

# IBM Storage Defender: Data Resiliency Service

Phil Gerrard

Camila Barrera

Ondrej Bláha

Christian Burns

Erin Farr

Meghan Grable

Juan Carlos Jimenez

Alexis Kojic

James Morassutti

Ranjith Rajagopalan Nair

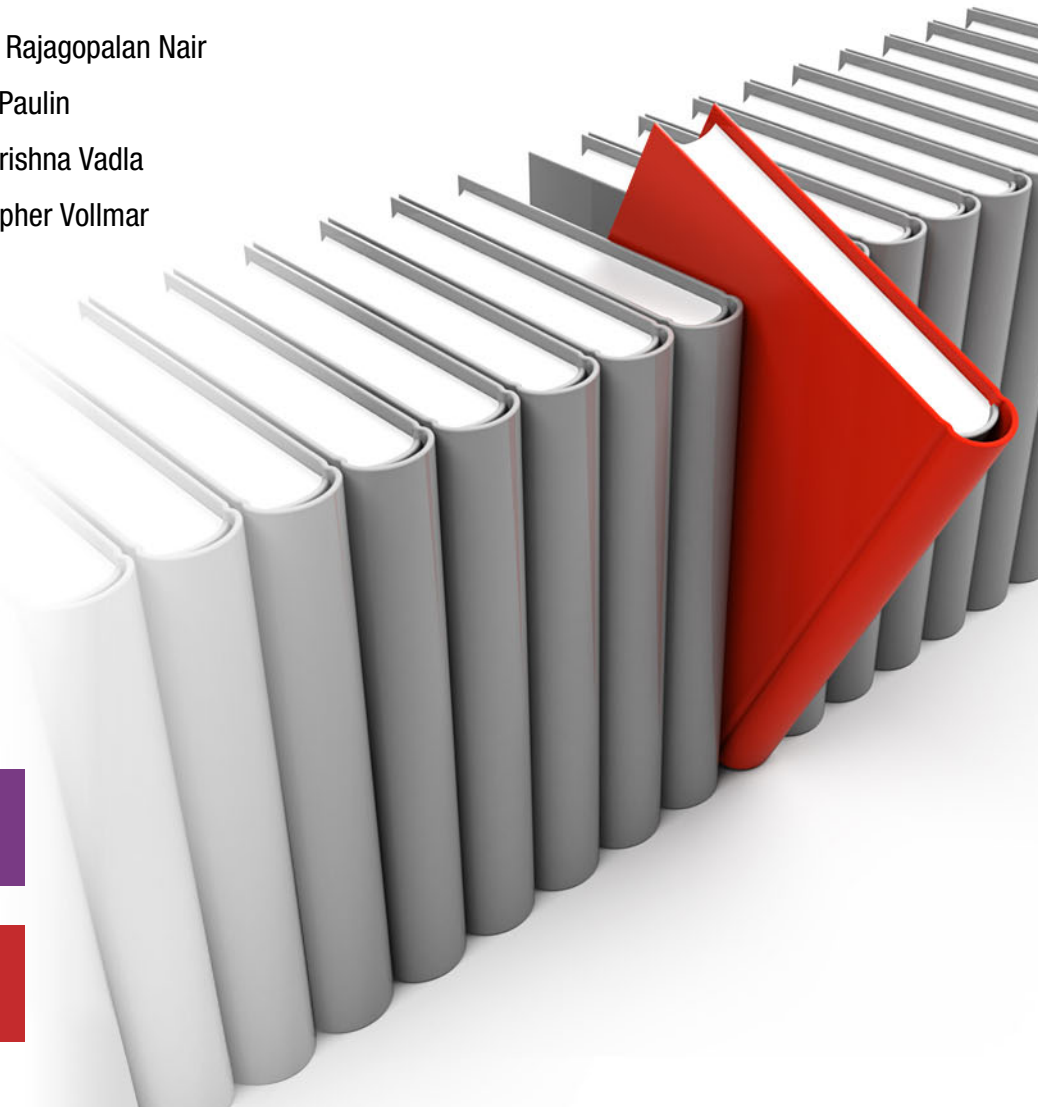
Daniel Paulin

Ramakrishna Vadla

Christopher Vollmar

**Storage**

 **Hybrid Cloud**







IBM Redbooks

## **IBM Storage Defender: Data Resiliency Service**

August 2025

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

**Second Edition (August 2025)**

This edition applies to IBM Storage Defender Data Protect 7.1.1 and 7.1.2 and IBM Storage Defender Data Resiliency Service (DRS) 2.0.13.

© Copyright International Business Machines Corporation 2025. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Notices</b> .....	v
Trademarks .....	vi
<b>Preface</b> .....	vii
Authors .....	vii
Now you can become a published author, too! .....	ix
Comments welcome .....	ix
Stay connected to IBM Redbooks .....	ix
<b>Summary of changes</b> .....	xi
August 2025, Second Edition .....	xi
<b>Chapter 1. Introduction</b> .....	1
1.1 IBM Storage Defender Data Resiliency Service .....	2
1.2 IBM Storage Defender overview and vision .....	3
1.3 IBM Storage Defender components and functions .....	4
<b>Chapter 2. IBM Storage Defender DRS and architecture overview</b> .....	7
2.1 DRS architecture and elements overview .....	8
2.1.1 The IBM Storage Defender mission .....	8
2.1.2 Architecture overview .....	9
2.2 IBM Storage Defender connection manager .....	10
2.2.1 Data sources .....	11
2.2.2 Recovery Locations .....	13
2.2.3 Sensor control nodes .....	14
2.2.4 IBM Storage Defender Sensors .....	15
2.3 Recovery groups .....	16
2.3.1 Auto grouping .....	17
2.3.2 Recovery groups and IBM Storage FlashSystem volume groups .....	18
2.3.3 Supported Configurations for IBM Storage FlashSystem .....	19
2.3.4 Recovery group and IBM Storage Protect .....	21
2.3.5 Supported configurations for Dell PowerMax .....	22
2.3.6 Recovery group and Oracle database .....	23
2.3.7 Recovery group and SAP HANA database .....	23
2.3.8 Recovery group and Microsoft Active Directory .....	24
2.3.9 Recovery group and IBM Db2 database .....	24
2.3.10 Recovery group and IBM Fusion .....	24
2.4 Profiles and IBM Clean Room .....	25
2.4.1 Governance and Clean Room profiles .....	25
2.4.2 IBM Clean Room .....	28
2.5 Adding resources in the Connection Manager and creating profiles in DRS .....	30
2.5.1 Adding resources in Connection Manager .....	30
2.5.2 Creating profiles in DRS .....	37
2.6 Auto-forwarding IBM Storage FlashSystem ransomware threat alerts to IBM Storage Defender .....	45
2.6.1 IBM FlashCore Module .....	45
2.6.2 Integration between DRS and IBM Storage Insights Pro .....	46
<b>Chapter 3. IBM Defender sensors</b> .....	49

3.1 What do sensors do . . . . .	50
3.2 Installing sensors . . . . .	53
3.2.1 Installing the sensor control software . . . . .	53
3.2.2 Adding a sensor control node . . . . .	54
3.2.3 Removing a sensor control node . . . . .	55
3.2.4 Installing an IBM Storage Defender sensor by using the Defender UI . . . . .	55
3.2.5 Installing an IBM Storage Defender sensor by using the CLI . . . . .	57
3.2.6 Uninstalling an IBM Storage Defender sensor by using the Defender UI . . . . .	58
3.2.7 Uninstalling an IBM Storage Defender sensor by using the CLI . . . . .	59
3.2.8 Requirements for IBM Storage Defender sensors . . . . .	60
<b>Chapter 4. Daily administration, alerting, testing, and validation . . . . .</b>	<b>63</b>
4.1 IBM Storage Defender DRS dashboard . . . . .	64
4.1.1 Resiliency Monitoring in the dashboard . . . . .	64
4.1.2 Recovery group status . . . . .	66
4.1.3 Governance profile status . . . . .	67
4.1.4 Recovery posture . . . . .	68
4.2 User management profiles . . . . .	69
4.2.1 Users . . . . .	72
4.2.2 User Groups . . . . .	74
4.3 Integrations for alerting . . . . .	75
4.4 Recovery testing and validation . . . . .	76
4.5 Activating the recovery plan . . . . .	80
<b>Abbreviations and acronyms . . . . .</b>	<b>83</b>

# 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

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

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

Db2®	IBM FlashSystem®	Redbooks®
Enterprise Design Thinking®	IBM Research®	Redbooks (logo)  ®
IBM®	IBM Spectrum®	X-Force®
IBM Cloud®	IBM Z®	z/OS®
IBM FlashCore®	QRadar®	

The following terms are trademarks of other companies:

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

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

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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

# Preface

This IBM Redpaper publication describes IBM's new cyber resiliency solution, IBM Storage Defender Data Resiliency Service (DRS). By using DRS, you can use new detection mechanisms for your environment to detect threats early and get a full view of the infrastructure by connecting primary storage arrays like IBM FlashSystem® and auxiliary storage solutions for backup, such as IBM Defender Data Protect and IBM Storage Protect. Also, users can set up Governance profiles to help ensure that their data is meeting internal or regulatory standards.

This IBM Redpaper publication is designed to help users and administrators to better understand how to set up, tailor, and configure this offering for their environments.

## Authors

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

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

**Camila Barrera** is a Storage Technical Sales Specialist based in Dallas, Texas. She holds a bachelor's degree in Management Information Systems from the University of Texas at Austin.

**Ondrej Bláha** works as a Technology EMEA subject matter expert (SME) and Architect focusing on IBM Storage Software with a specialization in data resilience (the Storage Defender strategy for primary and secondary workloads). He has been with IBM for more than 17 years, and for the last 10 years, he has served in several regional roles as an SME and Customer Technical Support or Technical Advisor for key IBM customers. Ondrej is an official IBM instructor for external IBM Software Training organizations who creates technical hands-on IBM Storage Defender courses in the EMEA region. In 2016, he received the "Best of IBM" award due to the delivery of key projects that still act as public references today. Ondrej is originally from the Czech Republic and lives in Prague.

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

**Erin Farr** is a Senior Technical Staff Member (STSM) who is based in the IBM Storage CTO Office, where she explores new technology for future products and shapes strategy in anticipation of industry trends. Her areas of focus are cybersecurity and cyber resiliency. She was instrumental in forming the vision for IBM Storage Defender, and she is passionate about helping customers prevent and recover from cyberattacks. Before joining IBM Storage in 2021, she was the team lead for the IBM Z® Center for Secure Engineering for z/OS®. She worked on product development for most of her career, in areas such as IBM z/OS UNIX, analytics, virtualization management, and open source.

**Meghan Grable** is a global Growth Product Manager who specializes in data management and resilience solutions, both Software as a Service (SaaS) and software-based, with a strong focus on Product-Led Growth (PLG) strategies. With over 5 years of experience, she has led cross-functional teams to develop cutting-edge technologies that empower organizations to exceed their compliance goals and enhance their cyber resilience against threats like cyberattacks, natural disasters, and human errors. Based in Raleigh, North Carolina, Meghan holds a degree in Service Design from the Savannah College of Art and Design. Her expertise in Service Design, enterprise design thinking, and PLG enables her to create innovative, customer-focused products that drive business success and growth directly through user engagement and product experience.

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

**Alexis Kojic** is a Storage Technical Sales Specialist who is based in Canada. He has 2 years of experience in the IT storage and cyber resilience field. He holds a BEng degree in Computer Engineering from Toronto Metropolitan University.

**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 as 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.

**Ranjith Rajagopalan Nair** is a Software Architect who is based at IBM India. He has worked at IBM for 20 years, which includes working on IBM Systems Storage for the past 10 years. Ranjith's current responsibility includes the development and delivery of IBM Storage Insights. Ranjith holds a master's degree in Computer Science from the University of Kerala.

**Daniel Paulin** is a Storage Software Architect who is based at IBM Croatia. An IT professional since 1997, he has worked as a systems engineer for two financial companies in Croatia. In 2003, he joined IBM, where he gained comprehensive experience in designing, developing, and deploying architectures and infrastructure for various storage and server solutions. Currently, Daniel is focused on IBM Storage solutions, particularly IBM Storage Defender. His work is part of IBM's broader initiative to enhance cyber resiliency and storage security, which helps ensure data protection across diverse IT infrastructures. Daniel plays a crucial role in promoting these innovations within the NCEE region, especially in storage management and safeguarding against data breaches.

**Ramakrishna Vadla** is an STSM and Lead Architect for IBM Storage Insights and IBM Spectrum Control. He is responsible for developing and designing the IBM Storage

Insights product, which monitors storage systems. With over 20 years of experience, he has worked on large-scale distributed systems across various technologies, including AIOps, microservices architecture, storage management, cloud-native services, and middleware systems. He has spoken at multiple technical forums, including the SNIA Storage Developer Conference and IBM global conferences, and has contributed to the open-source community. He holds a Master of Technology degree in Computer Science from the International Institute of Information Technology, Hyderabad, India.

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

## Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience by using leading-edge technologies. Your efforts 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](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

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

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

[ibm.com/redbooks](http://ibm.com/redbooks)

- Send your comments in an email to:

[redbooks@us.ibm.com](mailto: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>

# Summary of changes

This section describes the technical changes that were made in this edition of the paper and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

## August 2025, Second Edition

This revision reflects the addition, deletion, or modification of new and changed information:

- ▶ General updates: Updated publication screenshots and some language to reflect new GUI and UI changes included in the release of IBM Storage Defender 7.2 and DRS version 2.0.13.
- ▶ Added information related to Connection Manager updates and changes
- ▶ Updated supported data source information
- ▶ Updated supported workload information for VM protection
- ▶ Updated information related to sensors
- ▶ Added and updated information related to Recovery Groups and Auto Grouping functionality
- ▶ Updated information related to Defender integration with Storage Insights Pro and FlashCore/FlashSystem





# Introduction

This chapter introduces the IBM Storage Defender Data Resiliency Service (DRS). It describes this solution's overview and vision for the future. Later chapters describe the different functions of the solution, how to set them up, configure them, and run them to drive the most value.

This chapter describes the following topics:

- ▶ 1.1, “IBM Storage Defender Data Resiliency Service” on page 2
- ▶ 1.2, “IBM Storage Defender overview and vision” on page 3
- ▶ 1.3, “IBM Storage Defender components and functions” on page 4

## 1.1 IBM Storage Defender Data Resiliency Service

Today, organizations face severe threats to their data as the number of cyberattacks increases and malicious actors become more sophisticated. According to the [IBM X-Force® Threat Intelligence Index 2025 report](#), 42% of all reported incidents involved malware, making it the most common threat, and 28% of reported incidents were attributed to ransomware attacks. In addition to malware, IT organizations are threatened by natural disasters, system failures, human errors, and even sabotage. These events and others like them might result in any number of outcomes, including financial losses and harm to customer trust if sensitive data is compromised.

IBM Storage Defender (DRS) is a cloud-based data resilience platform to help organizations restart essential business operations if there are cyberattacks or other unforeseen catastrophic events. DRS provides data resilience and compliance, early threat detection, and safe and fast recovery orchestration for data that is stored across primary and auxiliary storage. The Software as a Service (SaaS) plane helps to detect and respond to cyberthreats, such as malware and ransomware attacks, and enables rapid recovery of data if there is a security breach or data loss. Administrators can take quick and effective action to minimize the risks of massive financial losses or damage to a company's reputation.

DRS offers features such as Governance and Compliance, Early Threat Detection, and Safe and Fast Recovery. It provides rapid recovery of data and applications if there is a cyber incident or data loss, minimizing downtime and helping ensure business continuity.

**Note:** The term *auxiliary storage* denotes a secondary or backup storage location that enables you to use copies of data in place before the data is recovered.

DRS provides the following benefits:

- ▶ Data resilience and compliance  
Set your resiliency standard to meet compliance across a data estate.
- ▶ Early threat detection  
Provides near real-time file system monitoring, backup anomaly analysis, IBM FlashSystem inline detection, and recovery time scanning.
- ▶ Safe and fast recovery  
By using air-gapped data, immutable snapshots, and Clean Room recovery, you can confidently and quickly recover your business operations.
- ▶ Connect storage ops and security operations (SecOps)  
Collect storage and security events to send alerts to support staff and other security tools. Deep integration with IBM QRadar® and Splunk.
- ▶ IBM FlashSystem integration  
Understand threats down to the IBM FlashSystem volume and virtual machine (VM) level, which helps speed up identification and the time to initiate remediation. Automatically trigger proactive Safeguarded Copy snapshots to limit damage and automatically recover to a Clean Room for testing.
- ▶ IBM Storage Defender Data Protect integration  
Catalog IBM Storage Defender Data Protect recovery points, understand how your policies align with your Governance goals, and automatically recover to a Clean Room for testing.

With the combination of SecOps, storage, and infrastructure tools, DRS can monitor end to end data movement and quickly supply critical information so that teams can make the most intelligent decision about recovery strategies. DRS presents data resilience and recoverability options across primary and auxiliary storage, bringing internal teams together with a comprehensive single pane of glass view and simplifying the orchestration of business recovery processes.

## 1.2 IBM Storage Defender overview and vision

This section describes the vision behind IBM Storage Defender, along with a high-level overview of the functions that IBM Storage Defender provides to meet those goals and customer needs.

Originally, backup solutions were focused on protecting against accidental data loss (user mishaps or data corruption), hardware failures, or natural disasters (such as hurricanes). As cyberattacks became more prevalent, the industry adapted to meet the growing needs of prevention and mitigation against bad actors attempting to cause harm. Many enterprises assert that having a good DR plan in place means that they are covered for responding to cyberattacks. However, cyber recovery has many different characteristics beyond simple DR:

- ▶ The impact of a natural disaster is regional, but a cyberattack can be global or can be targeted at certain datasets.
- ▶ Depending on the location of the backup data, you might expect that natural disasters cannot affect your data, but with cyberattacks, backups can be targeted first. Targeting backup or data copies further impacts recoverability, forcing a victim into paying the ransom. According to IDC, over 30% of data backups are successfully destroyed, and 55% in North America.<sup>1</sup>
- ▶ The probability of a natural disaster is relatively low compared to a cyberattack.

In addition, although an enterprise might practice their DR recovery, research<sup>2</sup> shows that few enterprises are practicing cyber recovery, which includes aspects outside of traditional DR:

- ▶ Playbooks to help ensure seamless interaction with incident responders.
- ▶ Antivirus scanning during recovery to avoid the reintroduction of dormant malware.
- ▶ Practicing identification of good data copies at scale. Cyberattacks are not as instantaneous as a power outage, for example, and the points of impact might vary across available recovery points.

The industry has responded to some of these threats with solutions that provided extra protection, such as air gaps or immutability. Initially, many auxiliary storage vendors also added threat detection to their solutions, which led Gartner to coin the term *cyberstorage*. Solutions with threat detection can be pure software or a dedicated appliance, but the trend is that threat detection and response capabilities are being added into storage across the industry. At first, only auxiliary storage vendors were providing this capability, but IBM saw a need for detection in primary storage.

However, as IBM investigated this trend, it was apparent that a series of concerns must be addressed:

---

<sup>1</sup> Source: 1. Ransomware 2024. If we have backups, why are we still paying a ransom?. IDC. March 2024. IDC Survey - Doc Document number:# US51941924

Source: 2. 2022 Gartner Hype Cycle report

<sup>2</sup> Ransomware 2024. If we have backups, why are we still paying a ransom?. IDC. March 2024. IDC Survey - Doc Document number:# US51941924

- ▶ If an attack is detected, who is expected to respond? Nobody expects storage administrators or data protection teams to become incident responders.
- ▶ Disparate solutions make it difficult to identify and find the last good copy. Recovery points can be primary storage snapshots or auxiliary storage backups. Often, they are managed by different tools and different teams. If an incident is actively occurring, a storage administrator or incident responder might not have a holistic view across both primary and secondary storage.
- ▶ If an enterprise takes backups once a day, is backup-based detection (only) fast enough to detect issues?
- ▶ How do you determine the scope of damage? Which systems were impacted? What is the timeline?
- ▶ Although storage-based threat detection is important, it is unlikely that someone would swap out their current solution only to get access to threat detection. How can this need for extra features be met?

IBM recognized the need for the following items:

- ▶ A way to provide these cyber recovery features that works with existing investments and current storage solutions.
- ▶ The ability to provide a holistic view across both primary and auxiliary storage for recovery and threat detection.
- ▶ Features that specifically address cyber recovery, such as Clean Rooms (isolated recovery environments) and antivirus scanning during recovery to avoid re-introduction of dormant malware.

The vision of IBM Storage Defender is one that meets all these needs. DRS is a Software as a Service (SaaS) management feature that is designed and intended to integrate with and sit above an enterprise's existing storage system investments. DRS enables a holistic view across primary and auxiliary storage and provides advanced ransomware detection and recovery features that address modern threats in storage environments.

## 1.3 IBM Storage Defender components and functions

DRS is a multi-faceted offering that has several functions that work together to stay ahead of data disruptions and attacks.

DRS has a centralized dashboard to promote cross-department visibility. Within the dashboard, Recovery Groups, resource summary, and usage monitoring are readily available in a simplified format. Also, Recovery Groups, Governance profiles, resources, and integrated configurations can be created and updated within this dashboard.

DRS deploys AI-powered sensors to quickly detect threats and anomalies from backup metadata, array snapshots, and other relevant threat indicators. Signals from all available sensors are aggregated to increase detection paths for a fast response.

IBM FlashSystem offers protection through immutable copies of data that are known as *Safeguarded copies*, which are isolated from production environments and cannot be modified or deleted through user error, malicious actions, or ransomware attacks. IBM Storage Defender includes IBM Storage hardware integration with IBM FlashSystem and SAN Volume Controller to include the usage of Safeguarded Copy as part of the DRS configuration.

IBM Storage Defender Data Protect offers an immutable auxiliary storage solution that incorporates backups with rapid recovery, policies to lock data even from administration removal, and two-person integrity checking. It features a scale-out, clustered architecture with deep integration into databases and hypervisors and a robust global management structure.

By integrating DRS with a security information and event management (SIEM) tools like Splunk or QRadar, you can use advanced notification aggregation so that crucial information is available and used for initiating the next steps between infrastructure and SecOps teams. This setup provides needed information use when deciding whether recovery plans should be implemented immediately or how best to address threats.

Clean Room isolation helps ensure that the backups are clean and malware-free before returning the data to a production environment. As a customer-managed resource, DRS provides guided testing workflows to recover, test, and isolate backups before pushing them to production systems, which help ensure that clean recovery data is present.

DRS also brings in data from various points to help organizations become proactive in their approach to data resilience. Identifying threats early helps ensure the availability of business operations, which is essential to building operational resilience and trust. DRS is an advanced solution that helps organizations build operational resilience by bringing together multiple levels of threat detection and data protection that serve as a base when building out advanced lines of defense across primary and auxiliary storage. This technology enables users to effectively detect and respond to cyberattacks and other unforeseen threats to storage environments. When put together, these features enable DRS to help navigate unpredictable events and help ensure the continuity of vital business operations and processes.





## IBM Storage Defender DRS and architecture overview

This chapter describes the architecture and elements of IBM Storage Defender Data Resiliency Service (DRS). It breaks down the core functions and elements, including the local and cloud-based elements, and alerting.

This chapter describes the following topics:

- ▶ 2.1, “DRS architecture and elements overview” on page 8
- ▶ 2.2, “IBM Storage Defender connection manager” on page 10
- ▶ 2.3, “Recovery groups” on page 16
- ▶ 2.4, “Profiles and IBM Clean Room” on page 25
- ▶ 2.5, “Adding resources in the Connection Manager and creating profiles in DRS” on page 30
- ▶ 2.6, “Auto-forwarding IBM Storage FlashSystem ransomware threat alerts to IBM Storage Defender” on page 45

## 2.1 DRS architecture and elements overview

IBM Storage Defender provides end-to-end data resiliency. Understanding the DRS architecture and its elements help you to properly plan, test, and recover your critical data.

The architecture shows you how DRS fits into IBM Storage Defender and what are the elements that make up the architecture. This chapter describes the following elements:

- ▶ IBM Storage Defender Connection Manager
- ▶ Data sources
- ▶ Recovery locations
- ▶ Sensors
- ▶ Recovery groups
- ▶ Clean room

### 2.1.1 The IBM Storage Defender mission

DRS is an optional component within the IBM Storage Defender solution that provides cyber resiliency capabilities for managing primary and secondary data, and workloads. DRS concepts simplify the recovery of complex applications and automated recovery tests, and performs validation of primary and secondary data. Also, DRS can send notifications if anomalies are detected, and indicate when the trustworthiness of existing primary and secondary data sources has decreased.

