pm	Dec 15, 2023 8:01pm	32s	0 Bytes	0 Bytes	
☐	W2012_AD Size: 100 GiB	Dec 13, 2023 8:00pm	Dec 13, 2023 8:00pm	Dec 15, 2023 8:01pm	42s	113.4 MiB	7.8 MiB	

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

5. Select the VM for details task events

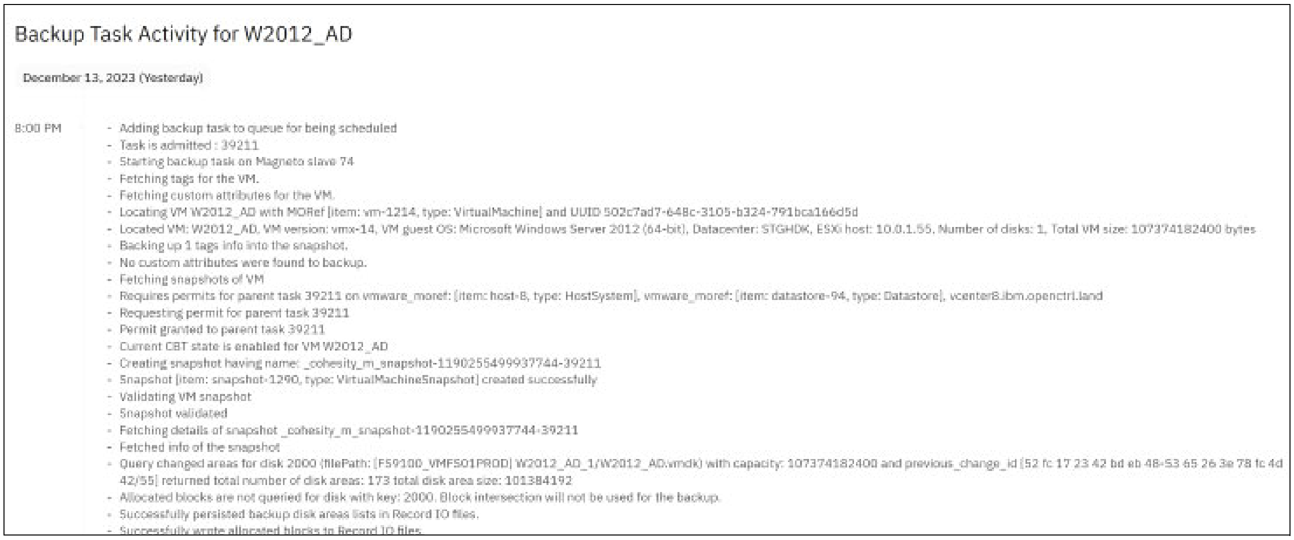

Backup Task Activity for W2012_AD

December 13, 2023 (Yesterday)

```
8:00 PM    - Adding backup task to queue for being scheduled
           - Task is admitted : 39211
           - Starting backup task on Magneto slave 74
           - Fetching tags for the VM.
           - Fetching custom attributes for the VM.
           - Locating VM W2012_AD with MORef [item: vm-1214, type: VirtualMachine] and UUID 502c7ad7-648c-3105-b324-791bca166d5d
           - Located VM: W2012_AD, VM version: vmx-14, VM guest OS: Microsoft Windows Server 2012 (64-bit), Datacenter: STGHDK, ESXi host: 10.0.1.55, Number of disks: 1, Total VM size: 107374182400 bytes
           - Backing up 1 tags info into the snapshot.
           - No custom attributes were found to backup.
           - Fetching snapshots of VM
           - Requires permits for parent task 39211 on vmware_moref: [item: host-8, type: HostSystem], vmware_moref: [item: datastore-94, type: Datastore], vcenter8.ibm.openctrl.land
           - Requesting permit for parent task 39211
           - Permit granted to parent task 39211
           - Current CBT state is enabled for VM W2012_AD
           - Creating snapshot having name: _cohesity_m_snapshot-1190255499937744-39211
           - Snapshot [item: snapshot-1290, type: VirtualMachineSnapshot] created successfully
           - Validating VM snapshot
           - Snapshot validated
           - Fetching details of snapshot _cohesity_m_snapshot-1190255499937744-39211
           - Fetched info of the snapshot
           - Query changed areas for disk 2000 (filePath: [FS9100_VMFS01PROD] W2012_AD_1/W2012_AD.vmdk) with capacity: 107374182400 and previous_change_id [52 fc 17 23 42 bd eb 48-53 65 26 3e 78 fc 4d
             42/55] returned total number of disk areas: 173 total disk area size: 101384192
           - Allocated blocks are not queried for disk with key: 2000. Block intersection will not be used for the backup.
           - Successfully persisted backup disk areas lists in Record IO files.
           - Successfully wrote allocated blocks to Record IO files.
```

Figure 3-14 Protection Group, backup task activity details

3.5 IBM Storage Ready Nodes

IBM Storage Defender Data Protect 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 Professional Services team following the network requirements and data resiliency parameters used in the sizing.

- Setting up nodes

- Creating the initial cluster
- Setting up data sources and protection policies

Note: The ISO software installation has to be done on each node. Downloading the ISO image is suggested as a time saver, and having multiple USB devices with the ISO image can save additional time when there are a larger number of nodes to be installed

3.6 IBM Storage Defender Ready Nodes

The IBM Storage Ready Node platform can be used as the foundation for an advanced scale out cluster that is highly secure, resilient, and provides exceptional performance using IBM Storage.

Defender. For detailed advice on designing the ideal solution for your enterprise data resiliency needs, contact your IBM seller or business partner.

Together they will review the workload types, the size, retention, and change rate of your data to suggest the necessary scale out cluster architecture needed to meet your enterprise data resiliency goals.

As stated in the previous sections, a minimal cluster can be made using a quantity of four nodes (8U total rack space).

Fore information on the latest IBM Storage Defender supported hardware see:
https://www.ibm.com/support/pages/node/6985577

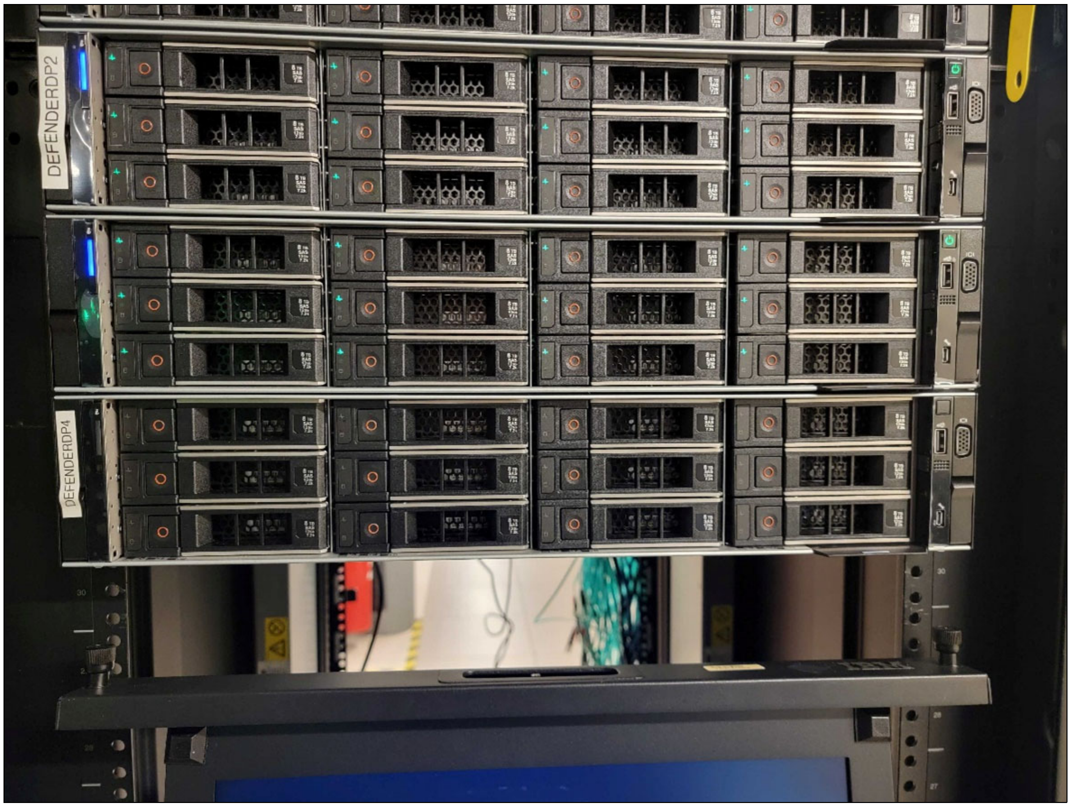

Figure 3-15 IBM Storage Ready Nodes storage bays and USB port access via the front of the cabinet

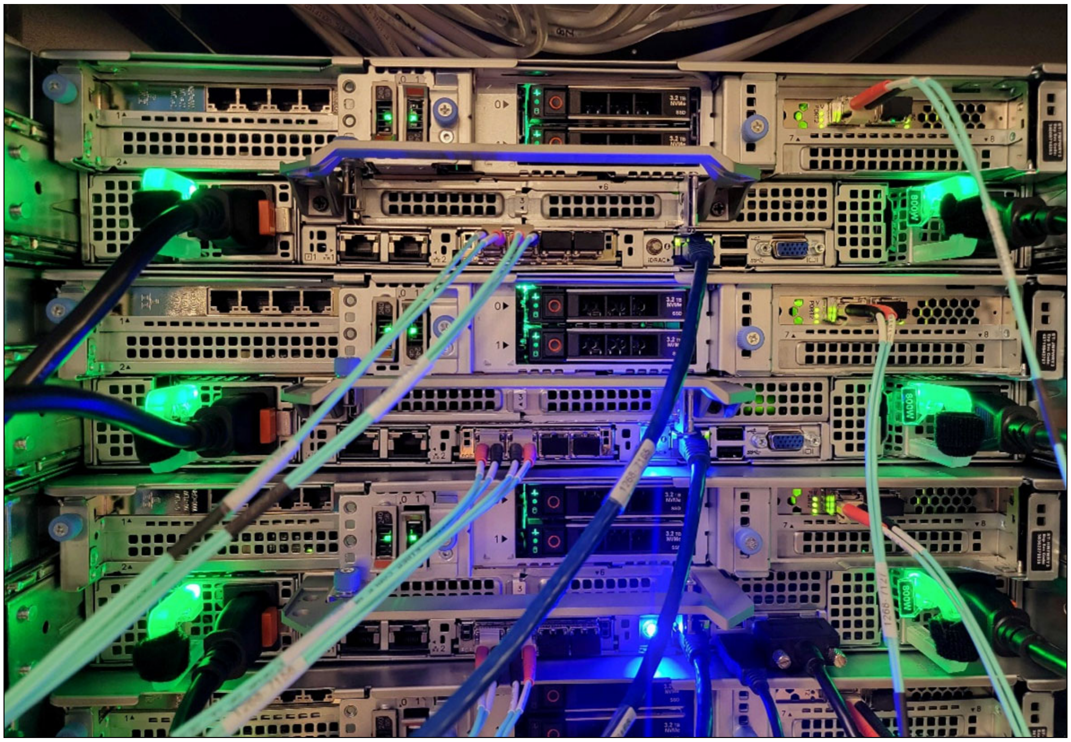

Figure 3-16 IBM Storage Ready Nodes management and network and optional FC adapter access via the rear of the cabinet

3.7 Registering with IBM Data Management Service

IBM Storage Defender provide a convenient way to manage the status of Storage Protect servers along with VMs and other protected workloads registered to the Data Management service (DMS) GUI. To register an IBM Storage Protect server with the IBM Storage Defender Data Protect (DDP) DMS, go to the DMS login and select Connection Agents from the left navigation menu. There you will see an option to download the connection agent for the Storage Protect server. Use the directions to download and extract the connection agent on the Storage Protect server.

1. Log into the Data Management service by navigating to the Data Management panel and clicking the Launch Data management button
2. From the Data Management service page, select the hamburger icon in the top left of the page and select Settings > Access Management
3. Select the Tokens tab
4. Select the Create box in the upper right corner of the screen. A Create Token input window appears 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 using 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 Data Protect v7.1.0 and beyond can be integrated with snapshots from the following SAN storage arrays:

- IBM FlashSystems (Includes immutable IBM Safeguarded copy snapshots)
- Nimble Storage, HPE Alletra 5000, and/or HPE Alletra 6000
- Pure Storage Array

Registering VM Hypervisors:

When creating 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 easier to create them prior to creating the Protection Group, as described in Chapter 6, “Integrating IBM Defender Data Protection with IBM Storage Protect” on page 95 of this document.

3.8 VMware Workload Use Case

One of the key capabilities of IBM Storage Defender Data Protect is the ability to perform Instant mass Restore to drastically reduce downtime after an incident like a large-scale malware attack or ransomware attack.

Recovering VMS doc reference: https://ibm.biz/BdGrAb

IBM Defender Data Protect can recover Protected Objects (such as VMs) from a Snapshot created earlier by a Protection Group. You can recover VMs from a cluster or a currently registered archive. You can recover VMs to their original location or a new location. The scale out architecture of the IBM Defender Data Protect cluster provides the resources needed to do large scale recovery of hundreds of VMs at a time. This greatly improves the RTO for groups of servers and applications.

For further information see Chapter 6, “Integrating IBM Defender Data Protection with IBM Storage Protect” on page 95 which will explore the VM backup process and the Instant Mass Restore capability using a lab environment.

4

Protecting VMware Workloads With IBM Data Protect

In this chapter, we will explore key concepts, different functions and features of IBM Data Protect as well as how to configure agentless backup workloads designated for a VMWare environment through vCenter or a single ESXi host. Backups are snapshot based incremental forever backup with the ability to recover a whole VM or restore a specific file/folder within a VM.

This chapter includes the following topics:

- 4.1, “IBM Data Protect Overview” on page 44
- 4.2, “Configuration of Data Protect Sources” on page 44
- 4.3, “Recovering Data with IBM Data Protect” on page 47
- 4.4, “Replicating Workloads with IBM Data Protect clusters” on page 53
- 4.5, “Monitoring the Protection job list” on page 56
- 4.6, “Recover from Replicated Copies” on page 57
- 4.7, “CloudArchive the Workloads with IBM Data Protect Cluster” on page 61

4.1 IBM Data Protect Overview

IBM Data Protect is a powerful tool in the hands of an administrator. An administrator will benefit from understanding the many features and functions that will help an organization realize the full potential of IBM Storage Defender Protect's ability to secure data from potential threats. To help draw visual clarity, below is a list of functions and definitions.

Definitions used in this chapter:

Protection - Protection is used to define parameters around a specific workload, such as a type of workload, objects to backup/protect, a policy for retention/frequency and the storage domain to which the data sets will be stored within.

Sources - Sources are used to register (attach) a source host.

Policies - Policies are parameters for retention, frequency, location to where backup isbeing stored as well as designating replication, archive location and cloud destination.

Storage Domains - Storage Domains are used to define how the data will be processed (dedup and/or compression), stored, quotas can be set as well as Fault tolerance & redundancy can be defined.

Remote Cluster – Remote Cluster option is used to configure replication targets for replicated protection

External Targets – External target allows for archival or tiering targets to be configured for additional copies of data sets.

The next few sections will walk through configuring each feature, starting with registering sources to configuring a complete protection policy, with the end goal of completing a protection policy for VMWare environments, as well as configuring additional copies for replication and/or archival.

4.2 Configuration of Data Protect Sources

When configuring IBM Data Protect sources, 'sources' refer to a host or hosts designated for protection. The focus of this Redbook and its examples use sources that contain Virtual Machines, specifically VMWare.

The following example (Figure 4-1 on page 45) shows the steps used to Register new virtual machines and configure a new source:

1. Expand the Data Protection section on the left side of the Defender GUI
2. Click on Sources, then select Register
3. Select Virtual Machines for registration then select the source type from the drop down
4. Select VMWare: vCenter or ESXi host
5. Fill out the IP/hostname, Username, password for the designated VMWare cluster/host
6. Enable any of the listed options that are desired for Defender to use with the VMware environment this will be applied to
7. Click Register to complete the registration process

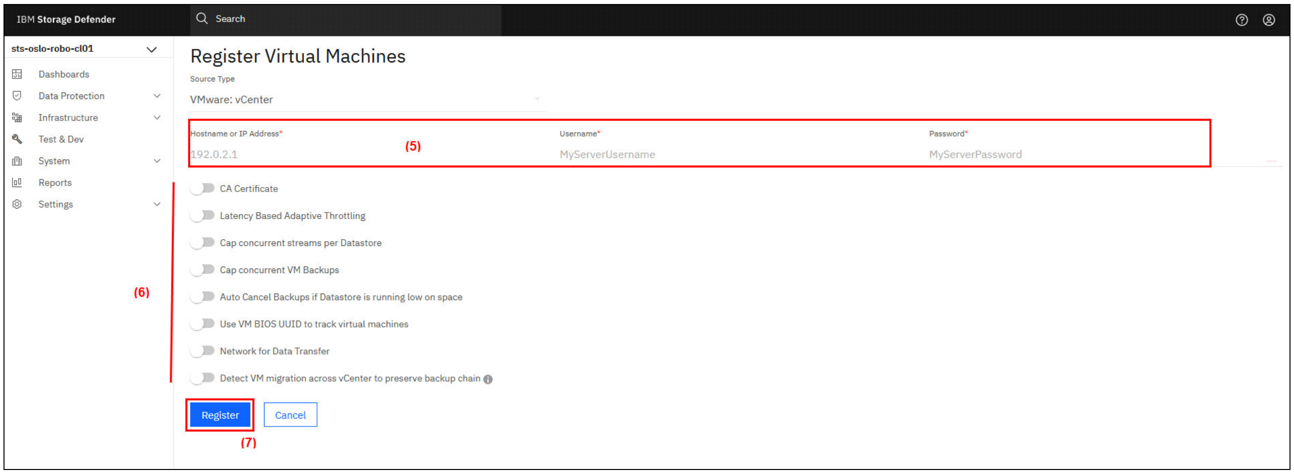

Register Virtual Machines

Source Type: VMware: vCenter

Hostname or IP Address*: 192.0.2.1

Username*: MyServerUsername

Password*: MyServerPassword

- CA Certificate
- Latency Based Adaptive Throttling
- Cap concurrent streams per Datastore
- Cap concurrent VM Backups
- Auto Cancel Backups if Datastore is running low on space
- Use VM BIOS UUID to track virtual machines
- Network for Data Transfer
- Detect VM migration across vCenter to preserve backup chain

Figure 4-1 Register Sources, steps 5-7

> **Note:** As a best practice, it is strongly recommended to have a minimum of 10GbE links, or higher, between IBM Data Protect cluster and the VMWare hosts to achieve the most effective RPO/RTO.

4.2.1 Configuration of a Data Protect Policy

Policies are the mechanism used by Defender for configuring the retention and frequency of recovery points, as well as location to where backup is being stored, replication target, archive location and cloud destination.

The following example shows the steps needed to configure a Data Protect policy for the previously configured VMware backup.

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

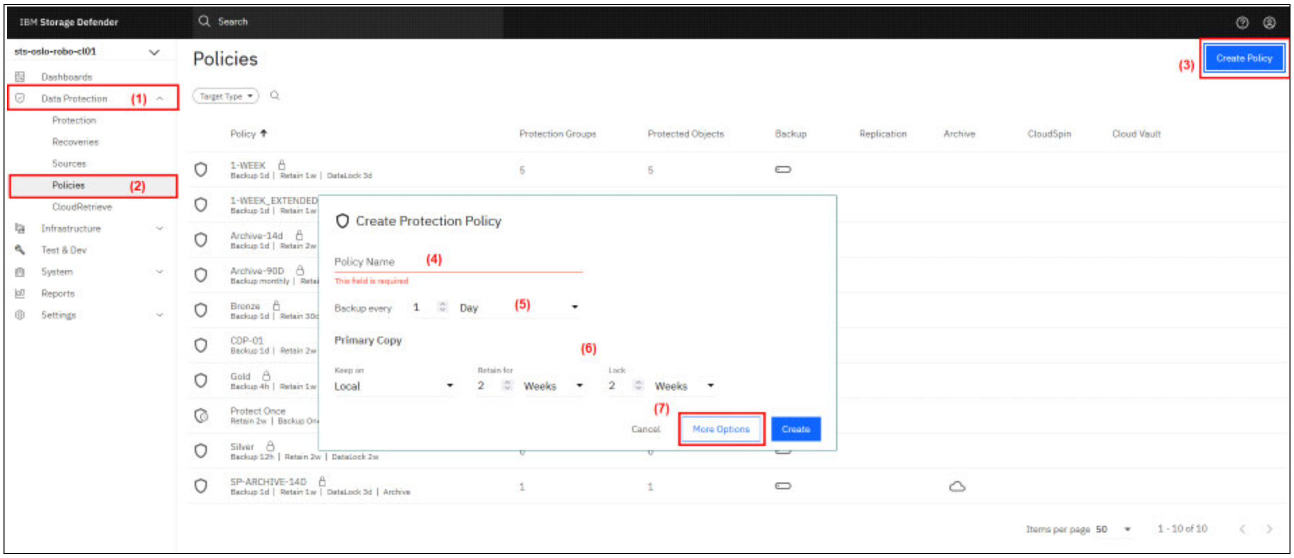

Figure 4-2 Policy Sources configuration steps 1-7

8. Configure Replication, Archive or CloudSpin as needed for the backup settings
9. Click create

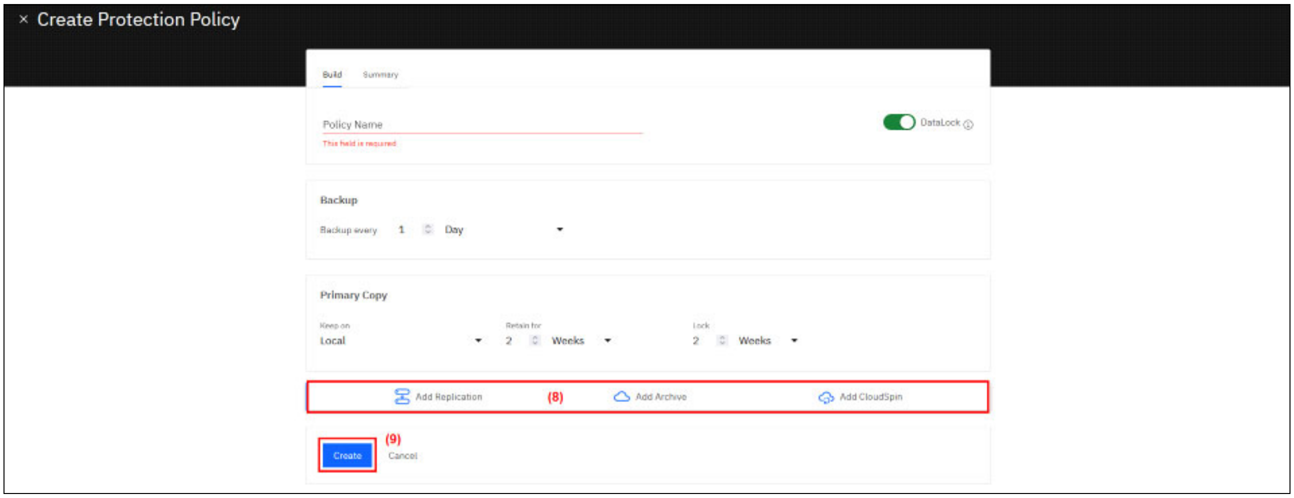

Figure 4-3 Policy Sources configuration steps 8 and 9.

> **Note:** Replication, Archive or Cloudspin can only be applied to a policy which will be attached to a data protection group. Replication, Archive and Cloudspin have their own frequency/retention when configuring. This will be discussed later in chapter.

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

- **Quality of Services (QoS) -** Policy determines the type of storage used and the latency factor in writing data to the view.
- **Periodic Full backup** - Periodic Full Backups are scheduled for: Day, Week, Month, Year. Schedules can also be made for multiple periodic full backups with different frequencies and dates.

- **Continuous Data Protection** - Continuous Data Protection defines the CDP schedule for capturing logs from the VMs as opposed to restore the VM applications to a point-in-time as instead of periodic snapshots available with regular backups.
- **Quiet Times** - Quiet Times define time periods when new protection runs are not started.
- **Customize Retires** - Customize Retries by default, Data Protect attempts to capture Snapshots three times before the protection run fails. The default time between retries is five minutes. Here you can customize the number of retries and how long to wait between each attempt.
- **BMR (Bare Metal Recovery) Backup** - BMR Backup provides a backup schedule and retention period for Bare Machine Recovery (BMR) system data on a physical server.
- **Log Backup** - Log Backup adds a database logs schedule if there are plans on restoring databases to a specific point in time between two full database server backups.

4.3 Recovering Data with IBM Data Protect

In this section we outline the recovery types available, tasks related to recovery and steps required when recovering data for a protected Data Protect cluster.

4.3.1 Recover from a Data Protect Cluster

When recovering from an IBM Data Protect cluster, the following tasks and operations will be created and run:

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

IBM Data Protect 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 will be instantly available after the recovery to the target location begins. Upon completion of the recovery, the VM data will be moved to the target storage location. The VM can be accessed when the data is copied to the VM from the IBM Defender Data Protect cluster. In these cases, the VM's performance may 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 where the recovered VM will reside.

Copy Recovery:

The VMs will be available in the target location only after all data is copied to the target storage from the source location (IBM Storage Defender Data Protect cluster or cloud). Once the restore is complete, the VM will be 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 the VM on. As a best practice, IBM Storage Defender Data Protect recommends for better performance, to choose Copy Recovery for powered-off VMs.

To perform a VM recovery:

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

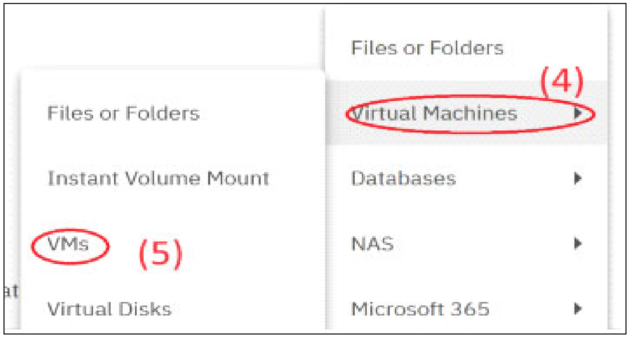

Figure 4-4 Recovering VM with Data Protect, steps 4-5

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

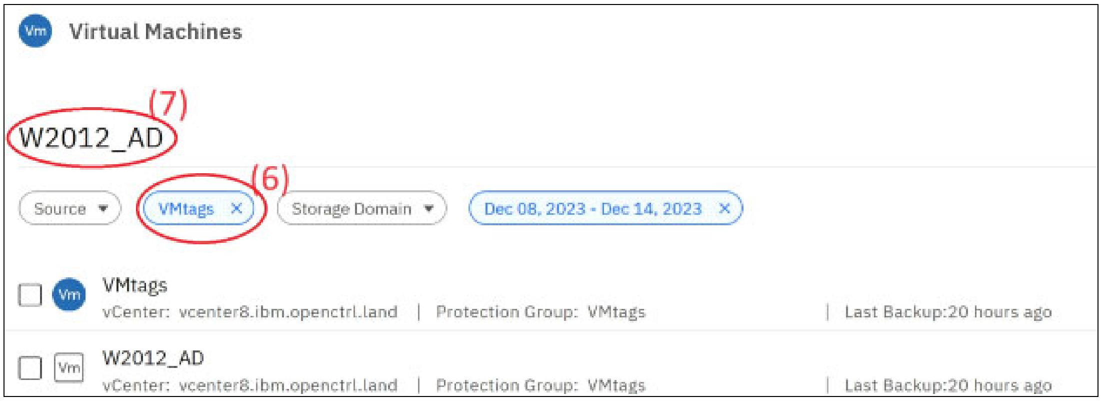

□ VMtags
vCenter: vcenter8.ibm.openctrl.land | Protection Group: VMtags | Last Backup:20 hours ago

□ W2012_AD
vCenter: vcenter8.ibm.openctrl.land | Protection Group: VMtags | Last Backup:20 hours ago

Figure 4-5 Recovering VM with Data Protect, steps 6-7

8. Select the VM to recover and Next for the recovery method (Instant Recovery)
9. Select the Existing VM Handling

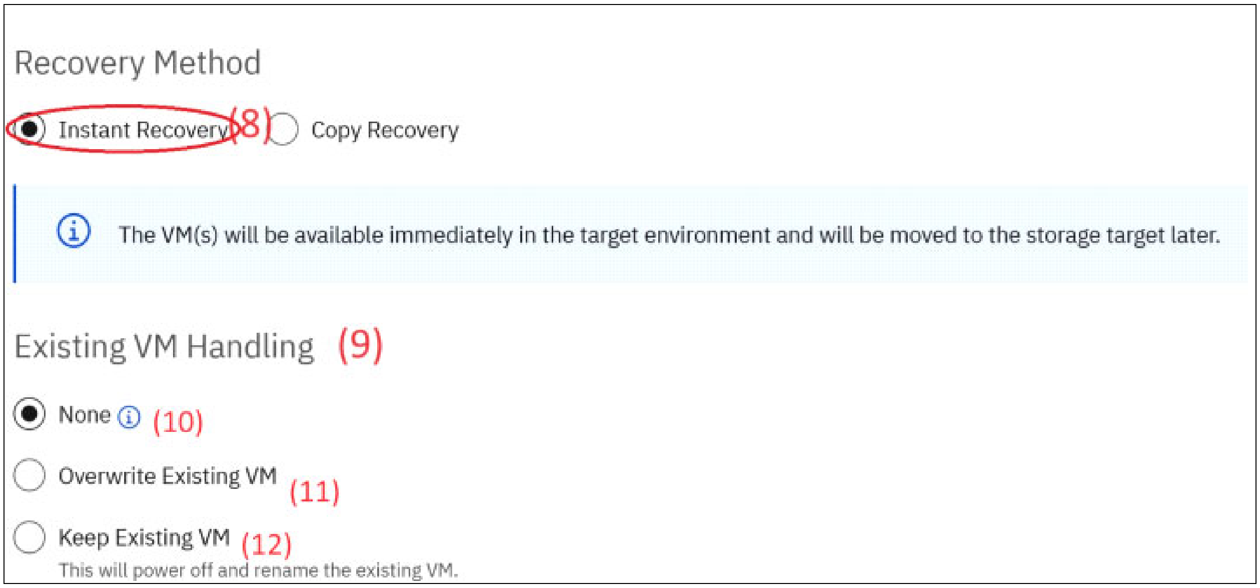

Recovery Method

● Instant Recovery ○ Copy Recovery

The VM(s) will be available immediately in the target environment and will be moved to the storage target later.

Existing VM Handling

● None

○ Overwrite Existing VM

○ Keep Existing VM
This will power off and rename the existing VM.

Figure 4-6 Recovering VM with Data Protect, steps 8-12

10. None:
    Recover will fail if there is an existing VM: Recovers the VM as a new VM and retains the original VM. In the Rename field available in the Recovery Options, specify the name of the recovered VM by adding prefix and/or suffix strings to the name of the original VM. By default, the recovered VM is named as copy-<original_VM_name>.
11. Overwrite Existing VM:
    Recovers the VM by deleting the original VM. The recovered VM will have the original VM name. When choosing this option, the original VM is deleted prior to the recovery taking place. A recovery failure will result in the loss of the original VM.

> **Note:** ‘Attempt Differential Recovery’: Attempts to recover the VM by overwriting only the difference between the original VM and the snapshot selected for recovery. Any newly added data in the original VM is deleted. If you want to reclaim free space for thin-provisioned disks, then IBM Storage Defender Data Protect recommends not to attempt differential recovery and only perform a recovery using Overwrite Existing VM. Differential recovery reduces the amount of data transfered during the recovery process. In the pulse log, you can view the amount of data transfer saved by selecting differential recovery.

12. Keep Existing VM: Recover the VM as a new VM with the original VM name. The original VM is retained powered off. The original VM will be renamed to obsolete_<time_stamp><Orginal_VM_Name>.
13. Finish the recover task and monitor the workflow. If needed review the recovery activity log associated with the recovery process.

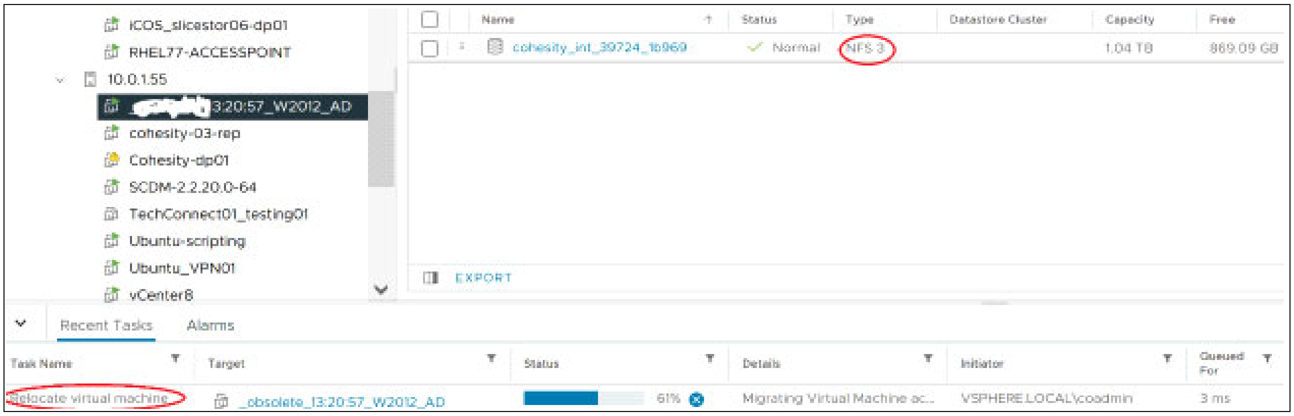

- ICOS_slicestor06-dp01
- RHEL77-ACCESSPOINT
- 10.0.1.55
  - [redacted]3:20:57_W2012_AD
  - cohesity-03-rep
  - Cohesity-dp01
  - SCDM-2.2.20.0-64
  - TechConnect01_testing01
  - Ubuntu-scripting
  - Ubuntu_VPN01
  - vCenter8

Name	Status	Type	Datastore Cluster	Capacity	Free
cohesity_int_39724_1b969	✓ Normal	NFS 3		1.04 TB	869.09 GB

Task Name	Target	Status	Details	Initiator	Queued For
Relocate virtual machine	_obsolete_13:20:57_W2012_AD	61%	Migrating Virtual Machine ac...	VSPHERE.LOCAL\coadmin	3 ms

Figure 4-7 Data Protect Instance Recovery vMotion from NFS to Production

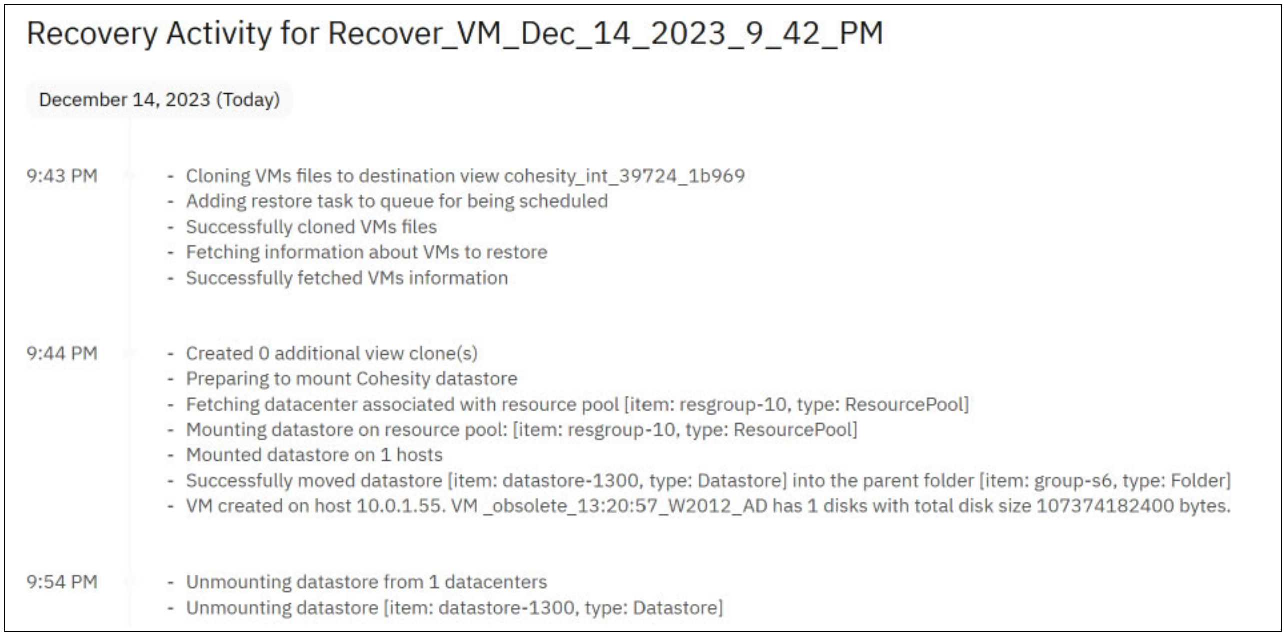

Recovery Activity for Recover_VM_Dec_14_2023_9_42_PM

December 14, 2023 (Today)

```
9:43 PM    - Cloning VMs files to destination view cohesity_int_39724_1b969
           - Adding restore task to queue for being scheduled
           - Successfully cloned VMs files
           - Fetching information about VMs to restore
           - Successfully fetched VMs information

9:44 PM    - Created 0 additional view clone(s)
           - Preparing to mount Cohesity datastore
           - Fetching datacenter associated with resource pool [item: resgroup-10, type: ResourcePool]
           - Mounting datastore on resource pool: [item: resgroup-10, type: ResourcePool]
           - Mounted datastore on 1 hosts
           - Successfully moved datastore [item: datastore-1300, type: Datastore] into the parent folder [item: group-s6, type: Folder]
           - VM created on host 10.0.1.55. VM _obsolete_13:20:57_W2012_AD has 1 disks with total disk size 107374182400 bytes.

9:54 PM    - Unmounting datastore from 1 datacenters
           - Unmounting datastore [item: datastore-1300, type: Datastore]
```

Figure 4-8 Data Protect Recover Activity Log

4.3.2 Instant Volume Mount recovery from Data Protect Cluster

The Instant Volume Mount operation restores an entire volume, to a running VM, and is the best way to restore data when a large data set needs to be recovered. This restoration method mounts the restored VMDK to a Windows VM allowing access to copy the required data wherever it is needed.

To recover a VMDK as Instant Volume Mount:

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

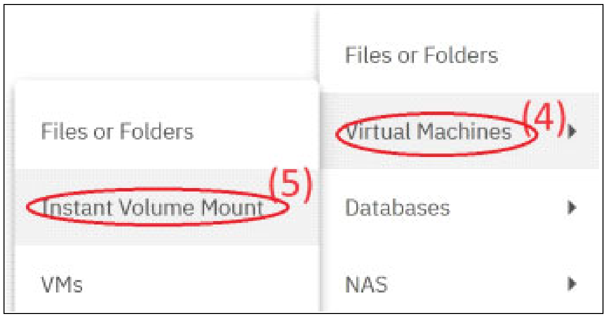

Figure 4-9 Instant Volume Mount recovery from Data Protect, steps 4-5

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

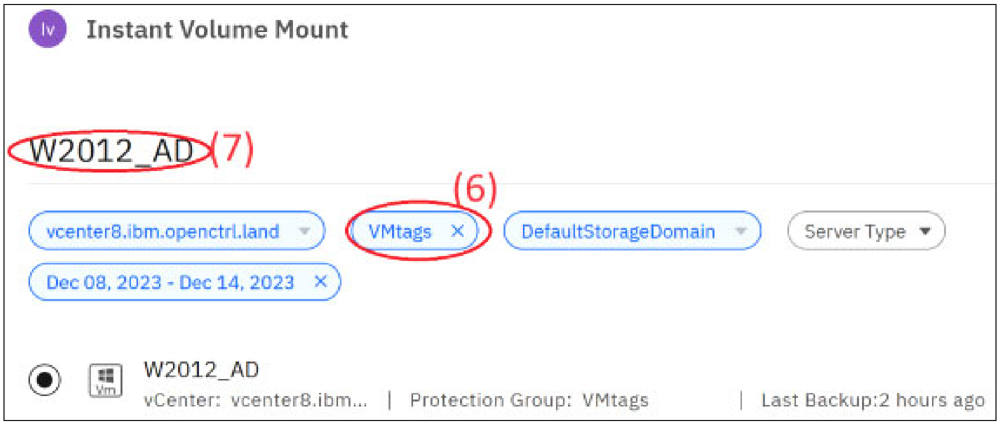

W2012_AD

● W2012_AD
vCenter: vcenter8.ibm... | Protection Group: VMtags | Last Backup:2 hours ago

Figure 4-10 Instant Volume Mount recovery from Data Protect, steps 6-7

8. Next select the recover options and the volume to recover
9. The default recover option will be the Original Location (Figure 4-11). If it is required that the vmdk should mount to another VM target, select New Location and browse to the vCenter and VM where the vmdk should be mounted (Figure 4-12 on page 52).

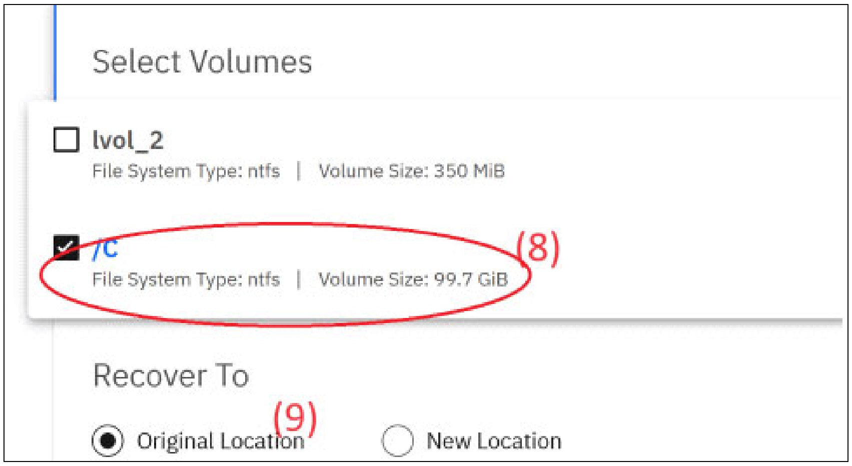

□ lvol_2
File System Type: ntfs | Volume Size: 350 MiB

☑ /C
File System Type: ntfs | Volume Size: 99.7 GiB

● Original Location ○ New Location

Figure 4-11 Instant Volume Mount recovery from Data Protect, steps 8-9

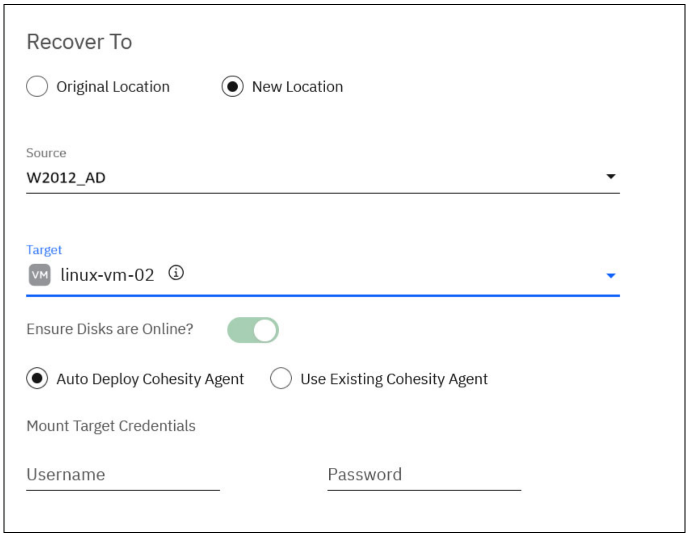

Recover To

○ Original Location ● New Location

Source
W2012_AD

Target
linux-vm-02

Ensure Disks are Online? (toggle on)

● Auto Deploy Cohesity Agent ○ Use Existing Cohesity Agent

Mount Target Credentials

Username

Password

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

10. Finish recover and monitor the workflow

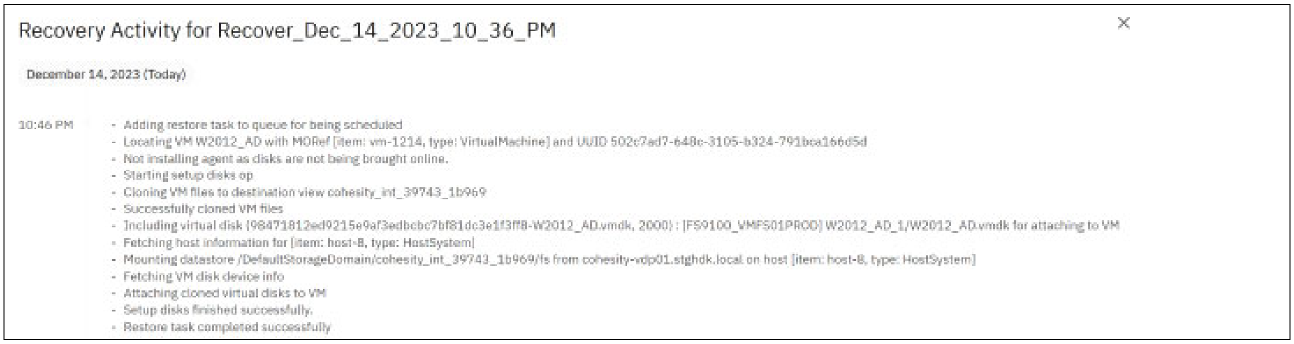

Recovery Activity for Recover_Dec_14_2023_10_36_PM

December 14, 2023 (Today)

```
10:46 PM   - Adding restore task to queue for being scheduled
           - Locating VM W2012_AD with MORef [item: vm-1214, type: VirtualMachine] and UUID 502c7ad7-648c-3105-b324-791bca166d5d
           - Not installing agent as disks are not being brought online.
           - Starting setup disks op
           - Cloning VM files to destination view cohesity_int_39743_1b969
           - Successfully cloned VM files
           - Including virtual disk (98471812ed9215e9af3edbcbc7bf81dc3e1f3ff8-W2012_AD.vmdk, 2000) : [FS9100_VMFS01PROD] W2012_AD_1/W2012_AD.vmdk for attaching to VM
           - Fetching host information for [item: host-8, type: HostSystem]
           - Mounting datastore /DefaultStorageDomain/cohesity_int_39743_1b969/fs from cohesity-vdp01.stghdk.local on host [item: host-8, type: HostSystem]
           - Fetching VM disk device info
           - Attaching cloned virtual disks to VM
           - Setup disks finished successfully.
           - Restore task completed successfully
```

Figure 4-13 Instant Volume Mount recovery from Data Protect activity log

The volume will be available in Windows disk manager and can be brought online via Windows disk management.

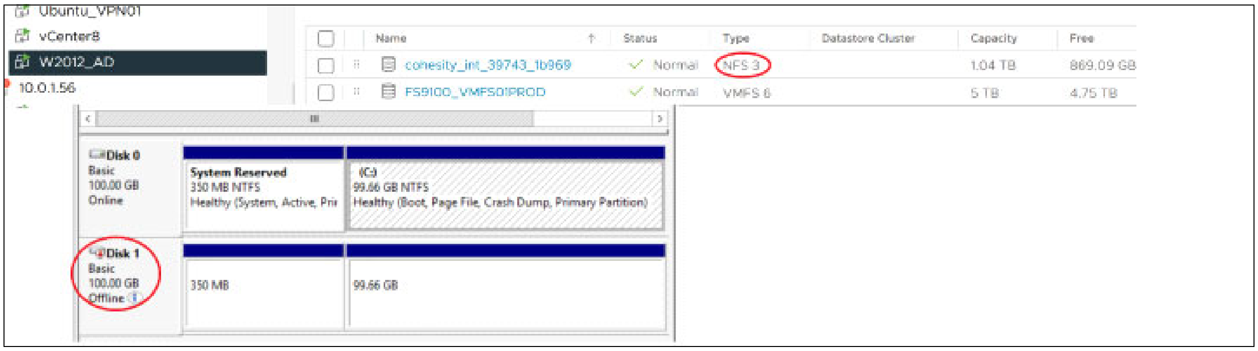

Name	Status	Type	Datastore Cluster	Capacity	Free
cohesity_int_39743_1b969	✓ Normal	NFS 3		1.04 TB	869.09 GB
FS9100_VMFS01PROD	✓ Normal	VMFS 6		5 TB	4.75 TB

Disk 0
Basic
100.00 GB
Online

System Reserved
350 MB NTFS
Healthy (System, Active, Pri

(C:)
99.66 GB NTFS
Healthy (Boot, Page File, Crash Dump, Primary Partition)

Disk 1
Basic
100.00 GB
Offline

350 MB

99.66 GB

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 Data Protect GUI to remove the volume from the target. The actions that occur during a teardown are described in the following procedure:

1. The Data Protect 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 acting as a datastore.
3. If no VMs are using the view (datastore), the view is umounted from all ESXi hosts.

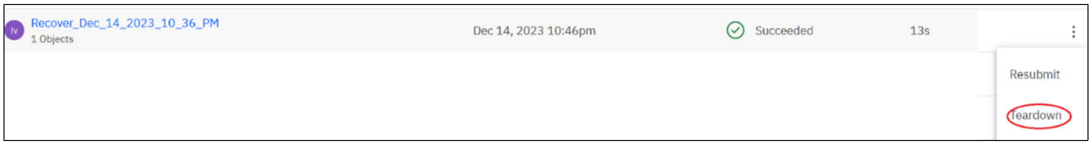

Recover_Dec_14_2023_10_36_PM
1 Objects
Dec 14, 2023 10:46pm
Succeeded
13s

Figure 4-15 Volume recovery Teardown option

4.4 Replicating Workloads with IBM Data Protect clusters

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

1. One-to-one model
   - A single production cluster can backup to a disaster recovery site
2. One-to-many model
   - A single production cluster can backup to multiple disaster recovery sites
3. Many-to-one model
   - Multiple production cluster can backup to remote cluster
4. Many-to-many model
   - Multiple production cluster can backup to multiple remote clusters across several sites

This technology works in conjunction with global data deduplication to greatly reduce storage requirements when data is spread across several sites. This also allows for a reduction in the network bandwidth required for replication of data for DR purposes.

4.4.1 VMware replication workflow one-to-one model

In this section, the replication workflow of VMware and IBM Data Protect Cluster according to the design principal One-To-One is described in further detail. The backup data originating from the source VMware vCenter on the primary host are ingested in the primary IBM Data Protect cluster DC1, which is then replicated to a secondary DR cluster (DC2).

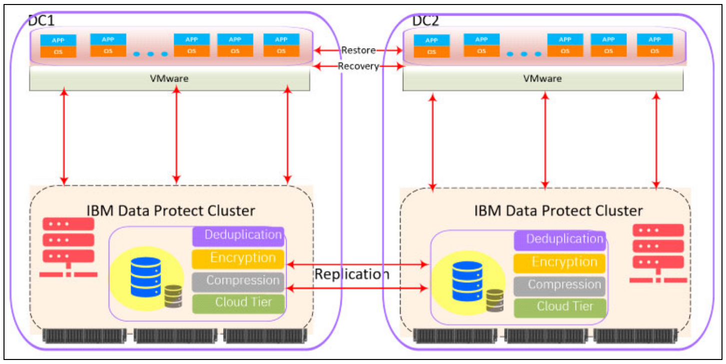

Figure 4-16 Replication workload One-To-One model

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

IP Failover and Load Balancing

Each node has a Virtual IP address (VIP), and clients can access any node on the cluster 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 allows the cluster to access client data that is physically present on any node on the cluster. The cluster depends on an external DNS server to do a round-robin on the VIP addresses so that client requests can be load-balanced.

The IBM Data Protect 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.

Register Remote Cluster

To register a remote cluster, the ip-address of the node, node-cluster, as well as the admin account and password are required. In a multiple node cluster always used the Virtual IP-address (VIP) to register the remote cluster.

> **Note:** The recommend best practice is to use one virtual IP for each cluster node

After connecting the clusters, it is possible to enable Remote Access [1] and Replication [2]. 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.

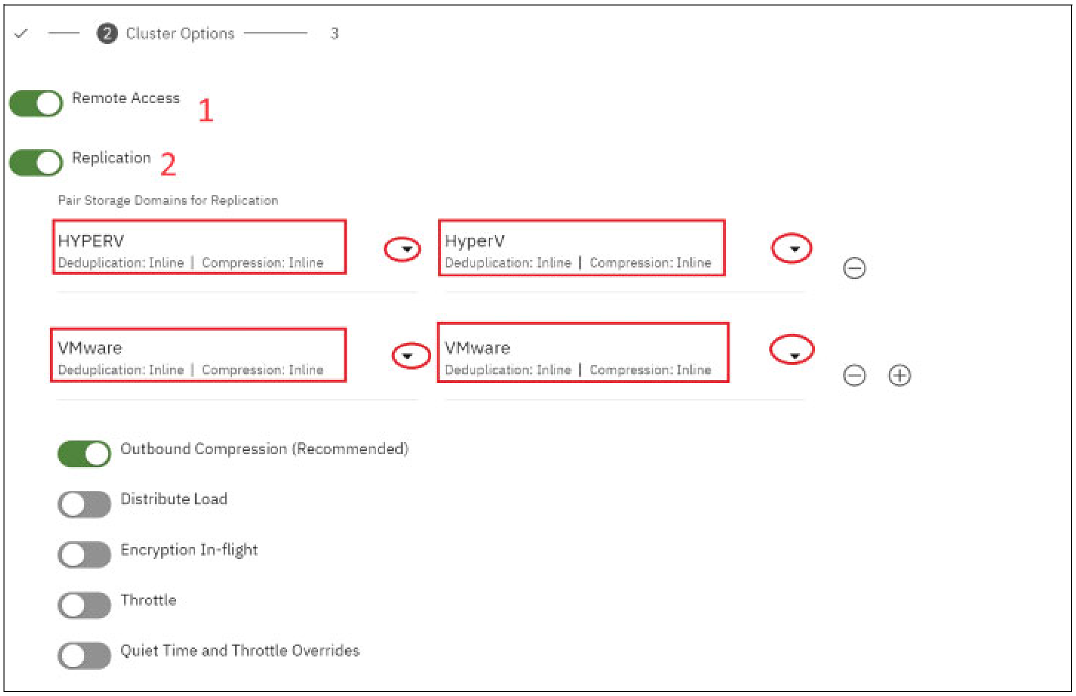

Remote Access (on) 1

Replication (on) 2

Pair Storage Domains for Replication

HYPERV Deduplication: Inline | Compression: Inline	HyperV Deduplication: Inline | Compression: Inline
VMware Deduplication: Inline | Compression: Inline	VMware Deduplication: Inline | Compression: Inline

Outbound Compression (Recommended) (on)

Distribute Load (off)

Encryption In-flight (off)

Throttle (off)

Quiet Time and Throttle Overrides (off)

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 on between 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’ then to configure:

1. In the replication field selected “replicate to”: Remote Cluster
2. Select the desired host name as the remote “Replication Target”

> **Note:** The following options may be toggled on/off as desired for the cluster
> - Outbound Compression
> - Distribute Load
> - Encryption In-flight
> - Throttle
> - Quiet time and Throttle Overrides

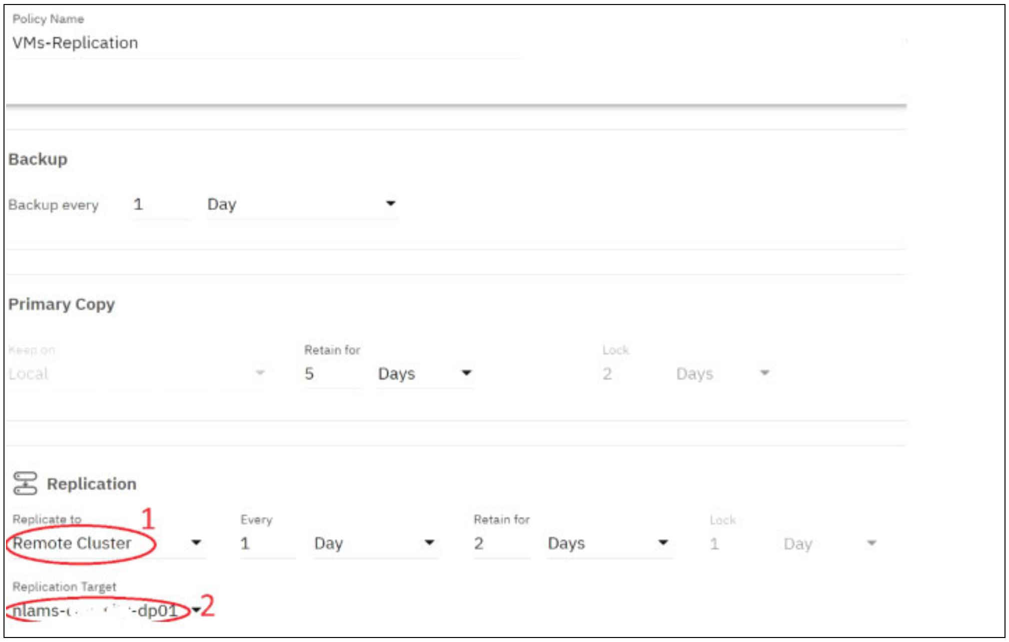

Policy Name
VMs-Replication

Backup
Backup every 1 Day

Primary Copy
Keep on Local
Retain for 5 Days
Lock 2 Days

Replication
Replicate to: Remote Cluster
Every 1 Day
Retain for 2 Days
Lock 1 Day
Replication Target: nlams-[redacted]-dp01

Figure 4-18 Modify the existing backup policy

After modifying the policies, double check the policies details view to ensure the configuration is correct.

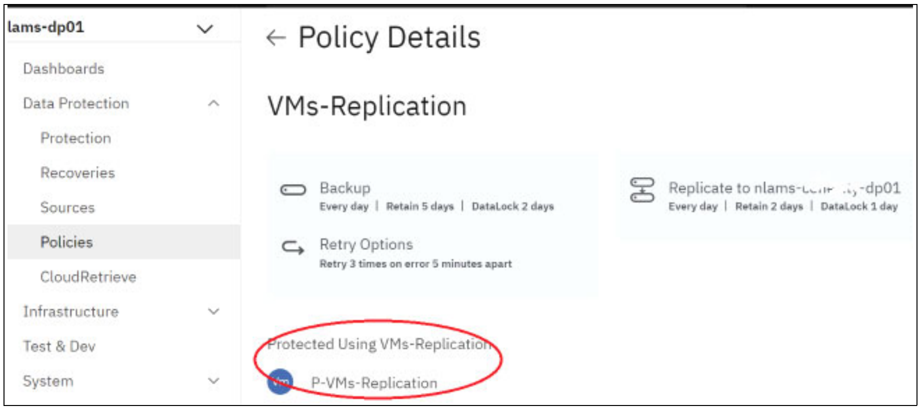

← Policy Details

VMs-Replication

Backup
Every day | Retain 5 days | DataLock 2 days

Retry Options
Retry 3 times on error 5 minutes apart

Replicate to nlams-[redacted]-dp01
Every day | Retain 2 days | DataLock 1 day

Protected Using VMs-Replication
P-VMs-Replication

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

4.5 Monitoring the Protection job list

To monitor the Protection jobs, use the left menu to navigate to the Data Protect section:

Under Data Protection select ‘Protection’ and Click on the protection group. For this example the group name is “DP01-P-VMs-Replication”.

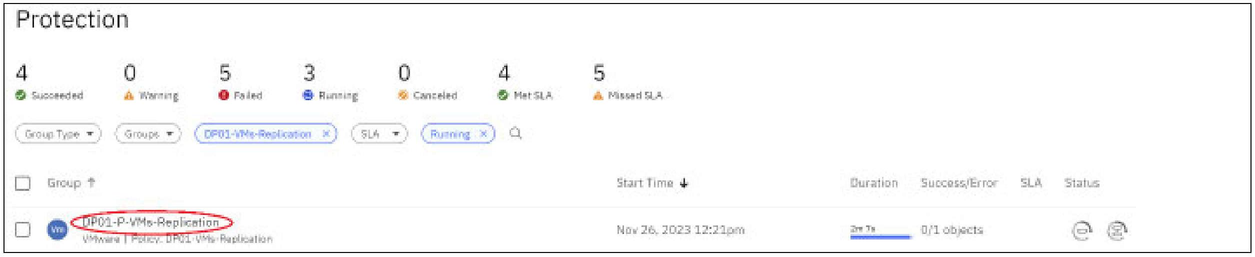

Protection

4 Succeeded · 0 Warning · 5 Failed · 3 Running · 0 Canceled · 4 Met SLA · 5 Missed SLA

□	Group	Start Time	Duration	Success/Error	SLA	Status
□	DP01-P-VMs-Replication VMware | Policy: DP01-VMs-Replication	Nov 26, 2023 12:21pm	2m 7s	0/1 objects		

Figure 4-20 Protection group status

Selecting the Protection Group will then display a list of the replication tasks associated with the group.

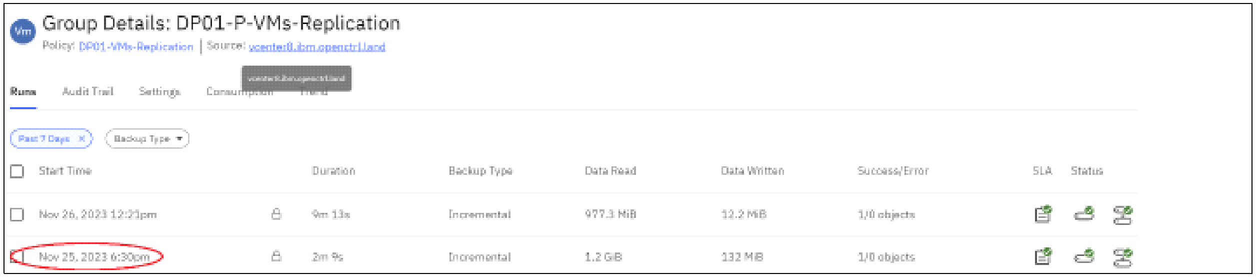

Group Details: DP01-P-VMs-Replication

Policy: DP01-VMs-Replication | Source: vcenter8.ibm.openctrl.land

□	Start Time	Duration	Backup Type	Data Read	Data Written	Success/Error	SLA	Status
□	Nov 26, 2023 12:21pm	9m 13s	Incremental	977.3 MiB	12.2 MiB	1/0 objects		
□	Nov 25, 2023 6:30pm	2m 9s	Incremental	1.2 GiB	132 MiB	1/0 objects		

Figure 4-21 Backup and replication task details

Once the list of jobs for the group is displayed, click on the desired job you wish to monitor, in this example the job list of “Nov 25, 2023 6:30pm” by clicking on it and selecting the tab “Replication” the details of the jobs replication process and reviewed.
In this example, the replication job did finish successfully with 16 GiB Logical Data backed up and 1 GiB of data transferred to the target replication group.

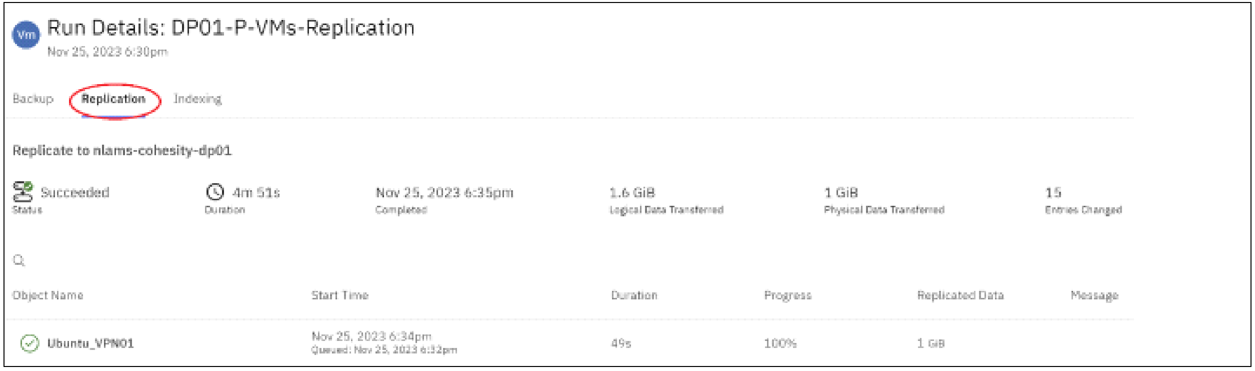

Run Details: DP01-P-VMs-Replication

Nov 25, 2023 6:30pm

Replicate to nlams-cohesity-dp01

Status	Duration	Completed	Logical Data Transferred	Physical Data Transferred	Entries Changed
Succeeded	4m 51s	Nov 25, 2023 6:35pm	1.6 GiB	1 GiB	15

Object Name	Start Time	Duration	Progress	Replicated Data	Message
Ubuntu_VPN01	Nov 25, 2023 6:34pm Queued: Nov 25, 2023 6:32pm	49s	100%	1 GiB	

Figure 4-22 Replication task details

4.6 Recover from Replicated Copies

In this section we will describe how to recover data from Replicated Copies in the event that a disaster strikes the Source Cluster. This can also be used when there is planned maintenance outage of the Source Cluster to maintain access to needed data and systems. When the IBM Defender Data Protect cluster or the Source is down or not available, Protection Groups can be failed over to the replication partner to maintain continuity.

After creating a Protection Group on a cluster with a replication schedule, once the first snapshot is replicated, the Protection Group is mark as ‘Failover Ready’ and the Protection Group become inactive. The Objects (VMs) and Policy are no longer associated with the inactive Protection Group.

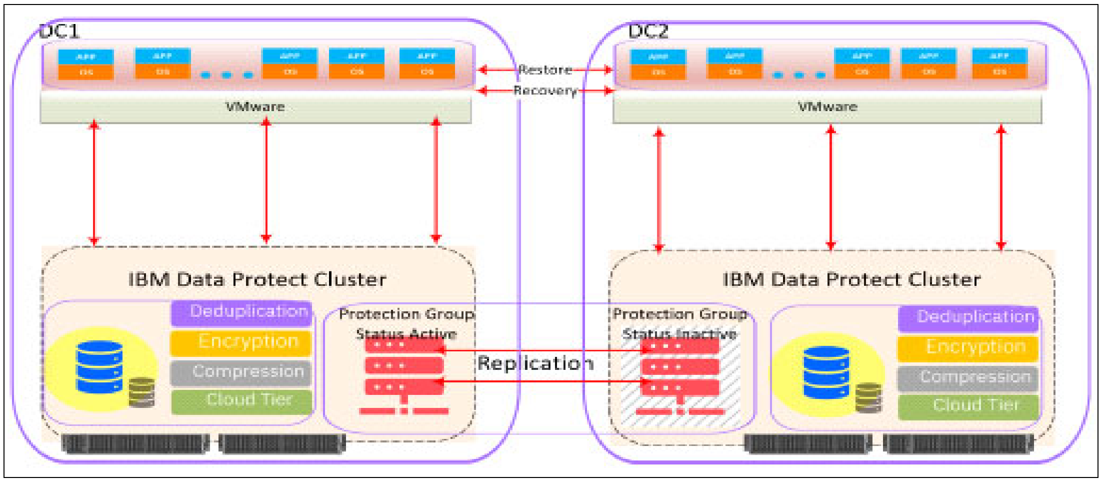

Figure 4-23 Protection Group Active / Inactive status

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 will no longer be captured by the original Protection Group on the original capturing cluster. However, the existing snapshots captured by the Protection Group are not deleted, they remain stored in the associated storage domain.

In the following 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 and Delete.

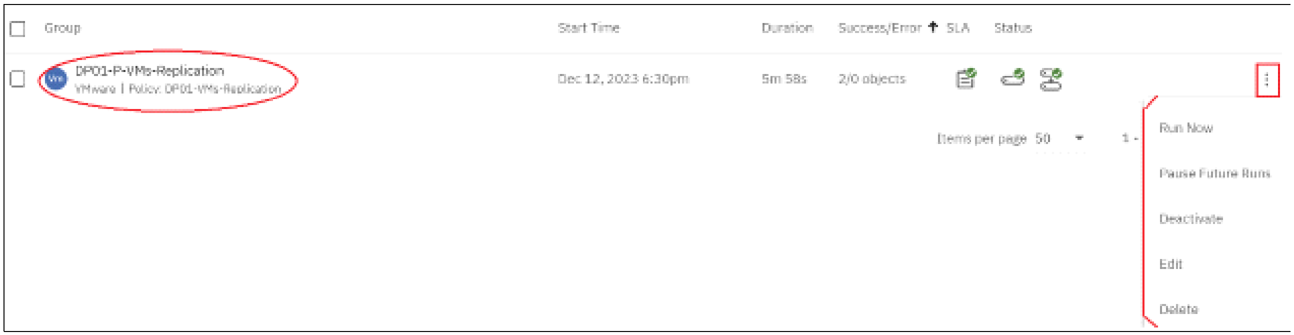

□	Group	Start Time	Duration	Success/Error ↑	SLA	Status
□	DP01-P-VMs-Replication VMware | Policy: DP01-VMs-Replication	Dec 12, 2023 6:30pm	5m 58s	2/0 objects		

Items per page 50 ▾ 1 -

Figure 4-24 Protection Group Actions

On the replicated cluster the Protection Group is marked as inactive and Failover Ready. This Protection Group will only have the options: Failover and Delete.

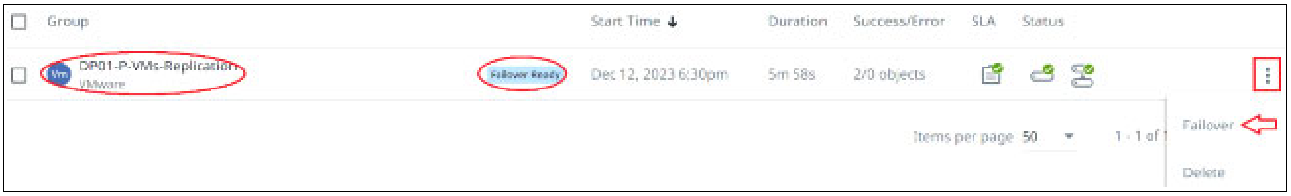

□	Group		Start Time ↓	Duration	Success/Error	SLA	Status
□	DP01-P-VMs-Replication VMware	Failover Ready	Dec 12, 2023 6:30pm	5m 58s	2/0 objects		

Items per page 50 1 - 1 of 1

Figure 4-25 Inactive Protection Group Actions

When activating the Failover feature specify the source vCenter server where recovered VMs will be placed. The example below shows vcenter8 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.

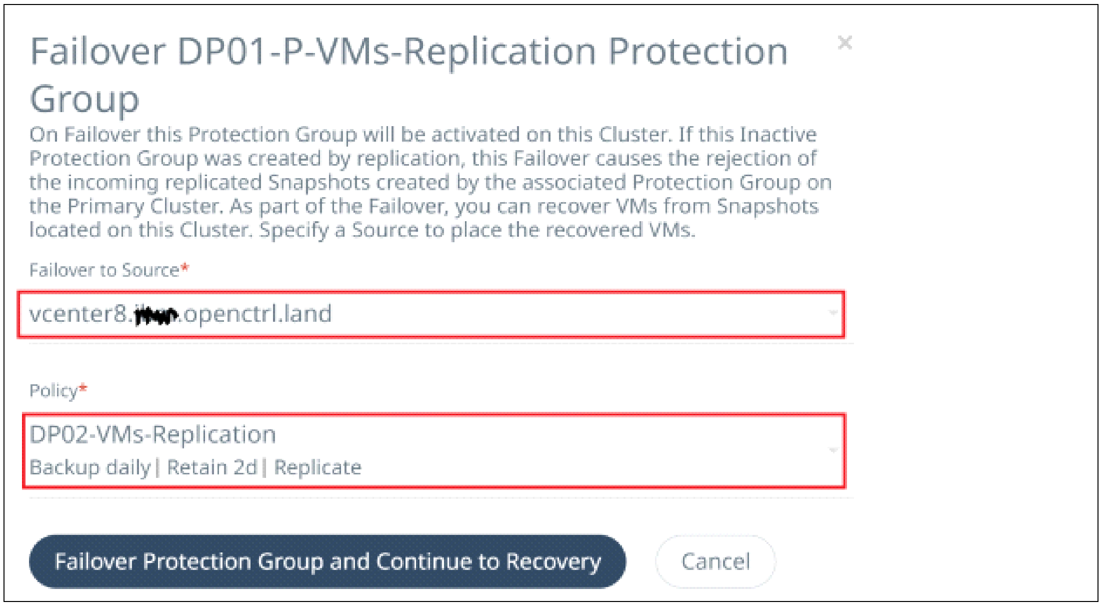

Failover DP01-P-VMs-Replication Protection Group

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.