DRS is a combination of cloud-based Software as a Service (SaaS) that is managed by IBM and an on-premises agent that manages communications from your data center. The data center agent is called the IBM Storage Defender Connection Manager, and collects telemetry data about your primary and secondary data, and data sources like VMware. The data stays on-premises. The telemetry data goes to the DRS, which helps secure and recover the data.

DRS can surface and aggregate the detection of operational threats on your production data. At the time of writing, this feature includes the following system-level detection:

- ▶ Detection on the file-system level by using IBM Storage Defender Sensor technology
- ▶ Detection on the storage block level by using IBM Storage FlashSystem and IBM FlashCore Module (FCM) technology and statistical analysis to identify threat patterns

DRS introduces the concept of recovery groups, which are used to group resources together within the DRS. The combination of resources enables DRS to perform automated test recoveries and to verify automatically whether the protection policies that are set up in the related data protection application meets the requirements for a cyber resilient environment. In DRS, multiple recovery groups can be defined. The key parts of a recovery group are the protected resources, for example, virtual machines (VMs); the Clean Room profile that defines the environment that can be used for automated test recoveries; and the Governance profile that specifies the requirements for cyber resiliency within each recovery group that is defined.

The DRS dashboard presents information that is relevant to cyber resiliency in a consolidated view. This dashboard displays the configured recovery groups and any potential informational or warning messages that are related to its cyber resiliency requirements being met for each of those groups. From the DRS dashboard, you can access all capabilities and configuration options for the service.

DRS is designed to enhance data cyber resiliency to help protect against events like hardware failures, human errors, sabotage, natural disasters, and ransomware. By consolidating key parts of the existing IBM Storage portfolio into a single solution, you can use new detection and protection capabilities on your data. DRS includes the following capabilities:

- ▶ Supports software protection for multiple operating systems inside a VMware environment.
- ▶ Deploys anomaly-based sensor agents on VMware VMs (IBM Storage Defender Sensors).
- ▶ Integrates and aggregates the hardware detection capabilities of IBM Storage FlashSystem to receive alerts from IBM Storage Insights and IBM Storage Defender. These alerts can be sent through integration with IBM QRadar and Splunk SIEM solutions.
- ▶ Can recover data from a more recent point in time by creating a SafeGuarded Copy (immutable hardware snapshot) on IBM Storage FlashSystem.
- ▶ Can recover IBM Defender Data Protect backups into the Clean Room for testing as part of the recovery group's collection of recovery points.
- ▶ Provides a dashboard that can help clients better understand inconsistencies between their primary storage copies and backup copies for the same workload or application.
- ▶ Additional dashboard features include the following items:
  - Can create and define recovery groups, which are a collection of data resources that should be backed up and recovered as a unit.
  - A summary of connected resources like VMs, data sources, Recovery Locations, and connection managers.
  - A license usage overview that highlights the number of recovery groups and deployed sensors.

## 2.1.2 Architecture overview

DRS is a component of IBM Storage Defender that runs in a cloud, and uses the on-premises IBM Storage Defender Connection Manager to inventory the available and important resources in a data center.

IBM Storage Defender Connection Manager provides on-premises data center connections to the following resources:

- ▶ Data Sources (IBM FlashSystem, IBM Storage Defender Data Protect, VMware vCenter, Dell EMC PowerMax, IBM Storage Protect & IBM Fusion)
- ▶ Recovery Locations
- ▶ Sensor Control Nodes and IBM Storage Defender Sensors

The data sources and Recovery Locations that are connected to the IBM Storage Defender Connection Manager are inventoried automatically, and IBM Storage Defender Sensors observe the systems on which they are installed.

After the inventory is done, you can create the following DRS elements:

- ▶ Recovery groups with resources
- ▶ Profiles:
  - Governance profiles
  - Clean Room profiles

Figure 2-1 shows a high-level overview of DRS.

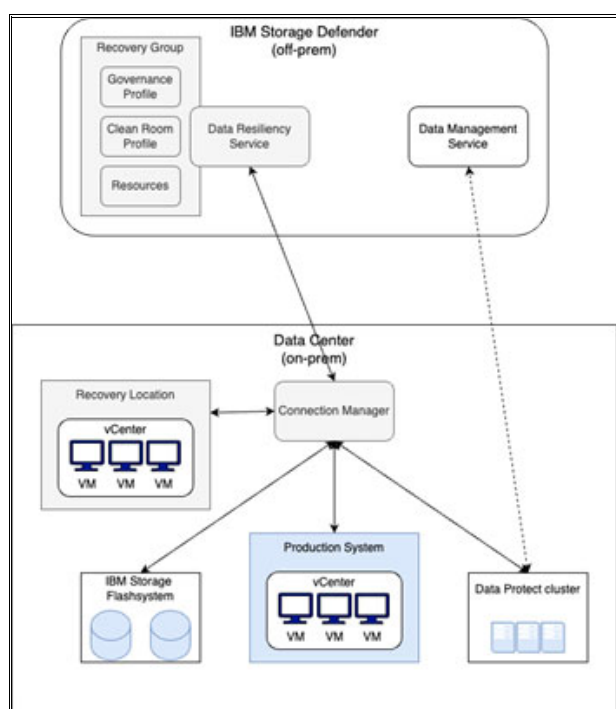


Figure 2-1 DRS: high-level overview

## 2.2 IBM Storage Defender connection manager

The connection manager is installed in your local data center and is used to maintain connections between your local storage systems, backup systems, and your IBM Storage Defender Data Protect cluster. It is also used to orchestrate and direct recovery operations across these resource groups.

You use a connection manager to access your local environment, carry out inventory operations, and perform test recovery and recovery operations by using Data Resiliency. After the connection manager is installed in your local environment, you can log in to the connection manager user interface.

DRS provides the connection manager image in OVA and ISO format, for deployment in an on-prem Data Center or cloud instance. The OVA can be deployed in your local VMware vCenter and the ISO can be deployed on a bare metal server. Inside the connection manager, Red Hat Enterprise Linux is used as the operating system. The Connection Manager software is built to become active and to quickly connect to your local resources and to the IBM Storage Defender DRS that runs in the cloud with less initial configuration before initial use.

Figure 2-2 shows the opening window of the Connection Manager.

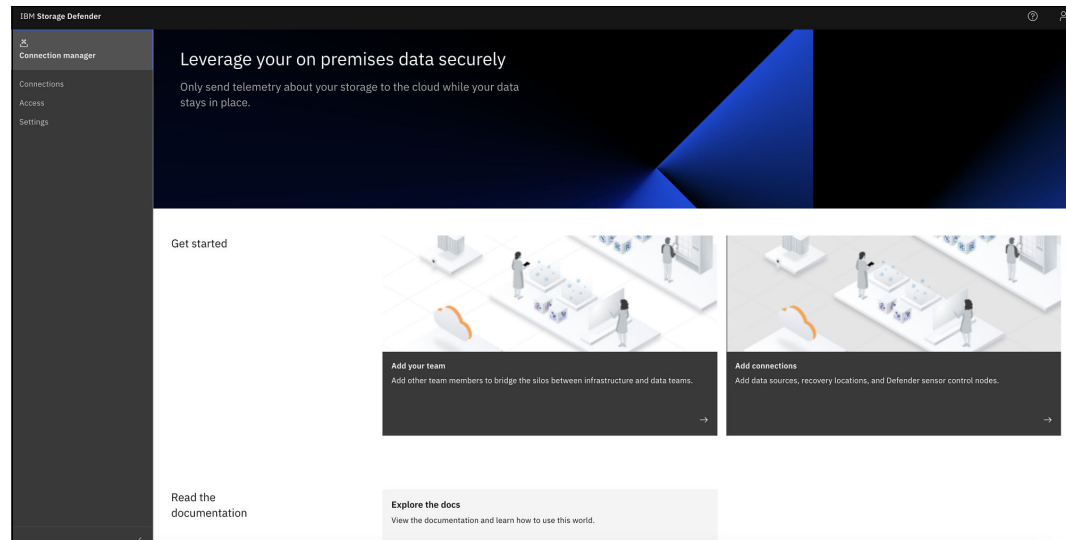


Figure 2-2 IBM Storage Defender Connection Manager

Name	Status	Type	Last Inventory Scan
c2f5200-01.ibmcloudlab.local	Healthy	IBM FlashSystem	Apr 10, 2025, 9:56 AM
c2f57200-01.ibmcloudlab.local	Healthy	IBM FlashSystem	Apr 10, 2025, 9:56 AM
c2f59150-01.ibmcloudlab.local	Healthy	IBM FlashSystem	Apr 10, 2025, 9:56 AM
g3vcenter.ibmcloudlab.local	Healthy	VMware vCenter	Apr 10, 2025, 9:56 AM
sts-10r-dpcluster.ibmcloudlab.local	Healthy	Defender Data Protect	Apr 10, 2025, 9:57 AM
sts-10r-robu.ibmcloudlab.local	Healthy	Defender Data Protect	Apr 10, 2025, 9:56 AM

Figure 2-3 Connection Manager connections list

Connections in Connection Manager include Data sources, Recovery Locations, and sensor control nodes. Typically, deploy only one Connection Manager at each physical location or in each designated pod or zone. You must register data sources to the Connection Manager instance that is in the same physical location.

Connection Manager also includes a job manager, which communicates internally with various workload agents that run in Connection Manager, and also catalogs Safeguarded copies for IBM FlashSystem.

By using a built-in SIEM agent, Connection Manager integrates with on-premises QRadar and Splunk installations to log security events from IBM Storage Defender.

## 2.2.1 Data sources

Data sources that you connect to the IBM Storage Defender Connection Manager are inventoried automatically. The inventory metadata is transferred to the DRS. Connection Manager supports the following data sources:

- ▶ IBM FlashSystem
- ▶ IBM Storage Defender Data Protect

- ▶ VMware vCenter, which includes extended support for IBM Storage Protect for Virtual Environments (SP4VE)
- ▶ IBM Storage Protect Fusion Applications
- ▶ Unisphere for Dell EMC PowerMax
- ▶ IBM Fusion

For IBM FlashSystem, Connection Manager gathers and maintains an inventory, catalogs safeguarded copies and recovery tasks, and restores from backup snapshots.

For IBM Storage Defender Data Protect clusters and VMware vCenters, Connection Manager scans for VMs and protected systems and sends the scan results to DRS. It also coordinates the recovery of VMs that are protected by IBM Storage Defender Data Protect and VMs that are protected by IBM Storage Protect for Virtual Environments (SP4VE).

For IBM Storage Protect, the Connection Manager gathers an inventory of the assets that are protected by IBM Storage Protect for Virtual Environments, so the user can recover VMs through DRS.

DRS support extends beyond traditional virtual machines and databases by adding inventory support for containers hosted on IBM Fusion. Figure 2-4 and shows examples of data sources in Connection Manager:

The screenshot shows the 'Connections' tab in the IBM Storage Defender Connection Manager. It displays a table of data sources with columns for Name, Status, Type, and Last Inventory Scan. There are 6 items listed, all with a status of 'Healthy'. The types include IBM FlashSystem, VMware vCenter, and Defender Data Protect. A search bar and a table control bar are also visible.

Name	Status	Type	Last Inventory Scan
c2b5200-01.ibmcdtlab.local	Healthy	IBM FlashSystem	Apr 10, 2025, 9:56 AM
c2b5700-01.ibmcdtlab.local	Healthy	IBM FlashSystem	Apr 10, 2025, 9:56 AM
c2b5910-01.ibmcdtlab.local	Healthy	IBM FlashSystem	Apr 10, 2025, 9:56 AM
g3vcenter.ibmcdtlab.local	Healthy	VMware vCenter	Apr 10, 2025, 9:56 AM
sts-tor-dpcluster.ibmcdtlab.local	Healthy	Defender Data Protect	Apr 10, 2025, 9:57 AM
sts-tor-robo.ibmcdtlab.local	Healthy	Defender Data Protect	Apr 10, 2025, 9:56 AM

Figure 2-4 IBM Storage Defender Connection Manager: data sources

DRS also supports additional workloads when protected by Defender Data Protect, allowing users to create recovery groups based on specific applications. DRS supports the following workloads when protect by Defender Data Protect:

- ▶ Oracle Databases
- ▶ SAP HANA Databases
- ▶ Microsoft Active Directory Domain Controllers
- ▶ IBM Db2® Databases

DRS collects the inventory of these resources from Defender Data Protect so that users can manually group together and view available recovery points to better understand their overall resiliency posture.

Figure 2-5 on page 13 shows examples of supported resources for creating recovery groups.

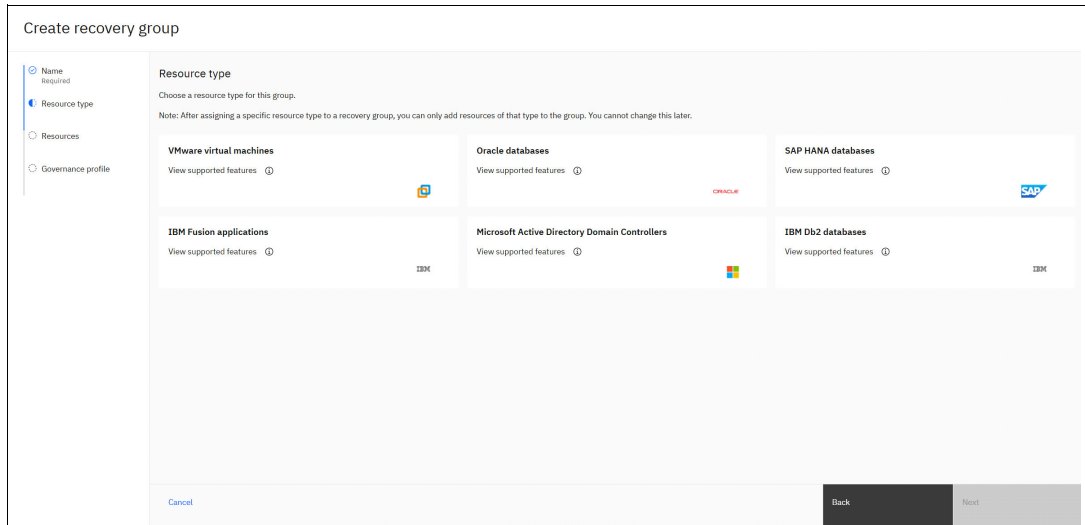


Figure 2-5 IBM Storage Data Resiliency Service: Recovery group resource types

## 2.2.2 Recovery Locations

The *Recovery Locations* concept is used to help recover workloads into an isolated environment. This concept introduces the ability to safely operate on resources that might be contaminated with viruses or other malware without the risk of infecting your production environment. Recovery Locations, like hypervisors that you connect to the IBM Storage Defender Connection Manager, are inventoried automatically. The inventory system metadata is transferred to the DRS.

Figure 2-6 shows the main window of the Recovery Locations.

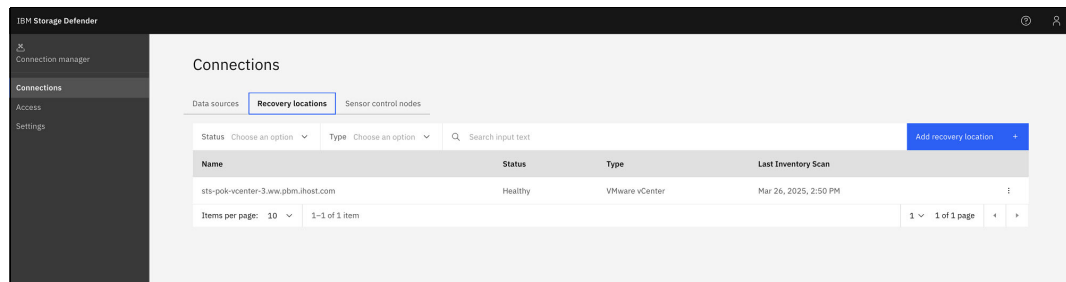


Figure 2-6 IBM Storage Defender Connection Manager: Recovery Locations

Figure 2-7 on page 14 shows an example of the relationship between your production environment and a Recovery Location using IBM FlashSystem and IBM Defender Data Protect as examples.

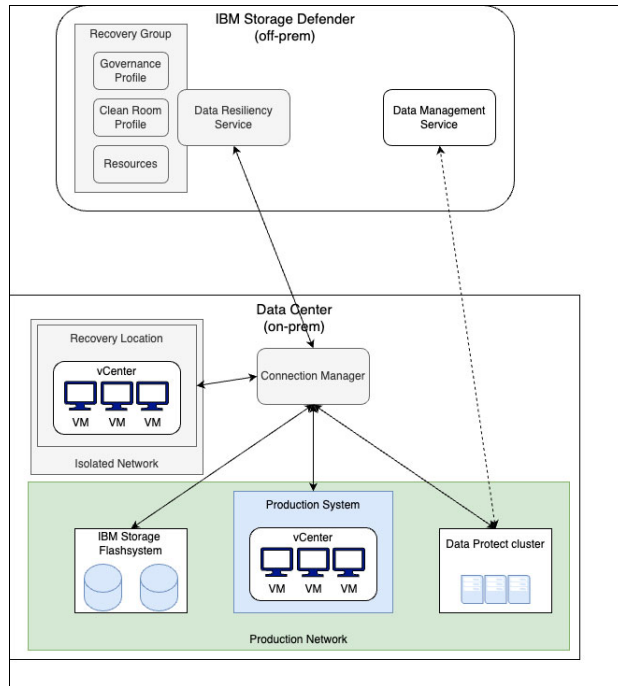


Figure 2-7 Recovery Location example diagram

### 2.2.3 Sensor control nodes

DRS implements the concept of sensor control nodes. The sensor control nodes are used to host the sensor management systems. The sensor management systems are used for sensors that are installed on resources like VMs.

The sensor control node hosts the sensor software and distributes it to the VMs that have sensors that are installed. These sensors observe the systems that they are installed on and can detect cyberattacks, like a ransomware attack, in real time. When the sensor detects a cyberattack, the sensor alerts you by sending messages to the on-premises Connection Manager and DRS.

The Connection Manager comes with a built-in sensor control node so you can start adding sensors right away. However, for larger environments that require deploying many sensors, the use of stand-alone control nodes is suggested. These can be added through the Connection Manager. Additionally, Ansible playbooks are provided to help users easily manage the sensors. Figure 2-8 on page 15 illustrates the sensor control architecture.

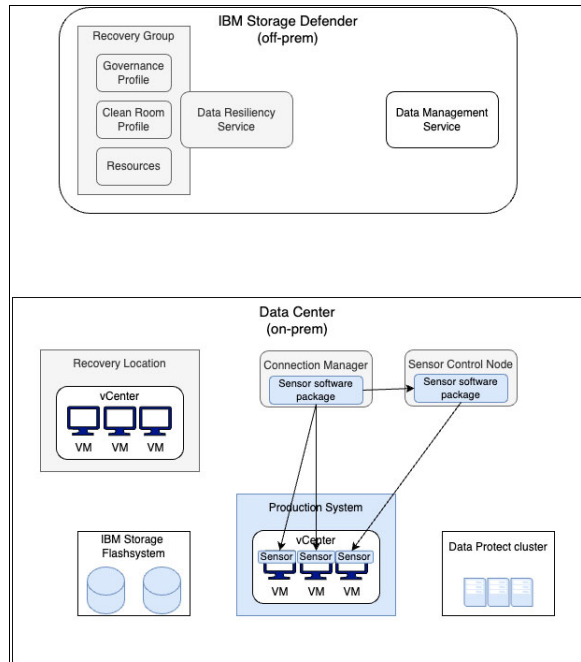


Figure 2-8 Sensor architecture overview

## 2.2.4 IBM Storage Defender Sensors

IBM Storage Defender Sensors implement a real-time detection mechanism for anomalous operations on file system objects for the hosts that they are installed on. IBM Storage Defender Sensors are part of the IBM Storage Defender product and can be deployed on VMs that are part of a recovery group. When the sensors are deployed, the sensors automatically send metadata to the DRS.

A high-level example of the workflow and data path for DRS sensors is shown in Figure 2-9.

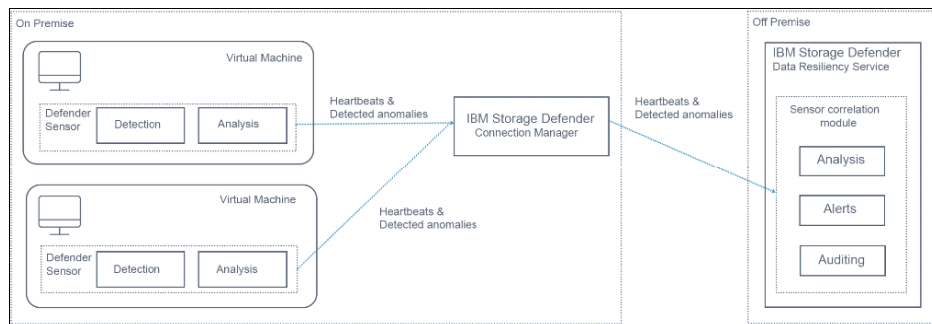


Figure 2-9 Defender sensor workflow

### IBM Storage Defender sensor operation

Storage Defender can collect sensor operation on premises and off premises then uses that information to correlate data, detect and report anomalies, and send notifications.

#### On premises

IBM Storage Defender sensors operate *on premises* by using the following information:

1. When installed, the sensors use file system and operating system interfaces to collect information about operations on file system objects.

2. While collecting this information, sensors analyze this information to identify anomalies for operations on file system objects.
3. Frequently, heartbeat information is sent to the IBM Storage Defender Connection Manager to signal that the sensor is active.
4. When anomalies are detected, the related information is sent to the IBM Storage Defender Connection Manager. A single Connection Manager can have many sensors that report data to it.

### ***Off premises***

IBM Storage Defender sensors operate *off premises* by using the following information:

1. The IBM Storage Defender connection manager reports the sensor data that is collected on premises to the IBM Storage Defender Data Resiliency Service.
2. The Data Resiliency Service correlates the information with recovery groups in your tenant.
3. When sensor heartbeat information is missing or when an anomaly is detected for file system object data operations, a case is opened for the related recovery group.
4. Depending on your notification settings, you are notified about the new case.

### ***Usage of sensor information***

IBM Storage Defender DRS uses the sensor information in the following ways:

1. When installed, the sensors use file system and operating system interfaces to collect information about operations on file system objects
2. The IBM Storage Defender Connection Manager reports the sensor data that is collected on premises to the IBM Storage Defender DRS.
3. The DRS correlates the information with recovery groups across all your Connection Managers.
4. When sensor heartbeat information is missing or when an anomaly is detected for a file system, a case is opened for the related recovery group.
5. Depending on your notification settings, notifications are sent out about the new case.

## **2.3 Recovery groups**

Recovery groups are a core concept within DRS. They include a combination of resources, Governance profiles, and Clean Room profiles. By prioritizing your data, you assign storage resources to a recovery group, which is assigned to a Governance profile and Clean Room profile. When creating the recovery group, DRS evaluates the assigned primary resources to determine whether the associated secondary resources contain corresponding information, such as data protection backups or snapshots of the primary resource.

For example, if a recovery group is assigned VM1, VM2, VM3, and VM4, then DRS determines whether it can find backup snapshots for these VMs in the secondary data sources within the same location (data center). When DRS correlates the primary and secondary resource data for the assigned VMs, it proceeds to test the recovery group based on the policy and Clean Room profile settings.

In the DRS dashboard, you can find the details about recovery groups. Figure 2-10 shows recovery group details.