Failover to Source*: vcenter8.[redacted].openctrl.land

Policy*: DP02-VMs-Replication
Backup daily | Retain 2d | Replicate

Figure 4-26 Protection Group configured for failover allowing continuous protection

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

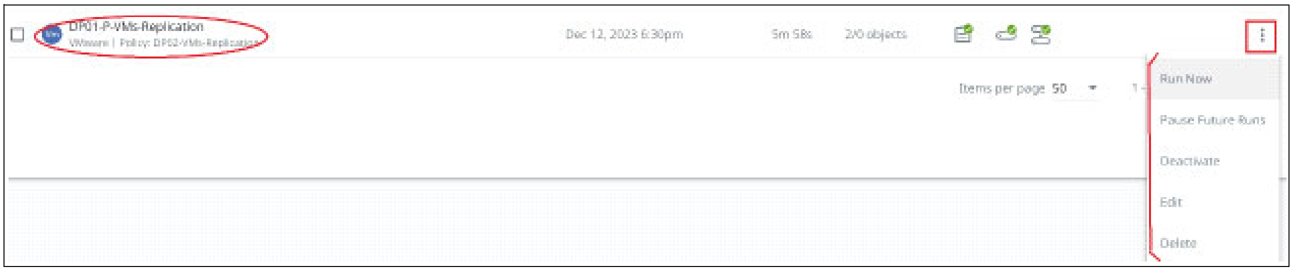

□	DP01-P-VMs-Replication VMware | Policy: DP02-VMs-Replication	Dec 12, 2023 6:30pm	5m 58s	2/0 objects		

Items per page 50

Figure 4-27 Protection Group Available Actions after Failover

The Protection Group is now ready to continue protecting the VMs on the replicated cluster and, there is now the ability to Edit the objects which require protection, Figure 4-28.

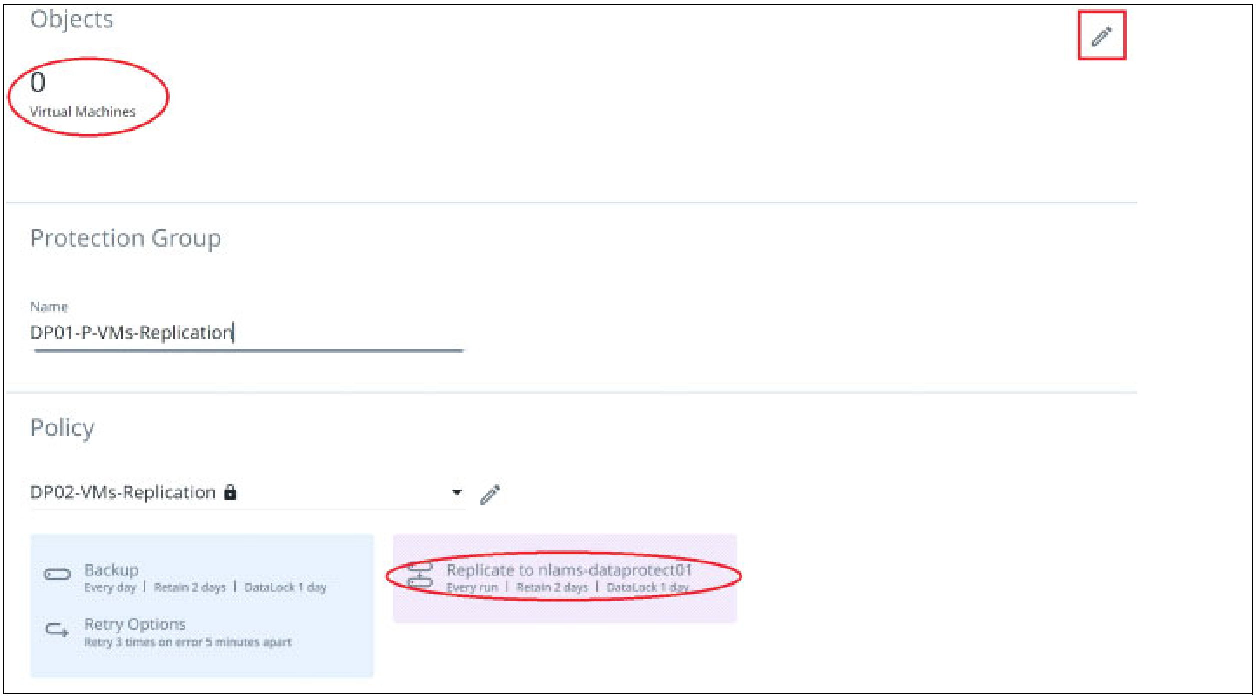

Objects

0
Virtual Machines

Protection Group

Name
DP01-P-VMs-Replication

Policy

DP02-VMs-Replication

Backup
Every day | Retain 2 days | DataLock 1 day

Retry Options
Retry 3 times on error 5 minutes apart

Replicate to nlams-dataprotect01
Every run | Retain 2 days | DataLock 1 day

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

> **Note:** *If a failover back to the original cluster is required, choose "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 has been returned to an available and accessible status, the system can be reverted back to running on the primary cluster by selecting "Deactivate" on the backup cluster. If the backup Protection Group is not Deactivated, the local backup job will continue to run successfully but the replication tasks will end in error state (Figure 4-29).

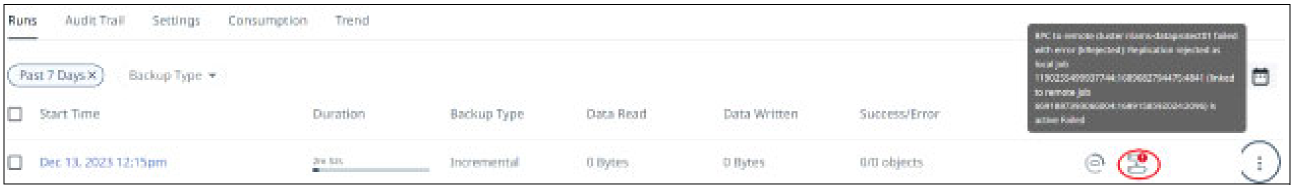

RPC to remote cluster nlams-dataprotect01 failed with error [kRejected]: Replication rejected as local job [illegible] (linked to remote job [illegible]) is active Failed

□	Start Time	Duration	Backup Type	Data Read	Data Written	Success/Error
□	Dec 13, 2023 12:15pm	[illegible]	Incremental	0 Bytes	0 Bytes	0/0 objects

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 on page 59).

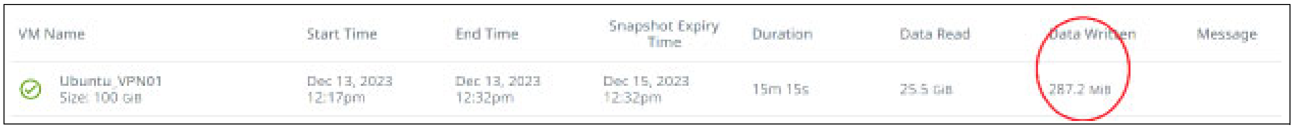

VM Name	Start Time	End Time	Snapshot Expiry Time	Duration	Data Read	Data Written	Message
✓ Ubuntu_VPN01 Size: 100 GiB	Dec 13, 2023 12:17pm	Dec 13, 2023 12:32pm	Dec 15, 2023 12:32pm	15m 15s	25.5 GiB	287.2 MiB	

Figure 4-30 Fallback Protection Group for continuous protection, local backup status

If the original source cluster is up and running, then “Deactivate Only” can be selected to make the Protection Group Failover Ready, Figure 4-31.

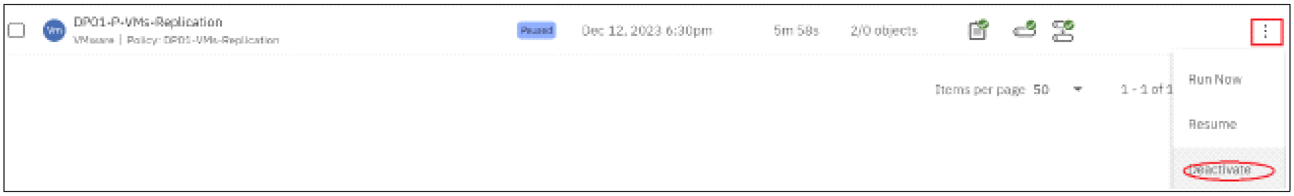

DP01-P-VMs-Replication
VMware | Policy: DP01-VMs-Replication
Paused Dec 12, 2023 6:30pm 5m 58s 2/0 objects

Items per page: 50 1 - 1 of 1

Figure 4-31 Fallback Protection Group deactivate option

Once ‘Deactivate’ is selected, the following options are presented for the Fallback cluster (Figure 4-32).

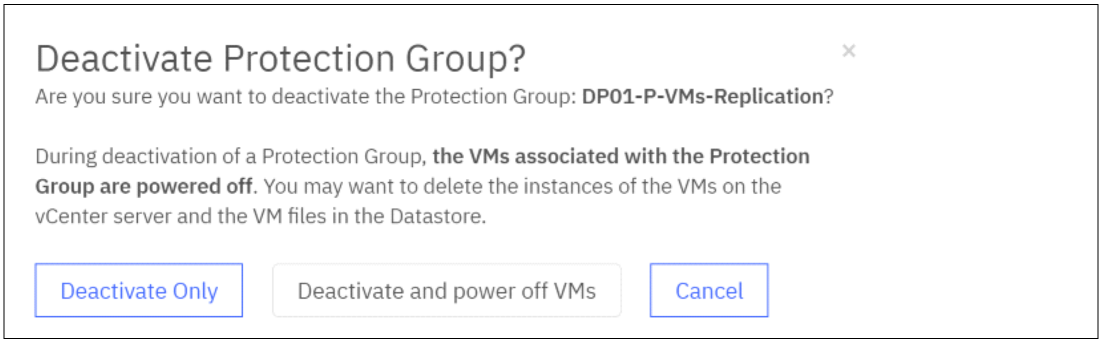

Deactivate Protection Group?

Are you sure you want to deactivate the Protection Group: **DP01-P-VMs-Replication**?

During deactivation of a Protection Group, **the VMs associated with the Protection Group are powered off**. You may want to delete the instances of the VMs on the vCenter server and the VM files in the Datastore.

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

Figure 4-33 ***Note:*** *The Deactivate and power off VMs options are only valid if there is an Active/Passive Datacenter or each vCenter on each Datacenter.*

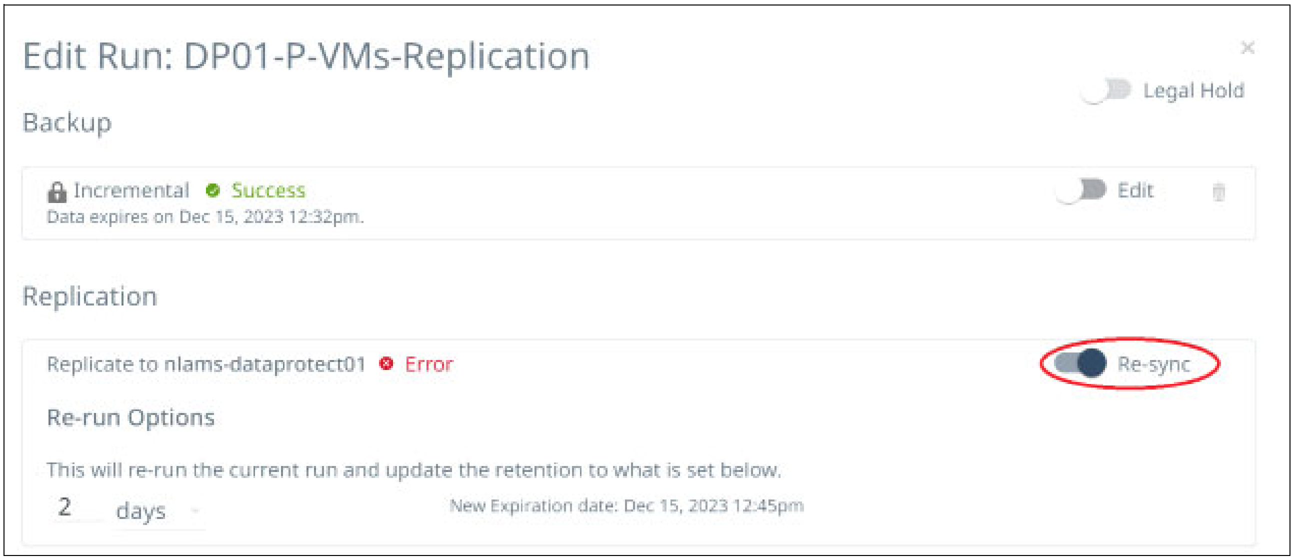

Edit Run: DP01-P-VMs-Replication

Backup

Incremental Success
Data expires on Dec 15, 2023 12:32pm.

Replication

Replicate to nlams-dataprotect01 Error

Re-run Options

This will re-run the current run and update the retention to what is set below.

2 days New Expiration date: Dec 15, 2023 12:45pm

Figure 4-34 Fallback Protection Group for continuous protection, Re-Sync after Deactivate Only

In this example, once the replication task finishes the Re-sync is completed. The run details show the replication task completed within 4m57s and replicated 1.7GB of data.

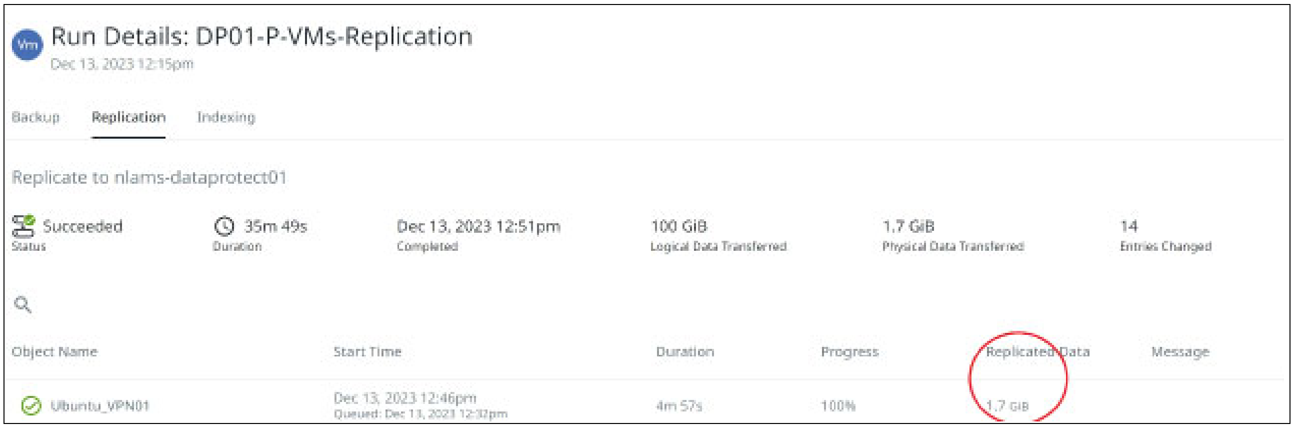

Run Details: DP01-P-VMs-Replication

Dec 13, 2023 12:15pm

Replicate to nlams-dataprotect01

Status	Duration	Completed	Logical Data Transferred	Physical Data Transferred	Entries Changed
Succeeded	35m 49s	Dec 13, 2023 12:51pm	100 GiB	1.7 GiB	14

Object Name	Start Time	Duration	Progress	Replicated Data	Message
Ubuntu_VPN01	Dec 13, 2023 12:46pm Queued: Dec 13, 2023 12:32pm	4m 57s	100%	1.7 GiB	

Figure 4-35 Fallback Protection Group for continuous protection, Re-Sync incremental Data Written

Now that replication is complete showing the full 1.7 GiB as Replicated Data.

4.7 CloudArchive the Workloads with IBM Data Protect Cluster

In today's technology-driven world, it is a major challenge having to store, protect and maintain availability of the ever-growing amount of data being created; think of everything from sensitive HR data and healthcare records to large media and entertainment files. As many organizations store their growing data collections on network-attached storage (GPFS/NAS) and run most of their workloads on VMWare infrastructure, it has never been more important to focus on NAS and VMWare data protection, how to store, protect and retain all that data.

CloudArchive Direct (CAD) is a policy-driven archival feature in IBM Storage Defender Data Protect that was built specifically to address these challenges by streaming data directly to lower-cost storage on an External Target without storing local backups. And While IBM Storage Defender Data Protect 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 Data Protect CloudArchive Direct can help to reliably archive this dataset securely to a supported S3 object store, one of which includes on premise IBM Cloud Object Storage.

4.7.1 VMware CloudArchive workflow S3 to IBM Cloud Object Storage (ICOS)

In this section, we describe the CloudArchive workflow of VMware according and provide an example of a One-to-One configuration using IBM Data Protect CAD. The backup data originates from the source VMware vCenter on the primary host and is then ingest to the primary IBM Data Protect cluster DC1. This data is then replicated to a secondary (DR) cluster DC2. In this configuration there is also an optional step of using the on-premises IBM iCOS erasure coding process solution for long-term data retention.

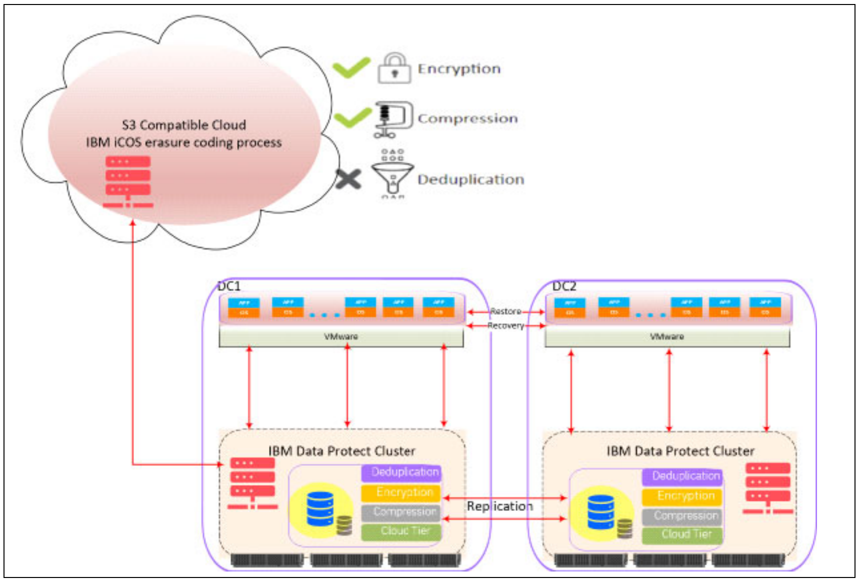

Figure 4-36 One to One CloudArchive workflow

Register External Targets as S3Compatible

To register an External Target the ip-address of the endpoint (iCOS Accessor) is needed, as is the Bucket name, admin Access Key ID and the Secret Access Key.

From the Infrastructure tab, select ‘External Targets’ then ‘Add external target’. When presented with the External Target fields, this example was configured using the following settings:

- select “Purpose”: Archival
- select “Storage Type”: S3Compatible
- select “Storage Class”: Regular
- fill-in “Bucket Name”: iCOSbucket
- fill-in “Access Key ID”: iCOSkeyID
- fill-in “Secret Access Key”: xxxxxxxxxxxxxx
- fill-in “Endpoint”: 10.0.2.x
- fill-in “AWS Signature Version”: Ver 4
- fill-in “External Target Name”: AnyName
- fill-in “Archival Format”: Incremental with Periodic Full

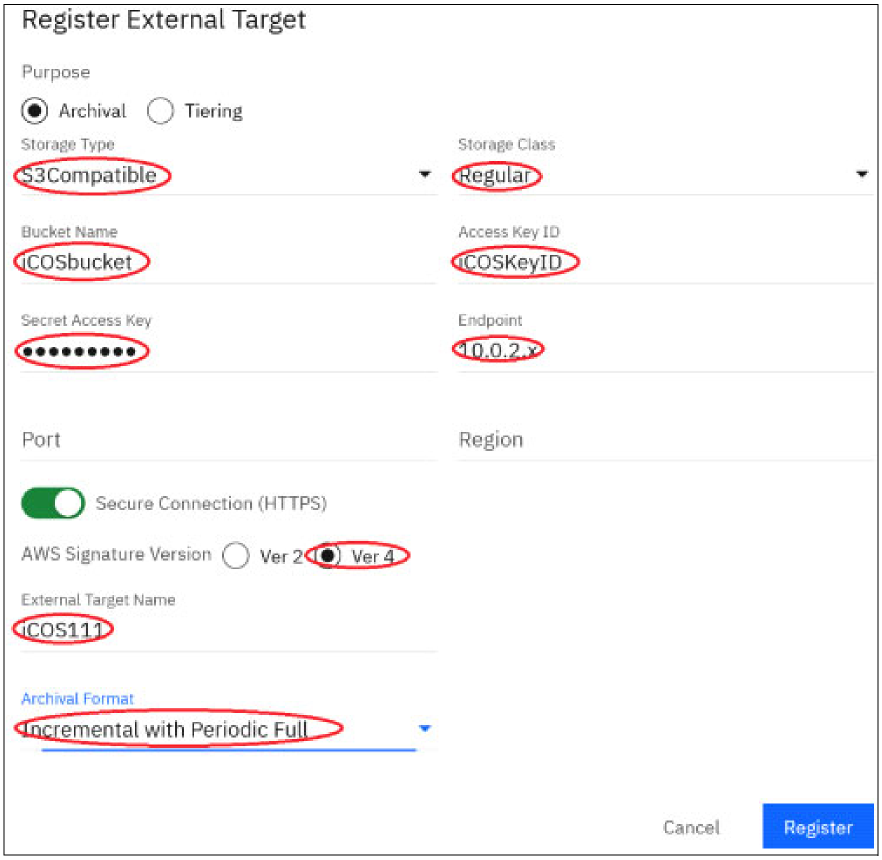

Register External Target

Purpose

● Archival ○ Tiering

Storage Type
S3Compatible

Storage Class
Regular

Bucket Name
iCOSbucket

Access Key ID
iCOSKeyID

Secret Access Key
●●●●●●●●●

Endpoint
10.0.2.x

Port

Region

Secure Connection (HTTPS)

AWS Signature Version ○ Ver 2 ● Ver 4

External Target Name
iCOS111

Archival Format
Incremental with Periodic Full

Figure 4-37 Register External Target IBM iCOS configuration example

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

Incremental with Periodic Full:

The first archive of a Protection Group is always a Reference Full Archive. A Reference Full forms the basis for Incremental Archives. With incremental archives, the next archive after a reference full is an Incremental Archive and 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 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 constitutes an increment on top of the previous incremental archive, i.e. it designates the previous Incremental Archive as its reference. In this way, incremental archives form an "archival chain" that starts at the selected archive and ends with the Reference Full.

In order 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 still refers to the Reference Full archive as its base, similar to the initial Incremental Archive.

The next Incremental Archive after a Full Archive will use the Full Archive as its a reference.

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

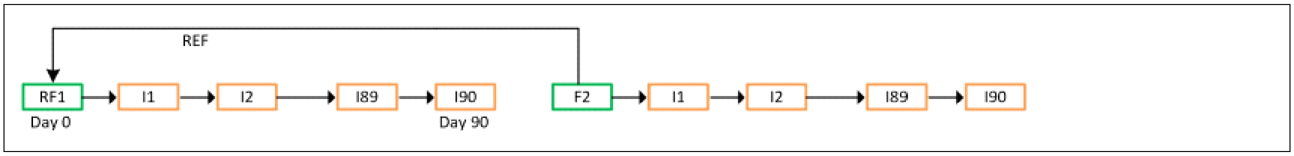

Figure 4-38 Incremental Archives reference workflow - using Reference Full Archive

During incremental archives, if the use of the current Reference Full Archive becomes too inefficient, mainly if its utilization as a reference drops below 50%, the Reference Full Archive is retired. At this point a new Reference Full Archive will be created.

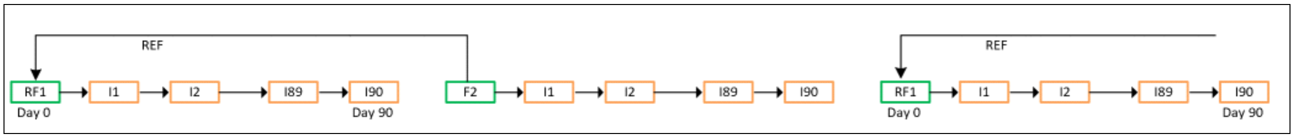

Figure 4-39 Reference Full Archive workflow - new base Reference Full Archive

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

CloudArchive Direct (CAD) incremental Forever:

Enabling this option will perform a first full archive and then continue with incremental forever archival of data from the Data Protect 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 will take place during the next reference archive creation.

> **Note:** IBM recommends using Incremental Forever and turning on Source Side Deduplication. This increases cloud performance and reduces the amount of data that the Data Protect cluster transfers to the external target.

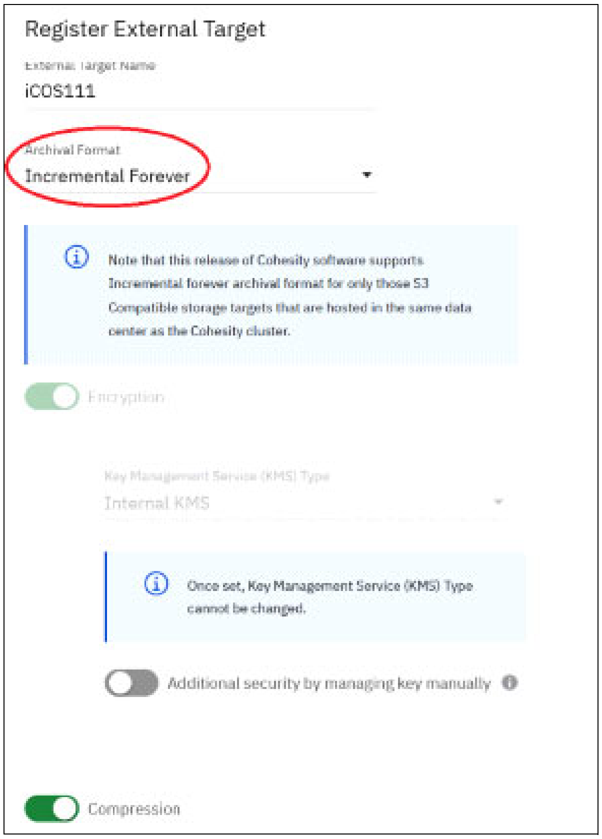

Register External Target

External Target Name
iCOS111

Archival Format
Incremental Forever

Note that this release of Cohesity software supports Incremental forever archival format for only those S3 Compatible storage targets that are hosted in the same data center as the Cohesity cluster.

Encryption

Key Management Service (KMS) Type
Internal KMS

Once set, Key Management Service (KMS) Type cannot be changed.

Additional security by managing key manually

Compression

Figure 4-40 Default settings when Registering an External Target

4.7.2 VMware CloudArchive restore IBM S3 IBM Cloud Object Storage

The following section provides an example of selecting and restoring an archive from IBM S3 iCOS. To perform the restore, select the Data Protection tab, then select ‘Recoveries’ and finally ‘Recover’. From here, filter the selections to find the desired VM archives:

Select the desired VM from the list to restore, then click on the pencil icon as shown in Figure 4-41 on page 65 to specify the location of the cloud and the recovery point.

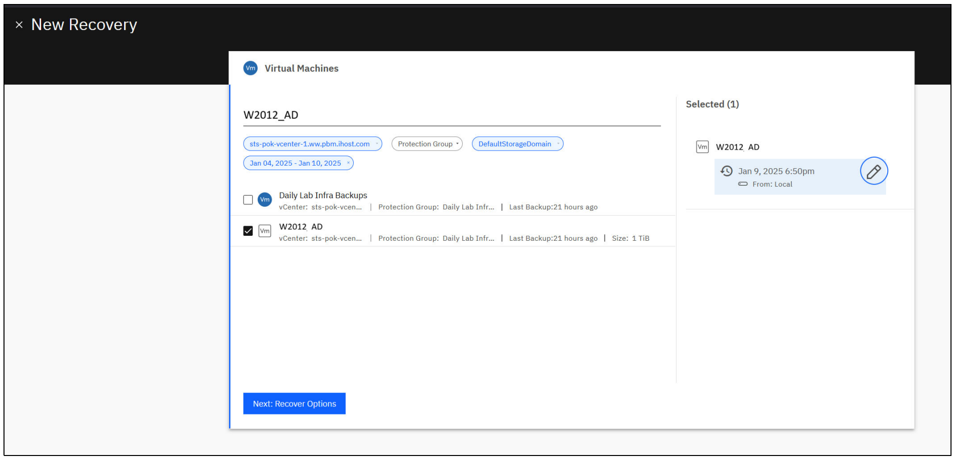

× New Recovery

Virtual Machines

W2012_AD

sts-pok-vcenter-1.ww.pbm.ihost.com
Protection Group
DefaultStorageDomain
Jan 04, 2025 - Jan 10, 2025

□ Daily Lab Infra Backups
vCenter: sts-pok-vcen... | Protection Group: Daily Lab Infr... | Last Backup:21 hours ago

☑ W2012_AD
vCenter: sts-pok-vcen... | Protection Group: Daily Lab Infr... | Last Backup:21 hours ago | Size: 1 TiB

Selected (1)

W2012_AD
Jan 9, 2025 6:50pm
From: Local

Figure 4-41 Virtual Machine restore settings pencil icon

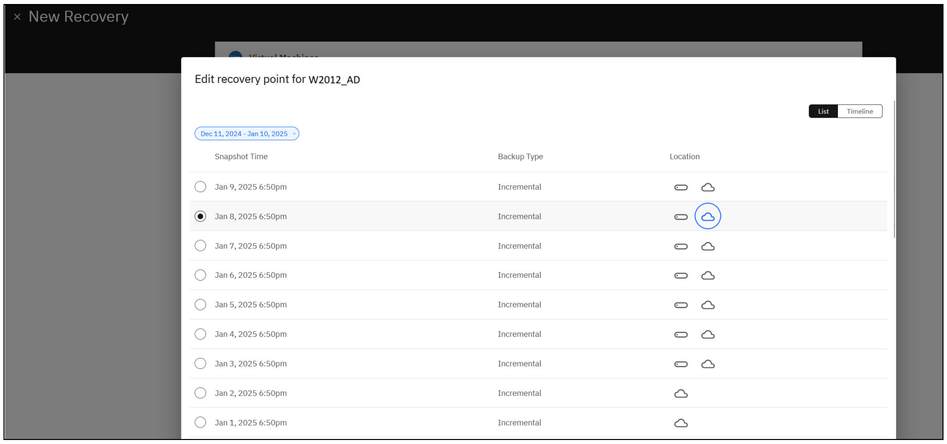

× New Recovery

Edit recovery point for W2012_AD

	Snapshot Time	Backup Type	Location
○	Jan 9, 2025 6:50pm	Incremental	▭ ☁
●	Jan 8, 2025 6:50pm	Incremental	▭ ☁
○	Jan 7, 2025 6:50pm	Incremental	▭ ☁
○	Jan 6, 2025 6:50pm	Incremental	▭ ☁
○	Jan 5, 2025 6:50pm	Incremental	▭ ☁
○	Jan 4, 2025 6:50pm	Incremental	▭ ☁
○	Jan 3, 2025 6:50pm	Incremental	▭ ☁
○	Jan 2, 2025 6:50pm	Incremental	☁
○	Jan 1, 2025 6:50pm	Incremental	☁

Figure 4-42 Point in time recover from cloud location

Once the archive is chosen and selected, the “Recovery Method” prompt will allow a choice of recovery methods and options. For this example the ‘Instance Recovery’ was selected and “Existing VM Handling”: Keep Existing VM option was used. Note that this option will power off and rename the existing VM as “_obsolete_xxx”. (Figure 4-43)

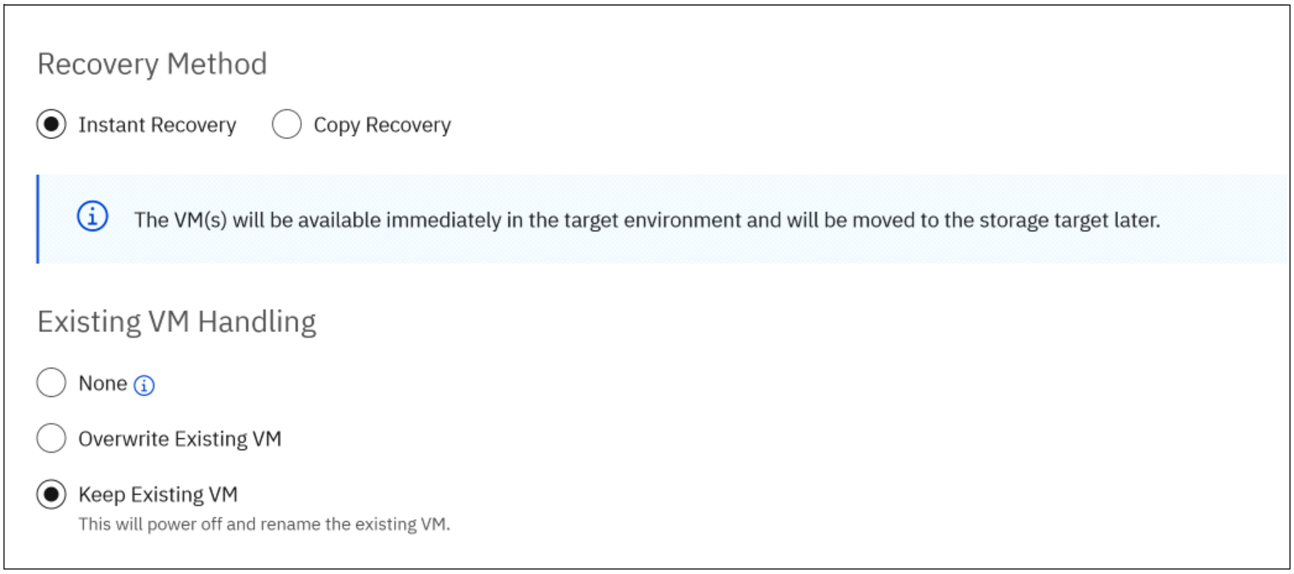

Recovery Method

● Instant Recovery ○ Copy Recovery

ⓘ The VM(s) will be available immediately in the target environment and will be moved to the storage target later.

Existing VM Handling

○ None ⓘ

○ Overwrite Existing VM

● Keep Existing VM
This will power off and rename the existing VM.

Figure 4-43 Recover method CloudArchive

In the activity log for the recovery process we can see the recovery is from the iCOS location (indicated by the kListVaultFSfiles task, Figure 4-44 on page 66).