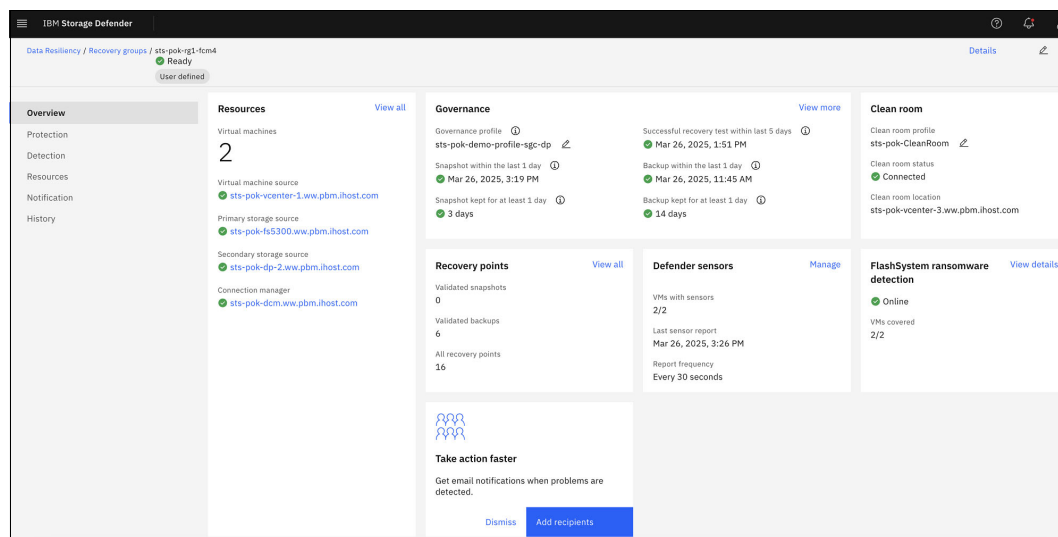


Figure 2-10 DRS recovery group details

## 2.3.1 Auto grouping

Defender can auto-generate recovery groups to organize asset copies for governance, posture analysis, testing, and recovery. The goal of Defender automation is to improve an organizations protection quality and recoverability. Defender is not designed to automate recovery. Instead, it assists the user in understanding what the best recovery points are, what system manages those copies, and how best to perform the needed recovery using that system.

Defender's auto-generated groups can be modified by the users. Additional Groups can also be generated by the user manually. The auto-created groups consider only data assets not otherwise grouped explicitly by users in Defender for inclusion.

Defender supports auto-grouping for the following products:

- ▶ VMware protected by Defender Data Protect
- ▶ Storage Protect for Virtual Environments
- ▶ Oracle
- ▶ SAP HANA
- ▶ Db2
- ▶ Active Directory protected by Defender Data Protect
- ▶ Volumes protected by FlashSystem and SVC

To automatically generate recovery groups navigate to the Recovery groups tab on the left (Figure 2-11 on page 18).

Select Generate Groups, and wait for the process to finish (Figure 2-12 on page 18).

When this process is completed, users can find the generated groups in the Recovery groups list.

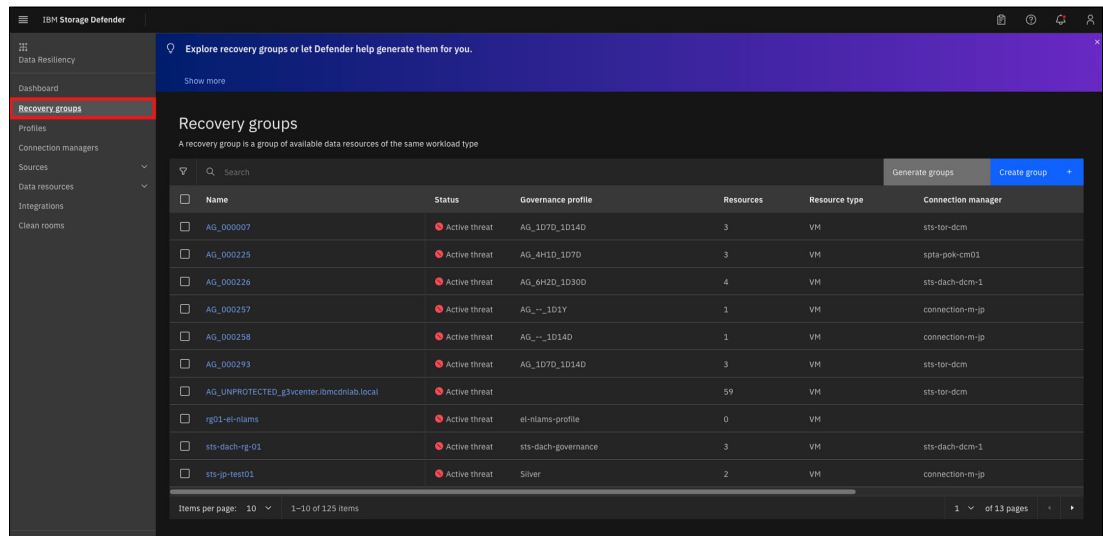


Figure 2-11 Recovery groups section

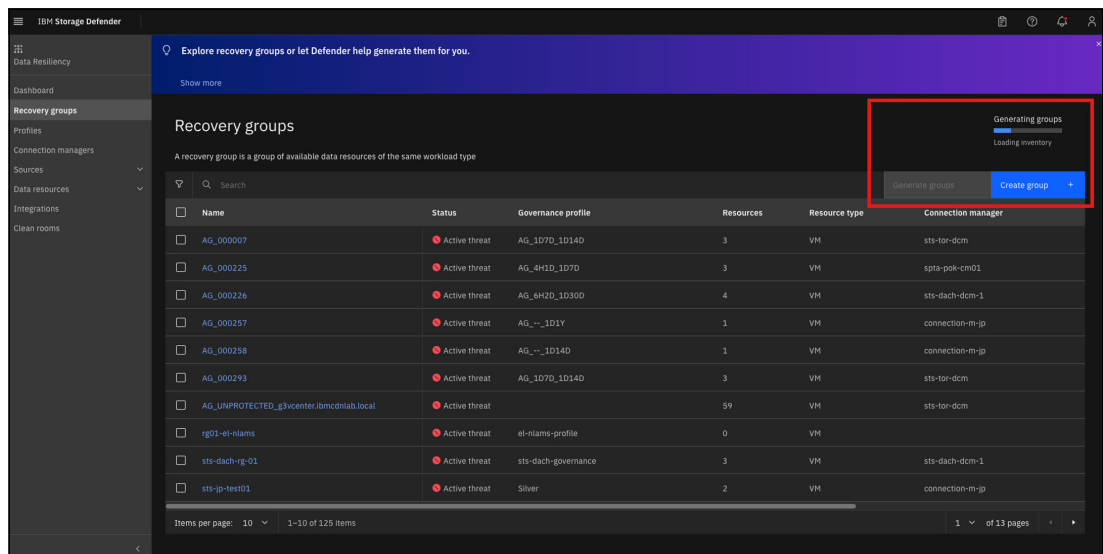


Figure 2-12 Create Group button and progress bar

## 2.3.2 Recovery groups and IBM Storage FlashSystem volume groups

Typically, a datastore in a vCenter relates to a single volume in the IBM Storage FlashSystem. Multiples of such single volumes can be combined in the same volume group of IBM Storage FlashSystem. IBM Storage Defender Data Resiliency Service supports recovery groups that contain virtual machines whose disks belong to the same datastore in the vCenter and a single volume in the IBM Storage FlashSystem volume group. Figure 2-13 on page 19 illustrates a typical setup of a recovery group that contains virtual machines that belong to the same vCenter. The image also illustrates that the datastore in vCenter uses a single volume that is part of a volume group in IBM Storage FlashSystem.

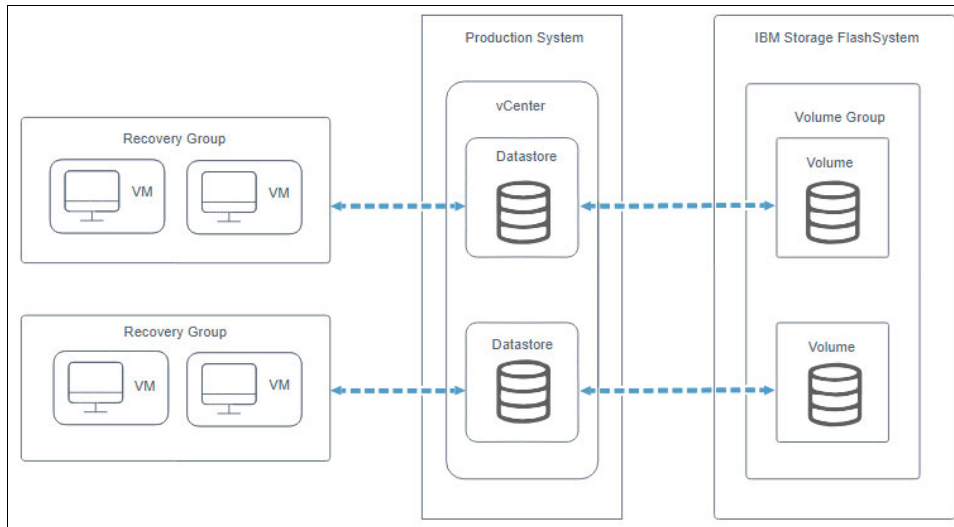


Figure 2-13 Datastore in vCenter uses a single volume, which is part of a volume group in IBM Storage FlashSystem

### 2.3.3 Supported Configurations for IBM Storage FlashSystem

IBM Storage Defender Data Resiliency Service supports various configurations of recovery groups, datastores, and volume groups in an IBM Storage FlashSystem. The following sections contains images that illustrate the supported configurations,

#### Recovery group resources in the same IBM FlashSystem volume

In this configuration (Figure 2-14), all the virtual machines that belong to the recovery group are located in the same VMware datastore. This datastore is located on a single volume in a volume group of the IBM Storage FlashSystem. This implies that a single snapshot policy is used to schedule and create snapshots of the virtual machines in the recovery group.

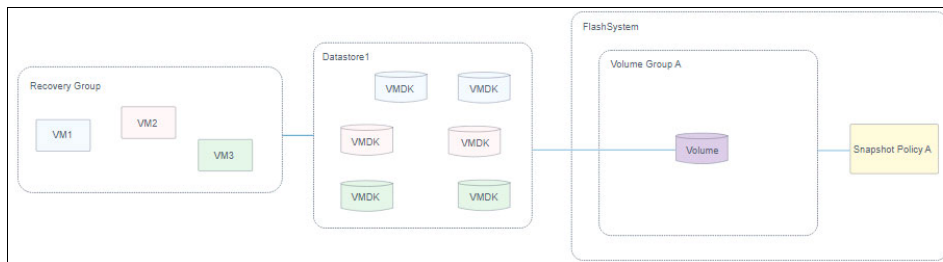
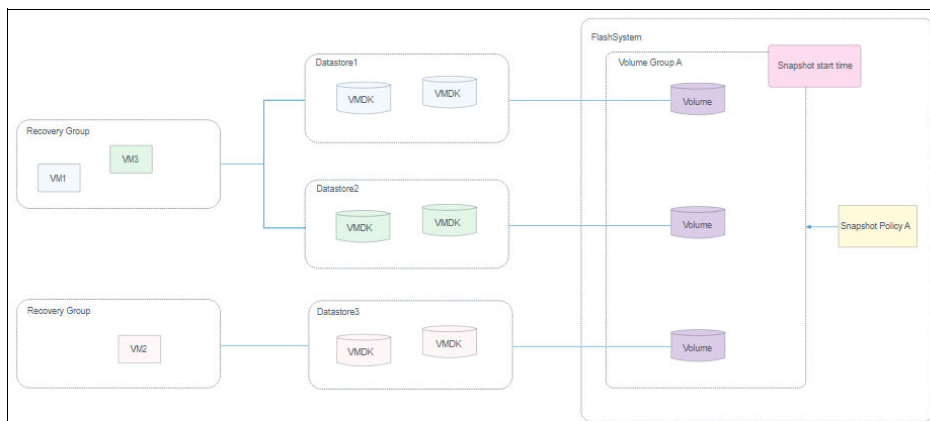


Figure 2-14 Resources of the recovery group are located in the same IBM Storage FlashSystem volume

#### Recovery group resources in the same IBM FlashSystem volume group

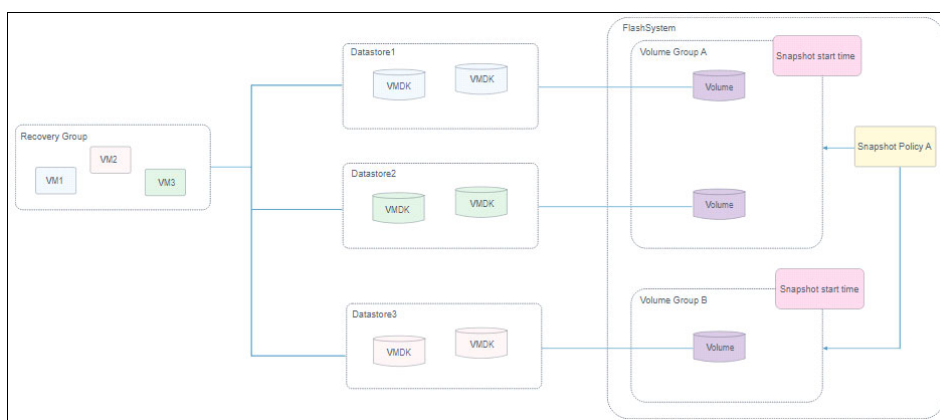
In this configuration (Figure 2-15 on page 20), the virtual machines that belong to the recovery group are located in more than one VMware datastore. The datastores are on single volumes in a volume group of the IBM Storage FlashSystem whereby each datastore has its own volume. This implies that a single snapshot policy is used to schedule and create snapshots of the virtual machines in the recovery group.



*Figure 2-15* Resources of the recovery group are located in the same IBM Storage FlashSystem volume group

### Recovery group resources located in different IBM FlashSystem volume groups

In this configuration (Figure 2-16), the virtual machines that belong to the recovery group are located in more than one VMware datastore. The datastores are on single volumes whereby each datastore has its own volume. The volumes belong to different volume groups of the IBM Storage FlashSystem. The same snapshot policy is used for all volume groups. This implies that a single snapshot policy is used to schedule and create snapshots of the virtual machines in the recovery group.



*Figure 2-16* Resources of the recovery group located in different IBM Storage FlashSystem volume groups

### Recovery group with single or multiple VMs in different IBM FlashSystem volume groups

In this configuration (Figure 2-17 on page 21), a single or multiple VMs can have vmdk files at multiple datastores that are located in different storage volumes and all storage volumes reside within the same IBM Storage FlashSystem volume group. The same snapshot policy is used for all volume groups. This implies that a single snapshot policy is used to schedule and create snapshots of the virtual machines in the recovery group.

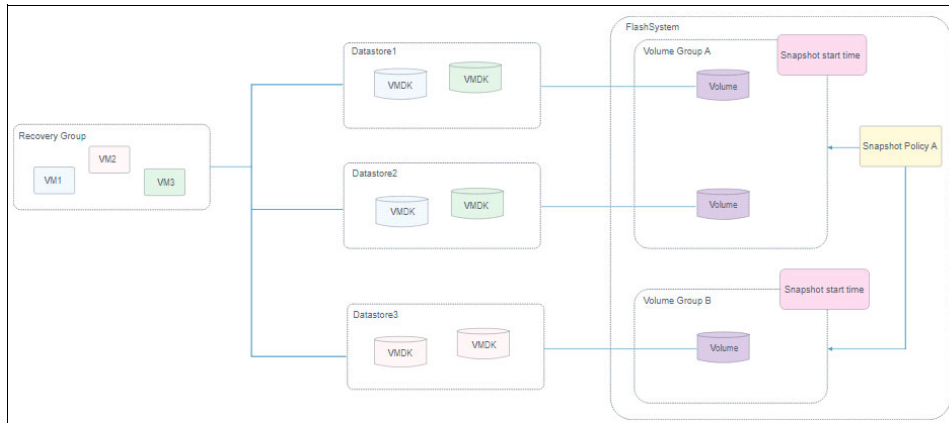


Figure 2-17 Resources of the recovery group with single or multiple VMs located in different IBM Storage FlashSystem volume groups

### 2.3.4 Recovery group and IBM Storage Protect

Data Resiliency Service identifies connected virtual machines (VMs) and relevant data resources of IBM Storage Protect to automatically generate recovery groups. You can also manually assign VMs to a specific recovery group. A recovery group is assigned to a governance profile and clean room profile to recover the assigned resources if there is a cyberattack.

IBM Storage Protect regularly backs up VMs in the production vCenter and stores the backups as snapshots in its storage pools. Each snapshot represents an available recovery point that Data Resiliency Service can recover to a clean room based on the configuration of recovery groups, governance profiles, and clean room profiles.

Figure 2-18 illustrates a typical setup of a recovery group that contains virtual machines that belong to the same vCenter. The image also illustrates that the VM in vCenter uses a single VM that is part of a storage pool in IBM Storage Protect.

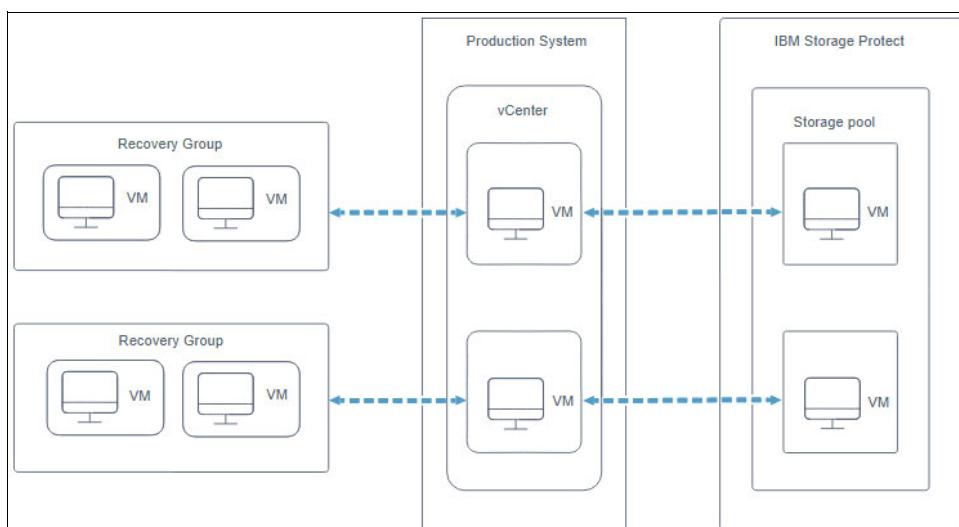


Figure 2-18 Typical setup of a recovery group that contains virtual machines that belong to the same vCenter

## Recovery group and Dell PowerMax storage groups

Typically, a datastore in a vCenter relates to a single volume in the Dell PowerMax storage. Multiple of such single volumes can be combined in the same storage group of Dell PowerMax storage. IBM Storage Defender Data Resiliency Service supports recovery groups that contain virtual machines, where the virtual machines belong to the same datastore in a vCenter and a single volume in the Dell PowerMax storage group. Figure 2-19 illustrates a typical setup of a recovery group for virtual machines that belong to a same vCenter. The image also illustrates that the datastore in vCenter uses a single volume that is part of a storage group in Dell PowerMax storage.

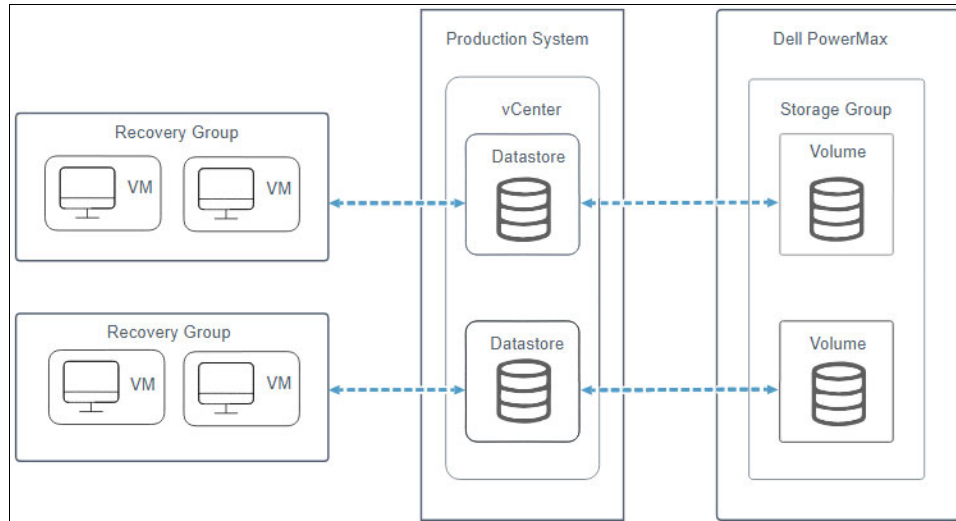


Figure 2-19 Datastore in vCenter uses a single volume that is part of a storage group in Dell PowerMax storage

### 2.3.5 Supported configurations for Dell PowerMax

IBM Storage Defender Data Resiliency Service supports various configurations of recovery groups, datastore, and storage groups in a Dell PowerMax.

#### Recovery group resources in the same Dell PowerMax volume

In this configuration (Figure 2-20), all the virtual machines that belong to the recovery group are located in the same VMware datastore. This datastore is on a single volume in a storage group of the Dell PowerMax. This implies that a single snapshot policy is used to schedule and create snapshots of the virtual machines in the recovery group.

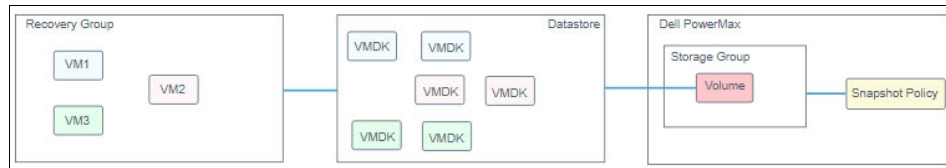


Figure 2-20 Resources of the recovery group located in the same Dell PowerMax volume

#### Recovery group resources distributed to different datastores but in the same Dell PowerMax storage group

In this configuration (Figure 2-21 on page 23), all the virtual machines that belong to the recovery group are located in multiple VMware datastores. Each datastore is on a single volume. The volumes are in the same storage group of the Dell PowerMax. This implies that

a single snapshot policy is used to schedule and create snapshots of the virtual machines in the recovery group.

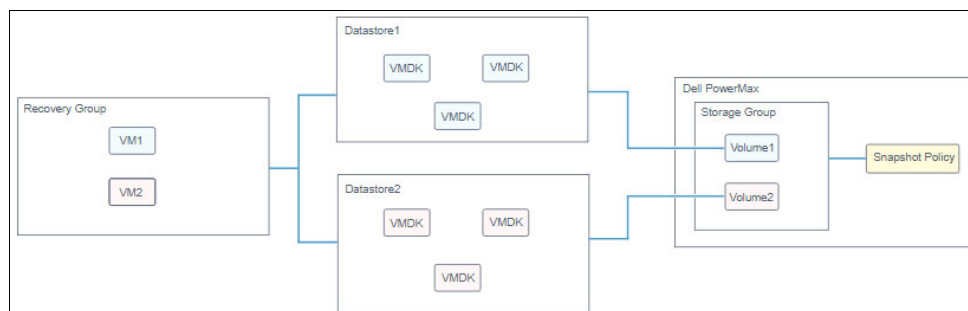


Figure 2-21 Recovery group resources distributed to different datastores but located in the same Dell PowerMax storage group

### Recovery group resources distributed to different datastores located in different Dell PowerMax storage groups

In this configuration (Figure 2-22), all the virtual machines that belong to the recovery group are located in multiple VMware datastores. Each datastore is on a single volume. The volumes are in different storage groups of the Dell PowerMax. All related storage groups share the same snapshot policy. This implies that a single snapshot policy is used to schedule and create snapshots of the virtual machines in the recovery group.

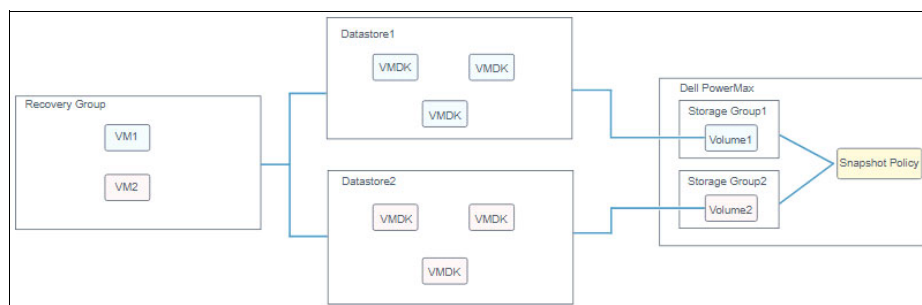


Figure 2-22 Recovery group resources are distributed to different datastores that are located in different Dell PowerMax storage groups

## 2.3.6 Recovery group and Oracle database

To prioritize data for Oracle databases, you must assign an Oracle database resource type to a recovery group that is assigned to a governance profile.

Typically, an Oracle database being protected in a Data Protect protection group creates a new Oracle database copy each time at a protection run. The connection manager gathers data that can be viewed in the available resources panel, the available copies panel, and in the recovery posture graph in the dashboard. You can use IBM Storage Defender Data Resiliency Service to create a recovery group for an Oracle resource type. You can search and select an Oracle database that belongs to a data protect protection group to be added in the recovery group.

## 2.3.7 Recovery group and SAP HANA database

To prioritize data for SAP HANA databases, you must assign an SAP HANA database resource type to a recovery group that is assigned to a governance profile.

SAP HANA system comprises multiple isolated databases and can consist of one host or a cluster of several hosts. SAP HANA system is identified by a single system ID (SID) and contains one or more tenant databases and one system database. Databases are identified by an SID and a database name. From the administration perspective, there is a distinction between tasks performed at system level and those performed at database level. Database clients, such as the SAP HANA cockpit, connect to specific databases.

**Note:** The support is for SAP HANA database with single host.

Typically, you create a backup of SAP HANA into Data Protect, which generates a new copy each time. The connection manager gathers data that can be viewed in the available resources panel, the available copies panel, and in the recovery posture graph in the dashboard. You can use IBM Storage Defender Data Resiliency Service to create a recovery group for SAP HANA resource type. You can search and select an SAP HANA database that belongs to a data protect protection group to be added in the recovery group.

### 2.3.8 Recovery group and Microsoft Active Directory

To prioritize the data for Microsoft Active Directory, you must assign an Active Directory Domain Controller resource type to a recovery group that is assigned to a governance profile.

Typically, you create a backup of Active Directory into IBM Storage Defender Data Protect that generates a new copy at every backup run. The connection manager gathers data that can be viewed in the available resources panel, the available copies panel, and in the recovery posture graph in the dashboard. IBM Storage Defender Data Resiliency Service helps you to create a recovery group for Microsoft Active Directory resource type. You can search and select Microsoft Active Directory that belongs to the data protect protection group to be added in the recovery group.

### 2.3.9 Recovery group and IBM Db2 database

To prioritize data for Db2 databases, you must assign a Db2 database resource type to a recovery group that is assigned to a governance profile.

Typically, you create a backup of the IBM Db2 database into IBM Storage Defender Data Protect that generates a new copy at every backup run. The connection manager gathers data that can be viewed in the available resources panel, the available copies panel, and in the recovery posture graph in the dashboard. IBM Storage Defender Data Resiliency Service helps you to create a recovery group for the IBM Db2 database resource type. You can search and select the IBM Db2 database that belongs to the data protect protection group to be added in the recovery group.

### 2.3.10 Recovery group and IBM Fusion

To prioritize data for IBM Fusion, you must assign a Fusion Application resource type to a recovery group that is assigned to a governance profile.

The IBM Fusion can protect Red Hat OpenShift applications on one or more connected Red Hat OpenShift clusters. The applications are protected by a policy that runs on a defined schedule and retains backups according to the retention configuration settings.

The connection manager collects information about the clusters, applications, backups, policies, and backup locations from the Fusion instance. This inventory data is then sent to

IBM Storage Defender Data Resiliency Service where it can be viewed in the recovery posture graph in the dashboard, available resources, and available copies. A recovery group can be created for the Fusion Application type. This recovery group can be associated with a governance profile to help ensure that the backups are occurring on the defined schedule and are kept for long enough.

The connection manager collects the information about the clusters, applications, backups, policies, and backup locations from the IBM Fusion instance. The collected inventory data is sent to IBM Storage Defender Data Resiliency Service, where a user can view the data in the available resources panel, the available copies panel, and in the recovery posture graph in the dashboard. The user can create a recovery group for the Fusion Application type. To take backups and retain the backup copies as needed, the user can associate the recovery group with an appropriate governance profile.

## 2.4 Profiles and IBM Clean Room

Profiles are used to define and set the recovery objectives. IBM Clean Room provides an isolated environment for the recovery of workloads.

### 2.4.1 Governance and Clean Room profiles

Use the Governance and Clean Room profiles to define and set the recovery objectives of recovery groups and recovery target environments.

#### **Governance profiles**

Governance profiles are created and applied to the recovery group so that specific recovery objectives can be defined and associated with one or more groups. These recovery objectives are composed of preset points in time for the recovery points and the preset minimum retention time for the recovery points. The Governance profile can specify a threshold time that must elapse before the next recovery test is performed for the recovery group. Separate recovery objectives can be defined for IBM Storage FlashSystem and IBM Storage Defender Data Protect, or IBM Storage Protect for Virtual Environments (SP4VE) independently.

The Governance profile definition enables one of the following use case definitions:

- ▶ Observation of the recovery objectives for IBM Storage FlashSystem recovery points (Safeguarded snapshot copies)
- ▶ Observation of the recovery objectives for IBM Storage Defender Data Protect or SP4VE recovery points
- ▶ Observation of both the recovery objectives for IBM Storage FlashSystem recovery points and IBM Storage Defender Data Protect / SP4VE recovery points

The test frequency objective is optional for all use cases.

Figure 2-23 on page 26 shows an overview of recovery objectives that are configured in the Governance profile.

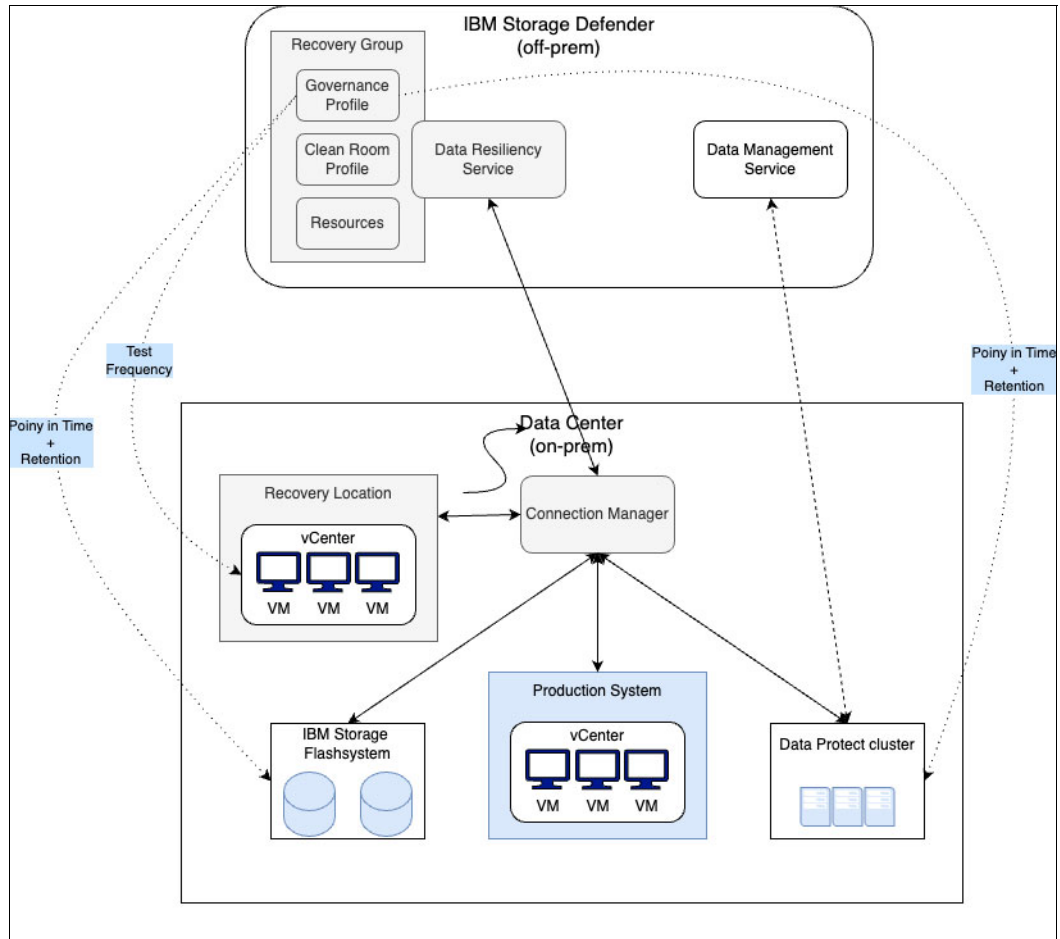


Figure 2-23 Recovery objectives that are configured in the Governance profile

## Clean Room profiles

The Clean Room profiles connect the recovery groups that belong to resources in the production environment with configuration and resources that are defined in DRS. The connected resources are IBM Storage FlashSystem, IBM Storage Defender Data Protect, and the Clean Room environment. This resource configuration defines how IBM Storage Defender behaves during a recovery event.

Figure 2-24 illustrates the logical connection between the different components.

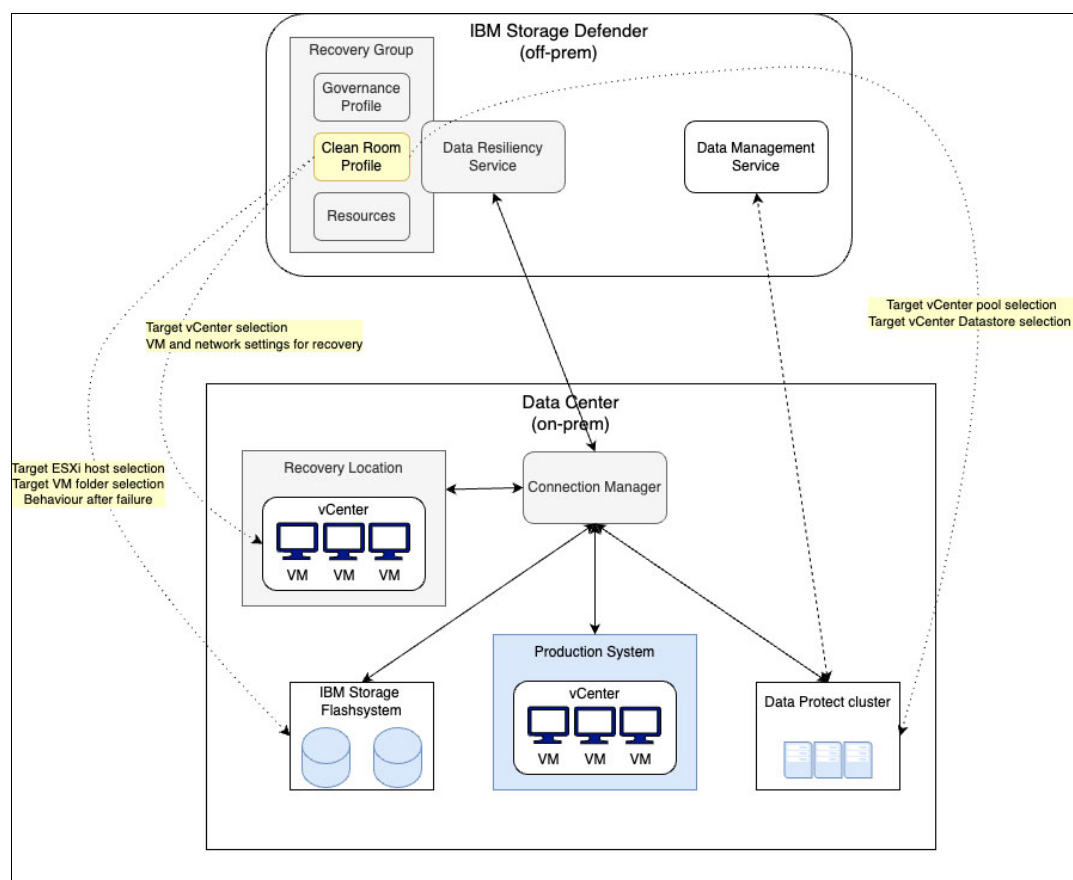


Figure 2-24 Clean Room objectives that are configured in the Clean Room profile

To help ensure the successful recovery of the recovery group that is assigned to the specific Clean Room profile, configuration requirements must be met. After you configure the Clean Room profile, you can use it for one of the following three use cases:

1. Recovery from IBM Storage FlashSystem Safeguarded snapshots
2. Recoveries from IBM Storage Defender Data Protect backup copies
3. Recovery from both IBM Storage FlashSystem Safeguarded snapshots and IBM Storage Defender Data Protect / SP4VE backup copies

**Important:** If these requirements are not met, the recovery of the VMs that belong to the specific recovery group fail for Clean Room recoveries.

In addition to the conceptual dependencies between the Clean Room profile and other IBM Storage Defender components, consider that the same Clean Room profile can be reused for different recovery groups. In cases where a Clean Room is associated with multiple recovery groups, the different recovery groups might have different requirements for their recovery. This situation is important when recovering from an IBM Storage FlashSystem because the requirements for network infrastructure, mapping of volumes, or SAN zoning might be different. Therefore, it might be beneficial to implement multiple Clean Room profiles with different configurations to provide more flexibility for the recovery scenarios that you want to implement for different recovery groups.

## Resources

All available resources that are managed by DRS and inventoried with Connection Manager are shown in the DRS GUI. DRS supports the following resources:

- ▶ VMs
- ▶ Connection managers
- ▶ Data sources
- ▶ Clean Rooms

Resources are added to a recovery group during its creation, and are checked during inventories by the Connection Manager.

### 2.4.2 IBM Clean Room

IBM Clean Room plays an important role in the IBM Storage Defender solution by enabling the recovery of workloads into an isolated environment. By using it, you can safely restore and investigate resources that might be contaminated with viruses or other malware without the risk of infecting your production environment.

Protected VMs can be recovered into an associated Clean Room for verification before their recovery into a production environment. IBM Storage Defender is connected to each VM instance and provides observation and assistance with this process.

A Clean Room environment setup has several similarities with a standard vCenter setup. Apart from the recovery groups that are restored by using data stores that are mapped from data protection solutions, a DMZ is implemented to enable access to the isolated portions of a Clean Room.

Figure 2-25 displays the high-level structure of a Clean Room environment.

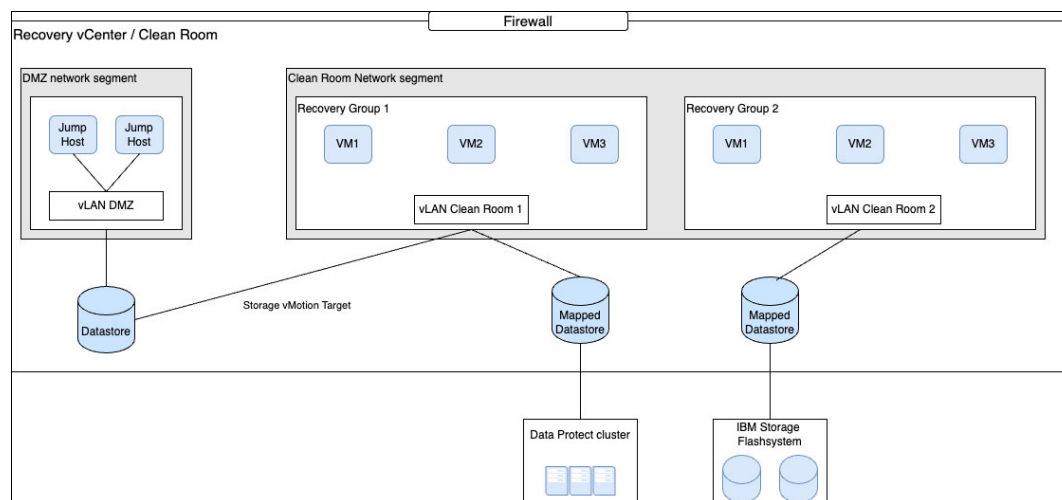


Figure 2-25 Clean Room environment schema

Isolation is an important aspect to consider when implementing Clean Room functions. There are multiple dimensions, such as isolation of infrastructure, network, and access management. In addition to isolation, you must monitor and log a Clean Room environment. The following sections describe the different aspects of isolation for a Clean Room environment.

## Infrastructure isolation

Isolating the infrastructure is an important aspect of a Clean Room environment. Isolation for physical resources refers to physical separation, where you have computer hardware that is used for a hypervisor and is independent of any production environment. When you use a cloud service provider, isolation refers to a logical separation that is configured by using different cloud accounts.

## Network segmentation and monitoring

Network segmentation comprises multiple aspects:

- ▶ **Logical separation and subnetting:** In addition to the recovered VMs, the Clean Room environment contains systems that are used for tools and management. Separate groups of systems into network segments to prevent the breakout of malware from infected systems. If multiple recovery groups are recovered into the same Clean Room to establish a temporary production environment, use a dedicated VLAN for each recovery group. Apart from the breakout prevention, the positive impact of the administrative separation of duty is another important benefit to this planning step.
- ▶ **Access control and firewalls:** Use firewalls and access control lists (ACLs) to control and monitor traffic between network segments. Also, enhance security by enforcing rules that are based on source, destination, and port.
- ▶ **Security zones and critical infrastructure protection:** Establish security zones, including a DMZ to separate public-facing servers and protect critical infrastructure components by limiting potential attack vectors.
- ▶ **Monitoring, encryption, and regular auditing:** Implement network monitoring tools and centralized logging to help ensure visibility and timely detection of security incidents. Also, implement secure communication between recovery groups in the same Clean Room. If applications require interaction, you can use VPNs and encryption. If the Clean Room is used for temporary production, conduct regular security audits to confirm all security measures are still valid and providing the expected protection.

## Identity management and logging

Implementing administrative separation of a Clean Room environment from a production environment helps provide an extra layer of security. This implementation can range from using a different set of administrative identities to a total separation of identity management in a separate directory service.

The logical separation of administrative roles for the production system and the Clean Room environment and strict limits on a user's permissions prevent a user from influencing both environments.

The implementation of auditable logging for all operations in the Clean Room helps ensure that any operation on the recovered data is traceable. This implementation includes the creation and configuration of the Clean Room; Clean Room operations, such as recovery, data masking, and anonymization; or temporary production usage of the data.

## Compliance and legal compliance

The bounded usage scope of a Clean Room environment enables comprehensive documentation of all operations in the Clean Room. The addition of the auditable logging in the configuration of the Clean Room environment allows for an event chain to be present and maintained to help ensure that proper procedures were followed or evaluated during a post-event review. With logging, the usage scope expands to include actions such as temporary production use or test recovery on the data in the Clean Room. These operation logs enable analysis or development, and you can use them to document events or actions

that are taken. A comprehensive review of this documentation can help you audit the regulatory compliance status of a company and confirm whether the requirements are being met.

For more information, see [IBM Storage Defender: Clean Room environments](#).

## 2.5 Adding resources in the Connection Manager and creating profiles in DRS

The following section describes how you can add resources in the IBM Storage Defender Connection Manager and how you can create profiles in DRS with these resources.

### 2.5.1 Adding resources in Connection Manager

After you deploy Connection Manager, you can log in to the Connection Manager GUI and add resources that you manage through DRS.

To log in to Connection Manager, enter the following link into a browser:

`https://<ConnectionManager IP or hostname>/login`

A login page opens, where you enter your username and password (Figure 2-26). Confirm the login by entering the confirmation code from your authentication application.

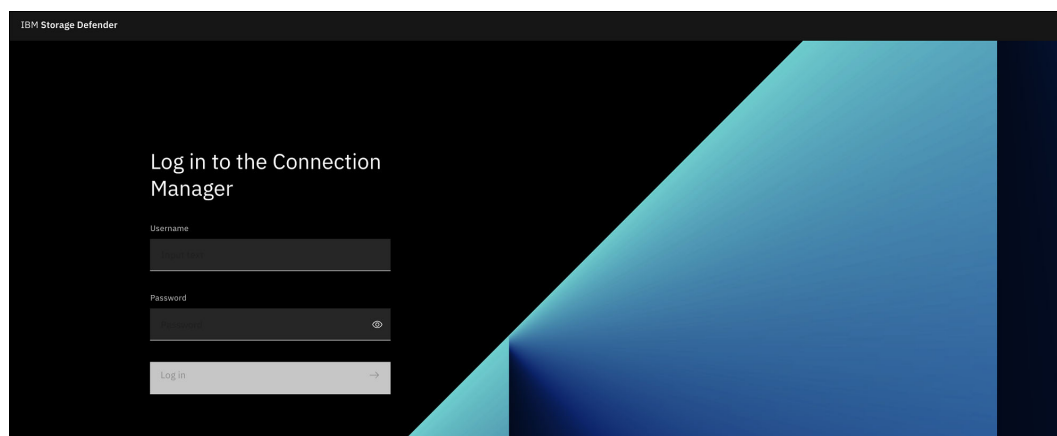


Figure 2-26 Connection Manager login

Figure 2-27 on page 31 shows the Connection Manager dashboard. You can use this dashboard to add resources from the **Connections** menu.

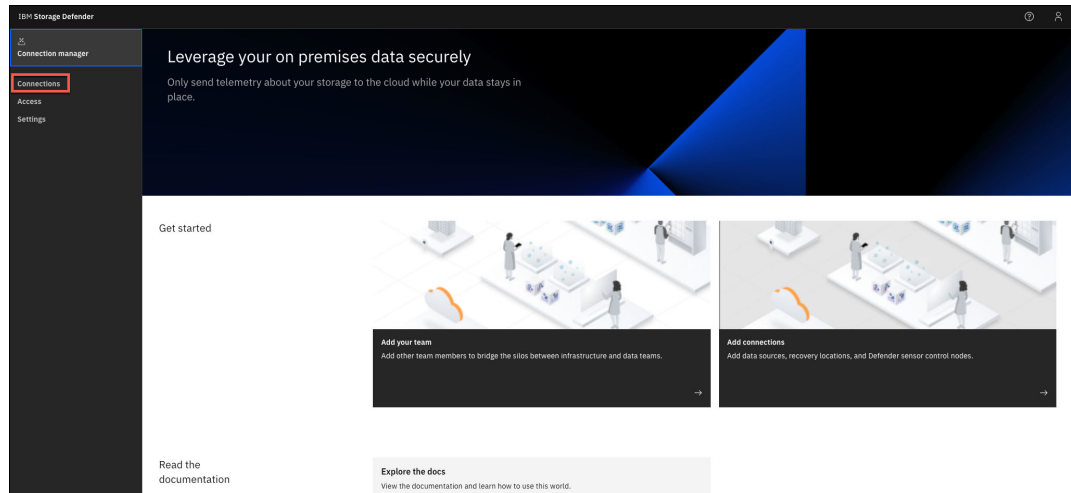


Figure 2-27 Connection Manager: Connections

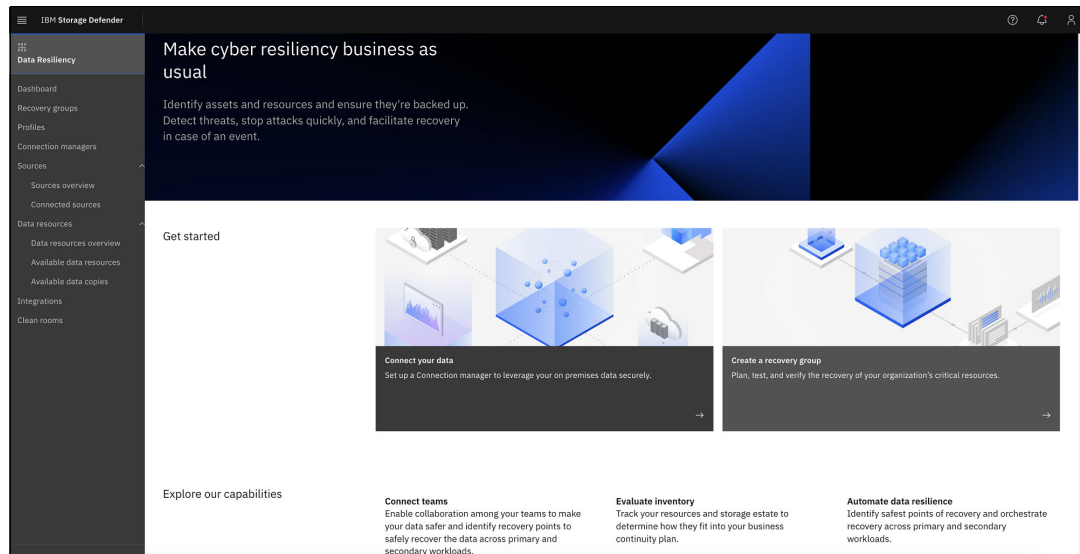


Figure 2-28 Connection Manager: Connections list

From the Connections dashboard data sources, you can add Recovery Locations and sensor control nodes to the DRS configuration.

## Adding data sources

To add data sources in the Connections dashboard, complete the following steps:

1. Select the **Data sources** tab and click **Add a data source**. A wizard opens in the right pane. It guides you through the process (Figure 2-29 on page 32). Select the type of data source that you want to add and click **Next**.