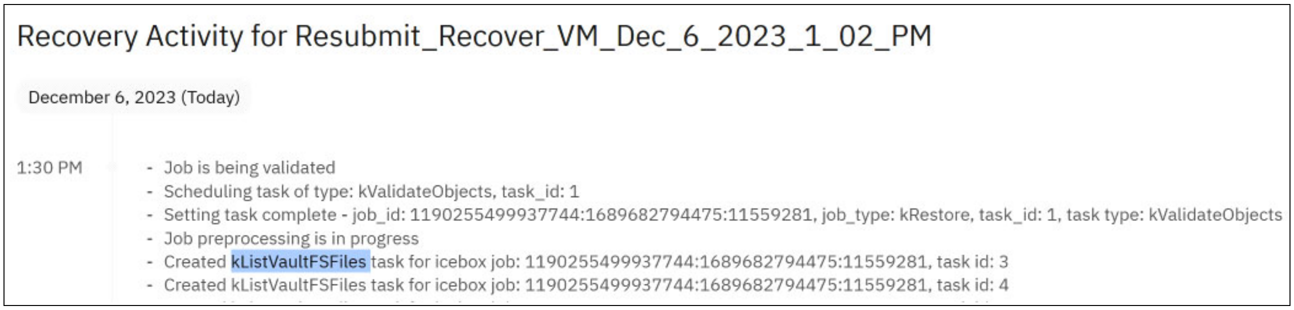

Recovery Activity for Resubmit_Recover_VM_Dec_6_2023_1_02_PM

December 6, 2023 (Today)

```
1:30 PM   - Job is being validated
          - Scheduling task of type: kValidateObjects, task_id: 1
          - Setting task complete - job_id: 1190255499937744:1689682794475:11559281, job_type: kRestore, task_id: 1, task type: kValidateObjects
          - Job preprocessing is in progress
          - Created kListVaultFSFiles task for icebox job: 1190255499937744:1689682794475:11559281, task id: 3
          - Created kListVaultFSFiles task for icebox job: 1190255499937744:1689682794475:11559281, task id: 4
```

Figure 4-44 Recovery Activity with CloudArchive task creation details

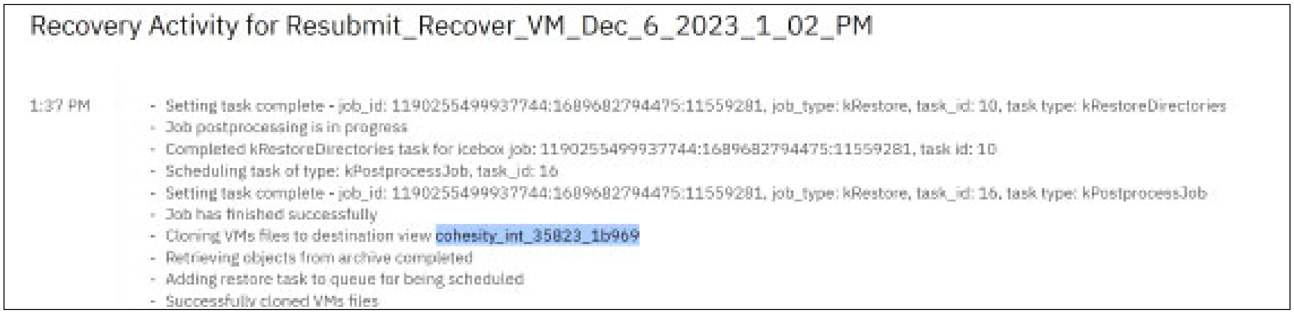

Recovery Activity for Resubmit_Recover_VM_Dec_6_2023_1_02_PM

```
1:37 PM   - Setting task complete - job_id: 1190255499937744:1689682794475:11559281, job_type: kRestore, task_id: 10, task type: kRestoreDirectories
          - Job postprocessing is in progress
          - Completed kRestoreDirectories task for icebox job: 1190255499937744:1689682794475:11559281, task id: 10
          - Scheduling task of type: kPostprocessJob, task_id: 16
          - Setting task complete - job_id: 1190255499937744:1689682794475:11559281, job_type: kRestore, task_id: 16, task type: kPostprocessJob
          - Job has finished successfully
          - Cloning VMs files to destination view cohesity_int_35823_1b969
          - Retrieving objects from archive completed
          - Adding restore task to queue for being scheduled
          - Successfully cloned VMs files
```

Figure 4-45 Recovery Activity with CloudArchive full action summary

The recovery snapshot is exported and can be seen in the NFS view for the vCenter DataCenter. The VM is then relocated to the original datastore. After the VM is successful cloned the NFS view is removed (Figure 4-45).

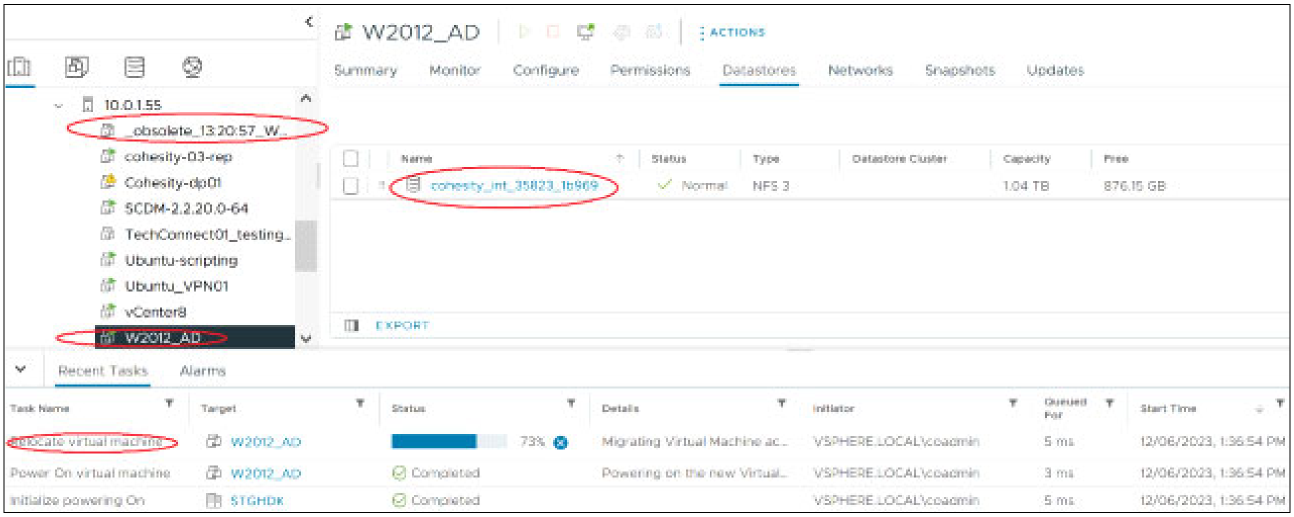

Name	Status	Type	Datastore Cluster	Capacity	Free
cohesity_int_35823_1b969	✓ Normal	NFS 3		1.04 TB	876.15 GB

Task Name	Target	Status	Details	Initiator	Queued For	Start Time
Relocate virtual machine	W2012_AD	73%	Migrating Virtual Machine ac...	VSPHERE.LOCAL\coadmin	5 ms	12/06/2023, 1:36:54 PM
Power On virtual machine	W2012_AD	Completed	Powering on the new Virtual...	VSPHERE.LOCAL\coadmin	3 ms	12/06/2023, 1:36:54 PM
Initialize powering On	STGHDK	Completed		VSPHERE.LOCAL\coadmin	5 ms	12/06/2023, 1:36:54 PM

Figure 4-46 Instant recovery CloudArchive to vCenter

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

5

Rapid Restore of Virtual Machines

In Chapter 4, “Protecting VMware Workloads With IBM Data Protect” on page 43 we learned about the various features and capabilities that Defender Data Protect offers for the protection of VMware virtual machine workloads. Chapter 5 will explore the different VMware recovery options available with Defender Data Protect:

- 5.1, “VM Recovery” on page 70
- 5.2, “File or Folder Recovery” on page 77
- 5.3, “Virtual Disk Recovery” on page 85
- 5.4, “Instant Volume Mount” on page 90

5.1 VM Recovery

Defender Data Protect allows for the recovery of one or more protected virtual machines from snapshots created previously by a Defender Data Protect Protection Group. See chapter 4 for more details on protecting VMware workloads. VMs can be recovered directly from a Data Protect cluster, or alternatively, 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 virtual machine from a Defender Data Protect backup:

1. Use the **DMS** left navigation menu to go to **Data Protection > Recoveries**, as shown in Figure 5-1:

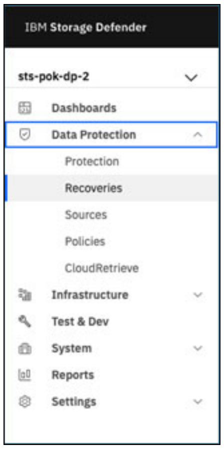

Figure 5-1 Storage Defender Data Protection menu

2. Then use the **Recover** menu on the top right of the **Recoveries** page to select **Recover > Virtual Machines > VMs**, as shown in Figure 5-2:

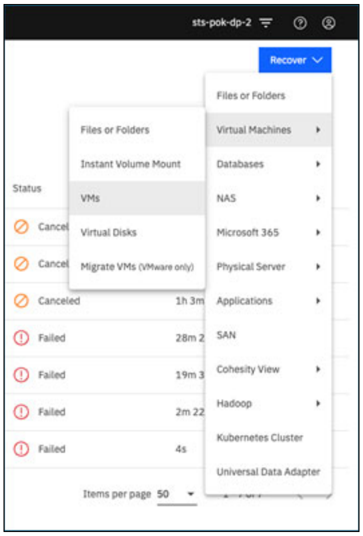

Figure 5-2 Virtual Machines Recovery Panel

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:

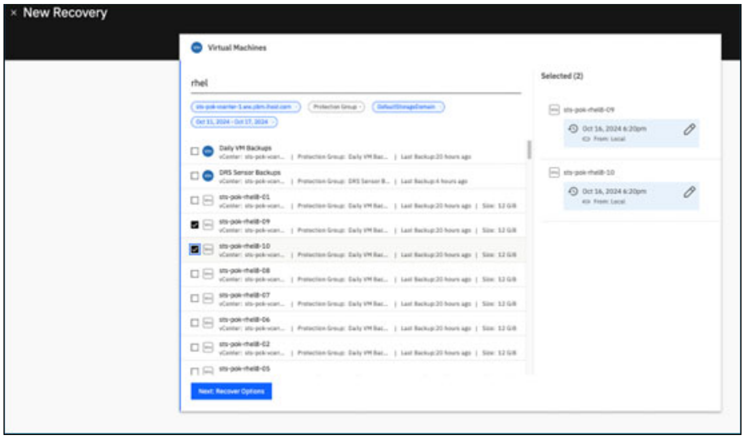

rhel

- ☐ Daily VM Backups
- ☐ DRS Sensor Backups
- ☐ sts-pok-rhel8-01
- ☑ sts-pok-rhel8-09
- ☑ sts-pok-rhel8-10
- ☐ sts-pok-rhel8-08
- ☐ sts-pok-rhel8-07
- ☐ sts-pok-rhel8-06
- ☐ sts-pok-rhel8-02
- ☐ sts-pok-rhel8-05

Selected (2):

- sts-pok-rhel8-09 — Oct 16, 2024 6:20pm, From: Local
- sts-pok-rhel8-10 — Oct 16, 2024 6:20pm, From: Local

Figure 5-3 New recovery VM search panel

> **Note:** For each virtual machine selected, Data Protect will automatically select 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-hand column, then use the pop-up window to select the desired snapshot, as shown in Figure 5-4.

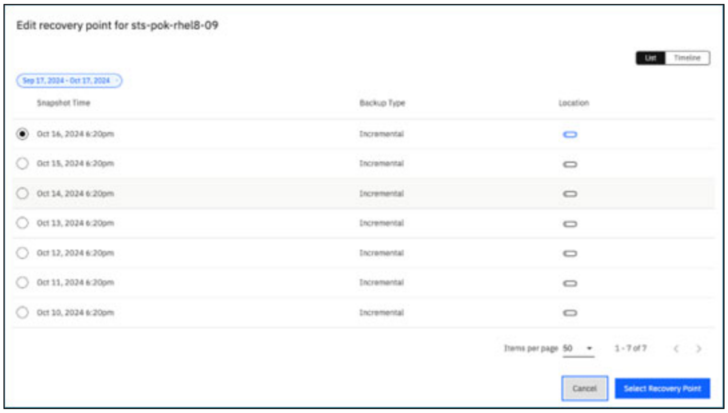

Edit recovery point for sts-pok-rhel8-09

	Snapshot Time	Backup Type	Location
●	Oct 16, 2024 6:20pm	Incremental	▭
○	Oct 15, 2024 6:20pm	Incremental	▭
○	Oct 14, 2024 6:20pm	Incremental	▭
○	Oct 13, 2024 6:20pm	Incremental	▭
○	Oct 12, 2024 6:20pm	Incremental	▭
○	Oct 11, 2024 6:20pm	Incremental	▭
○	Oct 10, 2024 6:20pm	Incremental	▭

Figure 5-4 Alternate restore point selection panel

> **Note:**
>
> - Selecting a Protection Group (In our example, "Daily VM Backups" or "DRS Sensor Backups") will restore all the VMs 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 available for recovery. If the example Data Protect cluster had created archive copies on external targets like cloud storage or Spectrum Protect, those snapshots would be indicated by cloud icons under the Location column. Simply click on the cloud icon to recover from an archival snapshot.

4. Once you have selected the desired snapshot(s), click the blue **Next: Recover Option**s button to proceed.
5. On the next screen, note the summary of the selected VMs at the top of page, as show, in Figure 5-5. If necessary, click the **pencil** icon to edit your selection.

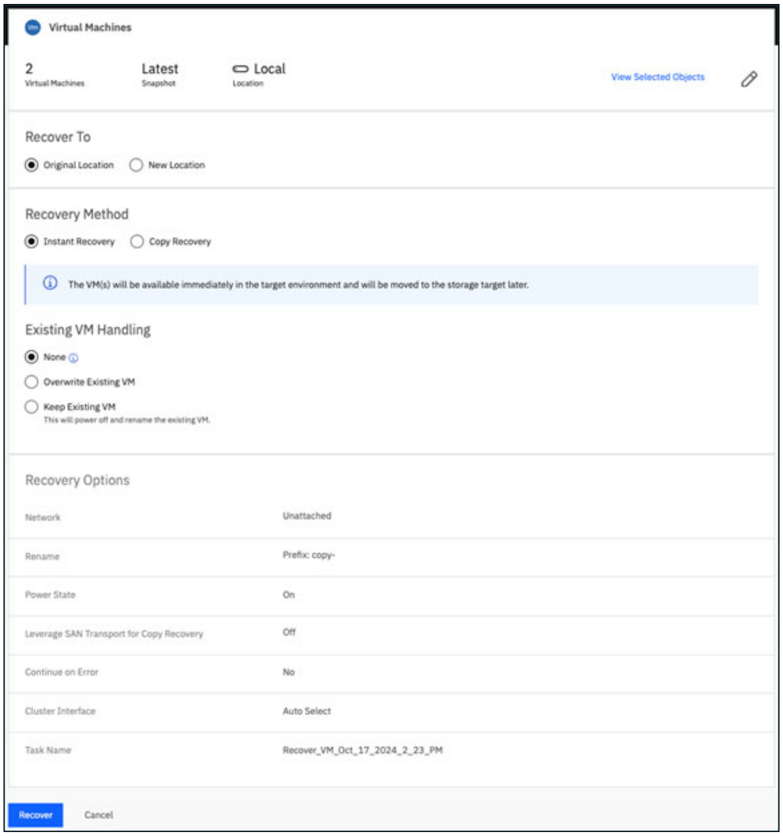

Virtual Machines

2 Virtual Machines | Latest Snapshot | Local Location | View Selected Objects

Recover To

● Original Location ○ New Location

Recovery Method

● Instant Recovery ○ Copy Recovery

The VM(s) will be available immediately in the target environment and will be moved to the storage target later.

Existing VM Handling

● None

○ Overwrite Existing VM

○ Keep Existing VM
This will power off and rename the existing VM.

Recovery Options

Network	Unattached
Rename	Prefix: copy-
Power State	On
Leverage SAN Transport for Copy Recovery	Off
Continue on Error	No
Cluster Interface	Auto Select
Task Name	Recover_VM_Oct_17_2024_2_23_PM

Figure 5-5 Virtual machine recovery summary panel

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

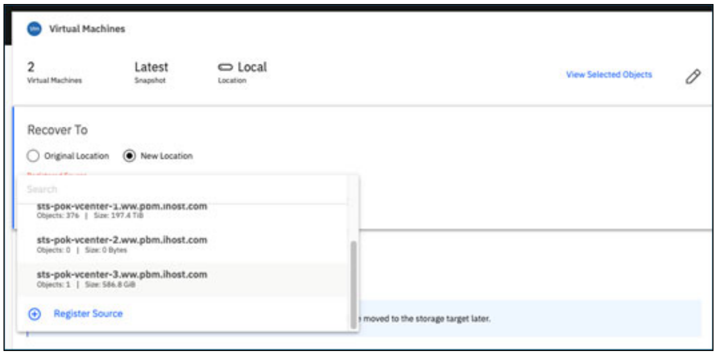

Virtual Machines

2 Virtual Machines | Latest Snapshot | Local Location | View Selected Objects

Recover To

○ Original Location ● New Location

Search

sts-pok-vcenter-1.ww.pbm.ihost.com
Objects: 376 | Size: 197.4 TiB

sts-pok-vcenter-2.ww.pbm.ihost.com
Objects: 0 | Size: 0 Bytes

sts-pok-vcenter-3.ww.pbm.ihost.com
Objects: 1 | Size: 586.8 GiB

Register Source

moved to the storage target later.

Figure 5-6 Virtual Machine recovery setting panel

Further, note that the **Recovery Method** is set to **Instant Recovery**.

> **Note:** When recovering VMs, two recovery methods are supported:
>
> - **Instant Recovery**: **Instant recovery** is the fastest method to get the VM(s) up and running. In **Instant Recovery**, the Data Protect cluster creates a writeable snapshot of the immutable snapshot(s) selected in step 3, and mounts it as a temporary NFS datastore(s) in the target VMware environment. The VMs present in this datastore are then registered to the target vCenter and powered on. After the last VM in the recovery job is powered on, Data Protect then orchestrates the storage vMotion of the VMs from the NFS datastore to the target datastore. Once the vMotion tasks complete, the Data Protect cluster unmounts the temporary datastore. The benefit of **Instant Recovery** is that the VMs being recovered can be powered on before any data is moved from the backup to the target datastore. Instant Recovery provides the foundation for **Instant Mass Restore**, which we will cover in more detail in section 5.1.2, “Instant Mass Restore” on page 76.
> - **Copy Recovery**: **Copy Recovery** is the fastest method to get the VM files out of the backup system and back to the VMware datastore. It will first copy the VM files from the Data Protect cluster to the VMware target datastore, and only once the VM files are fully restored to the datastore can the VM(s) be powered on. To reiterate, **Copy Recovery** provides a faster method to move the files back to the datastore than **Instant Recovery**, but the VMs cannot be powered up until the VM files are fully restored.

When performing a VM restore and **Original Location** is selected, Data Protect offers multiple options for how to handle existing VMs in the original location:

- **None**: VM recovery will fail if there is an existing VM in the recovery location.
- **Overwrite Existing VM**:
  - For an **Instant Recovery**, the existing VM will first be deleted from the target system prior to the restoration
  - For a **Copy Recovery**, Data Protect will attempt a differential restore, where only the blocks that differ between the existing VM and the backup copy are restored from the backup. This 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 will be powered off and renamed prior to the restoration.

6. Finally, make note of the available **Recovery Options** shown in Figure 5-7:

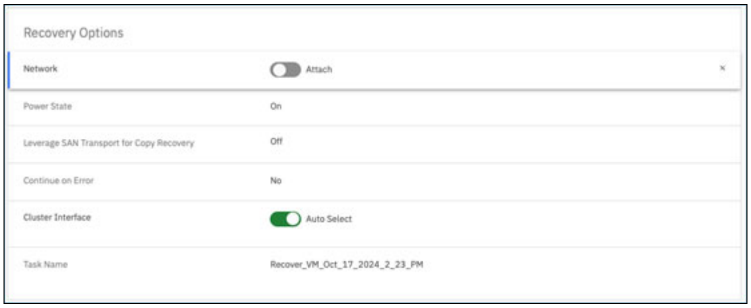

Recovery Options

Network	Attach
Power State	On
Leverage SAN Transport for Copy Recovery	Off
Continue on Error	No
Cluster Interface	Auto Select
Task Name	Recover_VM_Oct_17_2024_2_23_PM

Figure 5-7 VM Recovery Options panel

Network:

- For recovery to the **Original Location**, default value is **Detached** - all VMs will be powered on detached from any network. If **Attached** is selected, all VMs will be powered on attached to their original virtual network. Additionally, selecting **Attached** unlocks the

Start Connected option. If **Start Connected** is selected, the VMs will connect to the original network(s) when the VMs reboot. If unselected, the VMS will not be connected to any virtual networks up reboots.

> **Note:** When recovering to the **Original Location**, Data Protect will recover 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 will have the same setting.

- For recovery to a **New Location**, default value is **Detached** - all VMs will be powered on detached from any network. If **Attached** is selected, a drop-down menu is displayed, allowing you to select the destination virtual network to attach the VMs to. This network will be attached to all the VMs in the recovery job. To override this, you can click the **Add Network Override** link and select separate network settings for each VM in the recovery job.
- Additionally, selecting **Attached** unlocks the **Start Connected** and **Preserve MAC Address** options. If **Start Connected** is selected, the VMs will be connect to the selected network when the VMs reboot. If unselected, the VMs will not be connected to any virtual networks up reboots. If **Preserve MAC Address** is selected, VMs that were set to have manual MAC addresses will be 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, the recovered VMs will have manual MAC addresses.

Rename:

- Allows you to enter a prefix and/or suffix to the names of the recovered VMs.

Power State:

- Default value is **Power On** - VMs will be powered on upon restoration.

Leverage SAN Transport for Copy Recovery:

- Instruct Data Protect to leverage SAN transport for datastores that reside on Fibre Channel (FC) or iSCSI-based storage arrays. Enable this option to enable SAN transport mode over FC or iSCSI for recovery of VMware VMs.

> **Note:** SAN transport over FC requires that FC is configured and zoned between the Data Protect cluster and the target ESX host(s).

Continue on Error:

- Default value is **No** - the recovery job will fail upon the first error recovering a VM. If **Yes** is selected, the protection job will continue to run even if errors are encountered.

Cluster Interface:

- Only available for **Instant Recoveries**. Default value is **Auto Select** - the Data Protect cluster will automatically select the correct interface group to use for the recovery. To manually select the desired **Interface Group**, click the **Cluster Interface** row, disable the **Auto Select** slider and use the **Interface Group** drop-down menu to select the desired **Interface Group** to use for the recovery job.

Task Name:

- Optionally, change the default name of the recovery job.

7. Click **Recover** to initiate the restore job. You will be presented with a **Recovery Summary** screen, as shown in Figure 5-8:

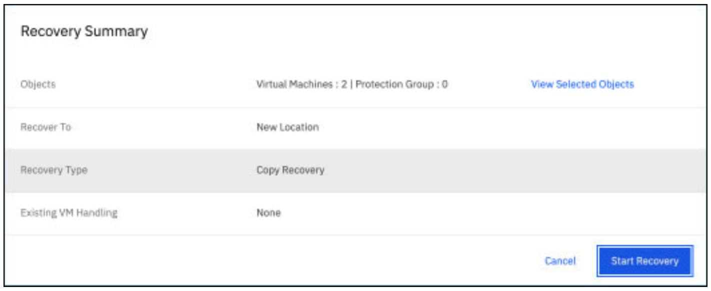

Recovery Summary

Objects	Virtual Machines : 2 | Protection Group : 0	View Selected Objects
Recover To	New Location	
Recovery Type	Copy Recovery	
Existing VM Handling	None	

Figure 5-8 Recovery summary panel

8. If the details look correct, click the **Start Recovery** button to initiate the restore. You will be redirected to the **Data Protection > Recoveries** page, where you will see the restore job running, as shown in Figure 5-9:

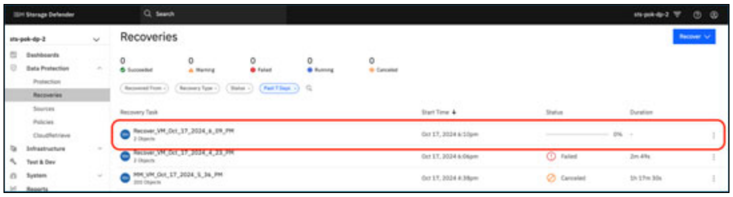

Recoveries

0 Succeeded · 0 Warning · 0 Failed · 0 Running · 0 Canceled

Recovery Task	Start Time	Status	Duration
Recover_VM_Oct_17_2024_6_09_PM 2 Objects	Oct 17, 2024 6:10pm	0%	-
Recover_VM_Oct_17_2024_4_23_PM 2 Objects	Oct 17, 2024 6:06pm	Failed	2m 49s
MM_VM_Oct_17_2024_5_36_PM 200 Objects	Oct 17, 2024 4:38pm	Canceled	1h 17m 30s

Figure 5-9 Recoveries details panel

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

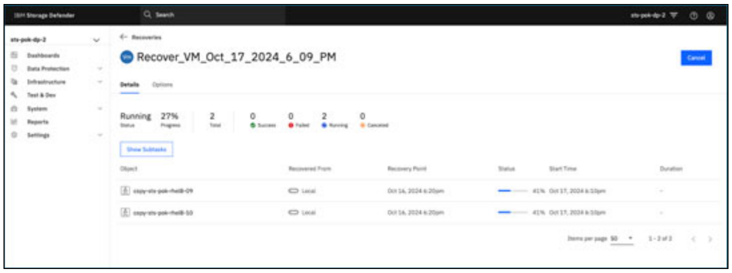

Recover_VM_Oct_17_2024_6_09_PM

Running (Status) · 27% (Progress) · 2 Total · 0 Success · 0 Failed · 2 Running · 0 Canceled

Object	Recovered From	Recovery Point	Status	Start Time	Duration
copy-sts-pok-rhel8-09	Local	Oct 16, 2024 6:20pm	41%	Oct 17, 2024 6:10pm	-
copy-sts-pok-rhel8-10	Local	Oct 16, 2024 6:20pm	41%	Oct 17, 2024 6:10pm	-

Items per page: 50 · 1 - 2 of 2

Figure 5-10 Recovery job details panel

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

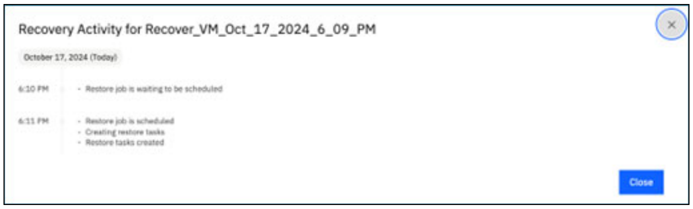

Recovery Activity for Recover_VM_Oct_17_2024_6_09_PM

October 17, 2024 (Today)

6:10 PM
- Restore job is waiting to be scheduled

6:11 PM
- Restore job is scheduled
- Creating restore tasks
- Restore tasks created

Figure 5-11 Recovery job log details

11. The figure Figure 5-12 on page 76 shows the successful completion status of the recovery job:

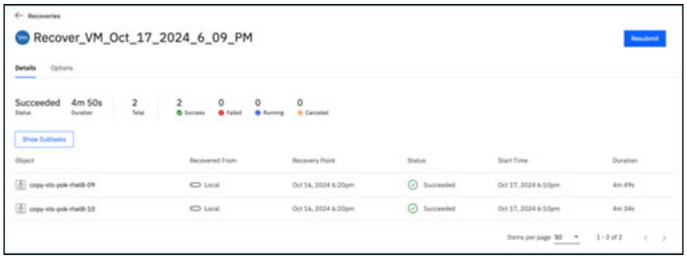

← Recoveries

Recover_VM_Oct_17_2024_6_09_PM

Succeeded (Status) · 4m 50s (Duration) · 2 (Total) · 2 Success · 0 Failed · 0 Running · 0 Canceled

Object	Recovered From	Recovery Point	Status	Start Time	Duration
copy-sts-pok-rhel8-09	Local	Oct 16, 2024 6:20pm	Succeeded	Oct 17, 2024 6:10pm	4m 49s
copy-sts-pok-rhel8-10	Local	Oct 16, 2024 6:20pm	Succeeded	Oct 17, 2024 6:10pm	4m 34s

Items per page 50 · 1 - 2 of 2

Figure 5-12 Recovery job completion status details

12. Further, as shown in Figure 5-13, we can confirm in vCenter that our VMs have been successfully restored:

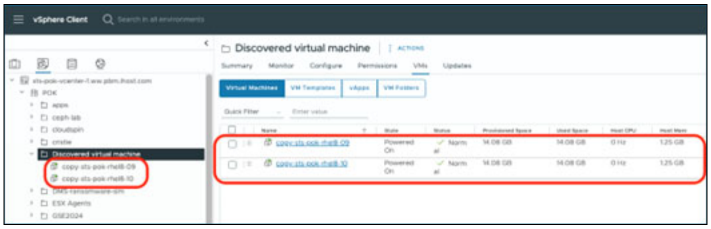

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem
copy-sts-pok-rhel8-09	Powered On	✓ Normal	14.08 GB	14.08 GB	0 Hz	1.25 GB
copy-sts-pok-rhel8-10	Powered On	✓ Normal	14.08 GB	14.08 GB	0 Hz	1.25 GB

Figure 5-13 Restored VMs displayed in Vcenter

Now that we have explored the fundamental capabilities of VMware virtual machine restoration, we will examine how Data Protect leverages these capabilities, along with its scale-out architecture, to support large scale, raid restore of virtual machines.

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 Instant Mass Restore (IMR). With IMR, you can use the Cohesity NFS views as the datastore, allowing the recovery process to happen much faster or instantly and circumventing the requirement of restoring the VMs to the original datastore before booting.

Before you can recover Protected Objects (VMs), a Snapshot must exist.

Key to recovery from a large-scale data corruption incident or other major intrusion is the ability to respond to the need to bring the business services back online quickly. The ability to select multiple VMs in a recovery task, and to recover them leveraging the Instant Recovery method discussed above in <link to Recovery Method description above>, is called Instant Mass Restore (IMR). In an Instant Mass Restore, the IBM Defender Data Protect cluster performing the recovery follows the same process as described above, but allows for Instant Recovery at scale. IMR facilitates the rapid recovery of hundreds or even thousands of VMs that is often required in cyber incidents.

The chart below details the use cases tested during a Data Protect Proof of Concept. It illustrates a variety of different recovery scenarios, including multiple types of large-scale restoration in a VMware environment. This set of use cases highlight the power of IBM Storage Defender Data Protect to drive workload restoration at scale and speed.

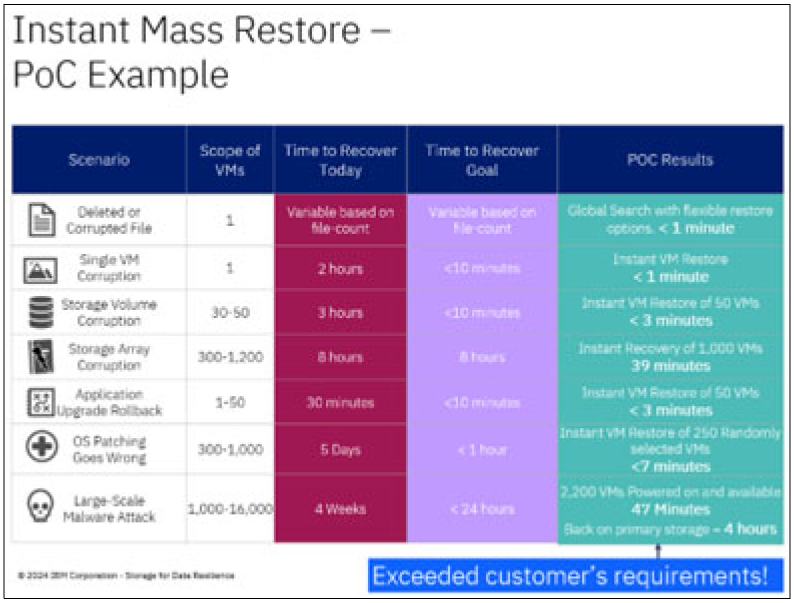

Instant Mass Restore – PoC Example

Scenario	Scope of VMs	Time to Recover Today	Time to Recover Goal	POC Results
Deleted or Corrupted File	1	Variable based on file-count	Variable based on file-count	Global Search with flexible restore options. **< 1 minute**
Single VM Corruption	1	2 hours	<10 minutes	Instant VM Restore **< 1 minute**
Storage Volume Corruption	30-50	3 hours	<10 minutes	Instant VM Restore of 50 VMs **< 3 minutes**
Storage Array Corruption	300-1,200	8 hours	8 hours	Instant Recovery of 1,000 VMs **39 minutes**
Application Upgrade Rollback	1-50	30 minutes	<10 minutes	Instant VM Restore of 50 VMs **< 3 minutes**
OS Patching Goes Wrong	300-1,000	5 Days	< 1 hour	Instant VM Restore of 250 Randomly selected VMs **<7 minutes**
Large-Scale Malware Attack	1,000-16,000	4 Weeks	< 24 hours	2,200 VMs Powered on and available **47 Minutes** Back on primary storage – **4 hours**

© 2024 IBM Corporation - Storage for Data Resilience

Exceeded customer's requirements!

Figure 5-14 Instant Mass Restore PoC details table

The example depicted here highlights the recovery of 200 virtual machines.

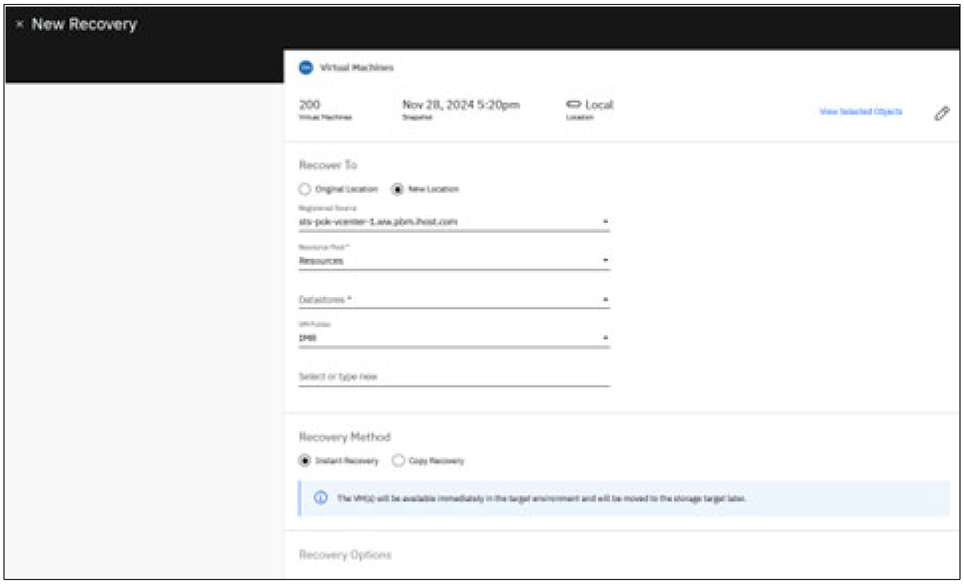

× New Recovery

Virtual Machines

200 Virtual Machines
Nov 28, 2024 5:20pm Snapshot
Local Location

View Selected Objects

Recover To

○ Original Location ● New Location

Registered Source: sts-pok-vcenter-1.aw.pbm.ihost.com

Resource Pool*: Resources

Datastores *

VM Folder: DMR

Select or type new

Recovery Method

● Instant Recovery ○ Copy Recovery

The VM(s) will be available immediately in the target environment and will be moved to the storage target later.

Recovery Options

Figure 5-15 Instant Mass Restore finished restore example

5.2 File or Folder Recovery

In addition to full virtual machine restoration, Data Protect provides the ability to recover files and folders from a virtual machine backup, to either their original location or to a new location. Alternatively, you can download the desired files and or folders contained in a Data Protect backup directly via the Data Protect UI (DMS). Optionally, you can instruct Data Protect 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 Data Protect to offer granular, file and folder level restore capabilities from VMware virtual machine backups is filesystem indexing. Data Protect, by default, attempts to index the filesystems of the Windows and Linux VMs that it backs up. While indexing is an optional parameter for VM backups, most users will find great benefit in leaving it enabled. Doing so allows Data Protect to include this index data in a global catalog, allowing Defender users to leverage a single search interface for recovering all the file data across their estate.

> **Note:** Data protect supports recovering files and/or folders from Windows VMs to Windows VMs, and from Linux VMs to Linux VMs only.
>
> File recovery and instant volume mount from a Data Protect backup of a Windows VM with Windows deduplication enabled for one or more volumes is only supported when the target Windows machines had deduplication installed. (Deduplication does not need to be enabled on the target machined volume(s), only installed.)

5.2.2 Recovering Files or Folders

Follow these steps to recover files or folders with Defender Data Protect:

1. Navigate to **Data Protection > Recoveries**, as shown in Figure 5-16:

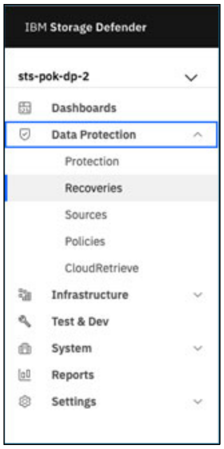

Figure 5-16 Storage Defender Data Protection menu

2. Then use the **Recover** menu on the top right of the **Recoveries** page to select **Recover > Virtual Machines > Files or Folders**, as shown in Figure 5-17:

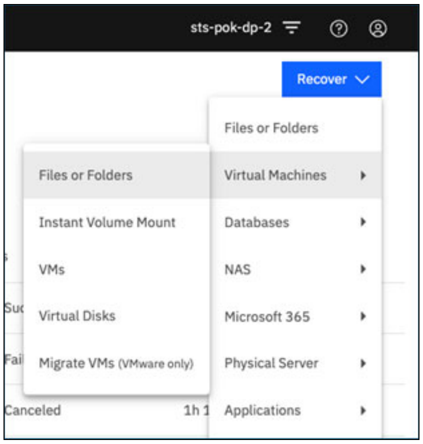

Figure 5-17 VM Recovery drop down menu

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

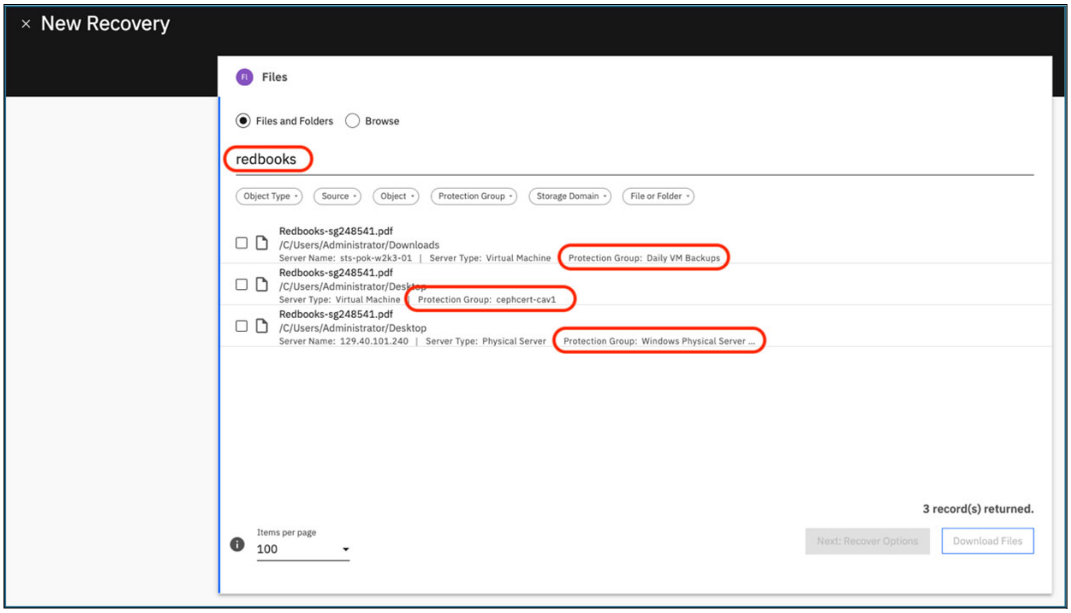
× New Recovery

Files

● Files and Folders ○ Browse

redbooks

□ Redbooks-sg248541.pdf
/C/Users/Administrator/Downloads
Server Name: sts-pok-w2k3-01 | Server Type: Virtual Machine | Protection Group: Daily VM Backups

□ Redbooks-sg248541.pdf
/C/Users/Administrator/Desktop
Server Type: Virtual Machine | Protection Group: cephcert-cav1

□ Redbooks-sg248541.pdf
/C/Users/Administrator/Desktop
Server Name: 129.40.101.240 | Server Type: Physical Server | Protection Group: Windows Physical Server ...

3 record(s) returned.

Items per page
100

Figure 5-18 VM New Recovery file search panel

> **Note:** The **Files and Folders** search feature allows you 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/Folder Recovery** job can restore multiple files and/or folders that were backed up as part of the same **Protection Group**. Use separate **File/Folder Recovery** jobs to recover files and/or folders backed up in separate **Protection Groups**.

4. Check the box for your desired file(s) or folder(s), to add the item(s) to the **Recovery Cart**. Note that selecting an item the search results will immediately filter the search results to only show items from the same **Protection Group and Recovery Point**, as shown in Figure 5-19 on page 80:

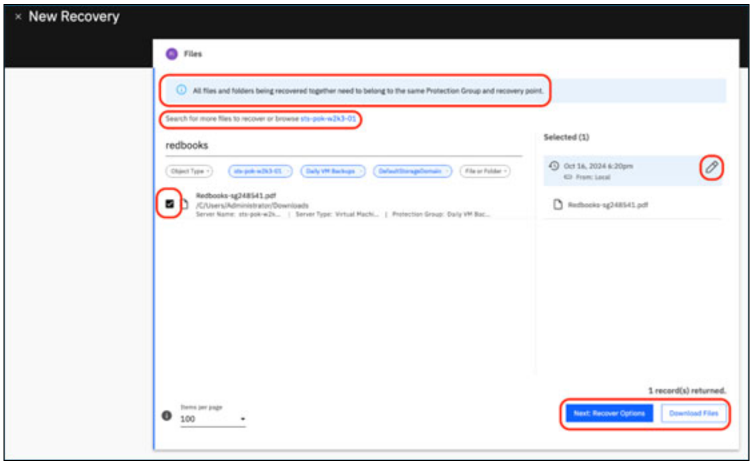
× New Recovery

Files

All files and folders being recovered together need to belong to the same Protection Group and recovery point.

Search for more files to recover or browse sts-pok-w2k3-01

redbooks

☑ Redbooks-sg248541.pdf
/C/Users/Administrator/Downloads
Server Name: sts-pok-w2k... | Server Type: Virtual Machi... | Protection Group: Daily VM Bac...

Selected (1)

Oct 16, 2024 6:20pm
From: Local

Redbooks-sg248541.pdf

1 record(s) returned.

Items per page: 100

Figure 5-19 Files selection for recovery example

5. Optionally, click the **pencil** icon to select an earlier version. (Again, Data Protect defaults to the latest backup copy.)

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

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

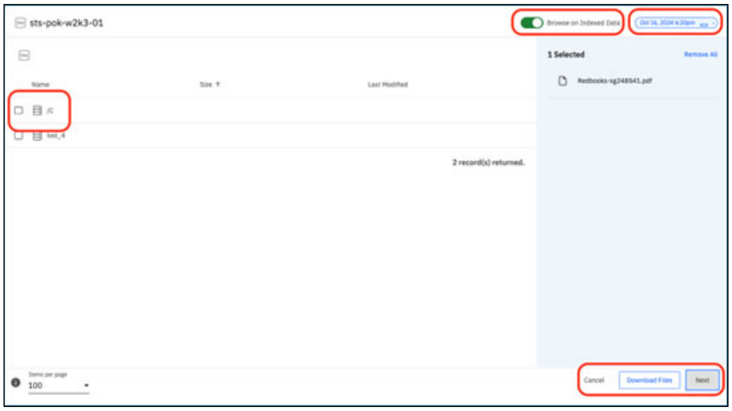
sts-pok-w2k3-01

Name	Size ↑	Last Modified
□ /C		
□ bml_4		

2 record(s) returned.

1 Selected

Redbooks-sg248541.pdf

Items per page: 100

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, Data Protect presents an indexed (cached) representation of the filesystem of this VM at the date and time of the selected **Recovery Point** shown at the top right. Browsing on indexed data allows for better search performance, only displaying files and folders that are indexed. Hidden paths, as well as manually excluded paths are not shown in this mode.

Flipping the slider off will trigger Data Protect 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**.

Clicking the **Recovery Point** timestamp allows you to select an alternate **Recovery Point**, as shown in Figure 5-21:

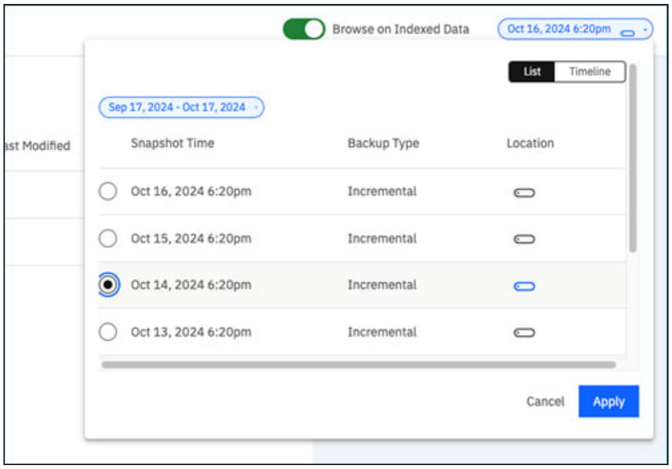

Browse on Indexed Data

Oct 16, 2024 6:20pm

Sep 17, 2024 - Oct 17, 2024

	Snapshot Time	Backup Type	Location
○	Oct 16, 2024 6:20pm	Incremental	⊂⊃
○	Oct 15, 2024 6:20pm	Incremental	⊂⊃
●	Oct 14, 2024 6:20pm	Incremental	⊂⊃
○	Oct 13, 2024 6:20pm	Incremental	⊂⊃

Figure 5-21 Alternative recovery point selection

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

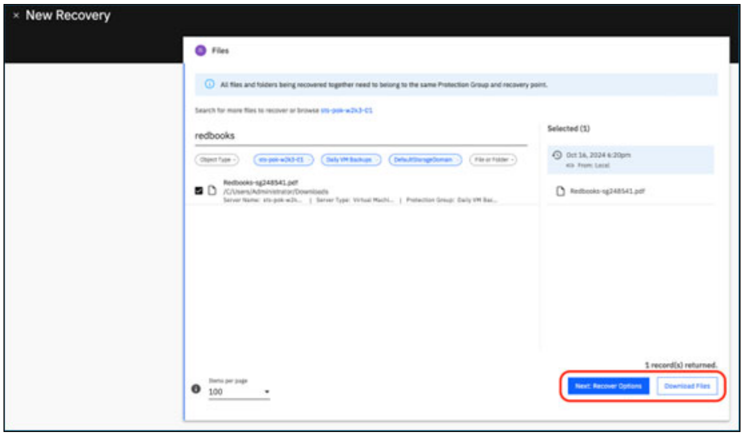

All files and folders being recovered together need to belong to the same Protection Group and recovery point.

Search for more files to recover or browse sts-pok-w2k3-01

redbooks

☑ Redbooks-sg248541.pdf
/C/Users/Administrator/Downloads
Server Name: sts-pok-w2k... | Server Type: Virtual Machi... | Protection Group: Daily VM Bac...

Selected (1)

Oct 16, 2024 6:20pm
From: Local

Redbooks-sg248541.pdf

1 record(s) returned.

Items per page 100

Figure 5-22 Return to recovery cart

8. From the original **Recovery Cart**, we can select either:

- **Next: Recover Options** to proceed with recovering the selected item(s)

or

- **Download Files** to download a zip archive of the selected items directly through the browser

9. In our example, we choose **Next: Recover Options**. The **New Recovery** page now displays a summary of the selected items at the top of the page (Again, click the **pencil** icon to edit the items selected for recovery) followed by additional recovery parameters and options, as show in Figure 5-23 on page 82:

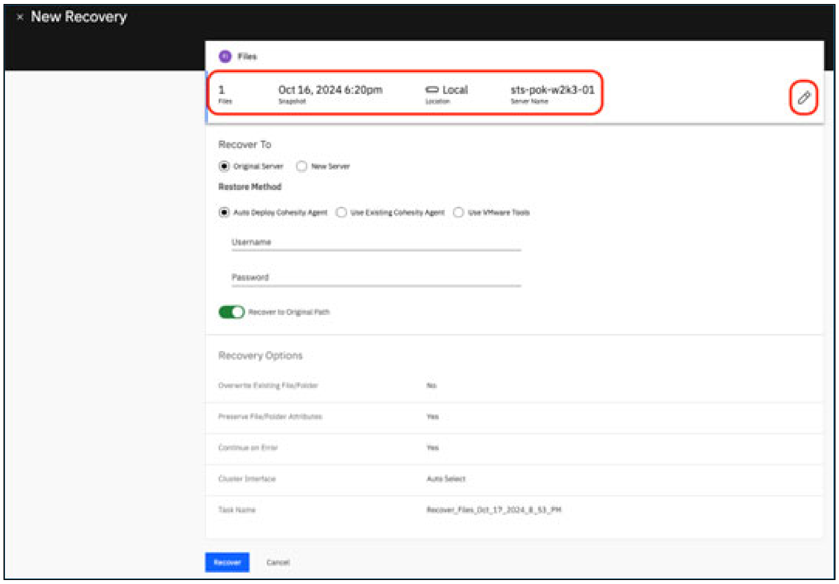

× New Recovery

Files

1 Files · Oct 16, 2024 6:20pm Snapshot · Local Location · sts-pok-w2k3-01 Server Name

Recover To

● Original Server ○ New Server

Restore Method

● Auto Deploy Cohesity Agent ○ Use Existing Cohesity Agent ○ Use VMware Tools

Username

Password

Recover to Original Path

Recovery Options

Overwrite Existing File/Folder	No
Preserve File/Folder Attributes	Yes
Continue on Error	Yes
Cluster Interface	Auto Select
Task Name	Recover_Files_Oct_17_2024_8_53_PM

Figure 5-23 New recovery summary page

The **Recover To** field allows you select either the **Original Server** or a **New Server** as the target for the file/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:

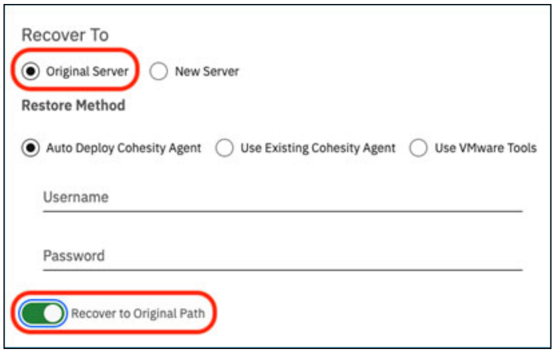

Recover To

● Original Server ○ New Server

Restore Method

● Auto Deploy Cohesity Agent ○ Use Existing Cohesity Agent ○ Use VMware Tools

Username

Password

Recover to Original Path

Figure 5-24 Recovery to original location settings panel

To recover the item(s) to an alternate path on the **Original Server**, flip the slider **Off** and enter the desired path for restoration.

To recover the item(s) to a **New Server**, select the **New Server** radio button, then use the **Source** and **Target** drop-down menus to select the destination ESX host or vCenter server and **Target** VM. Optionally, the **Source** drop-down menu allows for the registration of a new **Source** server "on the fly." In our example, we select the already-registered vCenter server **sts-pok-vcenter-1.www.pbm.ihost.com** as our **Source**, then select virtual machine **sts-pok-rhel8-10-8** as the **Target** for restore, as shown in Figure 5-25 on page 83:

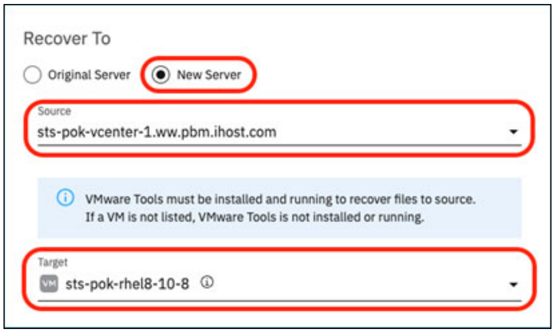

Recover To

○ Original Server ● New Server

Source: sts-pok-vcenter-1.ww.pbm.ihost.com

VMware Tools must be installed and running to recover files to source.
If a VM is not listed, VMware Tools is not installed or running.

Target: sts-pok-rhel8-10-8

Figure 5-25 Recovery to new server settings panel

> **Note:** The use of the term **Server** on this panel, whether **Original** or **New**, can be a bit misleading. Here **Server** denotes the **Virtual Server** (VM) to which you want to recover the item(s).
>
> The use of the term **Server** here, whether **Original** or **New**, can be a bit misleading. Here **Server** denotes the **Virtual Server** (VM) to which you want to recover the item(s).
>
> Additionally, **Source** is used here in the context of registered Data Protect protection sources. When **New Server** is selected, the **Source** drop-down menu loads the registered VMware data protection sources in this DMS instance (vCenter Servers and/or ESX hosts) and the **Target** drop-down menu dynamically loads based on the selected Source, allowing you to select the target VM to which you want to recover the item(s).

Regardless of whether the item(s) are recovered to the **Original Server** or a **New Server**, we must select the **Restore Method.** The available methods are:

- **Auto Deploy Cohesity Agent**: Data Protect will deploy the Cohesity agent using the supplied credentials, then recover the item(s) to the **Target** VM leveraging the agent.
- **Use Existing Cohesity Agent**: Data Protect will leverage the existing Cohesity agent on the **Target** VM to restore the items.

> **Note:** Recovering VMware files and/or folders with the IBM Storage Defender Data Protect agent requires that VMware tools be installed and running on the **Target** VM.

- **Use VMware Tools**: Data Protect will leverage VMware Tools to perform the recovery of the item(s). VMware Tools must be installed and running on the **Target** VM.

10. In our example, we know that VMware Tools are already installed on the original, source VM, so in Figure 5-26 on page 84 we select **Original Server**, Use **VMware Tools**, enter user credential to access the **Target** VM, and select **Recover to Original Path**:

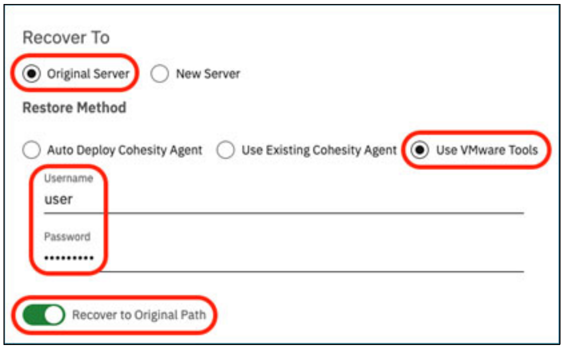

Recover To

● Original Server ○ New Server

Restore Method

○ Auto Deploy Cohesity Agent ○ Use Existing Cohesity Agent ● Use VMware Tools

Username: user

Password: •••••••••

Recover to Original Path

Figure 5-26 Recovery to original location settings panel

Data Protect offers the following additional Recovery Options for VMware file and folder recovery:

- **Overwrite Existing File/Folder**: Default value, **No**; If Yes, existing files may be overwritten as part of the recovery job.
- **Preserve File/Folder Attributes**: Default value, **Yes**; By default, Data Protect will preserve the ACLs, permissions, and timestamps for all recovered item(s). If **No**, the ACLs and permissions are not preserved.

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

- **Continue on Error**: Default value, Yes; The recovery job will upon the first error encountered. If set to No, the recovery will continue when errors are encountered.
- **Cluster Interface**: Default value is **Auto Select** - the Data Protect cluster will automatically select the correct interface group to use for the recovery. To manually select the desired **Interface Group**, click the **Cluster Interface** row, disable the **Auto Select** slider and use the **Interface Group** drop-down menu to select the desired **Interface Group** to use for the recovery job.
- **Task Name**: Optionally, change the default name of the recovery job.

11.To initiate the recovery job, click the Recover button. You will be redirected to the **Data Protection > Recoveries** page, where you will see the restore job running, as shown in Figure 5-27:

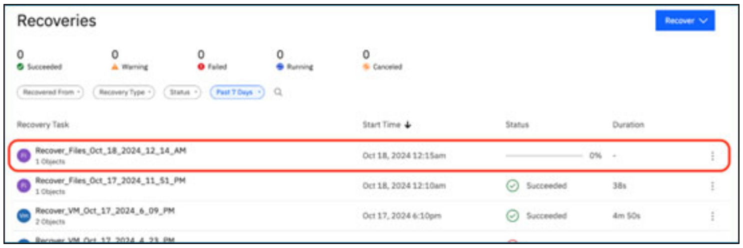

Recoveries

Recovery Task	Start Time	Status	Duration
Recover_Files_Oct_18_2024_12_14_AM 1 Objects	Oct 18, 2024 12:15am	0%	-
Recover_Files_Oct_17_2024_11_51_PM 1 Objects	Oct 18, 2024 12:10am	Succeeded	38s
Recover_VM_Oct_17_2024_6_09_PM 2 Objects	Oct 17, 2024 6:10pm	Succeeded	4m 50s
Recover_VM_Oct_17_2024_4_23_PM			

Figure 5-27 Recoveries status dashboard

12.Click on the recovery job to load its details, as shown, in Figure 5-28 on page 85:

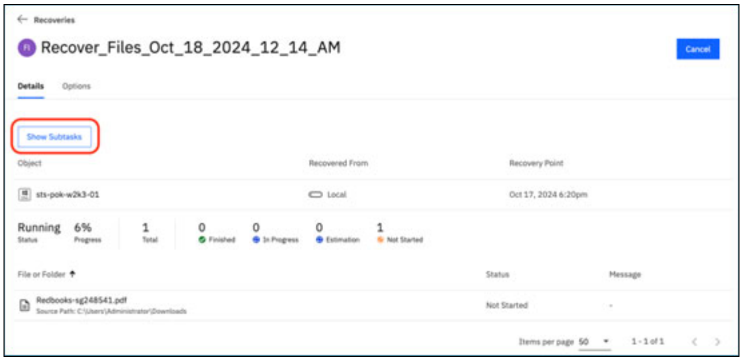

Recover_Files_Oct_18_2024_12_14_AM

Object	Recovered From	Recovery Point
sts-pok-w2k3-01	Local	Oct 17, 2024 6:20pm

Running Status · 6% Progress · 1 Total · 0 Finished · 0 In Progress · 0 Estimation · 1 Not Started

File or Folder	Status	Message
Redbooks-sg248541.pdf Source Path: C:\Users\Administrator\Downloads	Not Started	-

Items per page 50 · 1 - 1 of 1

Figure 5-28 File recovery job details panel

13. Click the **Show Subtasks** button to view the log of the recovery job:

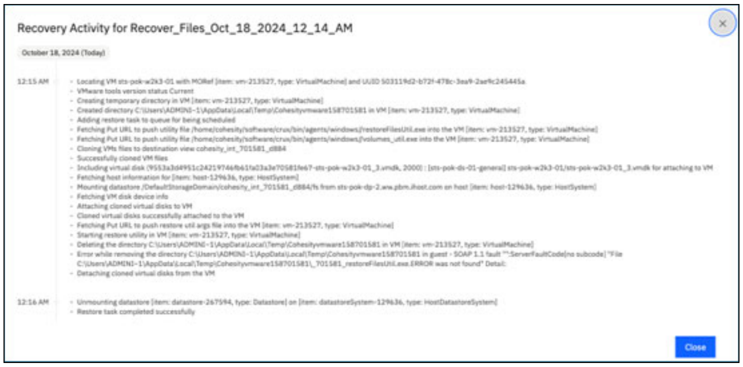

Recovery Activity for Recover_Files_Oct_18_2024_12_14_AM

October 18, 2024 (Today)

```
12:15 AM  - Locating VM sts-pok-w2k3-01 with MORef [item: vm-213527, type: VirtualMachine] and UUID 503119d2-b72f-478c-3ea9-2ae9c245445a
          - VMware tools version status Current
          - Creating temporary directory in VM [item: vm-213527, type: VirtualMachine]
          - Created directory C:\Users\ADMINI~1\AppData\Local\Temp\Cohesityvmware158701581 in VM [item: vm-213527, type: VirtualMachine]
          - Adding restore task to queue for being scheduled
          - Fetching Put URL to push utility file /home/cohesity/software/crux/bin/agents/windows//restoreFilesUtil.exe into the VM [item: vm-213527, type: VirtualMachine]
          - Fetching Put URL to push utility file /home/cohesity/software/crux/bin/agents/windows//volumes_util.exe into the VM [item: vm-213527, type: VirtualMachine]
          - Cloning VMs files to destination view cohesity_int_701581_d884
          - Successfully cloned VM files
          - Including virtual disk (9553a3d4951c24219746fb61fa03a3e70581fe67-sts-pok-w2k3-01_3.vmdk, 2000) : [sts-pok-ds-01-general] sts-pok-w2k3-01/sts-pok-w2k3-01_3.vmdk for attaching to VM
          - Fetching host information for [item: host-129636, type: HostSystem]
          - Mounting datastore /DefaultStorageDomain/cohesity_int_701581_d884/fs from sts-pok-dp-2.ww.pbm.ihost.com on host [item: host-129636, type: HostSystem]
          - Fetching VM disk device info
          - Attaching cloned virtual disks to VM
          - Cloned virtual disks successfully attached to the VM
          - Fetching Put URL to push restore util args file into the VM [item: vm-213527, type: VirtualMachine]
          - Starting restore utility in VM [item: vm-213527, type: VirtualMachine]
          - Deleting the directory C:\Users\ADMINI~1\AppData\Local\Temp\Cohesityvmware158701581 in VM [item: vm-213527, type: VirtualMachine]
          - Error while removing the directory C:\Users\ADMINI~1\AppData\Local\Temp\Cohesityvmware158701581 in guest - SOAP 1.1 fault "":ServerFaultCode[no subcode] "File
            C:\Users\ADMINI~1\AppData\Local\Temp\Cohesityvmware158701581\_701581_restoreFilesUtil.exe.ERROR was not found" Detail:
          - Detaching cloned virtual disks from the VM

12:16 AM  - Unmounting datastore [item: datastore-267594, type: Datastore] on [item: datastoreSystem-129636, type: HostDatastoreSystem]
          - Restore task completed successfully
```

Figure 5-29 File recovery log details

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

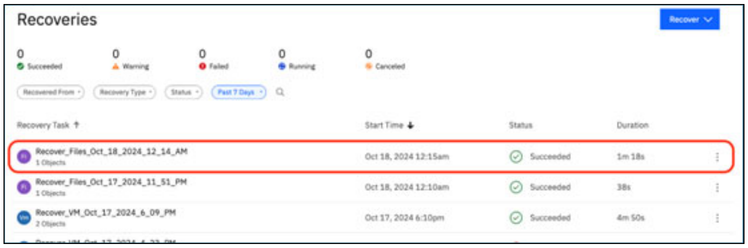

Recoveries

0 Succeeded · 0 Warning · 0 Failed · 0 Running · 0 Canceled

Recovery Task	Start Time	Status	Duration
Recover_Files_Oct_18_2024_12_14_AM 1 Objects	Oct 18, 2024 12:15am	Succeeded	1m 18s
Recover_Files_Oct_17_2024_11_51_PM 1 Objects	Oct 18, 2024 12:10am	Succeeded	38s
Recover_VM_Oct_17_2024_6_09_PM 2 Objects	Oct 17, 2024 6:10pm	Succeeded	4m 50s

Figure 5-30 Recovery job status panel

5.3 Virtual Disk Recovery

Defender Data Protect supports the recovery of VMware VM virtual disks from snapshot backups that it has created. For more information on creating VMware VM backups with Data Protect, see **4.2.3, Configuration of a Data Protect Protection Group**. Recovery can be to the original VM or to a different VM registered within the same vCenter.

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

1. Data Protect creates a writable clone of the selected VMDK file from the selected Recovery Point and mounts it to the target ESX host(s) as a temporary NFS datastore from the Data Protect cluster.
2. Data Protect attaches the new VMDK virtual disk(s) to the target VM and initiates a storage vMotion of the new VMDK from the temporary datastore to the target datastore.
3. The temporary datastore is automatically removed upon vMotion completion.

> **Note:** After virtual disk recovery competes, it may be necessary to perform additional, operating system-level actions to make the recovered disks online and available.

5.3.1 Recovering Virtual Disks

Follow these steps to recover VMware virtual disk with Defender Data Protect:

1. Navigate to **Data Protection > Recoveries**, as shown in Figure 5-31:

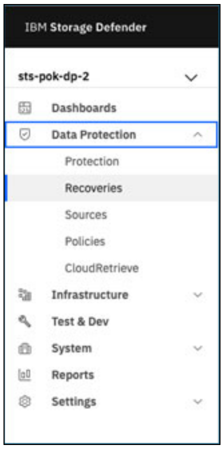

Figure 5-31 Storage Defender Data Protection menu

2. Then use the **Recover** menu on the top right of the **Recoveries** page to select **Recover > Virtual Machines > Virtual Disks** as shown in Figure 5-32:

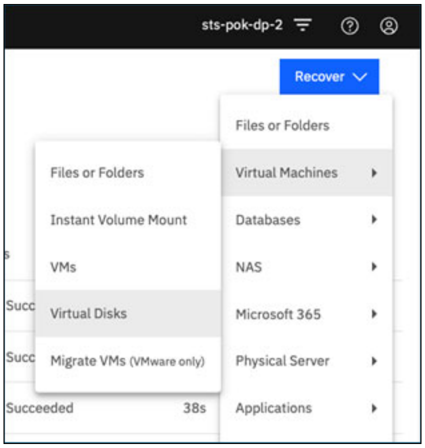

Figure 5-32 Virtual disks recovery option

3. In the **New Recovery** window, as shown in Figure 5-33, enter a wildcard search term to search for the virtual machine whose disks you want to recover, then select the desired

VM. Note that Data Protect defaults to the most recent **Recovery Point**. Again, optionally use the pencil icon to select an earlier **Recovery Point**. In our example, we leave the default **Recovery Point** selected and click the **Next: Recover Options** button:

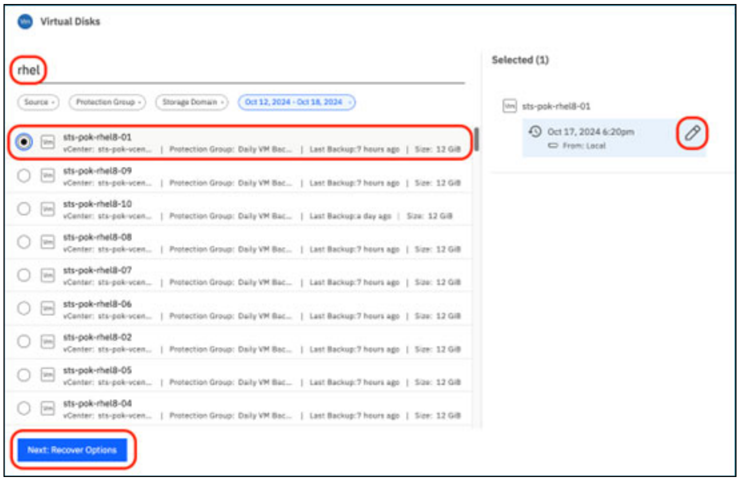

Virtual Disks

rhel

Source ▾ | Protection Group ▾ | Storage Domain ▾ | Oct 12, 2024 - Oct 18, 2024 ▾

● sts-pok-rhel8-01
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:7 hours ago | Size: 12 GiB

○ sts-pok-rhel8-09
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:7 hours ago | Size: 12 GiB

○ sts-pok-rhel8-10
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:a day ago | Size: 12 GiB

○ sts-pok-rhel8-08
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:7 hours ago | Size: 12 GiB

○ sts-pok-rhel8-07
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:7 hours ago | Size: 12 GiB

○ sts-pok-rhel8-06
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:7 hours ago | Size: 12 GiB

○ sts-pok-rhel8-02
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:7 hours ago | Size: 12 GiB

○ sts-pok-rhel8-05
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:7 hours ago | Size: 12 GiB

○ sts-pok-rhel8-04
vCenter: sts-pok-vcen... | Protection Group: Daily VM Bac... | Last Backup:7 hours ago | Size: 12 GiB

Next: Recover Options

Selected (1)

sts-pok-rhel8-01
Oct 17, 2024 6:20pm
From: Local

Figure 5-33 Virtual disk recovery search panel

The **New Recovery** screen (Figure 5-34) will now show a summary of the selected source virtual machine and associated **Recovery Point**. The **Recover To** option defaults to **Original Location**. Select **New Location** to restore the virtual disk(s) to a different VM:

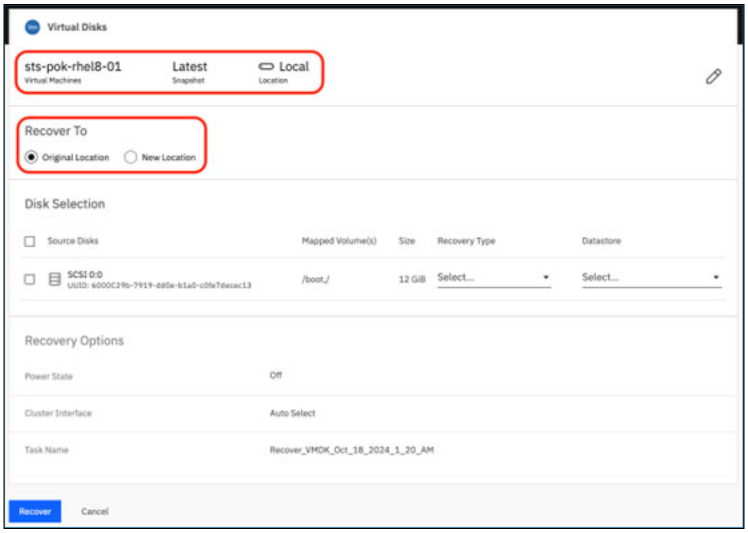

Virtual Disks

sts-pok-rhel8-01 Virtual Machines | Latest Snapshot | Local Location

Recover To
● Original Location ○ New Location

Disk Selection

□	Source Disks	Mapped Volume(s)	Size	Recovery Type	Datastore
□	SCSI 0:0 UUID: 6000C29b-7919-dd0e-b1a0-c0fe7decec13	/boot,/	12 GiB	Select... ▾	Select... ▾

Recovery Options

Power State	Off
Cluster Interface	Auto Select
Task Name	Recover_VMDK_Oct_18_2024_1_20_AM

Recover Cancel

Figure 5-34 Virtual disks recovery options panel

4. Next, check the box(es) for the disk(s) that you want to recover (Figure 5-35 on page 88), and for each disk, select the recovery type:

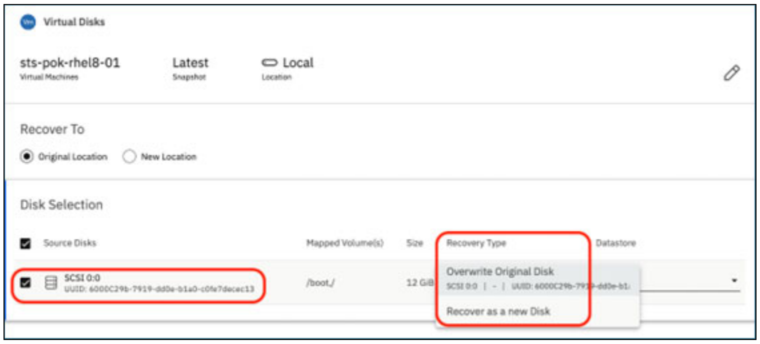

Virtual Disks

sts-pok-rhel8-01 Virtual Machines | Latest Snapshot | Local Location

Recover To

● Original Location ○ New Location

Disk Selection

☑ Source Disks	Mapped Volume(s)	Size	Recovery Type	Datastore
☑ SCSI 0:0 UUID: 6000C29b-7919-dd0e-b1a0-c0fe7decec13	/boot./	12 GiB	Overwrite Original Disk SCSI 0:0 | - | UUID: 6000C29b-7919-dd0e-b1 Recover as a new Disk	

Figure 5-35 Virtual disks recovery panel disk selection

The available Recovery Types are:

- **Overwrite Original Disk**: The original disk will be overwritten in the original datastore with the recovered virtual disk. It is recommended that VM disk IO be quiesced prior to performing the disk recovery.
- **Recover as a new Disk**: The recovered virtual disk will be attached to the target VM as a new virtual disk. You must chose this option if the virtual disk has been deleted from the VM since the time of the backup.