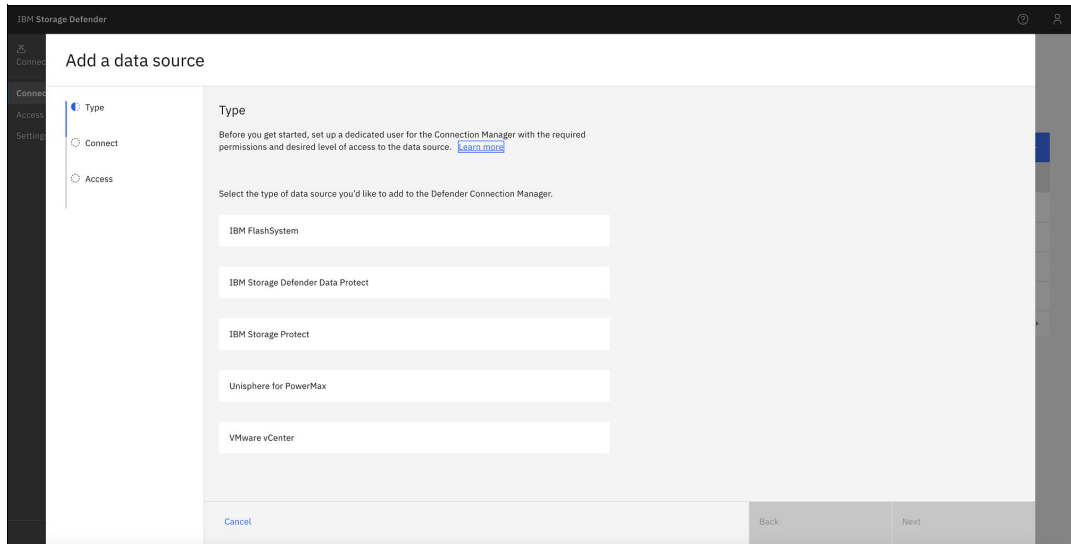


Figure 2-29 Connection Manager: Add a data source pane

2. Enter the hostname or IP address of the data source and click **Next**. (Figure 2-30).

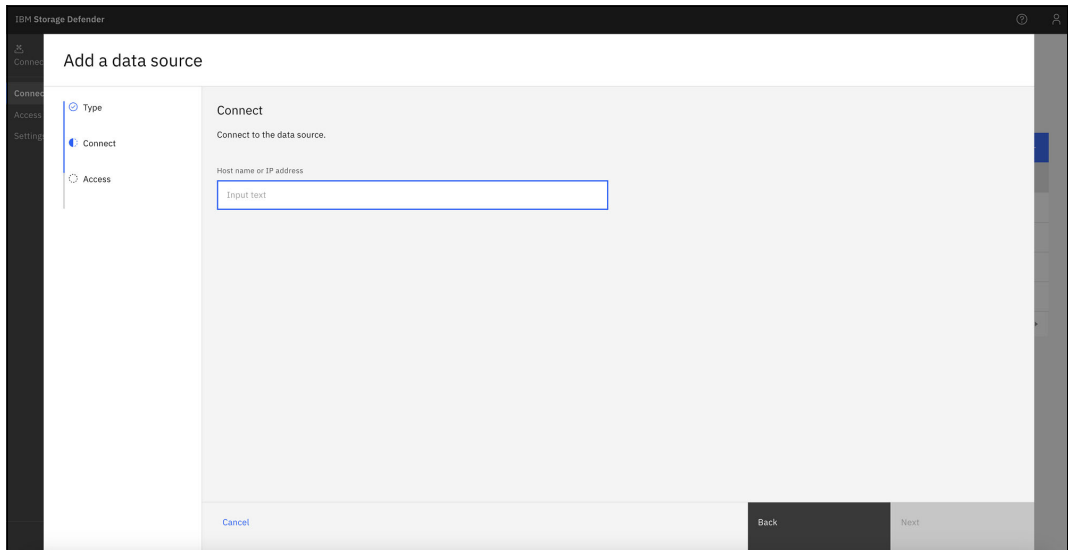


Figure 2-30 Connection Manager: Add a data source details pane

3. Review the certificate details and click **Next**. (Figure 2-31 on page 33)

**Add a data source** ✕

Type **Connect** Access

Connect to the data source.

Host name or IP address  
10.208.100.55

Issuer name  
Expiry  
Thu Oct 10 08:32:24 UTC 2024  
Common Name  
78E02HN  
Fingerprint (SHA256)  
90eb5683e2e838470601b6fa8743676302528dacc6bc9f  
1550065b2a64301444

**Self-signed certificate**  
This data source is using a self-signed certificate. Verify the certificate details before accepting by clicking **Next**.

Back Next

Figure 2-31 Connection Manager: Add a data source details pane

Enter the credentials to be used by Connection Manager to access this data source (Figure 2-32).

**Add a data source** ✕

Type Connect **Access**

These credentials will be used by the Connection Manager to access the data source. Make sure you're using dedicated credentials with the required permissions and desired level of access. [Learn more](#)

Username  
superuser

Password  
.....

Back Add

Figure 2-32 Connection Manager: Add a data source credentials pane

4. Click **Add**. Once the process completes, the new data source is added to Connection Manager, as shown in Figure 2-33.

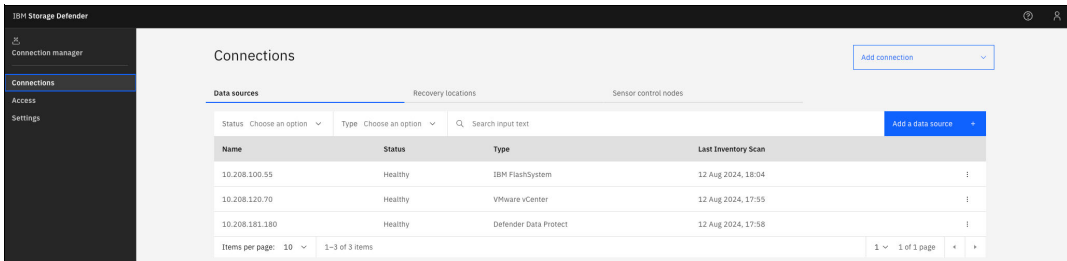


Figure 2-33 Connection Manager: New data source

## Adding Recovery Locations

To add Recovery Locations in the Connections dashboard, complete the following steps:

1. Select the **Recovery locations** tab and click **Add recovery location**. A wizard opens in the right pane. Enter the hostname or IP address of the VMware vCenter that you want to add and click **Next** (Figure 2-34).

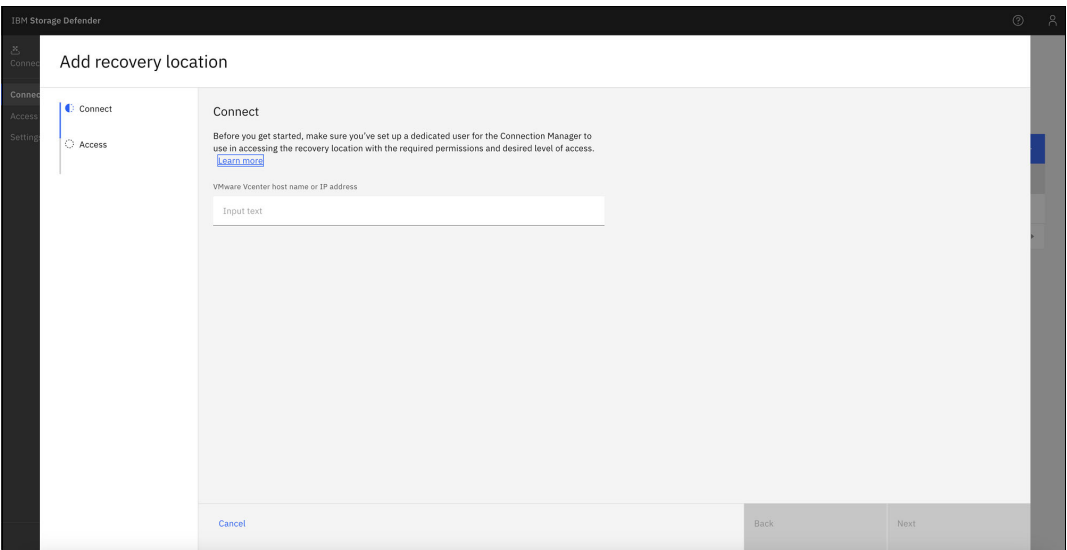


Figure 2-34 Recovery Location: Adding recovery location pane

2. Review the certificate details and click **Next** (Figure 2-35 on page 35).

**Add recovery location** ✕

**Connect** Access

Before you get started, make sure you've set up a dedicated user for the Connection Manager to use in accessing the recovery location with the required permissions and desired level of access. [Learn more](#)

VMware Vcenter host name or IP address

10.208.120.60

Issuer name

Expiry  
Sat Jun 27 20:39:25 UTC 2026

Common Name  
CA

Fingerprint (SHA256)  
9329cf4139a6fba00a25ac5282ecb182db2e12f95af058538e28ca8ed131ccb2

**Self-signed certificate** ✕  
This data source is using a self-signed certificate. Verify the certificate details before accepting by clicking **Next**.

Cancel **Next**

Figure 2-35 Recovery Location: Certificate details panel

Enter the dedicated credentials with the required permissions and the level of access to the environment for this Recovery Location and click **Add** (Figure 2-36).

**Add recovery location** ✕

**Connect** Access

Enter the dedicated credentials with the required permissions and desired level of access to this recovery location. [Learn more](#)

Username

administrator@ssc1ab.ibm.si

Password

.....

Back **Add**

Figure 2-36 Recovery Location: Add location panel

Figure 2-37 shows that the new Recovery Location was successfully added to the Connections list.

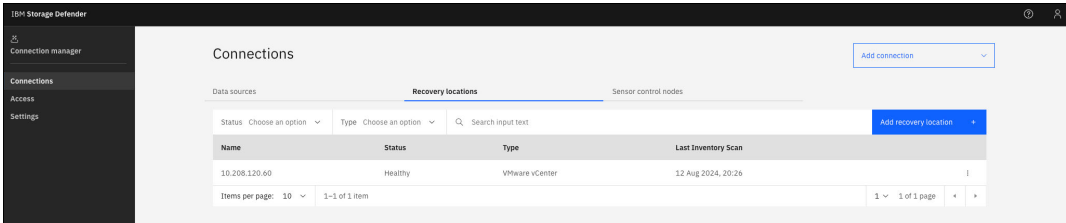


Figure 2-37 Add Recovery Location

Adding sensor control nodes

Connection Manager comes with a built-in control node. If you want to use your own control nodes, you can add them through the Connection Manager GUI and use the provided Ansible playbooks to manage the sensors.

To add a control node into the Connections dashboard, complete the following steps:

- 1. Select the **Sensor control nodes** tab and click **Add control node** (Figure 2-38). A wizard opens in the right pane. Enter the Ansible control node hostname, and click **Next**.

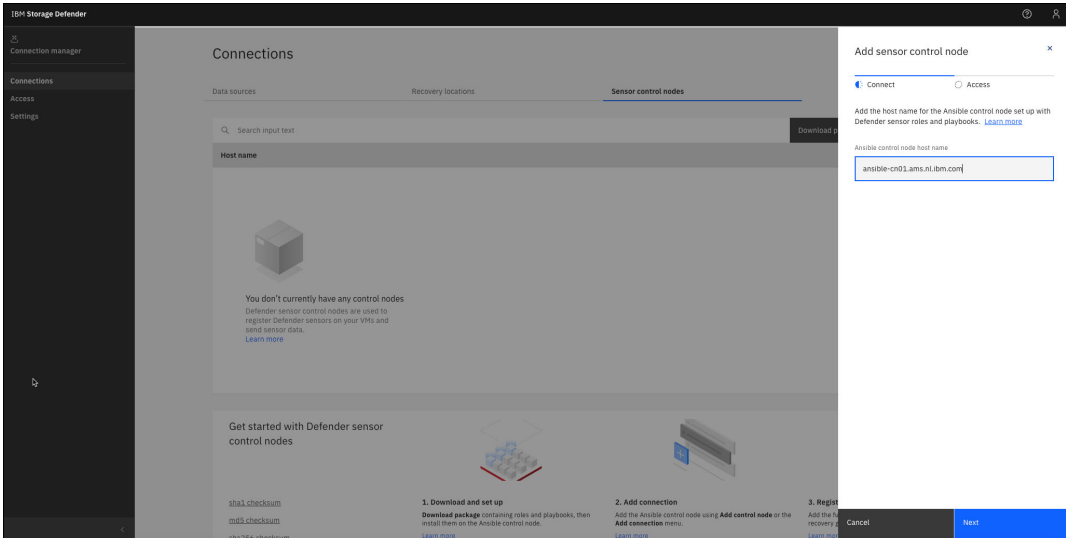


Figure 2-38 Add sensor control node pane

- 2. Enter the credentials that you created on the Ansible control node during the IBM Storage Defender sensor setup and click **Add** (Figure 2-39 on page 37).

**Add sensor control node** [X]

**Connect** [Access]

Use the credentials created on the Ansible control node during the Defender sensor setup. [Learn more](#)

Username  
admin

Password  
[Masked Password]

Back Add

Figure 2-39 Add sensor control node pane

The new sensor control node was added to Connection Manager (Figure 2-40).

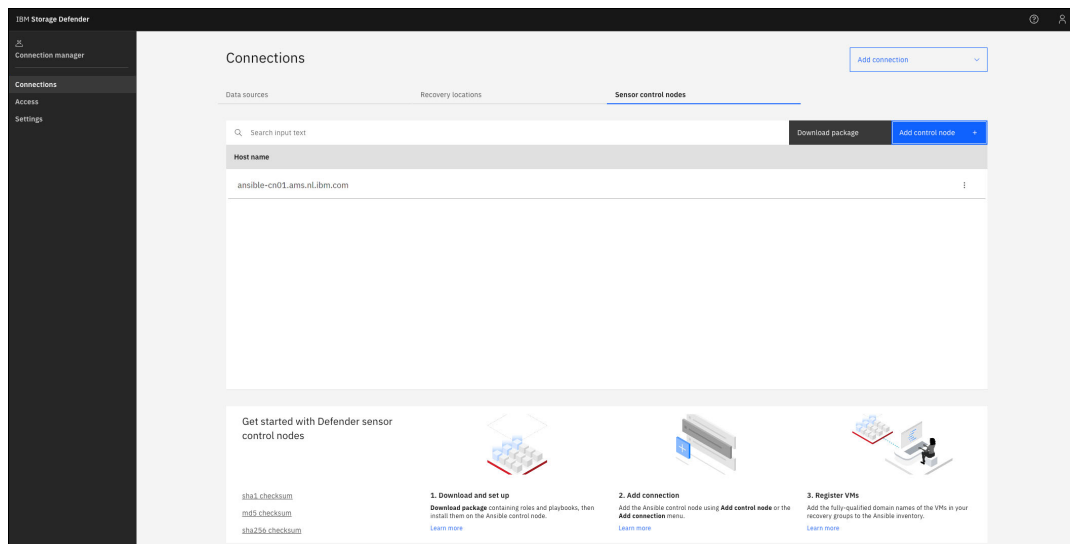


Figure 2-40 Sensor control node connections list

## 2.5.2 Creating profiles in DRS

Profiles that are used in DRS are assigned to recovery groups and consist of Governance and Clean Room components. Policy Governance profiles can be assigned to recovery groups to monitor alignment with your backup policies and recovery point objectives. Clean Room profiles are used by a recovery group to specify the Clean Room location and setting

that is needed to recover data. To create profiles in DRS, select **Profiles** from the menu, which opens the Profiles dashboard.

To create a Governance profile, complete the following steps:

1. Select the **Governance** tab and click **Create profile**. The Create governance profile window opens (Figure 2-41). Under the **Details** tab, enter the name for a Governance profile and its description, and then click **Next**.

The screenshot shows the 'Create governance profile' window with the 'Details' tab selected. The window has a title bar with a close button. Below the title bar are four tabs: 'Details' (active), 'Immutable snapshots Optional', 'Backups Optional', and 'Recovery testing Optional'. The main content area contains the instruction 'Give this profile a name and a brief description of when it should be used.' followed by a 'Name' field with the text 'Demo' and a 'Description' text area. At the bottom are 'Cancel' and 'Next' buttons.

Figure 2-41 Creating a Governance profile window

2. When creating a Governance profile, use the **Immutable Snapshots** tab to select thresholds for immutable snapshot recovery points that are available from the IBM FlashSystem server (Figure 2-42). Select the checkbox to enable point in time verification and retention time verification for the specified time interval. Click **Next**.

The screenshot shows the 'Create governance profile' window with the 'Immutable snapshots' tab selected. The window has a title bar with a close button. Below the title bar are four tabs: 'Details', 'Immutable snapshots' (active), 'Backups Optional', and 'Recovery testing Optional'. The main content area contains the instruction 'Select the thresholds for immutable snapshot recovery point objectives.' followed by two sections. The first section has a checked checkbox 'At least one immutable snapshot within the last' and a field with '1' and a 'Days' dropdown. The second section has a checked checkbox 'Immutable snapshots kept for at least' and a field with '1' and a 'Days' dropdown. At the bottom are 'Back' and 'Next' buttons.

Figure 2-42 Governance profile: Immutable snapshots verification window

- Under the **Backups** tab, select thresholds for backup copy recovery points that are available from IBM Storage Defender Data Protect (Figure 2-43). Select the checkbox to enable point in time verification and retention time verification for the specified time interval. Click **Next**.

The screenshot shows the 'Create governance profile' window with the 'Backups' tab selected. The window has a title bar with a close button. Below the title bar is a tab bar with four tabs: 'Details' (selected), 'Immutable snapshots' (Optional), 'Backups' (Optional), and 'Recovery testing' (Optional). The main content area is titled 'Select the thresholds for backup recovery point objectives.' and contains two sections. The first section is 'At least one backup within the last' with a checkbox checked, a numeric input field set to '1', and a unit dropdown menu set to 'Days'. The second section is 'Backups kept for at least' with a checkbox checked, a numeric input field set to '1', and a unit dropdown menu set to 'Days'. At the bottom of the window are two buttons: 'Back' and 'Next'.

Figure 2-43 Governance profile: Backups verification window

- In the **Recovery testing** tab, select the thresholds for successful recovery testing. Select the checkbox to enable test frequency verification and specify time interval (Figure 2-44).

The screenshot shows the 'Create governance profile' window with the 'Recovery testing' tab selected. The window has a title bar with a close button. Below the title bar is a tab bar with four tabs: 'Details' (selected), 'Immutable snapshots' (Optional), 'Backups' (Optional), and 'Recovery testing' (Optional). The main content area is titled 'Select the thresholds for successful recovery testing.' and contains one section: 'At least one successful test every' with a checkbox checked, a numeric input field set to '1', and a unit dropdown menu set to 'Days'. At the bottom of the window are two buttons: 'Back' and 'Create'.

Figure 2-44 Governance profile: Recovery Testing window

- Click **Create**. A new Governance profile is created.

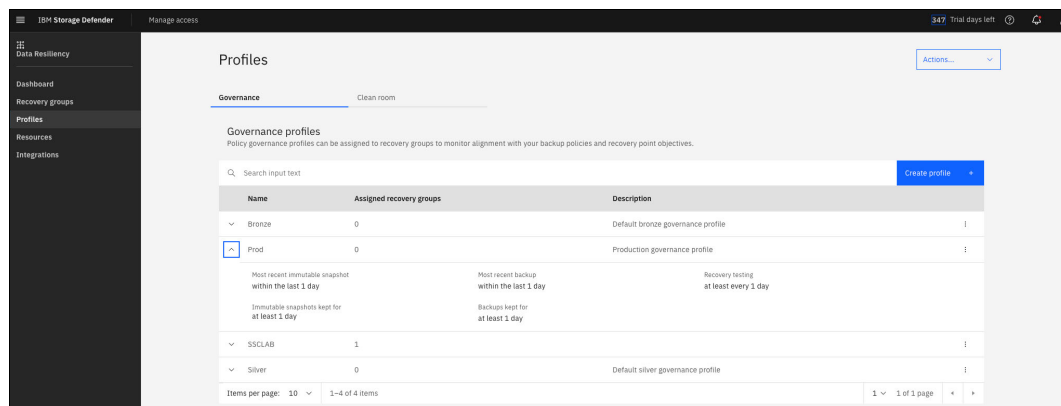


Figure 2-45 New Governance profile created in the profiles list

To create Clean Room profile, complete the following steps:

1. Select the **Clean Room** tab and click **Create profile** (Figure 2-46).

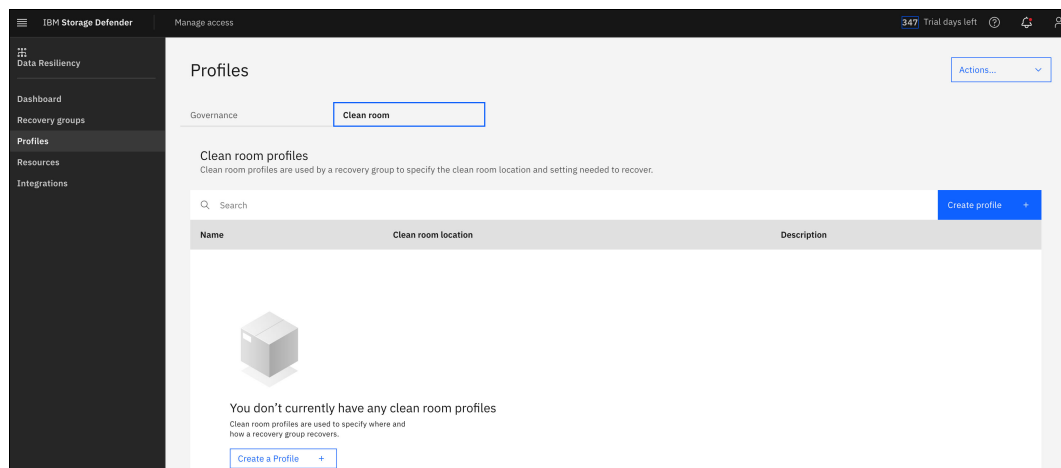


Figure 2-46 Creating a Clean Room profile

2. In the Create Clean Room profile window, under the **Details** tab, specify the name for a Clean Room profile and provide a description of the Clean Room profile (Figure 2-47). Click **Next**.

Create clean room profile

Details

Clean room settings

Immutable snapshot recovery  
Optional

Backup recovery  
Optional

Give this clean room profile a name a brief description of when it should be used.

Profile name

ssclab-CleanRoom

Description (optional)

ssclab clean room profile

Cancel

Next

Figure 2-47 Clean Room profile details window

3. Under the **Clean room settings** tab, enter the Clean Room location and recovery preferences (Figure 2-48). The settings under this tab are global in the context of the profile and influences the recovery from IBM Storage FlashSystem and IBM Storage Defender Data Protect. Click **Next**.

Create clean room profile

Details

Clean room settings

Immutable snapshot recovery  
Optional

Backup recovery  
Optional

Choose a clean room location and the preferences for recoveries using this profile

Clean room location

EN20-vcenter2.sscslab.ibm.si

Power state ⓘ

Off

Attach to network ⓘ

Off

Previous

Next

Figure 2-48 Clean Room settings window

4. Under the **Immutable snapshot recovery** tab, enter your recovery preferences when recovering from immutable snapshots with IBM Storage FlashSystem (Figure 2-49).

Create clean room profile

Details Clean room settings **Immutable snapshot recovery** Backup recovery

Optional Optional

Choose the settings to be used when recovering using an immutable snapshot with IBM FlashSystem.

Recovery settings ⓘ

☒ On

ESXi host ⓘ

10.200.120.61

vCenter folder ⓘ

CleanRoom

Clean up on failure ⓘ

☐ Off

Previous Next

Figure 2-49 Clean Room: Immutable snapshots settings window

5. Under the **Backup recovery** tab, specify your recovery preferences when recovering from IBM Storage Defender Data Protect (Figure 2-50 on page 44). If you plan to recover from a backup, select the vSphere resource pool from the drop-down list that you use for recovery. The default resource pool on each vCenter is the pool that is called Resources. You can have other resource pools that you created in your vCenter. All available resource pools can be selected for recovery. You can select the vCenter data store from the drop-down list that you want to use for recovery with this policy. Click **Create** to create a Clean Room profile with the specified values.

Create clean room profile

Details Clean room settings Immutable snapshot recovery Optional Backup recovery Optional

Choose the settings to be used when recovering using a backup with IBM Defender Data Protect.