If **Recover as a new Disk** is selected, you must select a target datastore for the new disk with the Datastore drop-down menu, as show in Figure 5-36:

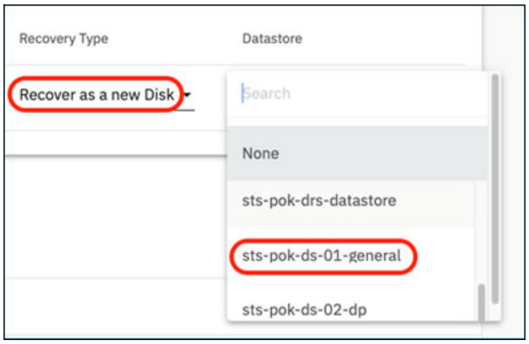

Recovery Type: Recover as a new Disk

Datastore

Figure 5-36 Virtual disk recovery panel ‘Recovery as a new disk option’ panel

> **Note:** If recovery fails with **Overwrite Original Disk** selected, the original disk will not be available. If this occurs, retry the recovery to restore the disk(s).
>
> For recoveries where one or more disks will be overwritten, Data Protect will automatically powered off the target VM prior to recovery. Optionally, the target VM can be powered on automatically after successful recovery.

5. In our example (Figure 5-37 on page 89), we select **Original Location** and choose **Recover as a new Disk**, then select **sts-pok-ds-01-general** as the target datastore:

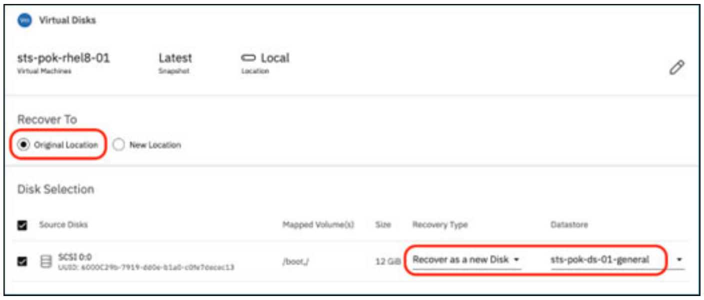

Virtual Disks

sts-pok-rhel8-01 Virtual Machines | Latest Snapshot | Local Location

Recover To

● Original Location ○ New Location

Disk Selection

☑	Source Disks	Mapped Volume(s)	Size	Recovery Type	Datastore
☑	SCSI 0:0 UUID: 6000C29b-7919-dd0e-b1a0-c0fe7decac13	/boot,/	12 GiB	Recover as a new Disk	sts-pok-ds-01-general

Figure 5-37 Virtual disk recovery options panel

Virtual Disk **Recovery Options** are:

- **Power State**: Default value is Off - VMs will not be powered off prior to restoration
- **Cluster Interface**: Default value is **Auto Select** - the Data Protect cluster will automatically select the correct interface group to use for the recovery. To manually select the desired **Interface Group**, click the **Cluster Interface** row, disable the **Auto Select** slider and use the **Interface Group** drop-down menu to select the desired **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**).** Note that the value changes from **Off** to a slider labeled **Power off VMs before Restore**:

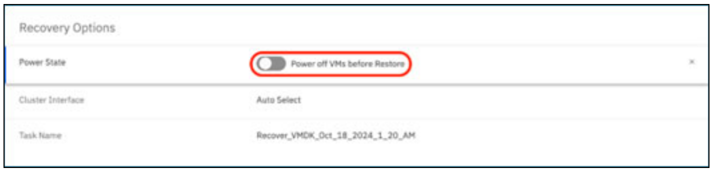

Recovery Options

Power State	Power off VMs before Restore
Cluster Interface	Auto Select
Task Name	Recover_VMDK_Oct_18_2024_1_20_AM

Figure 5-38 Recovery options panel power options

7. Flip the **Power off VMs before Restore** slider on (Figure 5-39). Note that there is now an option to **Power on VMs after Restore**:

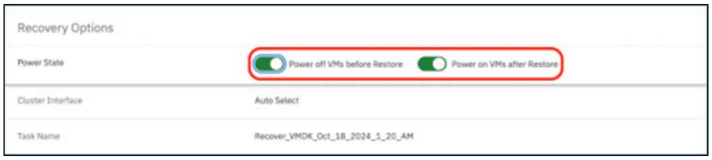

Recovery Options

Power State	Power off VMs before Restore Power on VMs after Restore
Cluster Interface	Auto Select
Task Name	Recover_VMDK_Oct_18_2024_1_20_AM

Figure 5-39 Recovery options panel 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 on page 90), we revert to the default **Power State** options and click **Recover**. The browser redirects to the **Data Protection > Recoveries** page, and we can see our virtual disk restore job running:

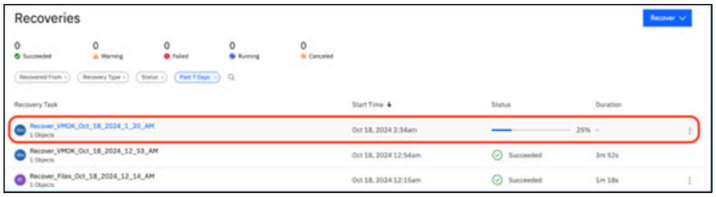

Recoveries

0 Succeeded 0 Warning 0 Failed 0 Running 0 Canceled

Recovery Task	Start Time	Status	Duration
Recover_VMDK_Oct_18_2024_1_20_AM 1 Objects	Oct 18, 2024 2:34am	25%	-
Recover_VMDK_Oct_18_2024_12_53_AM 1 Objects	Oct 18, 2024 12:54am	Succeeded	3m 52s
Recover_Files_Oct_18_2024_12_14_AM 1 Objects	Oct 18, 2024 12:15am	Succeeded	1m 18s

Figure 5-40 Recoveries status panel

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

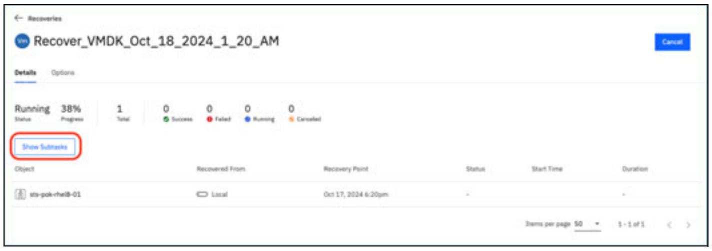

Recover_VMDK_Oct_18_2024_1_20_AM

Running (Status) 38% (Progress) 1 (Total) 0 Success 0 Failed 0 Running 0 Canceled

Object	Recovered From	Recovery Point	Status	Start Time	Duration
sts-pok-rhel8-01	Local	Oct 17, 2024 6:20pm	-		-

Items per page 50 1 - 1 of 1

Figure 5-41 Recovery details panel

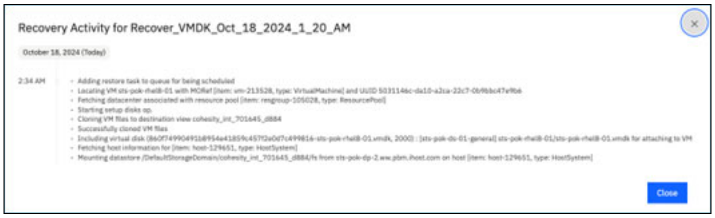

Recovery Activity for Recover_VMDK_Oct_18_2024_1_20_AM

October 18, 2024 (Today)

```
2:34 AM  - Adding restore task to queue for being scheduled
         - Locating VM sts-pok-rhel8-01 with MORef [item: vm-213528, type: VirtualMachine] and UUID 503114bc-da10-a2ca-22c7-0b9bbc47e9b6
         - Fetching datacenter associated with resource pool [item: resgroup-105028, type: ResourcePool]
         - Starting setup disks op.
         - Cloning VM files to destination view cohesity_int_701645_d884
         - Successfully cloned VM files
         - Including virtual disk (860f749904918954e41859c457f2a0d7c499816-sts-pok-rhel8-01.vmdk, 2000) : [sts-pok-ds-01-general] sts-pok-rhel8-01/sts-pok-rhel8-01.vmdk for attaching to VM
         - Fetching host information for [item: host-129651, type: HostSystem]
         - Mounting datastore /DefaultStorageDomain/cohesity_int_701645_d884/fs from sts-pok-dp-2.ww.pbm.ihost.com on host [item: host-129651, type: HostSystem]
```

Figure 5-42 Recovery job log details

10. Figure 5-43 shows our virtual disk recovery job has completed successfully:

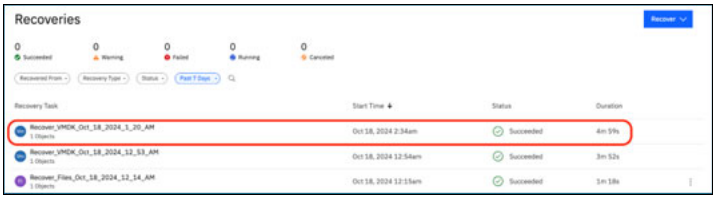

Recoveries

0 Succeeded 0 Warning 0 Failed 0 Running 0 Canceled

Recovery Task	Start Time	Status	Duration
Recover_VMDK_Oct_18_2024_1_20_AM 1 Objects	Oct 18, 2024 2:34am	Succeeded	4m 59s
Recover_VMDK_Oct_18_2024_12_53_AM 1 Objects	Oct 18, 2024 12:54am	Succeeded	3m 52s
Recover_Files_Oct_18_2024_12_14_AM 1 Objects	Oct 18, 2024 12:15am	Succeeded	1m 18s

Figure 5-43 Recoveries job summary panel

5.4 Instant Volume Mount

Instant Volume Mount (IVM) allows selected backup volume(s) to be made available to a target server target without the need for any automated movement of the backup data from the Data Protect system. Common use cases include granular recovery of Microsoft SQL, Exchange and Sharepoint databases, where the database/application admin can then perform the desired granular recovery operations manually or via scripting. Instant Volume Mount supports mounting to the original VM or a new VM.

Note: Instant Volume Mount is only available for backup volumes stored locally on the Data Protect cluster. In other words, **Instant Volume Mount** is not available from cloud/archive backup copies.

The default path of the recovered file or instantly mounted volume is *<Cluster_IP_address>/<Cloned view name>\<PATH>*. In order to access this location, both the target for the IVM and the Data Protect cluster must be in the same Active Directory domain.

5.4.1 Instant Volume Mounting a VMware virtual disk

Follow these steps to recover VMware virtual disk with Defender Data Protect:

1. Navigate to **Data Protection > Recoveries**, as shown in Figure 5-44:

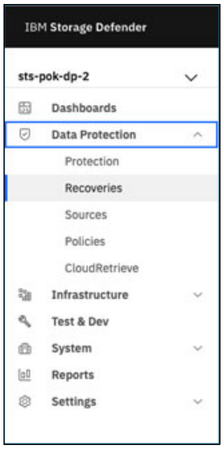

Figure 5-44 Storage Defender Data Protection menu

2. Then use the **Recover** menu on the top right of the **Recoveries** page to select **Recover > Virtual Machines > Instant Volume Mount** as shown in Figure 5-45:

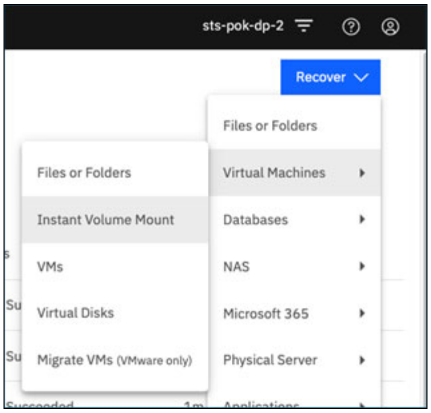

Figure 5-45 Instant Volume Mount option

3. In the **New Recovery** window, as shown in Figure 5-46, enter a wildcard search term to search for the virtual machine whose disks you want to instantly mount, then select the desired VM. Note that Data Protect defaults to the most recent **Recovery Point**. Again,

optionally use the pencil icon to select an earlier **Recovery Point**. In our example, we leave the default **Recovery Point** selected and click the **Next: Recover Options** button:

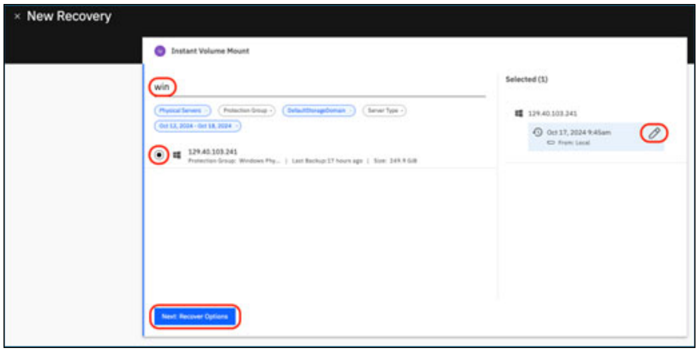

Figure 5-46 Recovery point options panel

4. The **New Recovery** screen will now show a summary of the selected source virtual machine and associated **Recovery Point**, as shown in Figure 5-47. The **Select Volumes** slider defaults to **Off**, allowing the selection of one or more specific volumes for instant mounting in the drop-down menu. Optionally, flip the slider to select all of the volumes that were backed up with this VM at the time of the selected **Recovery Point**:

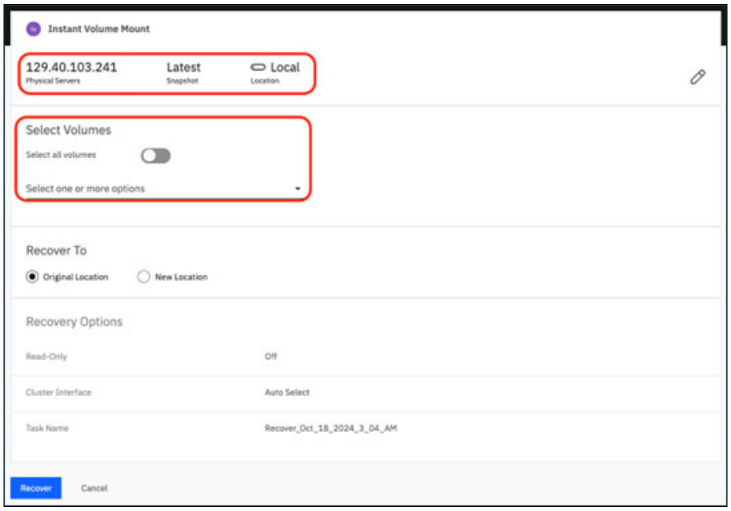

Read-Only	Off
Cluster Interface	Auto Select
Task Name	Recover_Oct_18_2024_3_04_AM

Figure 5-47 Instant Volume mount options panel

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

Additional **Recovery Options** for VMware **Instant Volume Mount** are:

- **Read-Only**: Default value is **Off** - the selected volumes will be mounted to the target VM in read-only mode.
- **Cluster Interface**: Default value is **Auto Select** - the Data Protect cluster will automatically select the correct interface group to use for the recovery. To manually select the desired **Interface Group**, click the **Cluster Interface** row, disable the **Auto Select** slider and use the **Interface Group** drop-down menu to select the desired **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**:

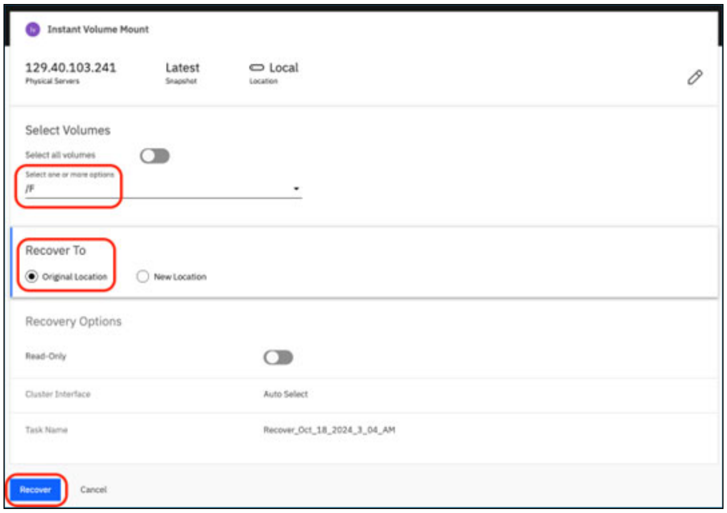

Instant Volume Mount

129.40.103.241 Physical Servers | Latest Snapshot | Local Location

Select Volumes

Select all volumes

Select one or more options: /F

Recover To

● Original Location ○ New Location

Recovery Options

Read-Only

Cluster Interface: Auto Select

Task Name: Recover_Oct_18_2024_3_04_AM

Figure 5-48 Instant volume mount recover to options

7. The Figure 5-49 below shows our **Instant Volume Mount** job has been successfully initiated:

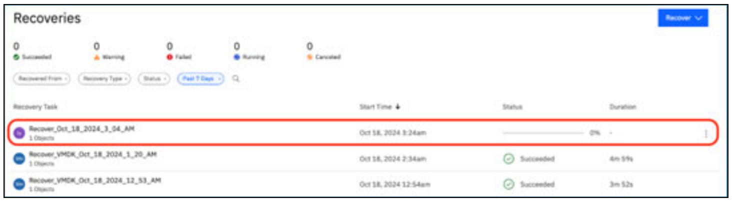

Recoveries

0 Succeeded | 0 Warning | 0 Failed | 0 Running | 0 Canceled

Recovery Task	Start Time	Status	Duration
Recover_Oct_18_2024_3_04_AM 1 Objects	Oct 18, 2024 3:24am	0%	-
Recover_VMDK_Oct_18_2024_1_20_AM 1 Objects	Oct 18, 2024 2:34am	Succeeded	4m 59s
Recover_VMDK_Oct_18_2024_12_53_AM 1 Objects	Oct 18, 2024 12:54am	Succeeded	3m 52s

Figure 5-49 Recoveries job log panel

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

5.4.2 Teardown of a Mounted Volume

Once the mounted volume is no longer needed, it can unmounted from the target host by selecting Teardown from the top-right menu, as shown below:

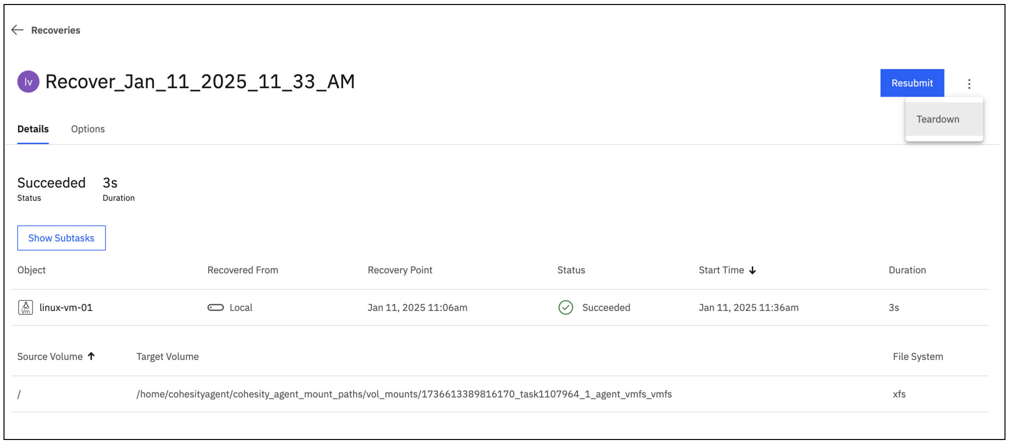

Recover_Jan_11_2025_11_33_AM

Succeeded
Status

3s
Duration

Object	Recovered From	Recovery Point	Status	Start Time ↓	Duration
linux-vm-01	Local	Jan 11, 2025 11:06am	✓ Succeeded	Jan 11, 2025 11:36am	3s

Source Volume ↑	Target Volume	File System
/	/home/cohesityagent/cohesity_agent_mount_paths/vol_mounts/1736613389816170_task1107964_1_agent_vmfs_vmfs	xfs

Figure 5-50 Recovery panel Teardown option selection

Then click to Teardown to confirm that you want to unmount the volume:

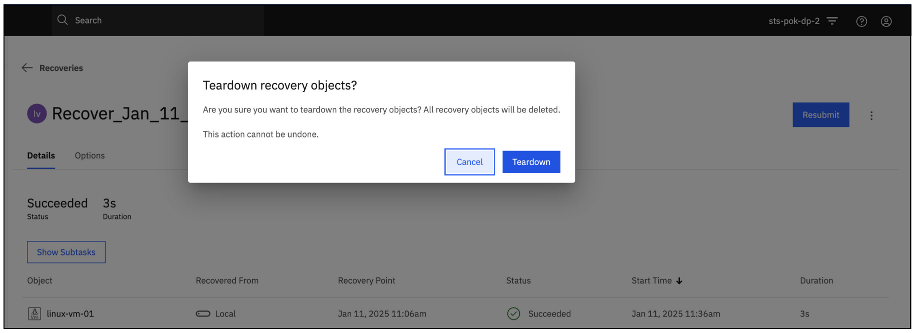

Teardown recovery objects?

Are you sure you want to teardown the recovery objects? All recovery objects will be deleted.

This action cannot be undone.

Succeeded
Status

3s
Duration

Object	Recovered From	Recovery Point	Status	Start Time ↓	Duration
linux-vm-01	Local	Jan 11, 2025 11:06am	✓ Succeeded	Jan 11, 2025 11:36am	3s

Figure 5-51 Teardown recovery object confirmation dialog

Integrating IBM Defender Data Protection with IBM Storage Protect

In this chapter, we demonstrate how to configure IBM Storage Defender Data Protect to send archives to an external S3 target using the S3 agent. The target in this example is a Storage Protect server. This type of configuration provides a Glacier-like S3 to tape archive for longer term retention use cases. We will show use cases, examples how to setup a DP Policy and a protection group along with an example data retrieval.

This chapter contains the following sections:

- 6.1, “Data Protection Overview” on page 96
- 6.2, “Deployment Exercise Overview” on page 97
- 6.3, “Create an S3 Bucket on the IBM Storage Protect Server” on page 97
- 6.4, “Capture Access Credentials on the SP server” on page 101
- 6.5, “Register the S3 Bucket as External Target on Data Protect Cluster” on page 101
- 6.6, “Create Polices to Archive Data to The S3 Bucket” on page 104
- 6.7, “Monitoring the Protection status” on page 109
- 6.8, “Retrieve from the S3 External Storage” on page 111
- 6.9, “Additional Process Monitoring Examples” on page 115

6.1 Data Protection Overview

IBM Storage Defender Data Protect (DP) provides support for sending data to external S3 targets. These S3 external target resources can be located in a public cloud, private cloud or on-premises S3 capable storage devices. For example an S3 compatible object storage on the cloud or S3 object storage on-prem (ICOS), allowing for configurations which fit a wide range of customer needs and requirements.

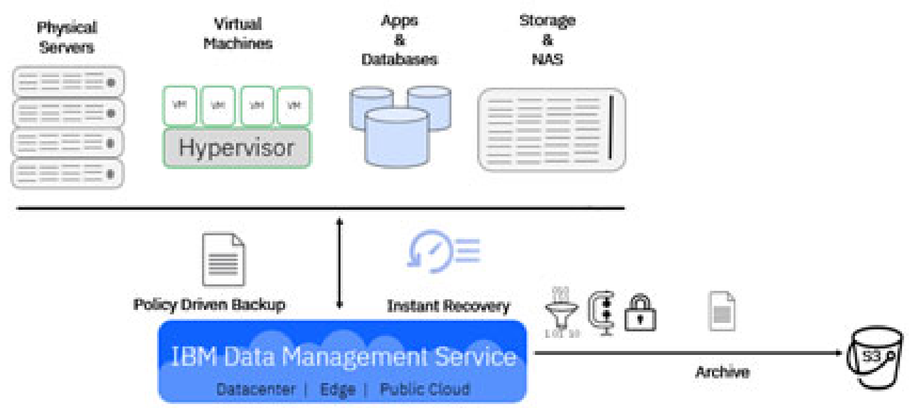

Figure 6-1 Workload protection with Archive to S3 Storage example

> **Note:** DP archive to Storage Protect (SP) tape currently is intended for weekly or monthly archives. Each archive is a full copy of the protection group from DP to SP. Incremental functionality for archives is not yet available.

S3 Offload Use Cases and Scenarios

- The IBM Storage Protect Server S3 agent can be set up for use by other components in the IBM Storage Defender stack. For example, it is possible to use the SP S3 Agent to provide an object storage bucket for Data Protect and IBM Storage Fusion.
- A requirement is in place to 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 copy/s 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 DP cluster with options to replicate the data so there are two copies for 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.
- A great use for existing tape libraries in the environment. When an archive copy is sent to the SP S3 bucket, the data is cached on the SP S3 cold data cache pool. It is then migrated to physical tape volumes that can be managed onsite or offsite for longer term retention needs. This cold data cache pool is helpful in reducing the number of tape drives required to meet the desired data throughput. Additionally, this disk cache will be used for restores.

Compatibility Details

A complete list of S3 storage compatibility for DP is available in the DP Support pages. At the writing of this book, the current release of DP is v7.1 and SP is 8.1.20. To find the latest information on IBM Storage Defender Data Protect see the following link:

https://www.ibm.com/docs/en/storage-defender/base?topic=storage-defender-data-protect

6.2 Deployment Exercise Overview

In this section, we take a closer look at how the IBM Storage Protect S3 Agent integration is set up for use with DP.

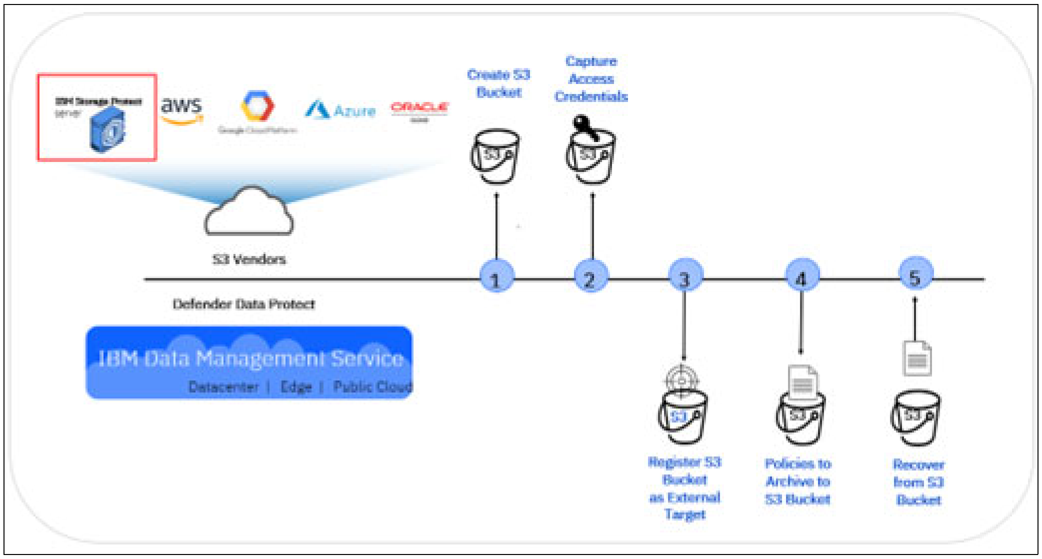
IBM Storage Protect server
aws
Google Cloud Platform
Azure
ORACLE
Create S3 Bucket
Capture Access Credentials
S3 Vendors
Defender Data Protect
1
2
3
4
5
IBM Data Management Service
Datacenter | Edge | Public Cloud
Register S3 Bucket as External Target
Policies to Archive to S3 Bucket
Recover from S3 Bucket

Figure 6-2 Leverage S3 compatible storage with IBM Storage Defender Data Protect

This diagram (Figure 6-2) shows S3 Storage options on the top left, with the IBM Data Management Server (DMS) at the bottom. In the middle is a timeline which shows the steps required to complete this integration.

These details for setup are covered in the sections below:

1. Create S3 Bucket on the SP Server
2. Capture the access credentials on the SP server
3. Register the S3 bucket as external target on DP cluster
4. Archive to the S3 bucket
5. Recover from the S3 bucket

6.3 Create an S3 Bucket on the IBM Storage Protect Server

These steps can be performed on a server with either IBM Storage Protect or IBM Defender Storage Protect licensing. Storage Protect servers need to be version 8.1.20 or higher.

How to Set up the SP server to receive DP archives

The command below will create 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 will be stored. In the command below the tape pool is defined as *coldtapestg*.

> **Note:** When creating a cache, ensure that it is big enough for the largest retrieve that may need to be done, as well as additional space for any additional incoming archives which may need to be processed.
>
> This Storage pool is used as a temporary cache location for the Data Protect (DP) archives before migrating them to the tape pool.
>
> For recovery, the IBM 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. Because of this, the size of DP groups and size of the VM’s protected will need to be taken into consideration when determining the amount storage pool space required for successful archive and retrieve operations.

Example 6-1 Define a coldcache storage pool on the SP 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)
```

Create the S3 Storage Agent:

Next, use the `define server` command to create the S3 agent. In the example below the server is defined as ‘s3-dp’.

Example 6-2 Define the Storage Agent which will connect to the S3 enabled storage

```
TSMSRC> def server s3-dp hla=<ip-address-of-SP-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 running the define server command, run the command in the resulting ANR4601I message as root from the host command line.

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 SP Operations Center you can see the results of this of this command execution, as well as view/copy the certificate information from the Servers/Object Agent drill down. (Figure 6-3 on page 99)

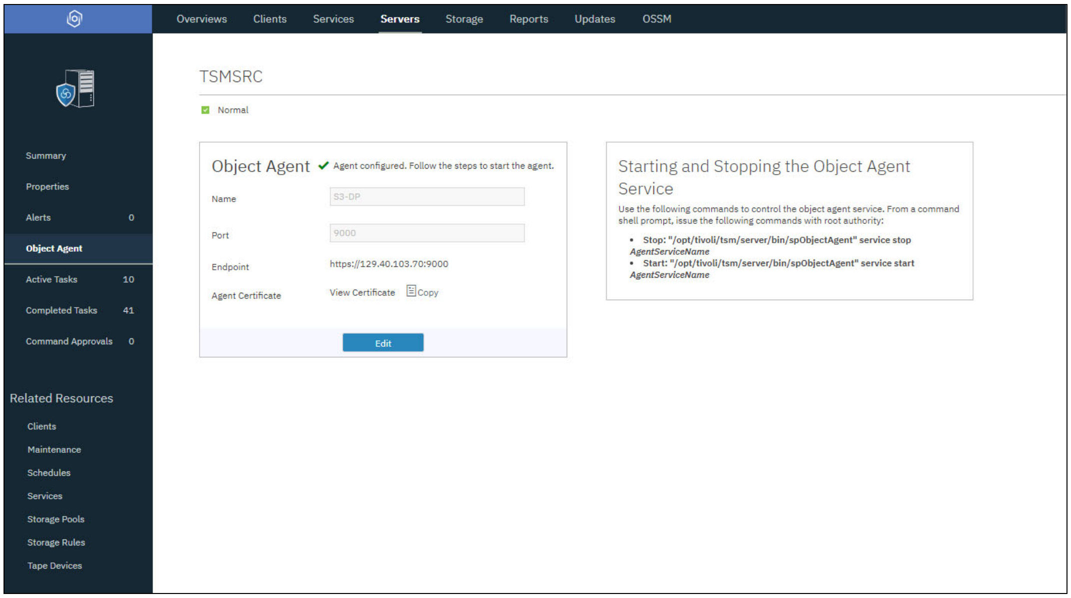

TSMSRC

☑ Normal

Object Agent ✓ Agent configured. Follow the steps to start the agent.

Name: S3-DP

Port: 9000

Endpoint: https://129.40.103.70:9000

Agent Certificate

Starting and Stopping the Object Agent Service

Use the following commands to control the object agent service. From a command shell prompt, issue the following commands with root authority:

- **Stop: "/opt/tivoli/tsm/server/bin/spObjectAgent" service stop *AgentServiceName***
- **Start: "/opt/tivoli/tsm/server/bin/spObjectAgent" service start *AgentServiceName***

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

From the Storage Protect command line, define the object domain.

In the example below, the standard pool ‘s3-fusion’ is a container storage pool used for Fusion backups and the coldcache pool is the name of the pool we created in the first Example 6-1 on page 98. The archives from Data Protect will first land in this coldcache pool then auto migrate to the nextstgpool which is set to ‘coldtapestg’.

Example 6-4 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.
```

Register the object client node on your SP server.

In the following example the `register node` command to register an node on the SP server to store the archive data. This command will use the same domain name that was used in the previous Example 6-4 on page 99.

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.
```

Create the S3 bucket used by Data Protect to send archives to Storage Protect.