Recovery settings ⓘ

☒ On

Data Protect resource pool

Resources

vCenter datastore

EN20DPVM1S

Previous Create

Figure 2-50 Clean Room: Backup recovery settings window

The Clean Room profile is created under the **Clean room** tab (Figure 2-51).

IBM Storage Defender Manage access 347 Trial days left

Data Resiliency

Dashboard Recovery groups Profiles Resources Integrations

Profiles

Governance Clean room

Clean room profiles

Clean room profiles are used by a recovery group to specify the clean room location and setting needed to recover.

Search Create profile

Name	Clean room location	Description
ssclab-CleanRoom	EN20-vcenter2:ssclab.ibm.si	ssclab clean room profile

Items per page: 10 1-1 of 1 item 1 1 of 1 page

Figure 2-51 Clean Room profile list

By creating the Governance and Clean Room profiles, you configured the recovery objectives of recovery groups and recovery target environments.

## 2.6 Auto-forwarding IBM Storage FlashSystem ransomware threat alerts to IBM Storage Defender

DRS integrates with IBM Storage Insights Pro and IBM FlashSystem to enable inline data anomaly detection on storage at the block level. IBM Storage FlashSystem offers new smart technology that is enabled by the fourth generation of IBM FlashCore Modules (FCM4), which are designed to continuously monitor statistics that are gathered from every I/O. IBM Storage FlashSystem uses machine learning models to detect anomalies like ransomware in less than a minute, which helps ensure that your business is protected before a cyberattack runs.

These ransomware alerts that are generated by IBM Storage Insights Pro for a monitored IBM FlashSystem can be auto-forwarded to IBM Storage Defender to trigger cyber resiliency workflows, and protect your systems as soon as possible. For customers subscribing to IBM Storage Insights Pro and IBM Storage Defender, this function enables enhanced protection from ransomware attacks with simple integration.

### 2.6.1 IBM FlashCore Module

IBM has been delivering high-performance, highly reliable customized flash modules for many years. With IBM FlashCore Modules (FCMs), the control path and the data path within the module are separated, which helps ensure that data can be accessed and transferred without any performance degradation that is caused by the control path. To enhance endurance and reliability, FCM modules have endurance features and RAID within the modules themselves. Numerous additional technologies and benefits are implemented.

For more information, see the [IBM Redbooks website](#) for publications about the IBM FlashSystem family, such as *IBM Storage FlashSystem 7300 Product Guide: Updated for IBM Storage Virtualize 8.7*, REDP-5741.

#### How IBM FCMs detect and report ransomware threats

In 2024, IBM introduced FCM4, which brought another industry-leading breakthrough that is called Ransomware Threat Detection. It is a process that identifies and responds to security threats before they can damage data or systems. FCM4 collects detailed statistics on every I/O operation (IOP) for each virtual disk (VDisk). This data is intelligently summarized for efficient processing. FCM4 transmits this summary to IBM Storage Virtualize, which relays it to an AI-powered inference engine. This engine can identify unusual activity, like potential ransomware attacks, in under a minute. On detection, an immediate alert is sent to IBM Storage Insights Pro, enabling swift action. Also, the information can be shared with IBM Storage Defender if it is available, which further strengthens your security posture.

With IBM Storage Virtualize 8.7 and FCMs with FCM firmware 4.1, the ransomware threat detection is further improved by the following process:

1. IBM FCMs collect and analyze detailed ransomware statistics from every I/O with no performance impact.
2. IBM Storage Virtualize runs an AI engine on every IBM FlashSystem, which is fed machine language (ML) models that are developed by IBM Research® and trained on real-world ransomware.
3. The AI engine learns what is normal for the system and detects threats by using data from FCM.

4. IBM Storage Insights Pro collects threat information from a connected IBM FlashSystem, and alerts trigger SIEM/SOAR software to initiate a response.
5. Statistics are fed back to IBM to improve ML models.

## 2.6.2 Integration between DRS and IBM Storage Insights Pro

IBM Storage Insights Pro is a subscription-based SaaS offering that provides enhanced monitoring, management, and optimization for storage environments. It is designed to help enterprises gain deeper insights into their storage infrastructure, improve operational efficiency, and proactively manage storage resources. IBM Storage Insights Pro has many AIOps capabilities that can help customers plan for the future and manage their infrastructure efficiently.

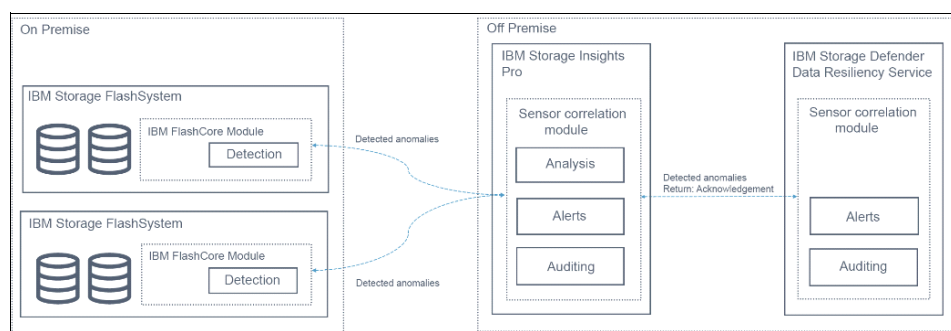


Figure 2-52 IBM Storage Defender integration with IBM FlashCore Module work flow

The following steps describe the working principle of the IBM Storage Defender integration with IBM FlashCore Module:

1. IBM FlashCore Module version 4 technology is built into the IBM Storage FlashSystem.
2. The IBM Storage FlashSystem is registered in IBM Storage Insights Pro. When the system is registered, the IBM FlashCore Module starts reporting the detected anomalies to your IBM Storage Insights Pro tenant.
3. IBM Storage Insights Pro correlates the data from multiple IBM FlashCore Modules and analyzes the data.
4. The IBM Storage FlashSystem must be registered in IBM Storage Defender Data Resiliency Service. This registration is done in the user interface of connection manager.
5. The IBM Storage Insights Pro communicates with Data Resiliency Service. The data that is related to your IBM Storage FlashSystem is sent to the Data Resiliency Service.
6. IBM Storage Defender correlates the information that is received from the storage system to recovery groups.
7. When the IBM FlashCore Module detects an anomaly for block level data operations, a case is opened for the related recovery group.
8. Depending on your notification settings you are notified about the new case.

For IBM FlashSystem running firmware 8.6.3 and later, IBM FCM (FCM4 with firmware 4.1) can detect ransomware threats in the data path and send threat details to IBM Cloud® Call Home. IBM Storage Insights Pro monitors ransomware threats that are detected on all monitored IBM FlashSystem systems and generates alerts. These alerts are sent to the storage administrator through email and are also displayed in the IBM Storage Insights Pro user interface. Also, IBM Storage Insights Pro identifies affected volumes, marking them as having detected ransomware threats.

## Enable DRS integration in Storage Insights Pro

Storage administrators who subscribe to both IBM Storage Insights Pro and DRS can direct ransomware alerts that are generated in IBM Storage Insights Pro to DRS. When an IBM FlashSystem is onboarded to both IBM Storage Insights Pro and IBM Storage Defender, IBM Storage Defender sends an integration request to IBM Storage Insights Pro. The IBM Storage Insights Pro administrator receives this request through the user interface and decides whether to forward ransomware alerts to IBM Storage Defender. After the administrator approves the request, the system is enabled to send ransomware alerts to IBM Storage Defender. When a ransomware threat is detected on any volume or volume group of the monitored IBM FlashSystem system, the alert is forwarded to the IBM Storage Defender webhook, including details such as storage system information, volume specifics, and the ransomware timestamp. On receiving and acknowledging the alert, IBM Storage Insights Pro notifies the user that IBM Storage Defender has acknowledged the alert and is actively addressing it. Then, the alert is available to be sent through DRS to any connected SIEM systems to notify the security operations (SecOpsClean Room) team. At the time of writing, QRadar and Splunk are supported.

The basic working principle of the integration between the two services is as follows:

- ▶ IBM FlashCore® Module 4 technology is built into the IBM Storage FlashSystem system that is used.
- ▶ The IBM Storage FlashSystem system is registered in IBM Storage Insights Pro. When the system is registered, the IBM FCM starts reporting the detected anomalies and ransomware threats to your IBM Storage Insights Pro tenant.
- ▶ The IBM Storage FlashSystem system must be registered in DRS. This registration is done in the user interface of Connection Manager.
- ▶ IBM Storage Insights Pro communicates with DRS. The health status of your IBM Storage FlashSystem system is sent to DRS so that if IBM Storage Insights Pro stops monitoring it, it can be made known to the users.
- ▶ IBM Storage Defender correlates the information that is received from the storage system to recovery groups.
- ▶ When the IBM FCM detects an anomaly for block-level data operations, a case is opened for the related recovery group.
- ▶ Depending on your notification settings you are notified about the new case. These notifications might include alerts that are sent to a connected SIEM.

## Viewing ransomware threats in DRS

IBM Storage Insights Pro reports the ransomware threats at the volume or volume group levels for an IBM FlashSystem system. An alert is sent and shown in DRS (Figure 2-53 on page 48). DRS uses recovery groups to group the related VMs. When a ransomware alert is received, DRS correlates the volume in the alert to the data store, and opens a case for the recovery group where the VM that uses the data store is. The newly opened cases can be viewed on DRS, and a recovery plan can be activated to recover to the last copy or last best copy that is available.

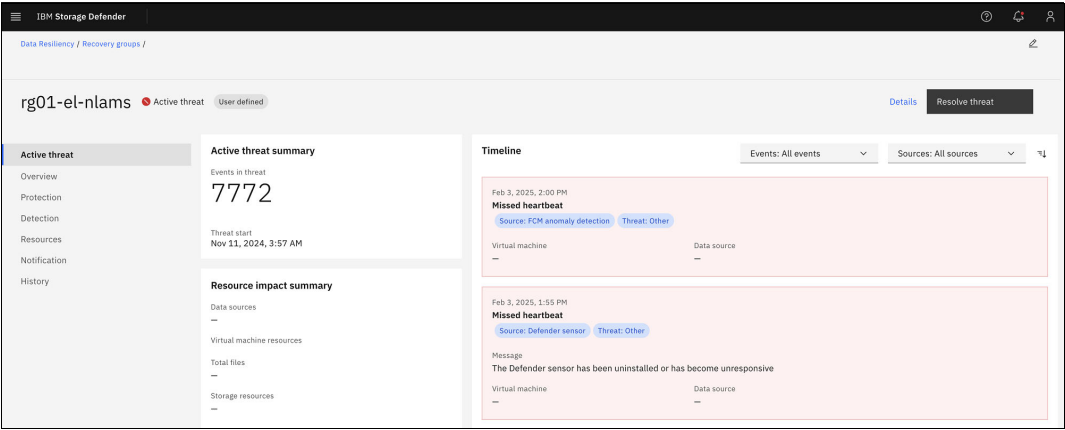


Figure 2-53 DRS showing a malware event notification message



## IBM Defender sensors

This section describes the IBM Defender sensors that are used for detecting threats against live data in near real time.

This chapter describes the following topics:

- ▶ 3.1, “What do sensors do” on page 50
- ▶ 3.2, “Installing sensors” on page 53

## 3.1 What do sensors do

IBM Storage Defender sensors are small, lightweight pieces of software that are installed into virtual machines (VMs) to monitor for file-pattern activities that resemble ransomware threats. Every 30 seconds, the sensor looks at Linux file-related event information to detect specific file patterns that ransomware variants tend to use. Sensors also use a pre-built machine learning model that trains on known ransomware variant patterns, which the sensors use to help identify similar patterns on the host where they are installed. If malicious activity is suspected based on these factors, a third check that uses file introspection is performed to determine whether the suspected victim files are encrypted.

If each of these criteria is met, an event is raised to IBM Storage Defender Data Resiliency Service (DRS) that indicates a possible malware event, and a case is opened, as shown in Figure 3-1.

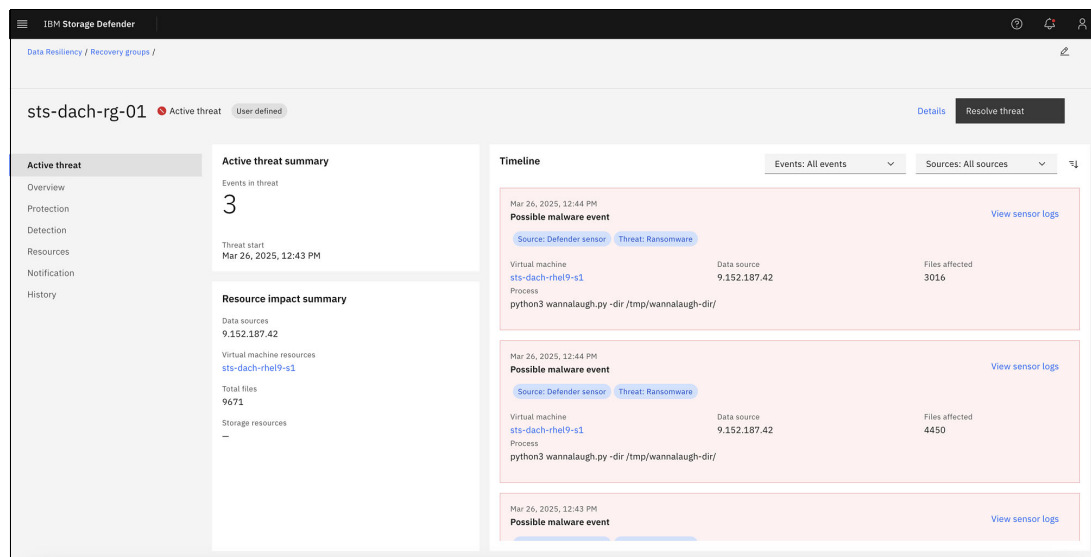


Figure 3-1 IBM Storage Defender timeline sensor event

In the details for the specific event (Figure 3-2 on page 51), both informational and actionable information is provided, which includes the following items (in the example):

- ▶ The type of event, which in this case is a “Possible malware event” of ransomware.
- ▶ The date and time that the event was detected, which can help with pinpointing clean copies for recovery and initial forensic analysis.
- ▶ The VM that is impacted (in this case, sts-pok-dsn-2-rhel) and its vCenter.
- ▶ The suspected malicious process: python3 ./filesEnc.py.
- ▶ The number of files that are affected (235) for this specific window of detection.
- ▶ The source (originator) of the event, in our case, an IBM Storage Defender sensor. IBM FlashSystem related events can also be raised from IBM Storage Insights Pro.



Log details

```
===== START Log Entry =====
Sensor ID: 4237e87b-1c82-1bb8-57a4-15912dfe3875
fqdn: sts-dach-rhel19-s1.boeblingen.de.ibm.com
Sensor version: 2.0.10-1733437154
Sensor commit ID: 26e7120
OS Name: Red Hat Enterprise Linux
OS Version: 9.5
Window ID: 2321
Time window seconds: 30
Collection start time: 2025-03-26_18:44:20
Analysis start time: 2025-03-26_18:44:51
----- START Processes involved in Malicious accesses -----
PID: 33883: command = 'python3 wannalaugh.py -dir /tmp/wannalaugh-dir/ '
----- END Processes involved in Malicious accesses -----
----- START Files involved in Malicious accesses -----
File: /tmp/wannalaugh-dir/9/024997.txt.WNNLGH: uid = 2147483647, size = 8432, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024996.txt.WNNLGH: uid = 2147483647, size = 2421296, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024994.txt.WNNLGH: uid = 2147483647, size = 57888, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024992.txt.WNNLGH: uid = 2147483647, size = 70912, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024990.txt.WNNLGH: uid = 2147483647, size = 3216, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024987.txt.WNNLGH: uid = 2147483647, size = 3904, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024986.txt.WNNLGH: uid = 2147483647, size = 6384, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024984.txt.WNNLGH: uid = 2147483647, size = 928, pids = [33883],
diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024983.txt.WNNLGH: uid = 2147483647, size = 3920, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024981.txt.WNNLGH: uid = 2147483647, size = 10560, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024980.doc.WNNLGH: uid = 2147483647, size = 44064, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024979.txt.WNNLGH: uid = 2147483647, size = 12512, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
File: /tmp/wannalaugh-dir/9/024978.txt.WNNLGH: uid = 2147483647, size = 73520, pids =
[33883], diagCode = 71.6483, encStatus = NotChecked
```

Figure 3-3 IBM Storage Defender sensor log details

Chapter 3. IBM Defender sensors 51

A summary of the impact is provided at the end of the log, including the total number of files (Figure 3-4 shows the suspected malicious accesses event details). Regardless of the number of files that are impacted, recovery happens at a volume level, and all files can be recovered to an earlier, unimpacted state.

**Note:** At the time of writing, the encryption detection identifies encryption only on files that are larger than 4 KB. Therefore, it is likely that if specific files are identified as impacted, it is probable that smaller files in these locations are, too.

Log details

File: /tmp/wannalaugh-dir/8/024430.jpg.WNNLGH: uid = 2147483647, size = 104672, pids = [33883], diagCode = 51.5, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024888.pdf.WNNLGH: uid = 2147483647, size = 31104, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024887.pdf.WNNLGH: uid = 2147483647, size = 5139200, pids = [33883], diagCode = 71.78659999999999, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024886.pdf.WNNLGH: uid = 2147483647, size = 14240, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/7/024028.doc.WNNLGH: uid = 2147483647, size = 89632, pids = [33883], diagCode = 51.5, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024885.pdf.WNNLGH: uid = 2147483647, size = 4864048, pids = [33883], diagCode = 71.78659999999999, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024884.pdf.WNNLGH: uid = 2147483647, size = 977008, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024883.pdf.WNNLGH: uid = 2147483647, size = 17536, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/9/024588.gif.WNNLGH: uid = 2147483647, size = 11376, pids = [33883], diagCode = 51.5, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024882.pdf.WNNLGH: uid = 2147483647, size = 50608, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024881.gif.WNNLGH: uid = 2147483647, size = 1747888, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024880.xls.WNNLGH: uid = 2147483647, size = 806944, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024879.xls.WNNLGH: uid = 2147483647, size = 67616, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024878.pdf.WNNLGH: uid = 2147483647, size = 48848, pids = [33883], diagCode = 71.6483, encStatus = NotChecked

File: /tmp/wannalaugh-dir/6/024877.xls.WNNLGH: uid = 2147483647, size = 27155488, pids = [33883], diagCode = 71.9107, encStatus = Detected

File: /tmp/wannalaugh-dir/6/024876.pdf.WNNLGH: uid = -1, size = 1955664, pids = [33883], diagCode = 60.9387, encStatus = NotChecked

----- END Files involved in Malicious accesses -----

Summary: totalMaliciousAccesses: 3016

Summary: totalMaliciousEncryptions: 30

Summary: isAlert: True

Summary: Analysis end time: 2025-03-26\_18:44:51

===== END Log Entry =====

Show less

Figure 3-4 Sensor log details (cont.)

The sensors also send regular heartbeats to the DRS to indicate that both the sensors and the dependent network connections are healthy. If a heartbeat is missed, an event is raised, as shown in Figure 3-5.

Oct 11, 2024, 2:16 AM

Missed heartbeat

Source: Defender sensor

Threat: Other

Message

The Defender sensor has been uninstalled or has become unresponsive

Virtual machine

Data source

DefenderDRSDemoVM2 DefenderDRSDemoVM3 10.3.69.10

Figure 3-5 Sensor heartbeat warning message

52

IBM Storage Defender: Data Resiliency Service

For any of these events, a case is opened so that actions can be reviewed and communicated between team members and teams. For example, after an event is analyzed by an admin or responder, information from the event is reviewed, and the cause can be addressed or confirmed. After the appropriate remediation is taken to resolve the issue, the case can be closed. When a case is closed, the corresponding event messages are cleared, as shown in Figure 3-5 on page 52. However, past events can still be viewed from the Detection window, as shown in Figure 3-6.

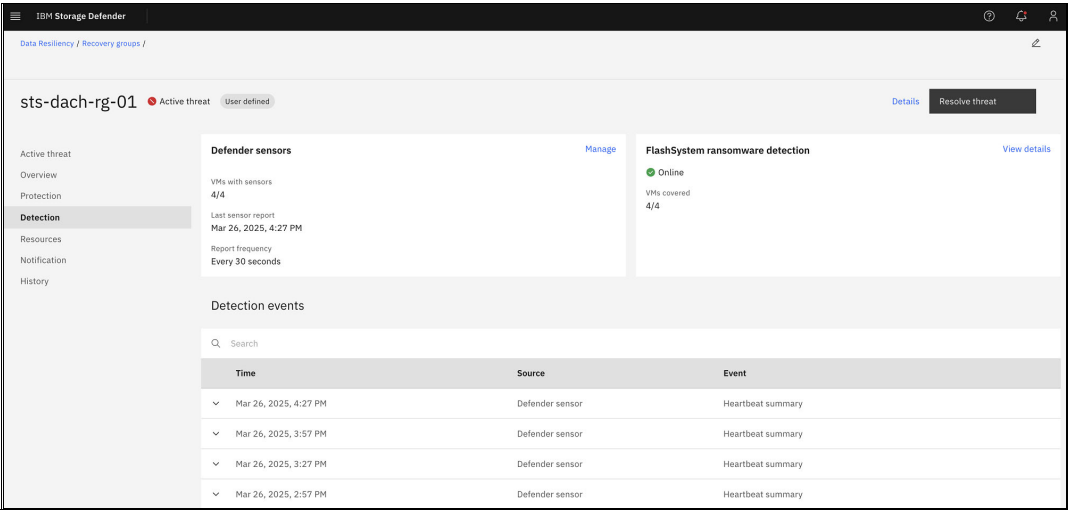


Figure 3-6 Recovery Group detection window

After closing the case, the DRS dashboard continues to allow access to historical events. You may search on previous threat events and drill down to review the details of those events.

## 3.2 Installing sensors

This section explains how to install an IBM Storage Defender sensor on one or more systems by using either the GUI or command-line interface (CLI). These sensors monitor the systems on which they are installed, enabling real-time detection of cyberthreats, such as ransomware attacks. There are two ways to deploy the sensors:

- ▶ Automatically from the Connection Manager by using the built-in facilities
- ▶ Through Ansible automation and deploying your own sensor control node

This section shows how to deploy a sensor control node outside of the IBM Storage Defender Connection Manager.

### 3.2.1 Installing the sensor control software

Download the IBM Storage Defender sensor control software by completing the following steps:

1. Log in to the system that you want to use as a sensor control node.
2. From that system, log in to a Connection Manager instance.
3. On the home page of the Connection Manager, click **Connections**.
4. Click **Sensor control nodes**.
5. Click **Download package**.

Install the sensor control software on the sensor control node by completing the following steps:

1. Log in to the system that you want to use as a sensor control node.
2. Copy the sensor download package to a working directory.
3. Unpack the compressed software package that you downloaded.
4. In the newly created directory, run the `setup.sh` shell script.

The script requires the following input values. Use unique names for each entity in the environment:

<b>Hostname:</b>	The FQDN of the Connection Manager.
<b>Username:</b>	Define a username that is to register IBM Storage Defender sensors that are installed on VMs for the sensor control node.
<b>Password:</b>	Define a password that is related to the username.
<b>Vault password:</b>	The username and password that were defined are stored and encrypted in a local Ansible vault. This password is used to protect the access to the vault.

### 3.2.2 Adding a sensor control node

To add a sensor control node to the Connection Manager, complete the following steps:

1. Log in to the Connection Manager instance.
2. On the home page of the Connection Manager, click **Connections**.
3. Click **Sensor control nodes**.
4. Click **Add control node**. This action opens a dialog box.
5. In the dialog box, enter the FQDN of the sensor control node.
6. Click **Next**.
7. Enter the username that was provided when you installed the sensor control software on the sensor control node.

**Note:** Multiple sensor control nodes can use the same username and password for sensor installation or registration. In this case, only one control node must be added. If you attempt to add more than one control node by using the same username, the following error occurs in the GUI:

Error getting source native ID: Username already in use. Select a different username.

8. Enter the password of the user.
9. Click **Add**.

Now, you see the registered sensor control node in the GUI.

### 3.2.3 Removing a sensor control node

To remove a sensor control node from the Connection Manager, complete the following steps:

1. Log in to the Connection Manager instance.
2. On the home page of the Connection Manager, select **Connections** → **Sensor control nodes**.
3. In the table that lists all the sensor control nodes, scroll to the relevant sensor control node.
4. In the row of the sensor control node, click the overflow menu (Figure 3-7), and then click **Remove**. This action opens a dialog box.