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

https://dl.min.io/client/mc/release/

As root, from the directory where the .rpm file is located, run the ‘`mcli alias set alias-name end-point access-key secret-key --insecure`’ command:

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`. Please 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.
```

Create the bucket.

In the following example s3-dp is the alias name from step Example 6-6 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 Capture Access Credentials on the SP server

The following information will be needed to register the SP bucket that was created to a new DP External Target.

Have the following items available to continue with the next configuration steps:

- Bucket Name
- Access Key ID
- Secret Access Key
- Endpoint: the IP address of your SP server
- Port number: the LLA used on the define server in Example 6-2 on page 98

6.5 Register the S3 Bucket as External Target on Data Protect Cluster

> **Note:** Data Protect cluster needs to be version 7.1 or higher.
>
> In DP, select AWS Signature version v4. For the External Target Name, the SP Node registered in step Example 6-5 on page 100 will be used. This will be the SP node under which the archives will be stored.

To register an External Target with your cluster, log in to the IBM Defender Dashboard and perform the following actions:

1. Select the Data Management to launch the Data Management Service

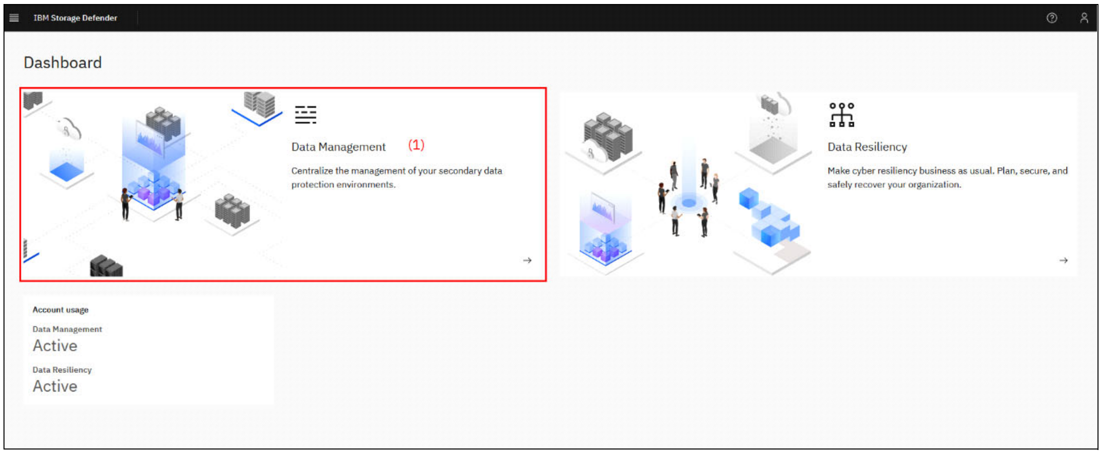

Centralize the management of your secondary data protection environments.

Make cyber resiliency business as usual. Plan, secure, and safely recover your organization.

Account usage

Data Management
Active

Data Resiliency
Active

Figure 6-4 Log 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 on page 102)

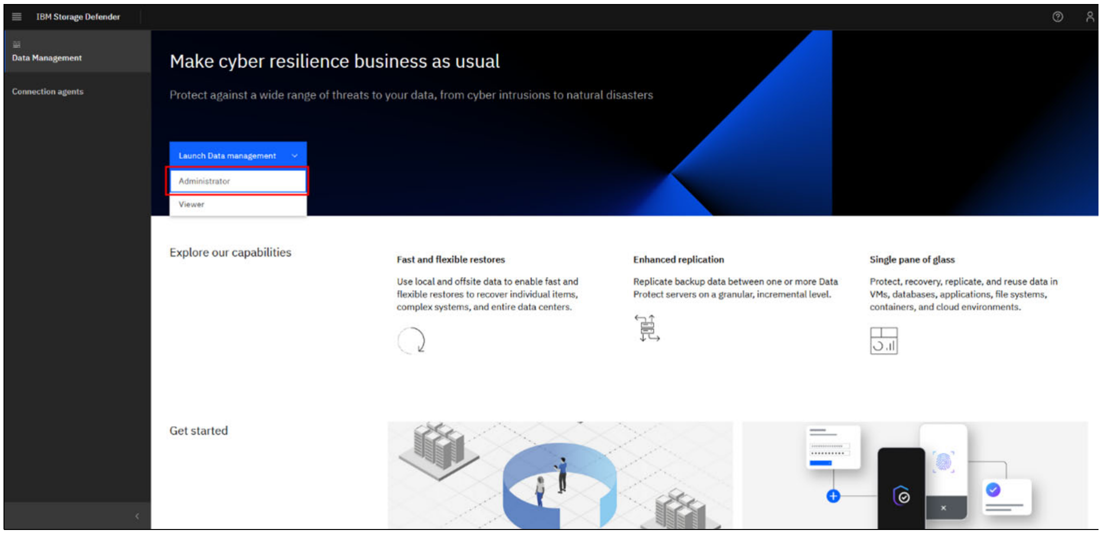

Make cyber resilience business as usual

Protect against a wide range of threats to your data, from cyber intrusions to natural disasters

Explore our capabilities

Fast and flexible restores

Use local and offsite data to enable fast and flexible restores to recover individual items, complex systems, and entire data centers.

Enhanced replication

Replicate backup data between one or more Data Protect servers on a granular, incremental level.

Single pane of glass

Protect, recovery, replicate, and reuse data in VMs, databases, applications, file systems, containers, and cloud environments.

Get started

Figure 6-5 Log in to IBM Defender Data Management Service, step 2

3. Select the cluster resource, in the drop down for the cluster select:
   a. Expand the Infrastructure option
   b. Select External Targets
   c. Click on the Add External Target button

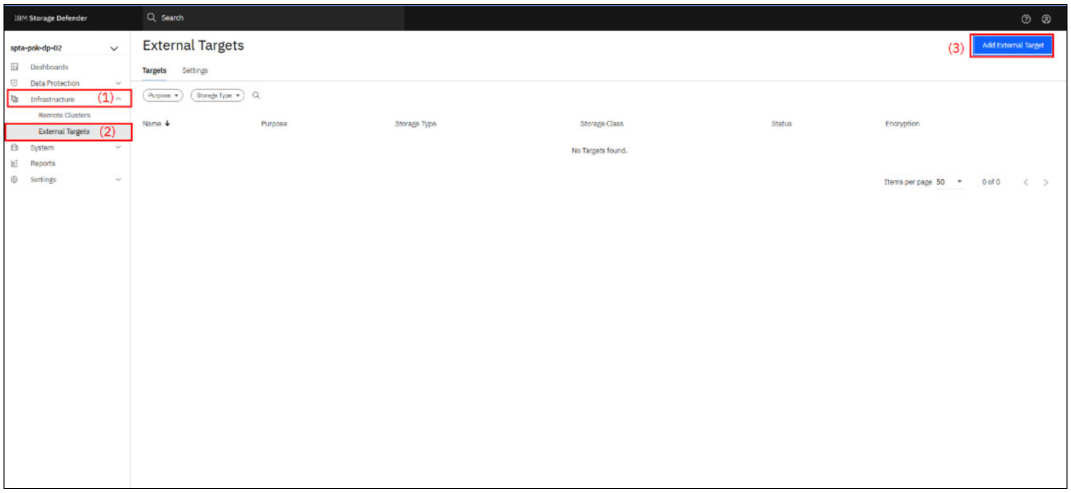

External Targets

Name ↓	Purpose	Storage Type	Storage Class	Status	Encryption
			No Targets found.		

Figure 6-6 Register an S3 External Target in Defender Data Protect, steps 1-3

In the Register External Target form:

4. Select Purpose (Archival).
5. Select Type (S3 Compatible)
6. Select Storage Class (Tape Based for Storage Protect S3 Agent to tape)
7. Enter Bucket Name.
8. Enter Access Key ID
9. 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 then, scroll down for additional options to continue the External Target Registration, do not click register at this time.

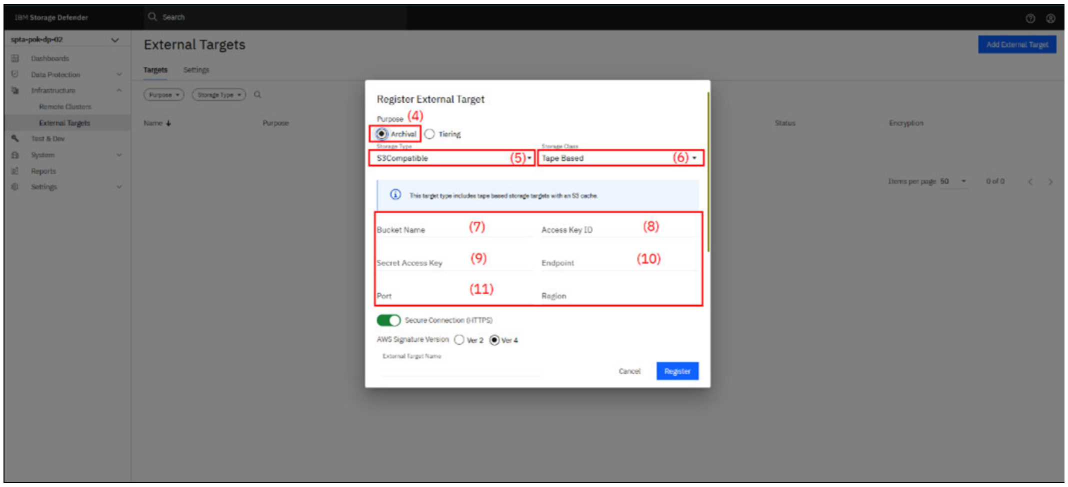

Register External Target

Purpose
● Archival ○ Tiering

Storage Type: S3Compatible

Storage Class: Tape Based

This target type includes tape based storage targets with an S3 cache.

Bucket Name

Access Key ID

Secret Access Key

Endpoint

Port

Region

Secure Connection (HTTPS)

AWS Signature Version ○ Ver 2 ● Ver 4

External Target Name

Figure 6-7 Register an S3 External Target in Defender Data Protect, 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 exposed via 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 to be the Storage Protect node name.
14. If desired, set additional options for Encryption. See the note about the Key Management Service type in the GUI panel. If you want to enable manual key management for extra security, turn it on 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 will still be able to recover data to its original cluster, but you will not be able to retrieve it onto a new cluster (for example, in a disaster-recovery scenario).

15. Edit the setting for Compression if desired
16. Enable and configure the Bandwidth Throttling settings if desired. 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.

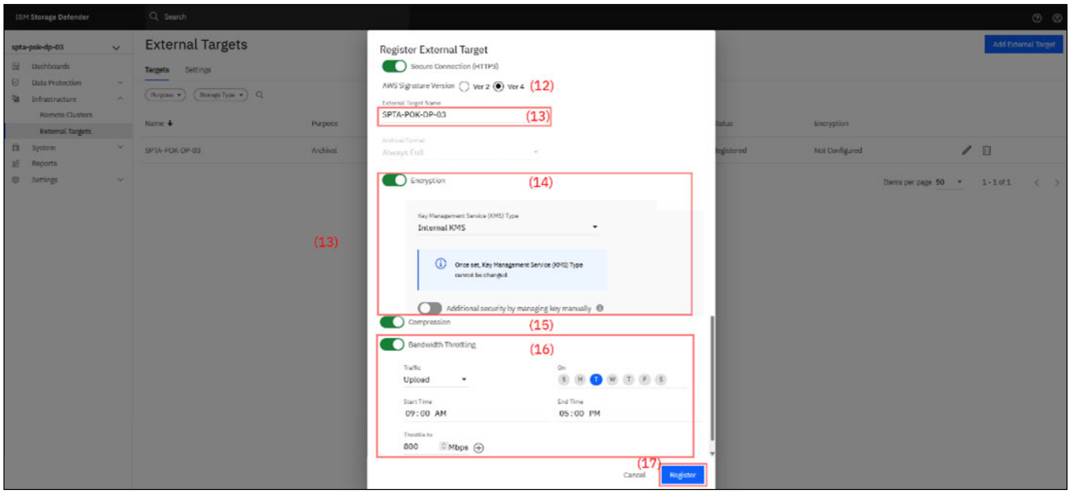

Name	Purpose	Status	Encryption
SPTA-POK-DP-03	Archival	Registered	Not Configured

Register External Target

Secure Connection (HTTPS): on

AWS Signature Version: ○ Ver 2 ● Ver 4 (12)

External Target Name: SPTA-POK-DP-03 (13)

Archival Format: Always Full

Encryption: on (14)

Key Management Service (KMS) Type: Internal KMS

Once set, Key Management Service (KMS) Type cannot be changed

Additional security by managing key manually: off

Compression: on (15)

Bandwidth Throttling: on (16)

Traffic: Upload

On: S M **T** W T F S

Start Time: 09:00 AM

End Time: 05:00 PM

Throttle to: 800 Mbps

(17) Cancel | Register

Figure 6-8 Register an S3 External Target in Defender Data Protect, steps 12-17

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

6.6 Create Polices to Archive Data to The S3 Bucket

This section is split into two sub-sections as there are multiple policies that need to be defined in Data Protect to correctly configure the system for this use case.

6.6.1 Create a Protection Policy for Archiving to the S3 Bucket

In IBM Data Protect, we will need to associate the external S3 target name from step 3 in the example above to a Data Protection policy in IBM Defender Data Protect. (see: 3 on page 102)

Create a Data Protection Policy by selecting the following GUI items:

1. Expand the Data Protection section
2. Select Policies
3. Click the button in the top right labeled ‘Create Policy’ to begin creating the new policy

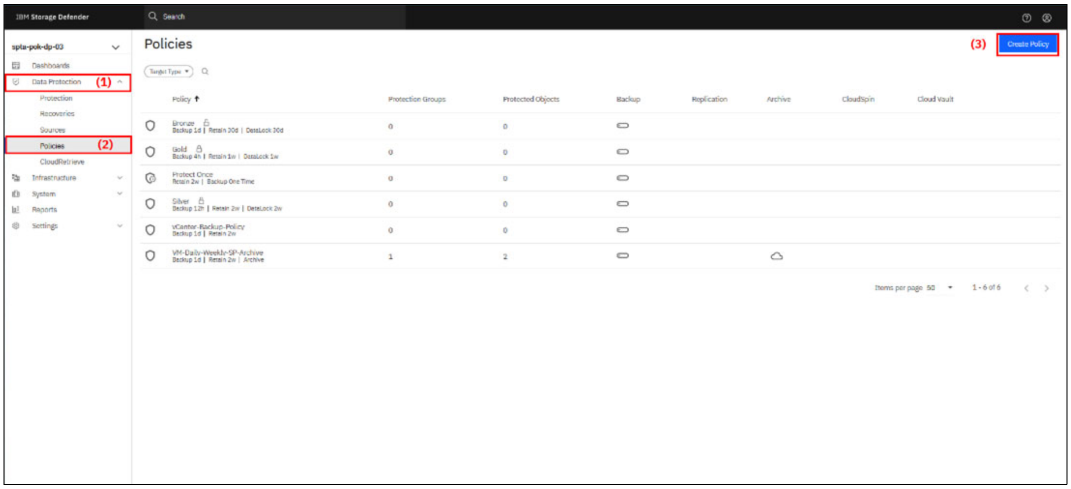

Policy	Protection Groups	Protected Objects	Backup	Replication	Archive	CloudSpin	Cloud Vault
Bronze Backup 1d | Retain 30d | DataLock 30d	0	0					
Gold Backup 4h | Retain 1w | DataLock 1w	0	0					
Protect Once Retain 2w | Backup One Time	0	0					
Silver Backup 12h | Retain 2w | DataLock 2w	0	0					
vCenter-Backup-Policy Backup 1d | Retain 2w	0	0					
VM-Daily-Weekly-SP-Archive Backup 1d | Retain 2w | Archive	1	2					

Figure 6-9 Create a Data Protection Policy for Archive to the S3 External target, steps 1-3

4. In the pop up dialog, select the ‘More Options’ button

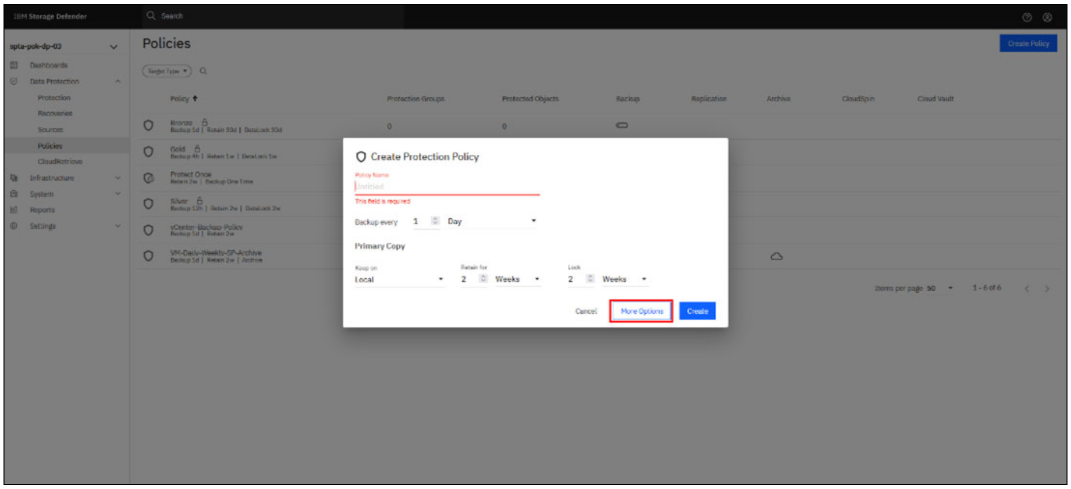

Create Protection Policy

Policy Name: Untitled

This field is required

Backup every 1 Day

Primary Copy

Keep on: Local
Retain for: 2 Weeks
Lock: 2 Weeks

Figure 6-10 Create a Data Protection Policy for Archive to the S3 External target, step 4

Provide and select the following items in the pop up dialog:

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

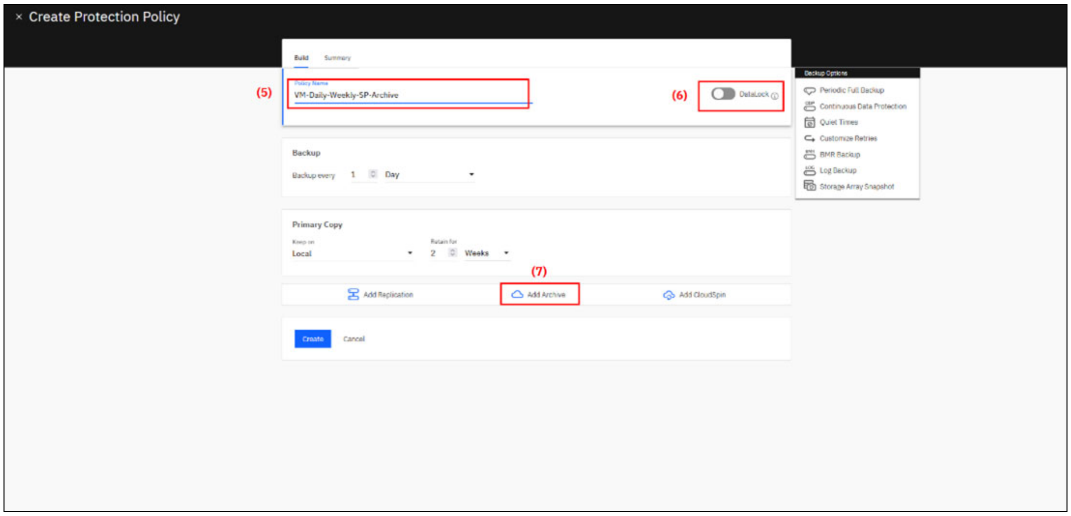

Create Protection Policy

Policy Name: VM-Daily-Weekly-SP-Archive

Backup

Backup every 1 Day

Primary Copy

Keep on: Local

Retain for: 2 Weeks

Figure 6-11 Data Protection Policy for Archive to the S3 External target, steps 5-7

8. locate the Archive section
9. Set the desired external target, frequency and retention settings

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

10. Click Create to create the policy

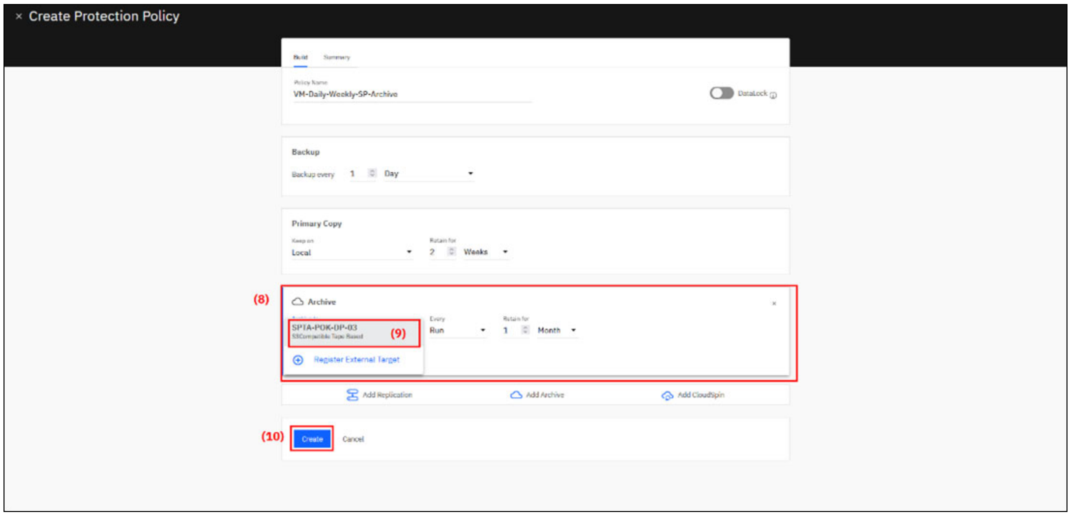

Create Protection Policy

Policy Name: VM-Daily-Weekly-SP-Archive

Backup

Backup every 1 Day

Primary Copy

Keep on: Local

Retain for: 2 Weeks

Archive

SPTA-POK-DP-03
S3Compatible Tape Based

Every: Run

Retain for: 1 Month

Figure 6-12 Create a Data Protection Policy for Archive to the S3 External target, steps 8-10

6.6.2 Configure the Data Protection Group to use the Data Protection Policy

The Data Protection group will include the definition of the workloads to be protected (VMs in this example) and the Data Protection Policy will then be applied to that group.

To configure the Data Protection Group select the following items in the GUI:

1. Expand Data Protection
2. Select the Protection section

3. In the top right of the GUI, select the Protect drop down
4. Select the workload type you want to protect (in this example Virtual Machines)

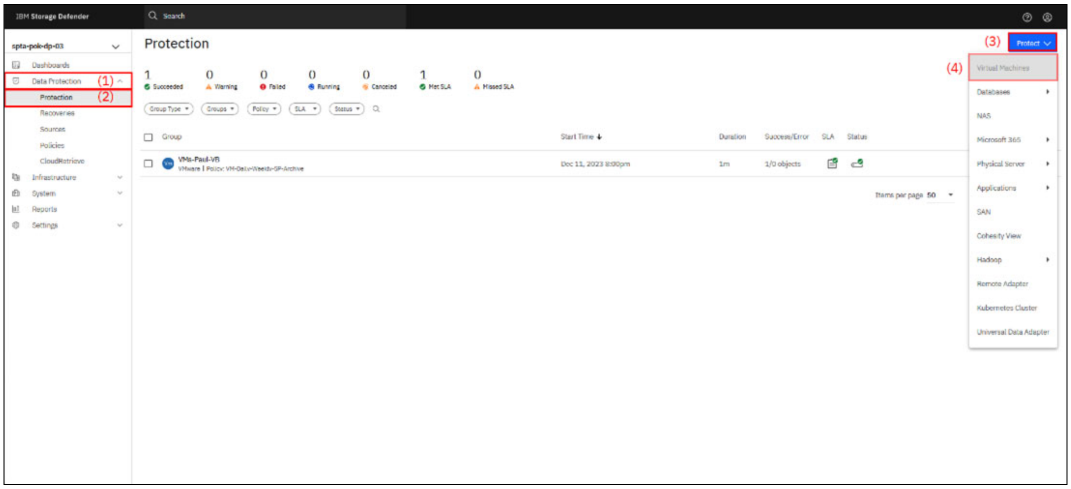

Figure 6-13 Create a new Data Protection Policy using the S3 External target, steps 1-4

5. In the New Protection dialogue window, select Add Objects.

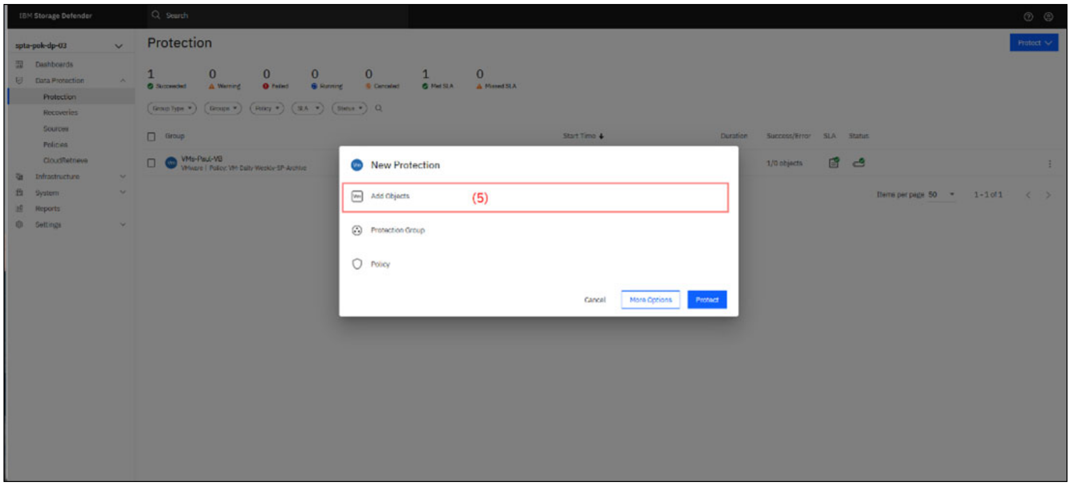

Figure 6-14 Virtual Machine New Protection options, step 5

6. Select the IP Address of the vCenter server
7. Search for VMs by using a name or pattern if desired
8. Select any desired ESX Hosts or individual VMs by name to be protected
9. Click Continue to complete this section

Note: Be aware of the number and size of VM's in this Protection Group.

The selected VMs will be managed in the Protection Group for backup and recovery operations. The sizing for the Storage Protect cold cache pool will be based on the ingest and the restore pattern.

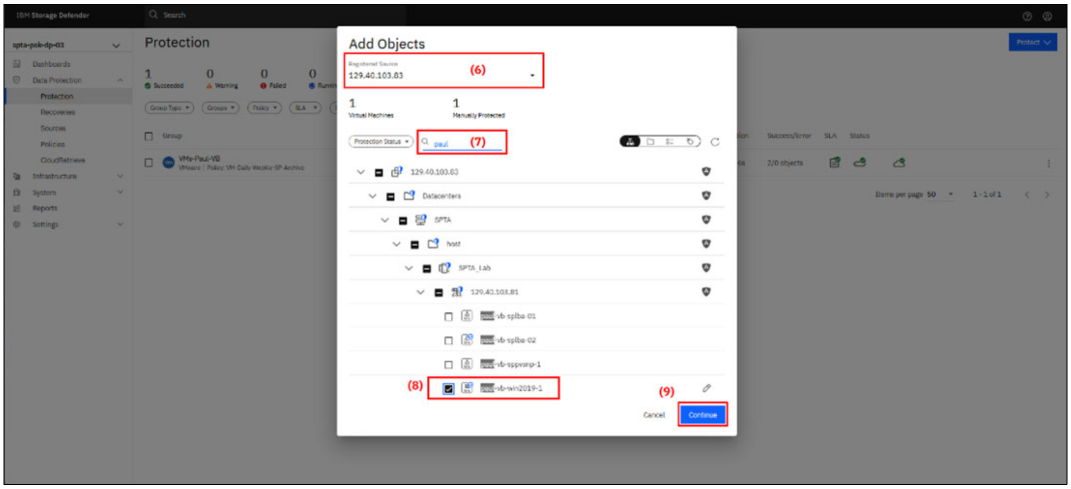

Registered Source: 129.40.103.83

Search: paul

- 129.40.103.83
  - Datacenters
    - SPTA
      - host
        - SPTA_Lab
          - 129.40.103.81
            - ☐ [redacted]-vb-splba-01
            - ☐ [redacted]-vb-splba-02
            - ☐ [redacted]-vb-sppvsnp-1
            - ☑ [redacted]-vb-win2019-1

Figure 6-15 Add Objects window steps 6-9

10. Enter a name for the new Protection Group
11. Select the Policy previously created for the weekly Archive to Storage Protect

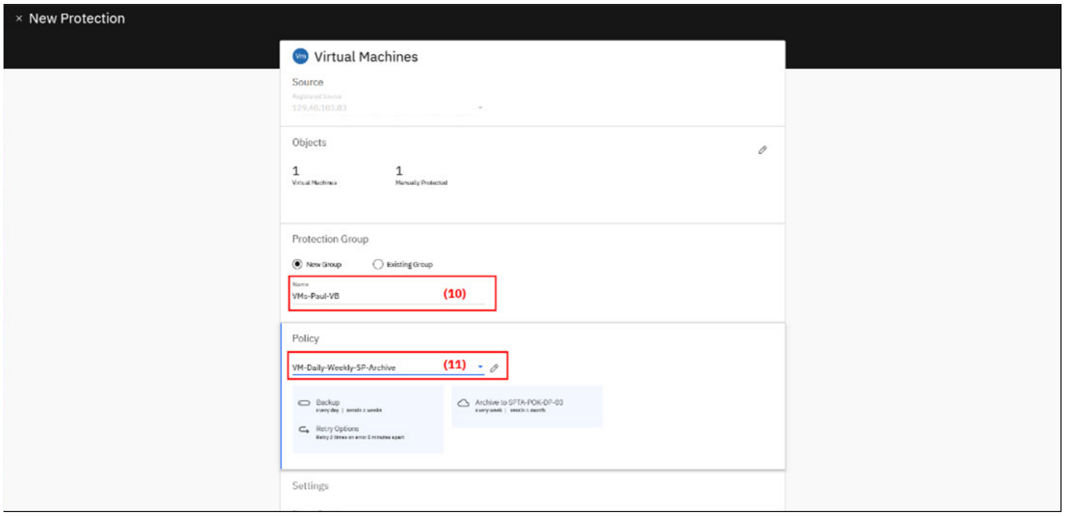

Registered Source: 129.40.103.83

Protection Group: ● New Group ○ Existing Group

Name: VMs-Paul-VB

Policy: VM-Daily-Weekly-SP-Archive

Backup: every day | retain 2 weeks

Retry Options: Retry 3 times on error 5 minutes apart

Archive to SPTA-POK-DP-03: every week | retain 1 month

Figure 6-16 Virtual Machines window steps 10-11

Scroll down to see more options and Additional Settings.

12. Review Settings for the Storage Domain
13. Review Additional Settings if needed
14. Click Protect to complete the configuration

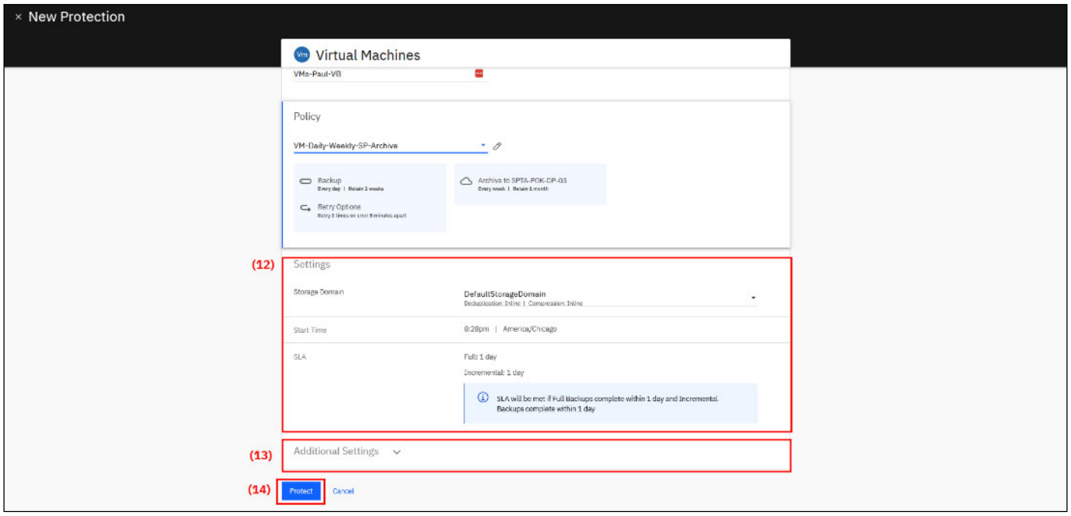

VMs-Paul-VB

Policy

VM-Daily-Weekly-SP-Archive

Backup
Every day | Retain 2 weeks

Retry Options
Retry 3 times on error 8 minutes apart

Archive to SPTA-POK-DP-03
Every week | Retain 1 month

Settings

Storage Domain: DefaultStorageDomain
Deduplication: Inline | Compression: Inline

Start Time: 8:28pm | America/Chicago

SLA: Full: 1 day
Incremental: 1 day

SLA will be met if Full Backups complete within 1 day and Incremental Backups complete within 1 day

Additional Settings

Figure 6-17 Virtual Machines window steps 12-14

Once this is completed, the Data Protection Group will be 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. The current recommendation is to limit the frequency to a weekly or monthly archive to meet retention needs.

6.7 Monitoring the Protection status

The status of Data Protection jobs are presented on the Protection panel. You can review the overall status for multiple protection jobs at a glance in this panel, as well as click on the job name to get additional details about the job.

The Data Protection Group Detail (Figure 6-18 on page 110) shows the daily VM backup statistics as well as:

1. The SLA policy
2. The job status
3. A weekly archive to cloud – in this case, our S3 External Target.

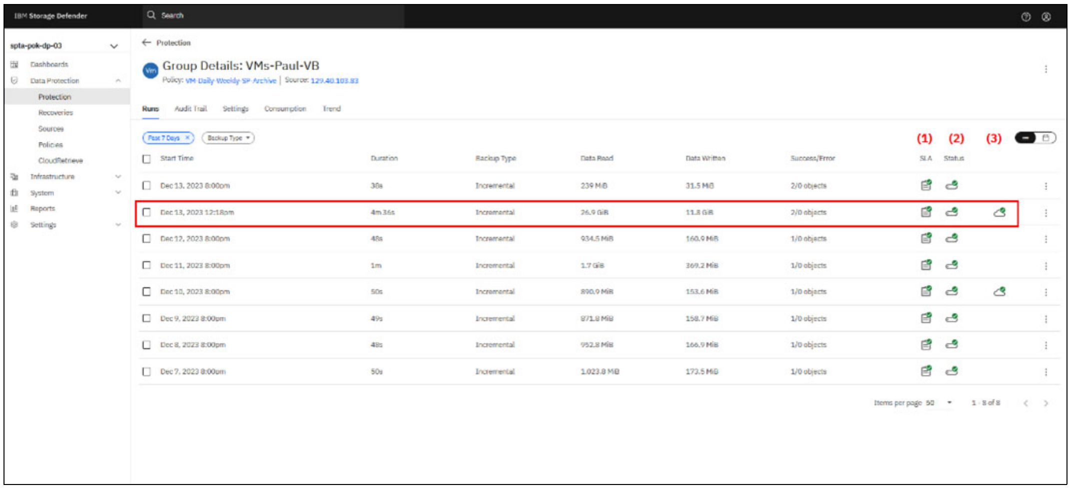

Group Details: VMs-Paul-VB

Policy: VM-Daily-Weekly-SP-Archive | Source: 129.40.103.83

☐	Start Time	Duration	Backup Type	Data Read	Data Written	Success/Error	SLA	Status	
☐	Dec 13, 2023 8:00pm	30s	Incremental	239 MiB	31.5 MiB	2/0 objects	✓	✓	
☐	Dec 13, 2023 12:18pm	4m 36s	Incremental	26.9 GiB	11.8 GiB	2/0 objects	✓	✓	✓
☐	Dec 12, 2023 8:00pm	48s	Incremental	934.5 MiB	160.9 MiB	1/0 objects	✓	✓	
☐	Dec 11, 2023 8:00pm	1m	Incremental	1.7 GiB	369.2 MiB	1/0 objects	✓	✓	
☐	Dec 10, 2023 8:00pm	50s	Incremental	890.9 MiB	153.6 MiB	1/0 objects	✓	✓	✓
☐	Dec 9, 2023 8:00pm	49s	Incremental	871.8 MiB	158.7 MiB	1/0 objects	✓	✓	
☐	Dec 8, 2023 8:00pm	48s	Incremental	952.8 MiB	166.9 MiB	1/0 objects	✓	✓	
☐	Dec 7, 2023 8:00pm	50s	Incremental	1,023.8 MiB	173.5 MiB	1/0 objects	✓	✓	

Items per page 50 ▾ 1 - 8 of 8

Figure 6-18 Reviewing the Protection Group details and status

In this example configuration, the daily backup with a weekly archive to S3 has been running for a couple of weeks. In the example below 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 new Virtual Machine to the Protection Group.

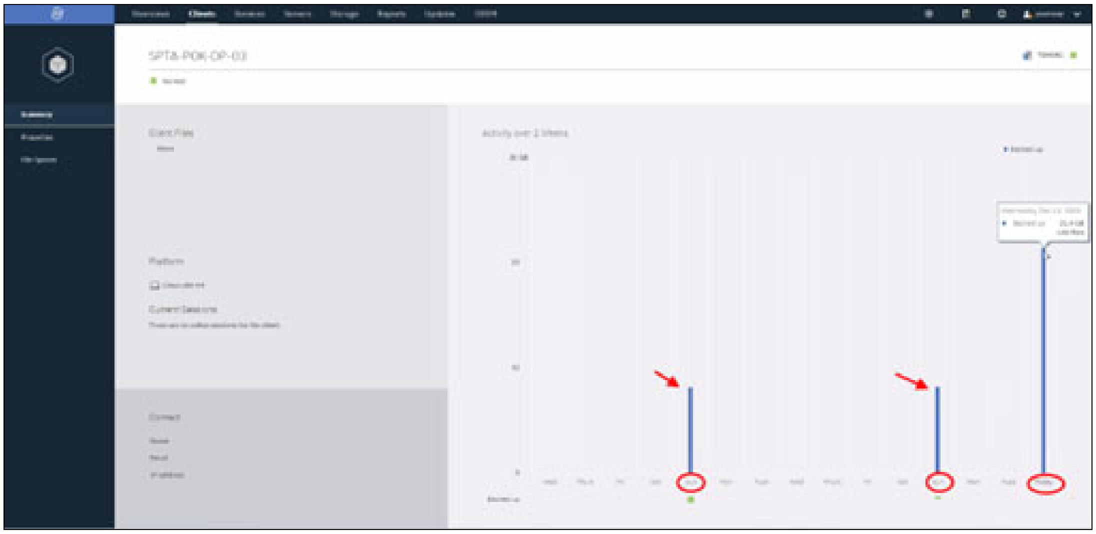

Figure 6-19 Reviewing the activity of the S3 Agent node on the Storage Protect server

Using the Operations Center GUI, we can monitor the weekly activity of the S3 Agent node on the Storage Protect server. A seen in Figure 6-19, there was an increase in the data being stored as a new VM was added to the Data Protection group which uses the S3 archive coldcache pool today. Additional details about coldcahce migration and volume status can be seen using the operations center to run queries on the Storage Protect server as shown in Figure 6-20 and Figure 6-21 on page 111.

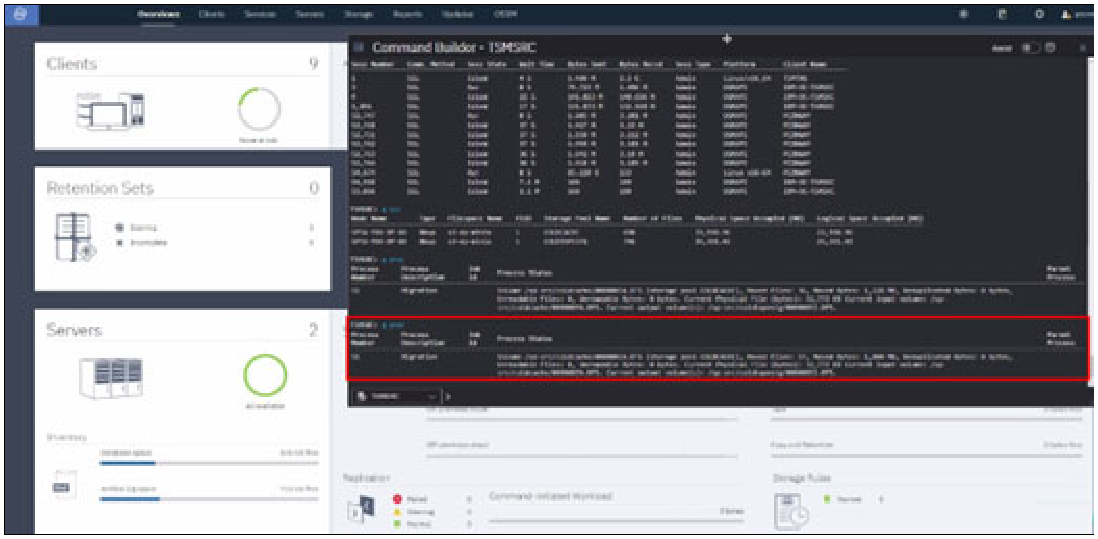

Figure 6-20 Screen shot from Storage Protect server showing migration from cold cache to cold tape

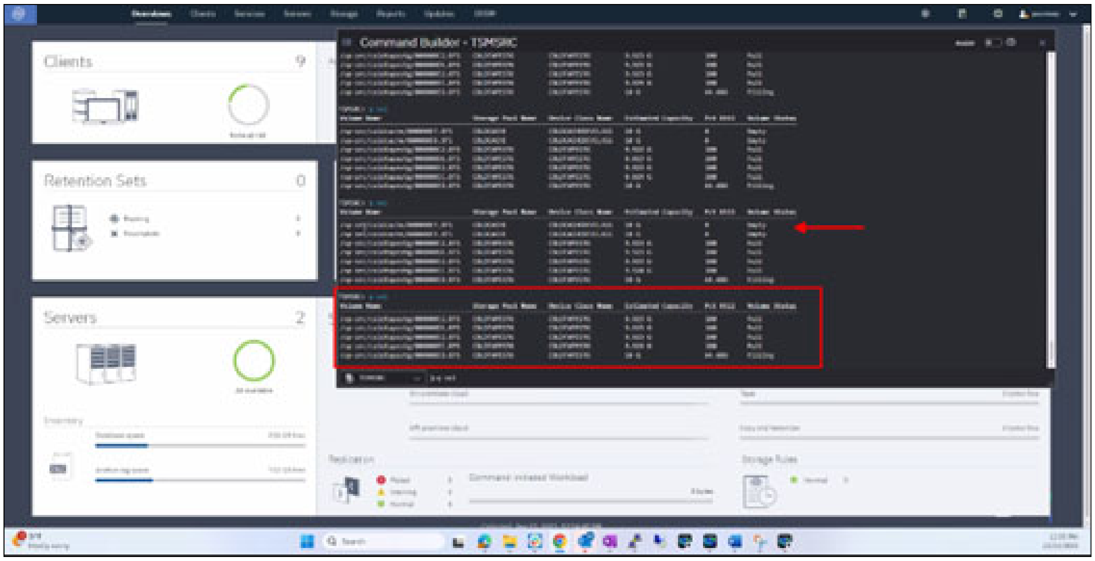

Figure 6-21 Screen shot from Storage Protect server showing volume status after migration

6.8 Retrieve from the S3 External Storage

With the flexible search tools available in IBM Storage Defender's Data Protection utility, there are multiple ways to find your desired data and create restore jobs for protected workloads.

6.8.1 Retrieve using Global Search

In the example below, we will use the Global Search bar to retrieve a VM from the S3 Archive.

1. Use the search bar and type the name of the VM
2. With the VM selected, you will see

- Search Results

- Virtual Machine information
- Protection Group information

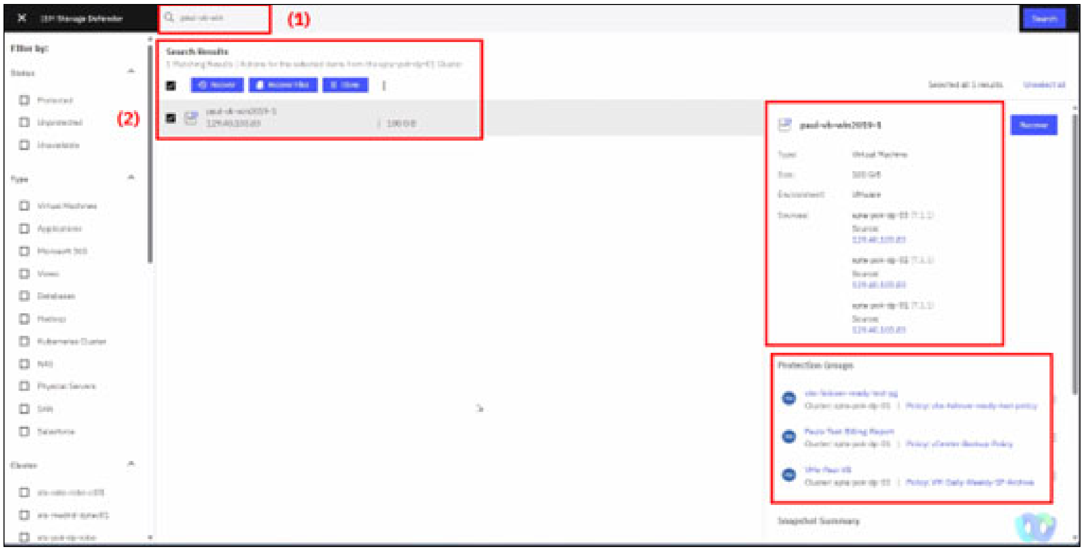

Figure 6-22 Recovery job steps 1-2

3. Use the Pencil icon to open the ‘Edit Recovery Point’ panel for the VM. The Edit Recovery Points table will display the selection of the Snapshots available by timestamp and the Location of the snapshot image.
4. Use the radio button to select the snapshot and click on the cloud to select the S3 location.
5. Click the Select Recovery Point button.

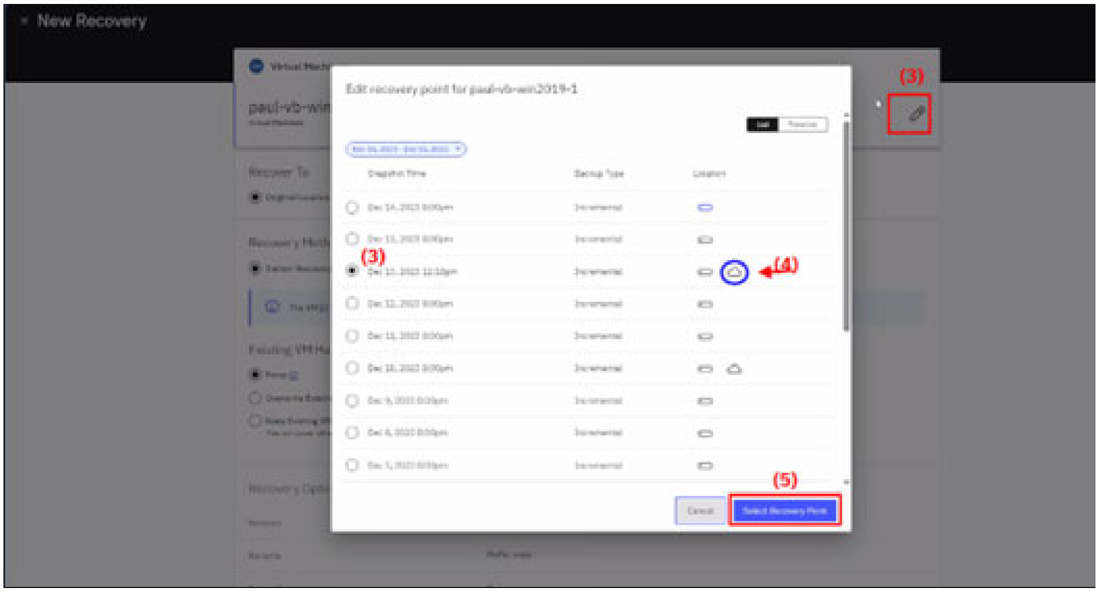

Figure 6-23 VM Recovery job steps 3-5

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

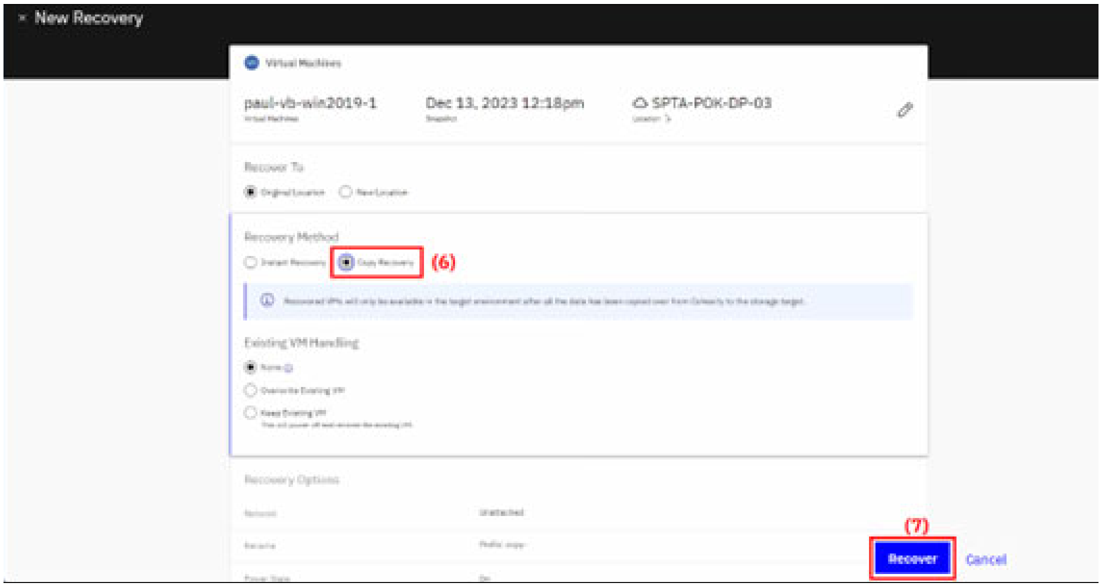

Figure 6-24 VM Recovery job steps 6-7

At this point the recovery of the selected backup data will proceed, and the VM data will be pulled back from the S3 target to the cache storage pool on the SP server. The backup data is then sent to the DP cluster and then recovered to the vCenter.

6.8.2 Retrieve using Recoveries Panel

Another way retrieve archived data is to use the Recoveries option under Data Protection in the Data Protect GUI.

To retrieve data using the recoveries panel use the following steps:

1. In the left had side of the GUI, select Data Protection
2. Then select Recoveries
3. Once selected, click the Recover button to expand the menu details
4. From the Recovery button menu, pick Virtual Machines
5. Finally select the VMs option

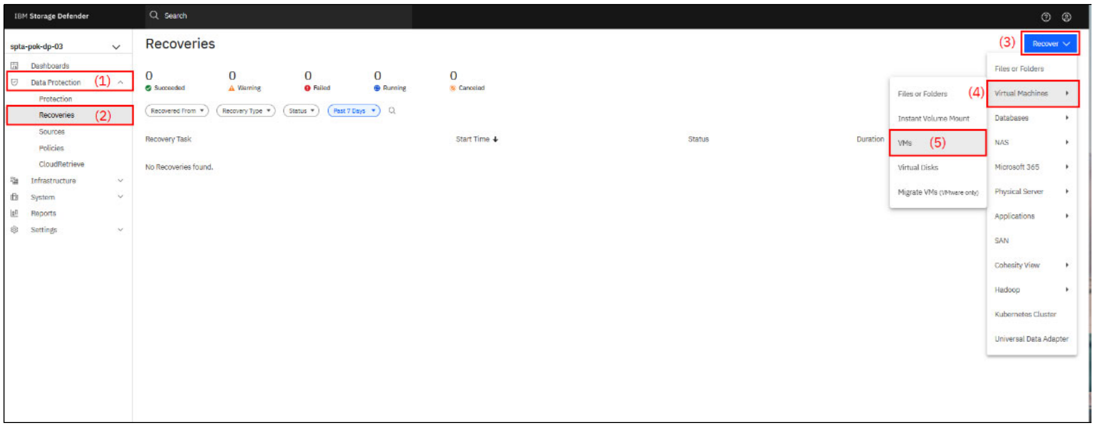

Figure 6-25 VM Recovery job via Recoveries UI, steps 1-5

6. Use the search bar in the panel to type the name of the desired VM
7. Select the VM using the check box
8. Use the pencil icon to select the Recovery Point and Location

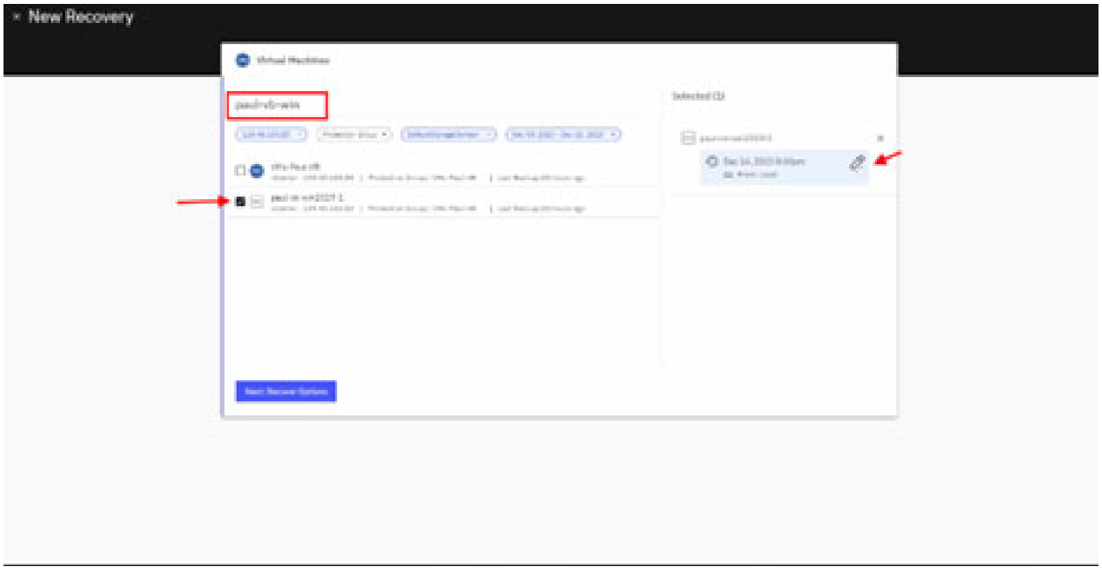

Figure 6-26 VM Recovery job via Recoveries UI, steps 6-8

9. Recovery point
10. Location – click on the cloud to select the External S3 target
11. Click the Select Recovery Point button

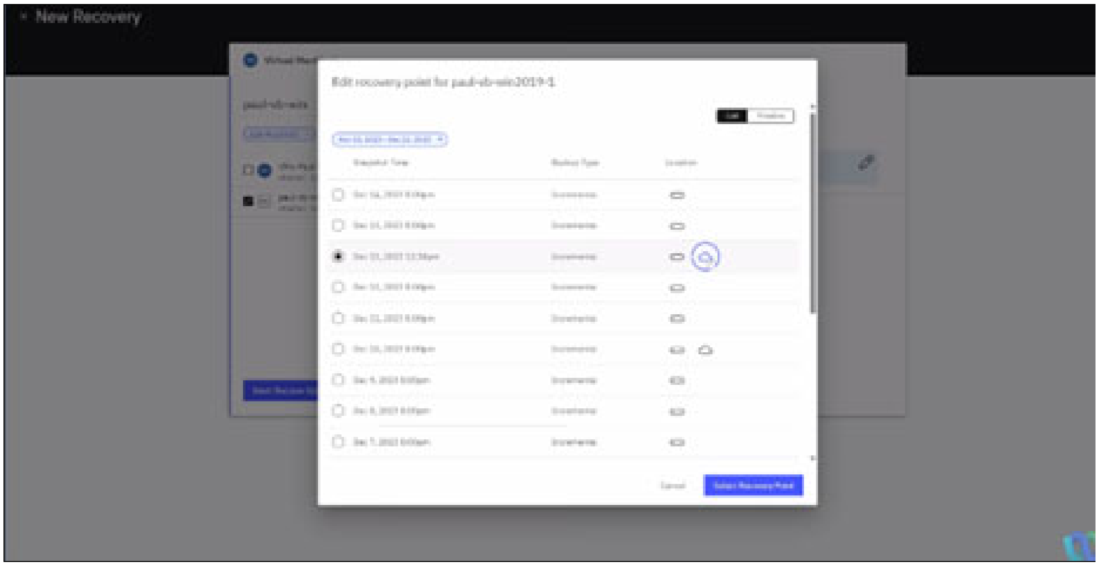

Figure 6-27 VM Recovery job via Recoveries UI, steps 9-11

12. In the Virtual Machine panel, select the Recovery method for Copy Recovery
13. Scroll down and click on Recover

Figure 6-28 on page 115

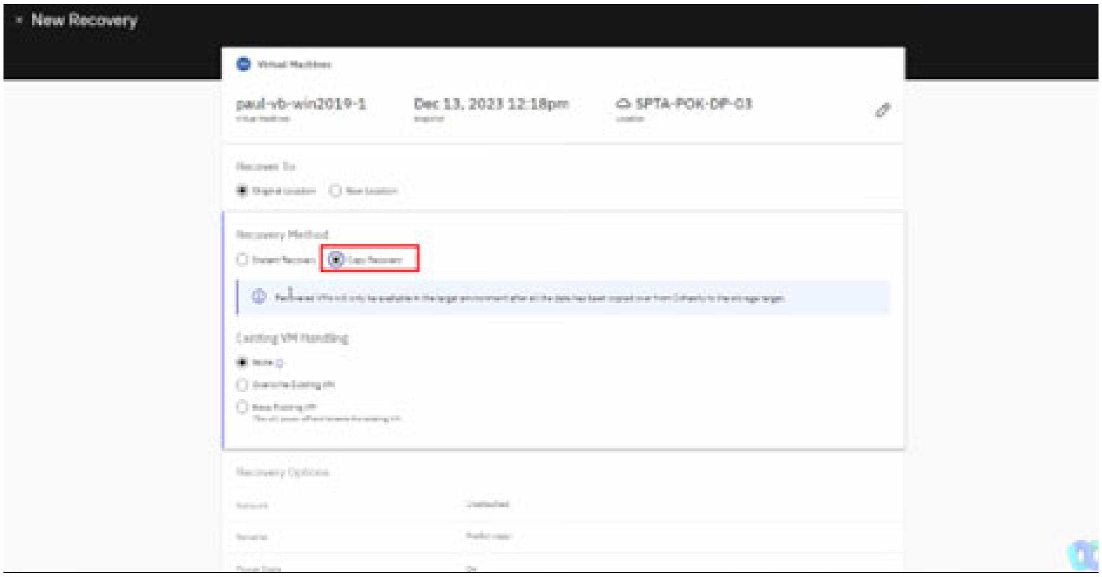

New Recovery

Virtual Machines

paul-vb-win2019-1 Dec 13, 2023 12:18pm SPTA-POK-DP-03

Recover To

● Original Location ○ New Location

Recovery Method

○ Instant Recovery ● Copy Recovery

Existing VM Handling

Recovery Options

Figure 6-28 VM Recovery job steps 12-13

Once complete the recovery will begin to retrieve the archive data.

6.9 Additional Process Monitoring Examples

The following screen shots are examples related to checking the status of archive and retrieval operations. This allows for users to conveniently monitor processes across the environment, from the ESX server to the Data Protect cluster and to the Storage Protect S3 Agent.

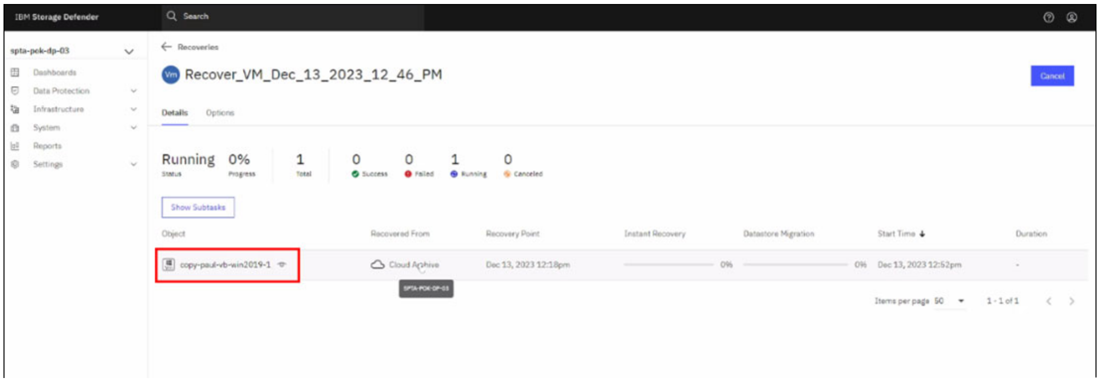

Recover_VM_Dec_13_2023_12_46_PM

Running 0% 1 0 0 1 0
Status Progress Total Success Failed Running Canceled

Object	Recovered From	Recovery Point	Instant Recovery	Datastore Migration	Start Time	Duration
copy-paul-vb-win2019-1	Cloud Archive	Dec 13, 2023 12:18pm	0%	0%	Dec 13, 2023 12:52pm	-

Figure 6-29 Virtual Machine Recovery job in progress

On the Storage Protect server, the Q Process command is used to monitor the retrieve from the cold tape pool to the cold cache pool. (Figure 6-30 on page 116)

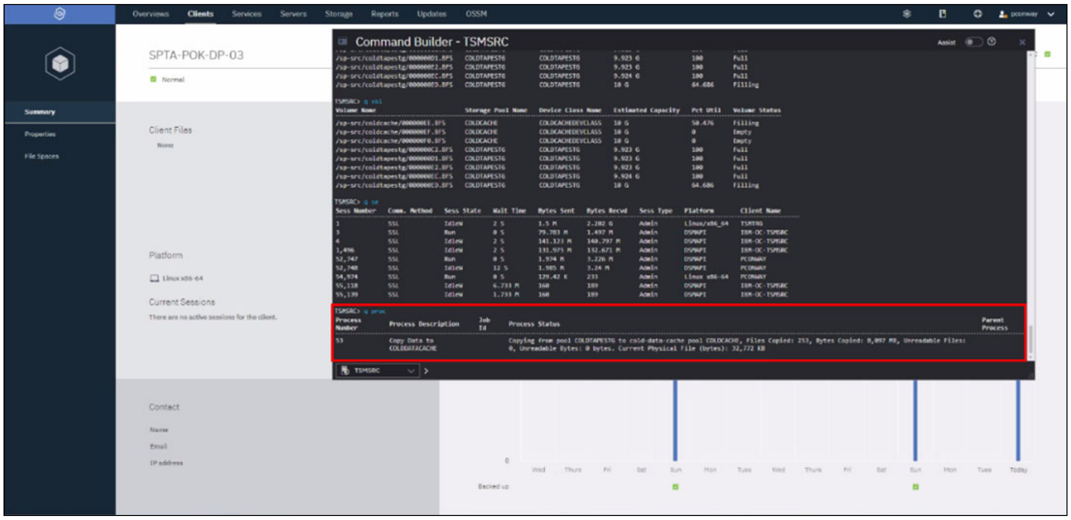

Command Builder - TSMSRC