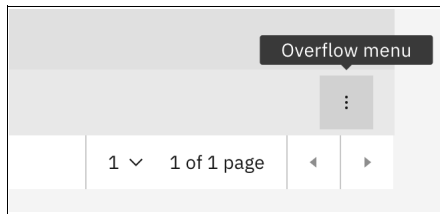


Figure 3-7 Overflow menu location

5. In the dialog box, click **Remove** to confirm that you want to remove the sensor control node from the Connection Manager.

### 3.2.4 Installing an IBM Storage Defender sensor by using the Defender UI

You can install the sensor on one or multiple systems directly through the IBM Storage Defender UI.

Before you begin, consider the following items:

- Review the system requirements.
- The procedure that is described in this topic covers adding a sensor to a server by using the Connection Manager embedded sensor control node feature.

To install an IBM Storage Defender sensor on one or more systems, complete the following steps:

1. Log in to IBM Storage Defender.
2. Click the hamburger menu (three horizontal lines) in the upper left of the window.
3. Select **Data Resiliency** → **Recovery Groups**.
4. From the list of Recovery Groups, select the row for the Recovery Group that you want to install the sensor on.

5. In the Overview window, find the Defender sensors tile and click **Get started** (see Figure 3-8).

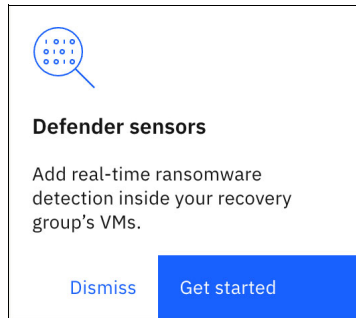


Figure 3-8 Defender sensors tile

**Note:** If you previously installed sensors, you see the **Manage** button on the Defender sensors tile.

6. In the Manage sensors window, select one or more VMs by checking the corresponding boxes.
7. Click **Add sensor +** in the title bar.

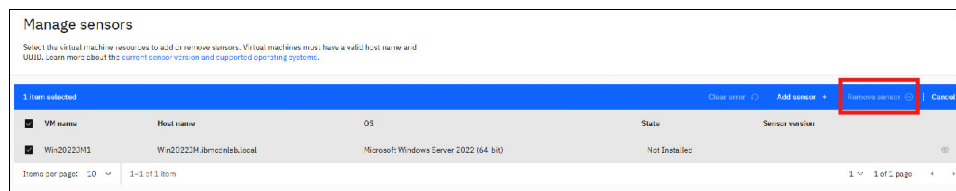


Figure 3-9 Sensor installation panel

8. Enter either the username and password or the SSH key for the VM.

**Note:** All the selected VMs must have the same login credentials.

9. Select **Add Sensor** to submit the installation request.

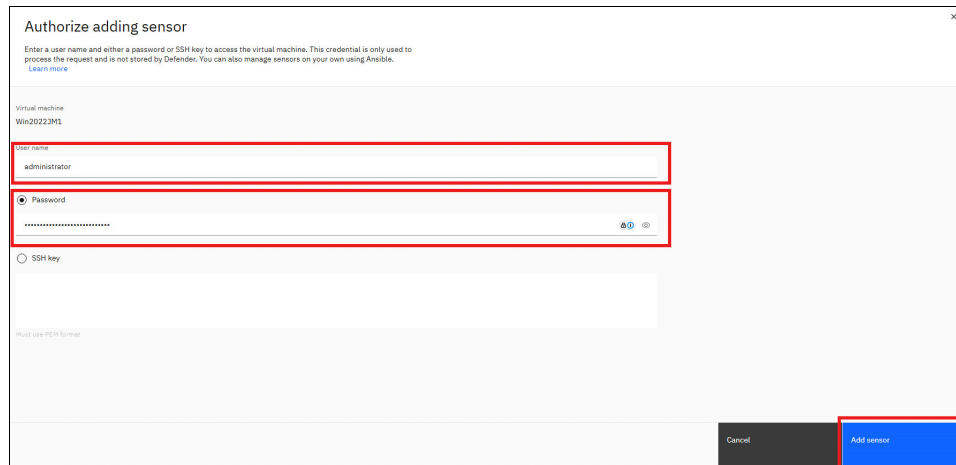


Figure 3-10 Sensor authorization panel

10. The selected VMs display the status “Installing” until the installation is complete.

**Note:** Monitor the **Notification** menu to check for completed or failed notifications for each sensor. If the status is TIMEOUT, the installation request was accepted but did not respond for 15 minutes. For the FAILED status, check the detailed error message in the notification.

After the installation completes, the sensor automatically begins monitoring file access activity on the system. If it detects any unusual access patterns, such as patterns that are associated with ransomware attacks, the sensor generates an alert. This alert is sent to the on-premises Connection Manager, which securely forwards it to the DRS. The sensor also periodically sends heartbeat messages through the Connection Manager to confirm that it is operating normally.

### 3.2.5 Installing an IBM Storage Defender sensor by using the CLI

You can install an IBM Storage Defender sensor on one or multiple systems by using the GUI or CLI. The sensors observe the systems that they are installed on and can detect cyberattacks like ransomware attacks in real time.

To install an IBM Storage Defender sensor on one or more systems, complete the following steps:

1. Log in to the system that is being used as the sensor control node.
2. Go to the working directory where the sensor control software is installed.

**Note:** This directory is the one that you specified when downloading and installing the sensor control software.

3. Create an inventory file containing the FQDNs of all the systems that you want to install the sensor on.
4. Modify the `/etc/ansible/hosts` file to include the FQDNs of the target systems.

**Note:** You can use a different file for the sensor inventory list. If so, use the `-i /your-directory/your-file` argument in step 5.

5. Add the FQDNs for the sensor hosts to the hosts file. Under the `[defender_sensor_hosts]` tag, list the FQDN of each system, one per line.

**Tip:** If you use a YAML inventory file, extend it with a `defender_sensor_hosts` group.

*Example 3-1 Ansible hosts file configuration*

```
[defender_sensor_hosts]
<FQDN1>
<FQDN2>
<FQDN3>

[defender_sensor_hosts:vars]
ansible_ssh_common_args='-o StrictHostKeyChecking=no'
ansible_connection=ssh
ansible_ssh_pass=<ssh password>
ansible_ssh_user=<ssh username>
```

```

all:
  vars:
    ansible_connection: ssh
    ansible_ssh_user: <ssh username>
    ansible_ssh_pass: <ssh password>
    ansible_ssh_common_args: '-o StrictHostKeyChecking=no'
  children:
    defender_sensor_hosts:
      hosts:
        <FQDN1>:
        <FQDN2>:
        <FQDN3>:

```

---

6. Run the Ansible playbook command in Example 3-2 to begin the installation.

*Example 3-2 Ansible playbook install command*

---

```

ansible-playbook sensor_install.yml --ask-vault-pass [-i
path_to_alternative_inventory_file>]

```

---

7. When prompted, enter the vault password that you created during the installation of the sensor control node software.

**Note:** To avoid saving passwords in the hosts file, use the arguments `--ask-pass` `--ask-become-pass` to provide the SSH and sudo passwords during the playbook run time.

After installation, the sensor automatically monitors file access activities on the system. If any unusual access patterns resembling ransomware attacks are detected, the sensor sends alert messages to the on-premises Connection Manager, which forwards these alerts securely to the DRS. The sensor also sends periodic heartbeat messages to the DRS through the Connection Manager, indicating normal operation.

### 3.2.6 Uninstalling an IBM Storage Defender sensor by using the Defender UI

To uninstall an IBM Storage Defender sensor from one or more systems by using the GUI, complete the following steps:

1. Log in to IBM Storage Defender and access the IBM Storage Defender dashboard.
2. Go to the Recovery Group:
  - a. Click the hamburger menu (three horizontal lines) in the upper left of the page.
  - b. Select **Data Resiliency** → **Recovery Groups**.
  - c. From the list of Recovery Groups, click the row corresponding to the Recovery Group where you want to uninstall sensors.
3. Manage sensors:
  - a. In the Recovery Group's Overview dashboard, find the Defender Sensors tile and click **Manage**.
  - b. Select the VMs from which you want to uninstall the sensor by checking the appropriate boxes.

**Note:** The Connection Manager uses FQDNs to perform sensor installation and uninstallation. You cannot select the following VMs for sensor uninstallation:

- ▶ VMs without an FQDN
- ▶ VMs that use localhost as the FQDN
- ▶ VMs with duplicate FQDNs

Any changes to VM network configurations are reflected in the GUI after the next inventory scan, which occurs automatically every hour or can be manually triggered.

#### 4. Uninstall the sensor:

- a. Click **Remove Sensor** in the title bar.
- b. Enter either the username and password, or the SSH key for the VMs.
- c. Click **Remove Sensor** to submit the uninstallation request.

**Note:** All selected VMs must share login credentials.

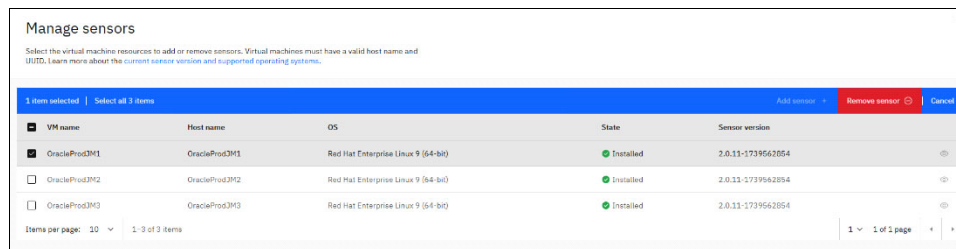


Figure 3-11 Sensor removal panel

#### 5. Monitor the process:

- a. The statuses of the selected VMs change to Uninstalling.
- b. You can monitor the Notification menu for updates about the success or failure of each sensor uninstallation.

**Tip:** If the status shows TIMEOUT, the request was accepted but did not receive a response for 15 minutes. For a FAILED status, check the detailed error message in the notification.

**Important:** If you are trying to uninstall a sensor that is associated with a Connection Manager that was destroyed or improperly backed up and restored during a Connection Manager OVA upgrade, the uninstallation fails. For troubleshooting, see [Resolving an IBM Storage Defender sensor uninstallation failure](#).

After the uninstallation completes, the IBM Storage Defender sensor service is removed from the selected VMs.

### 3.2.7 Uninstalling an IBM Storage Defender sensor by using the CLI

You can uninstall an IBM Storage Defender sensor from one or more systems by using either the GUI or CLI. To proceed with the CLI method, complete the following steps:

1. Log in to the sensor control node: Access the system you are using as the sensor control node.
2. Create an inventory file listing the FQDNs or IP addresses of all systems from which you want to uninstall the sensor.
3. Modify the `/etc/ansible/hosts` file to include the FQDN or IP address of each target system.

**Note:** You can use a different file for the sensor inventory list. If so, use the `-i /your-directory/your-file` argument in step 4.

4. Add the FQDN or IP address of all systems that you want to equip. Add one per line under the tag `[defender_sensor_hosts]`.
5. Run the Ansible playbook command that is shown in Example 3-3.

*Example 3-3 Ansible playbook uninstall command*

---

```
ansible-playbook sensor_uninstall.yml --ask-vault-pass [-I  
<path_to_alternative_inventory_file>]
```

---

6. Enter the Ansible vault password.

After the playbook runs, the sensor is removed from the host.

### 3.2.8 Requirements for IBM Storage Defender sensors

Before proceeding with the installation and registration of the IBM Storage Defender sensor, ensure that your system meets the following requirements in terms of supported operating systems and necessary software packages:

- ▶ Supported operating systems:
  - Red Hat Enterprise Linux Server 9 required packages:
    - `bash`
    - Kernel 5.9 or later
    - `libgomp`
    - `python3`
  - SUSE Linux Enterprise Server 15 SP5 required packages:
    - `bash`
    - Kernel 5.9 or later
    - `libgomp1`
    - `python311`

**Note:** To install `python311`, the Python3 module must be enabled. For details on enabling modules, refer to the SUSE Linux Enterprise Server documentation.

- Ubuntu 24.04 LTS required packages:
  - `bash`
  - `libgomp1`
  - `linux-image-generic 5.9` or later

- python3
- ▶ Supported file systems:
  - XFS
  - EXT4





## Daily administration, alerting, testing, and validation

This chapter provides an overview about how to bring together the elements for daily administration and test your recovery points.

This chapter describes the following topics:

- ▶ 4.1, “IBM Storage Defender DRS dashboard” on page 64
- ▶ 4.2, “User management profiles” on page 69
- ▶ 4.3, “Integrations for alerting” on page 75
- ▶ 4.4, “Recovery testing and validation” on page 76
- ▶ 4.5, “Activating the recovery plan” on page 80

## 4.1 IBM Storage Defender DRS dashboard

For daily administration and statuses at a glance, the IBM Storage Defender Data Resiliency Service (DRS) provides a dashboard landing area. This area is an at-a-glance perspective of the overall environment, including urgent issues, the Connection Managers, open cases, Recovery group status, and other information.

Figure 4-1 shows a view of the DRS dashboard.

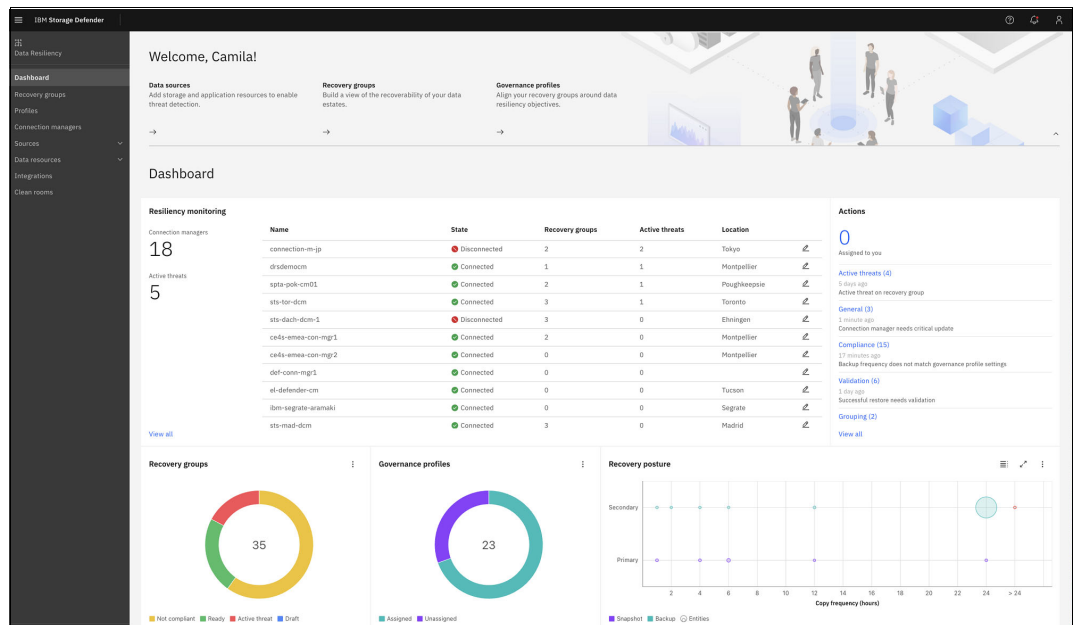


Figure 4-1 DRS Dashboard overview page

This dashboard contains several elements that enable users to view more information and context. These capabilities include the following items:

- ▶ Resiliency Monitoring through IBM Storage Defender Connection Managers.
- ▶ Actions that can be performed, which include open cases, assigned actions, required updates, and other issues.
- ▶ Recovery groups statuses.
- ▶ Governance Profiles statuses.
- ▶ Recovery Posture status.

### 4.1.1 Resiliency Monitoring in the dashboard

The DRS Dashboard provides an at-a-glance view of Resiliency Monitoring (Figure 4-2 on page 65), which highlights the status of the Connection Managers and any open cases.

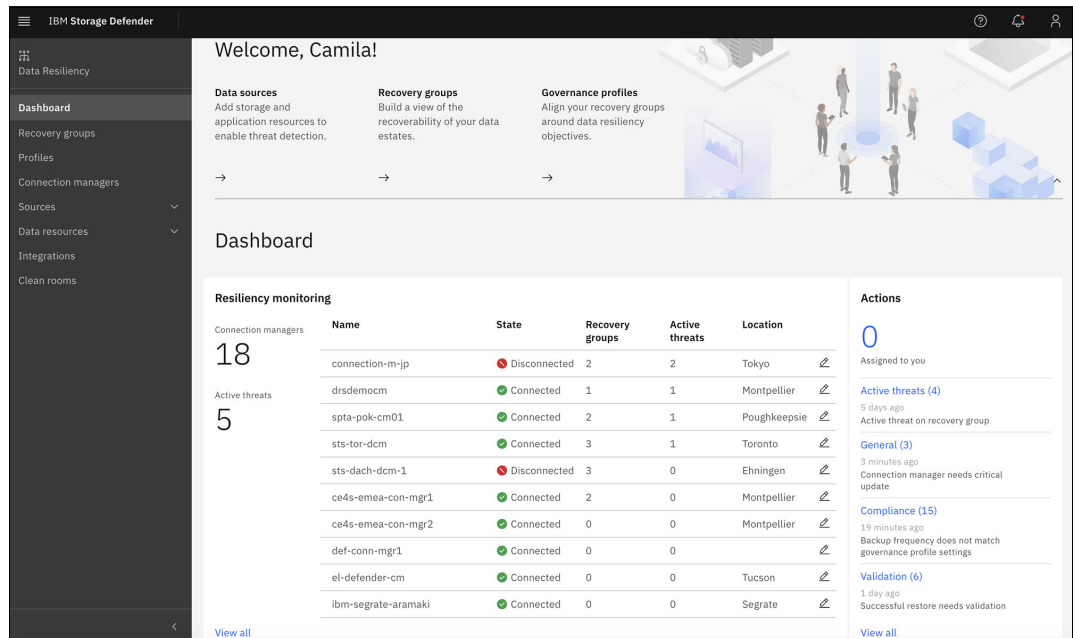


Figure 4-2 Resiliency Monitoring Dashboard window

With this dashboard, you can see locations, see the states or statuses of them, and drill in on the managed Connection Managers. You can use the **View All** link to see the **Connection Managers** tab of the Resources window. This window highlights the Connection Managers, their states, their types, hostnames, versions, and whether updates are required, as shown in Figure 4-3. It also allows administrators to push available upgrades to the Connection Managers remotely from the 'More Options icon' to the far right of each connection listed.

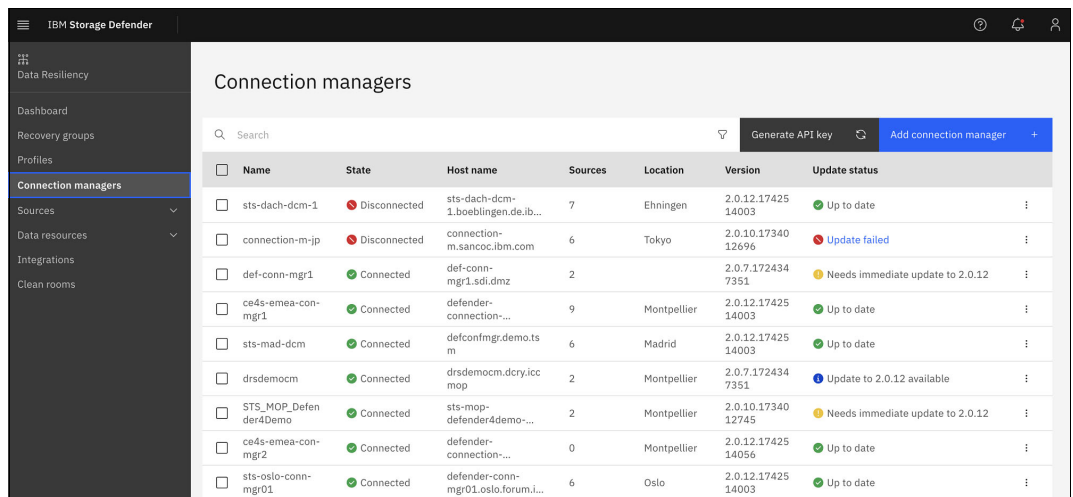


Figure 4-3 DRS Monitoring Resources overview window

Figure 4-4 on page 66 shows the actions that you can take, such as open cases, assigned actions, required updates, and others.

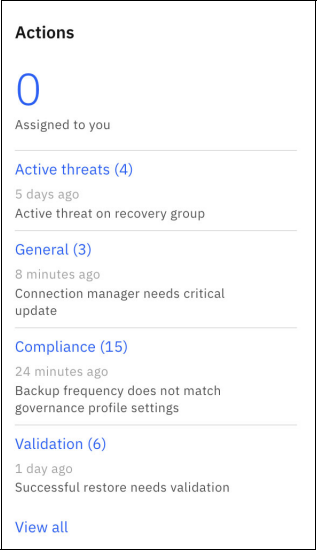


Figure 4-4 Actions summary window

Figure 4-5 shows the view that is available when you click the **Actions** menu. From here, you can see a deeper view of the actions to review recommendations and resolve issues, see pending actions, or view the history.

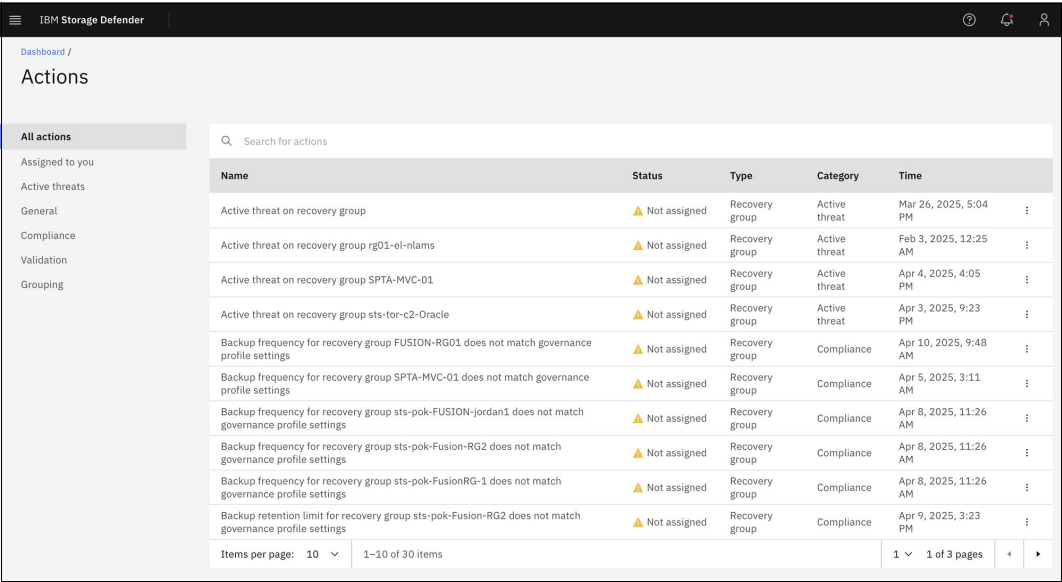


Figure 4-5 Actions window history details

### 4.1.2 Recovery group status

The Recovery groups pie chart (Figure 4-6 on page 67) shows a summarized view of all recovery groups across all the deployed environments linked by the various Connection Managers. This chart indicates the percentage of recovery groups that are ready, in draft, not compliant, or have an open case or a threat that is recorded on them that must be addressed.



Figure 4-6 Recovery groups pie chart on DRS dashboard

If you select Recovery groups on the left of the dashboard, you see the Recovery groups list (Figure 4-7), where you can drill down on any Recovery group or create one.

The screenshot shows the "Recovery groups" list in the IBM Storage Defender interface. A sidebar on the left contains navigation links: Data Resiliency, Dashboard, Recovery groups (selected), Profiles, Connection managers, Sources, Data resources, Integrations, and Clean rooms. The main area displays a table of recovery groups with columns for Name, Status, Governance profile, Resources, Resource type, and Connection manager. A search bar is at the top of the table.