```
/sp-src/coldtapestg/000000D1.BFS   COLDTAPESTG    COLDTAPESTG        9.923 G              100        Full
/sp-src/coldtapestg/000000E2.BFS   COLDTAPESTG    COLDTAPESTG        9.923 G              100        Full
/sp-src/coldtapestg/000000EC.BFS   COLDTAPESTG    COLDTAPESTG        9.924 G              100        Full
/sp-src/coldtapestg/000000ED.BFS   COLDTAPESTG    COLDTAPESTG        10 G                 64.686     Filling

TSMSRC> q vol
Volume Name                        Storage Pool Name  Device Class Name  Estimated Capacity  Pct Util  Volume Status
----------------------------------------------------------------------------------------------------------------------
/sp-src/coldcache/000000EE.BFS     COLDCACHE      COLDCACHEDEVCLASS  10 G                 50.476     Filling
/sp-src/coldcache/000000EF.BFS     COLDCACHE      COLDCACHEDEVCLASS  10 G                 0          Empty
/sp-src/coldcache/000000F0.BFS     COLDCACHE      COLDCACHEDEVCLASS  10 G                 0          Empty
/sp-src/coldtapestg/000000C2.BFS   COLDTAPESTG    COLDTAPESTG        9.923 G              100        Full
/sp-src/coldtapestg/000000D1.BFS   COLDTAPESTG    COLDTAPESTG        9.923 G              100        Full
/sp-src/coldtapestg/000000E2.BFS   COLDTAPESTG    COLDTAPESTG        9.923 G              100        Full
/sp-src/coldtapestg/000000EC.BFS   COLDTAPESTG    COLDTAPESTG        9.924 G              100        Full
/sp-src/coldtapestg/000000ED.BFS   COLDTAPESTG    COLDTAPESTG        10 G                 64.686     Filling

TSMSRC> q se
Sess Number  Comm. Method  Sess State  Wait Time  Bytes Sent  Bytes Recvd  Sess Type  Platform      Client Name
-----------------------------------------------------------------------------------------------------------------
1            SSL           IdleW       2 S        1.5 M       2.202 G      Admin      Linux/x86_64  TSMTRG
3            SSL           Run         0 S        79.703 M    1.497 M      Admin      DSMAPI        IBM-OC-TSMSRC
4            SSL           IdleW       2 S        141.123 M   140.797 M    Admin      DSMAPI        IBM-OC-TSMSRC
1,496        SSL           IdleW       2 S        131.975 M   132.671 M    Admin      DSMAPI        IBM-OC-TSMSRC
52,747       SSL           Run         0 S        1.974 M     3.226 M      Admin      DSMAPI        PCONWAY
52,748       SSL           IdleW       12 S       1.985 M     3.24 M       Admin      DSMAPI        PCONWAY
54,974       SSL           Run         0 S        129.42 K    233          Admin      Linux x86-64  PCONWAY
55,118       SSL           IdleW       6.733 M    160         189          Admin      DSMAPI        IBM-OC-TSMSRC
55,139       SSL           IdleW       1.733 M    160         189          Admin      DSMAPI        IBM-OC-TSMSRC

TSMSRC> q proc
Process   Process Description   Job   Process Status                                                                                      Parent
Number                          Id                                                                                                        Process
----------------------------------------------------------------------------------------------------------------------------------------------------
53        Copy Data to                Copying from pool COLDTAPESTG to cold-data-cache pool COLDCACHE, Files Copied: 253, Bytes Copied: 8,097 MB, Unreadable Files:
          COLDDATACACHE               0, Unreadable Bytes: 0 bytes. Current Physical File (bytes): 32,772 KB
```

TSMSRC

Figure 6-30 Query process output from the SP server showing the data copy to coldcache

The query volume command shows the Archived snapshot restored to the S3 cold cache pool. The query session command shows the data has been sent from the Storage Protect server to the Data Protect cluster. Figure 6-31

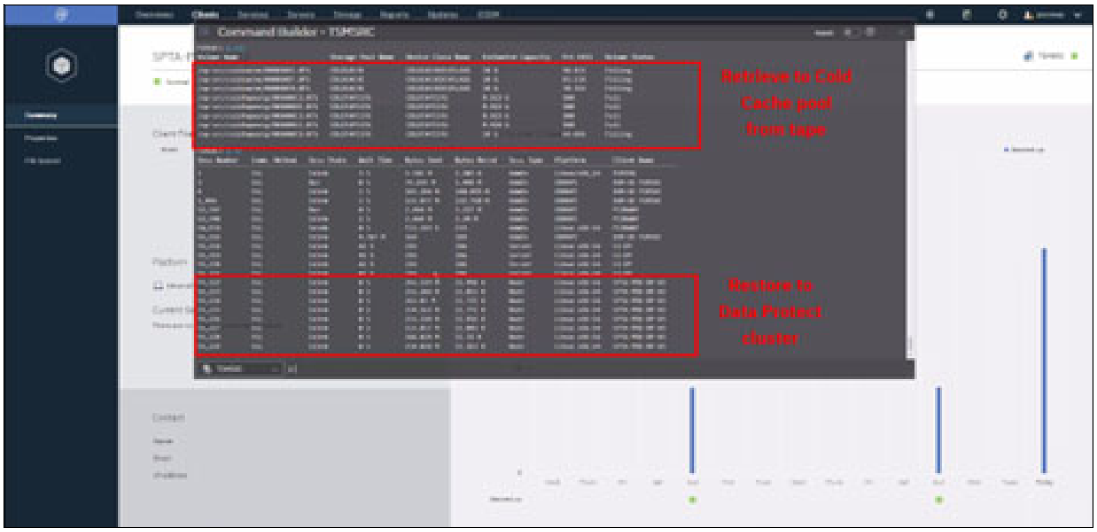

Figure 6-31 Storage Protect command line query volume and query session outputs for recovery related coldcache data movement

The screenshots Figure 6-32 on page 117 and Figure 6-33 show examples of the DP recovery log details for the recovery task.

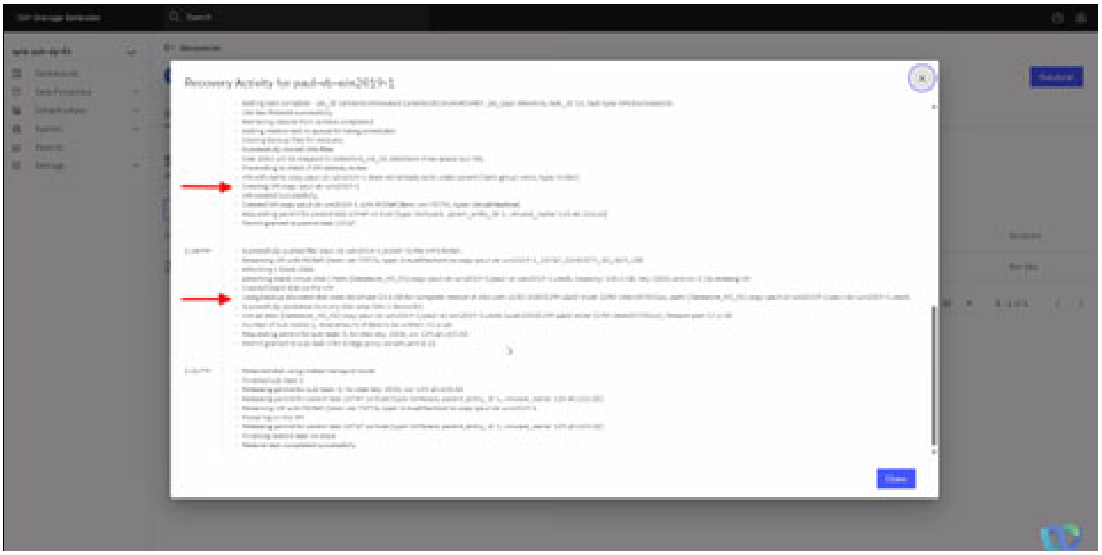

Recovery Activity for paul-vb-win2019-1

Figure 6-32 In Data Protect, click the Recovery job name to get the current status of the job with Recovery Activity details and completion status

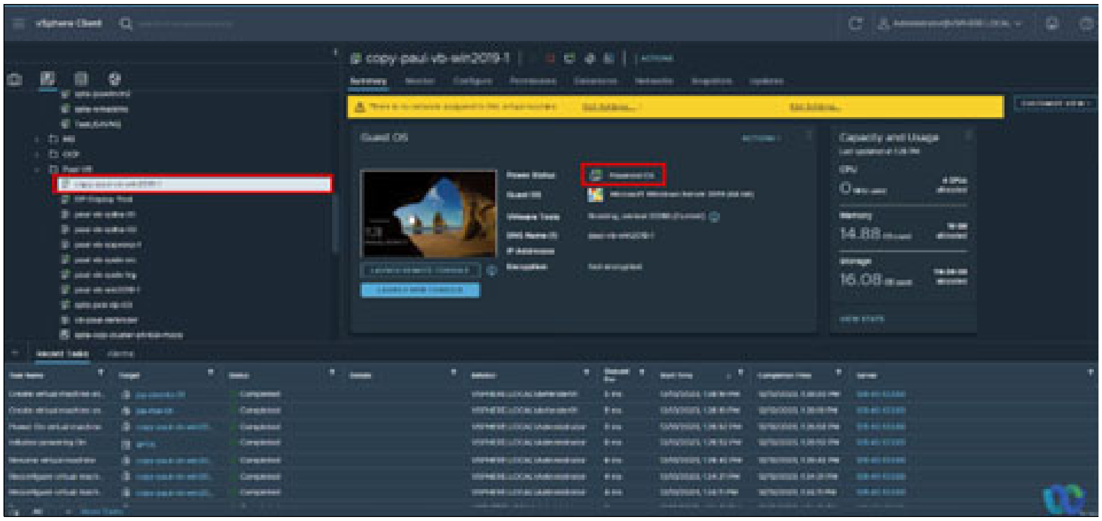

Figure 6-33 Use the vCenter Client to show the VM has been copied and powered on

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5” spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.

Redbooks

IBM Storage Defender: IBM Data Management Service and IBM

SG24-8554-00

ISBN DocISBN

(1.5” spine)
1.5”<-> 1.998”
789 <->1051 pages

Redbooks

IBM Storage Defender: IBM Data Management Service and IBM Data

SG24-8554-00

ISBN DocISBN

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

Redbooks

IBM Storage Defender: IBM Data Management Service and IBM

SG24-8554-00

ISBN DocISBN

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

Redbooks

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

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

(0.1”spine)
0.1”<->0.169”
53<->89 pages

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5” spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.

Redbooks

IBM Storage Defender: IBM Data Management

SG24-8554-00

ISBN DocISBN

(2.5” spine)
2.5”<->nnn.n”
1315<-> nnnn pages

Redbooks

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

SG24-8554-00

ISBN DocISBN

(2.0” spine)
2.0” <-> 2.498”
1052 <-> 1314 pages

SG24-8554-00

ISBN DocISBN

Printed in U.S.A.