Name	Status	Governance profile	Resources	Resource type	Connection manager
rg01-el-nlams	Active threat	el-nlams-profile	0	VM	
SPTA-MVC-01	Active threat	SPTA-Demo-Profile-MVC	3	VM	spta-pok-cm01
sts-jp-test01	Active threat	Silver	2	VM	connection-m-jp
sts-jp-test02	Active threat	Bronze	25	VM	connection-m-jp
sts-mop-demo	Active threat	sts-mop-demo	2	VM	drsdemocom
sts-tor-c2-Oracle	Active threat	C2Toronto_Production	3	VM	sts-tor-dcm
AD-RG01	Not compliant	GovProfile-MAD	1	AD Domain Controller	sts-mad-dcm
CE45-EMEA-primary-databases-RG	Not compliant	CE45-EMEA-primary-database-profile	3	VM	ce45-emea-con-mgr1
CE45APAC	Not compliant	Bronze	1	VM	
defender4Demo6930	Not compliant	Bronze	0	VM	
FUSION-RG01	Not compliant	GovProfile-MAD	1	Fusion application	sts-mad-dcm
MAD-RG02	Not compliant	GovProfile-MAD	2	VM	sts-mad-dcm

Figure 4-7 Recovery groups list

### 4.1.3 Governance profile status

This pie chart (Figure 4-8) highlights the number of Governance profiles and what percentage of them were assigned to a recovery group. These Governance profiles help users follow internal or regulatory compliance that is mandated around retention, frequency, and testing frequency of copies of data.

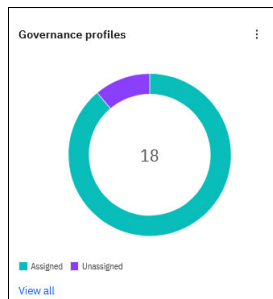


Figure 4-8 Governance profile pie chart on DRS dashboard

You can create and modify your existing Governance and Clean Room profiles within the **Profiles** tab (Figure 4-9).

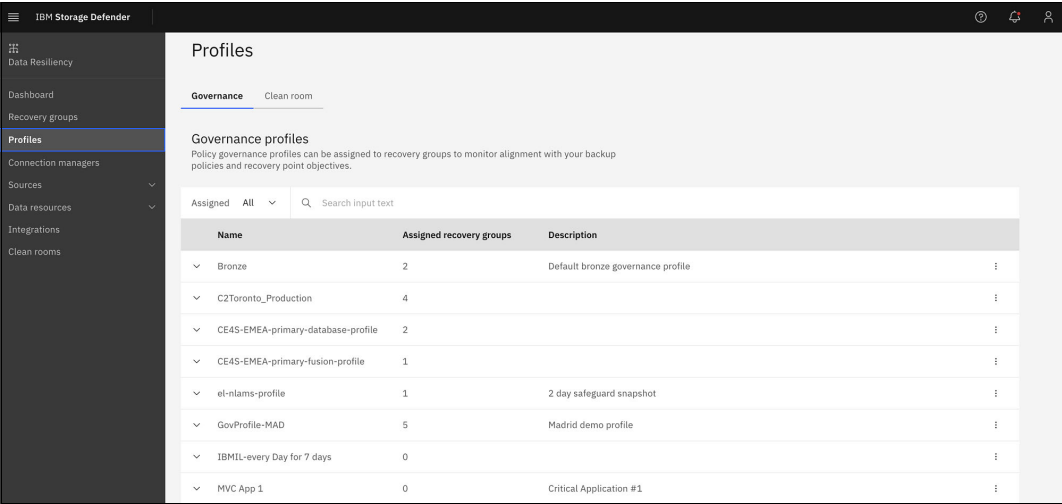


Figure 4-9 Profiles tab on DRS

### 4.1.4 Recovery posture

The recovery posture graphic (Figure 4-10) helps you quickly understand your recovery posture. On the Y axis, you see Secondary and Primary, which refers to auxiliary storage (for example, backups in IBM Storage Defender Data Protect) and primary storage (for example, an IBM FlashSystem system). On the X axis, you see the copy frequency, which is how often that your system creates copies. By combining these two axes, you can view what the frequency policies are for your environment for both primary and secondary copies.

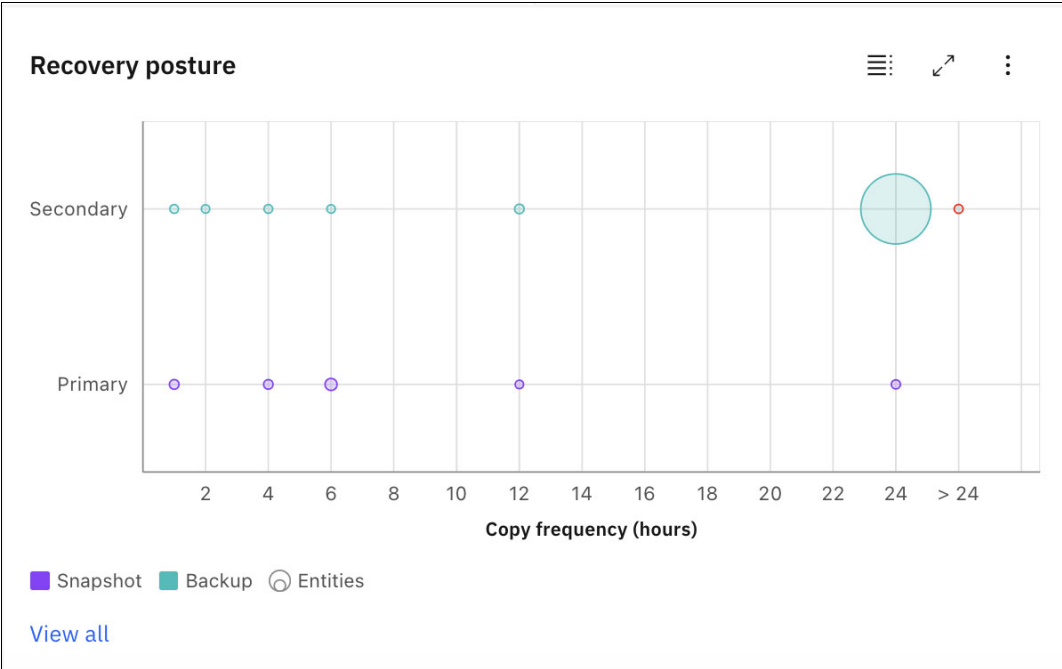


Figure 4-10 Recovery posture graphic in DRS dashboard

You can gather more information about your available resources, available copies, connections, and Connection Managers by clicking the **Resources** tab (Figure 4-11) and then clicking **Resources** in the left pane of the GUI.

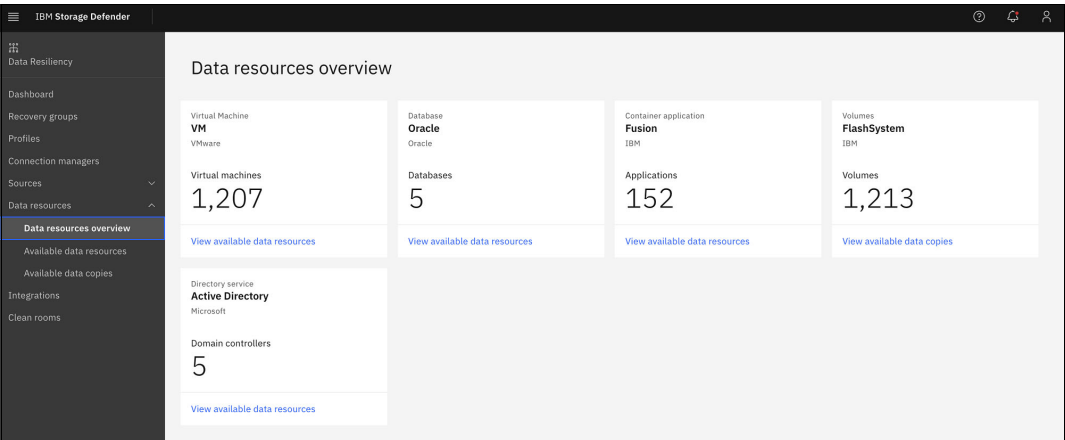


Figure 4-11 Resources tab on DRS

## 4.2 User management profiles

From the main Dashboard, Administrators can navigate to the Access section and select one of 3 options to manage user access (Figure 4-12):

- 1. Access Points
- 2. Users
- 3. User Groups

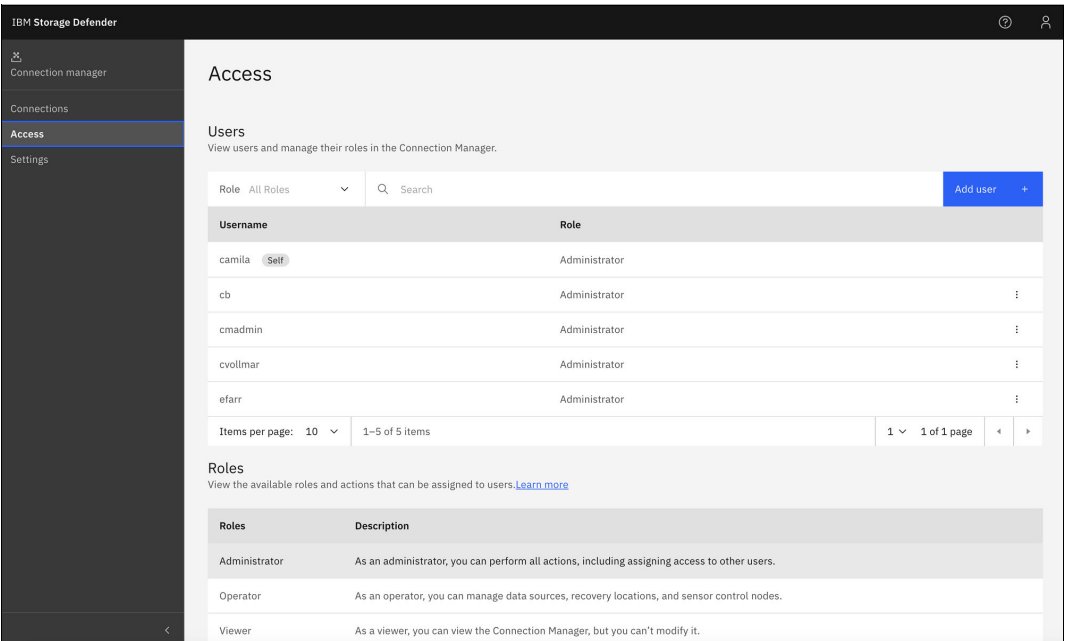


Figure 4-12 User management in DRS

For a user to appear in these options and be managed, they must first be enabled to see and work in DRS. From the main user management screen in IBM Storage Defender, ensure that the Data Resiliency Service option is enabled for the user's role.

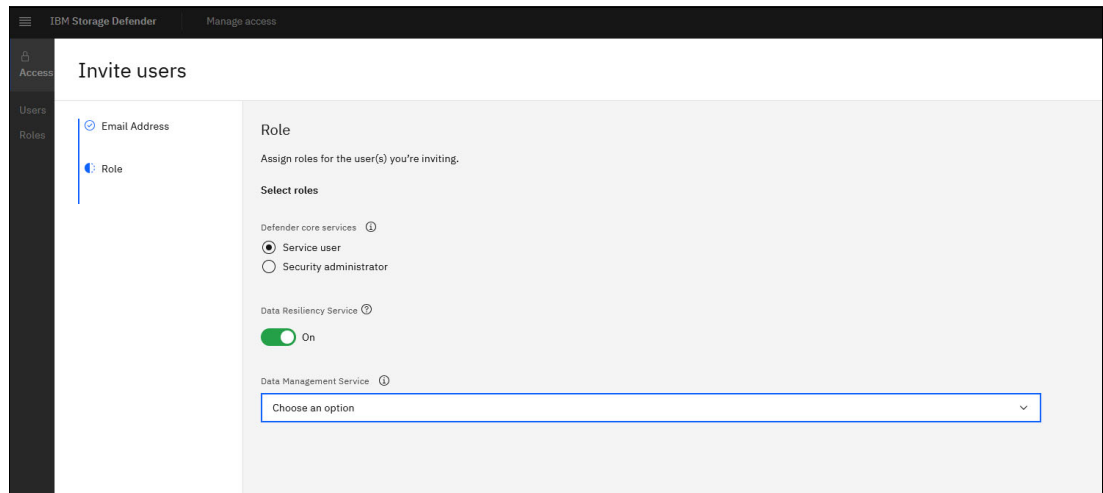


Figure 4-13 User Role options dialog

From there, administrators can navigate to the access points that they want to assign to their users.

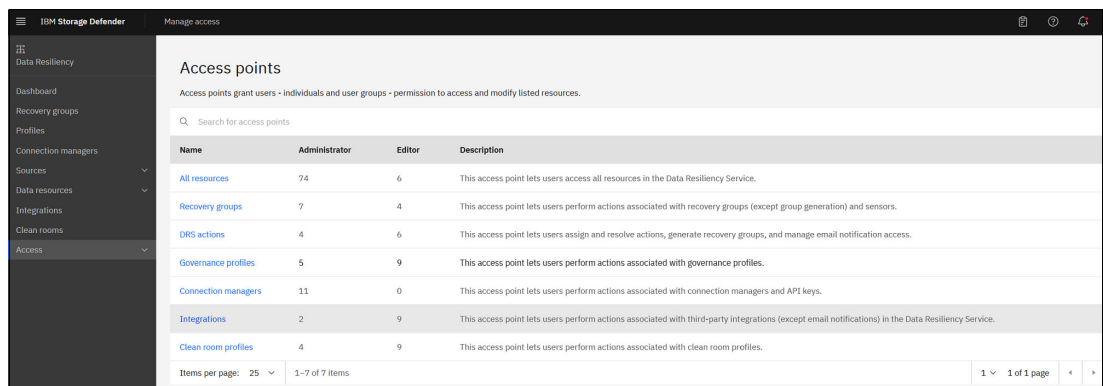


Figure 4-14 DRS Access points list

The following Access Points are available:

- ▶ All resources: Use this access point to access all resources in the Data Resiliency Service.
- ▶ Recovery groups: Use this access point to perform actions associated with recovery groups (except group generation) and sensors.
- ▶ DRS actions: Use this access point to assign and resolve actions, generate recovery groups, and manage email notification access.
- ▶ Governance profiles: Use this access point to perform actions associated with governance profiles.
- ▶ Connection managers: Use this access point to perform actions associated with connection managers and API keys.
- ▶ Integrations: Use this access point to perform actions associated with third-party integrations (except email notifications) in the Data Resiliency Service.
- ▶ Clean room profile: Use this access point to perform actions associated with clean room profiles.

By clicking any of the Access Points, the administrator can see which users are assigned which rights to each Access Point.

For example, each Access Point except for Connection Managers, has 2 options for role assignment; Administrator and Editor. The Connection Managers Access Point only has the Administrator role.

For example, each of the Access Point role assignment screens has a similar appearance, in that it lists the Users and User Groups assigned, as well as the capabilities of the two roles that can be assigned (see Figure 4-15). It also allows the Administrator to Manage user access for that specific function from that screen.

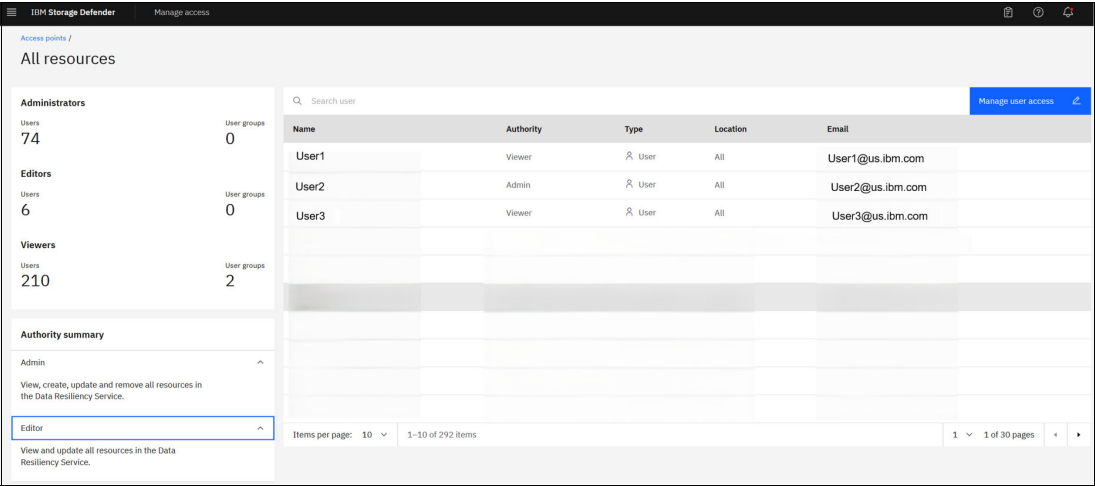


Figure 4-15 Access point user list

Table 4-1 describes the access point options and the rights provided to a user when admin or editor privileges are granted.

Table 4-1 Access point details

Access Point	Description	Admin	Editor
All resources	Use this access point to access all resources in the Data Resiliency Service.	View, create, update, and remove all resources in the Data Resiliency Service.	View and update all resources in the Data Resiliency Service.
Recovery groups	Use this access point to perform actions associated with recovery groups (except group generation) and sensors.	View, create, update, and remove recovery groups (except group generation) and perform all actions related to sensors (including remove).	View and update recovery groups (except group generation), and perform actions related to sensors (except remove).

Access Point	Description	Admin	Editor
DRS Actions	Use this access point to assign and resolve actions, generate recovery groups, and manage email notification access.	Generate recovery groups (locking/unlocking groups, including/excluding resources, resolving/dismissing events), assign and resolve actions, and manage email notification access.	Generate recovery groups (locking/unlocking groups, including/excluding resources, resolving/dismissing events) and assign or resolve actions.
Governance Profiles	Use this access point to perform actions associated with governance profiles.	View, create, update and remove governance profiles.	View and update governance profiles.
Connection Managers	Use this access point to perform actions associated with connection managers and API keys.	Perform all actions associated with API keys and connection managers (including update).	N/A
Integrations	Use this access point to perform actions associated with third-party integrations (except email notifications) in the Data Resiliency Service.	View, create, update, and remove third-party integrations (except email notifications) in the Data Resiliency Service.	View and update third-party integrations (except email notifications) in the Data Resiliency Service.
Clean room profile	Use this access point to perform actions associated with clean room profiles.	View, create, update, and remove clean room profiles.	View and update clean room profiles.

## 4.2.1 Users

Users provides a list of all the users available to have Access Points and roles assigned in DRS.

Name	Email	User groups	Access points
User1	User1@us.ibm.com	0	All resources
User2	User2@us.ibm.com	0	All resources
User3	User3@us.ibm.com	1	All resources, Clean room profiles, Integrations, Governance profiles
User4	User4@us.ibm.com	0	All resources
User5	User5@us.ibm.com	0	All resources
User6	User6@us.ibm.com	0	All resources
User7	User7@us.ibm.com	1	All resources, Clean room profiles, Integrations, Governance profiles

Figure 4-16 Users list view

Select a specific user to view their individual details (Figure 4-17). You can also select “View all details” to see the full view of that specific user and their current privileges (Figure 4-18).

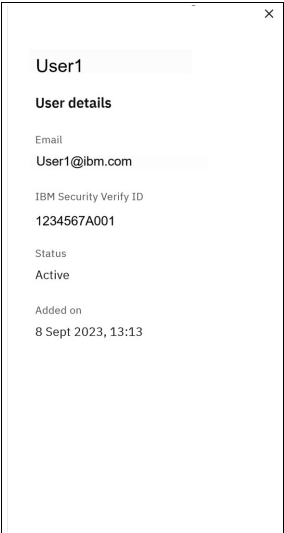


Figure 4-17 User Details

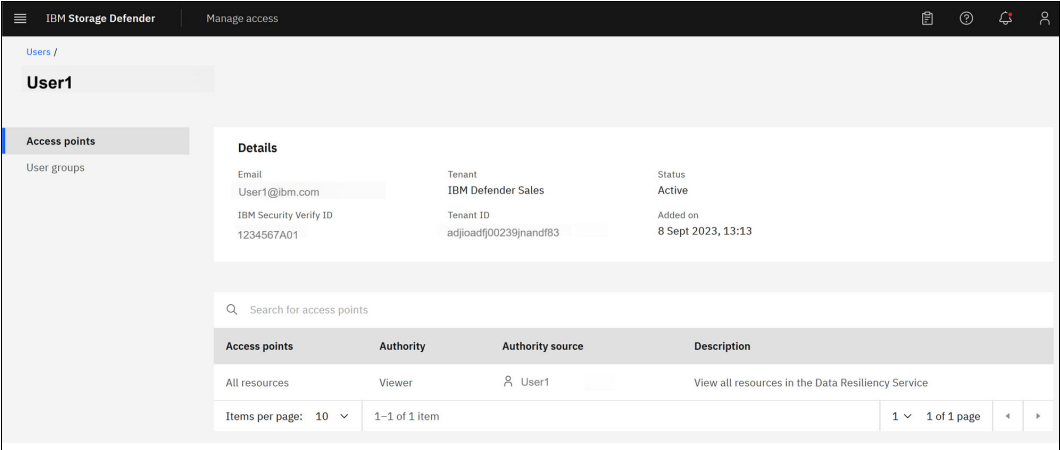


Figure 4-18 All User details

The administrator can click the Access Points to add further access points and responsibilities (Figure 4-19 on page 74).

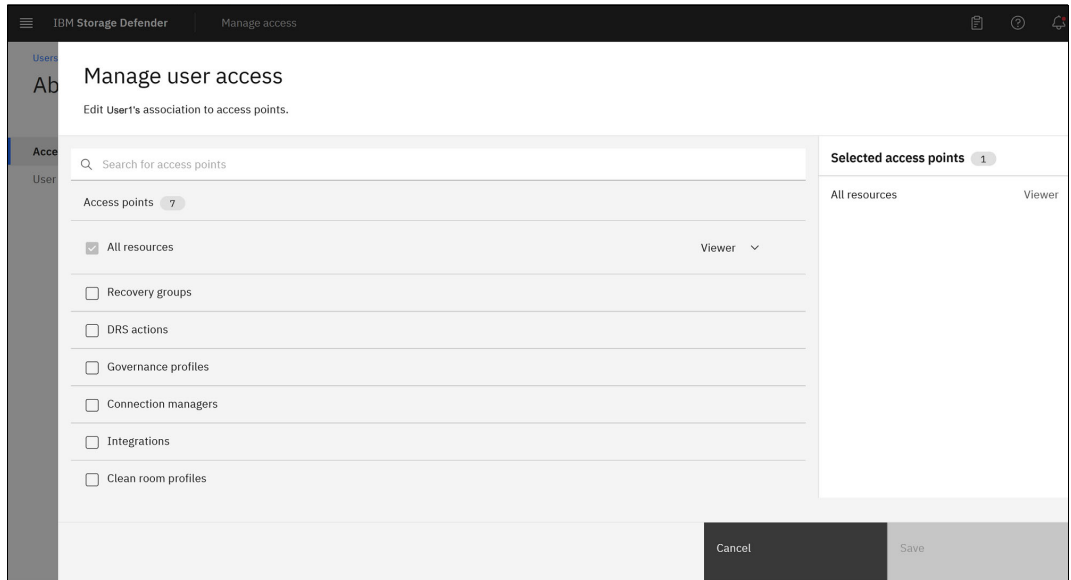


Figure 4-19 Manage user access options panel

This allows the administrator to create granular access points for users and divide roles across the environment. One example of this is providing only certain users with access to recovery groups and DRS actions without providing access to the Connection managers or Integrations because they might be supported by different administrators or editors in the organization.

## 4.2.2 User Groups

DRS also provides for the ability to create user groups that can have the same Access points and responsibilities assigned to them. After they are created, you can use the User groups to add or remove users from the group (Figure 4-20).

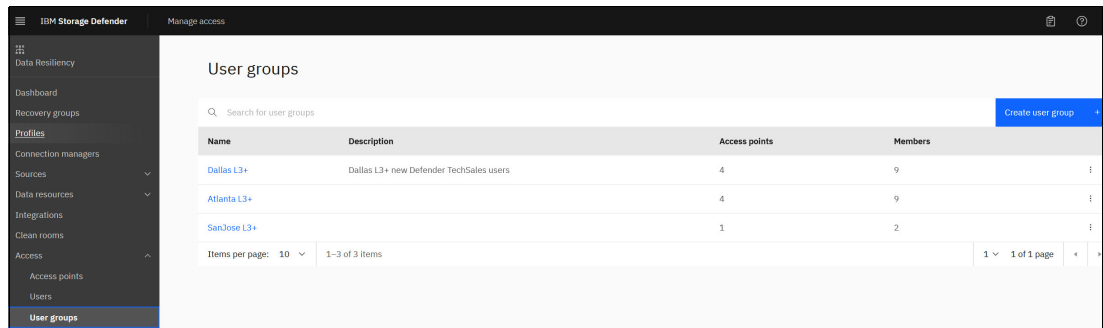


Figure 4-20 User groups management page

Select a User Group so that you can view the users and access points associated with that group.

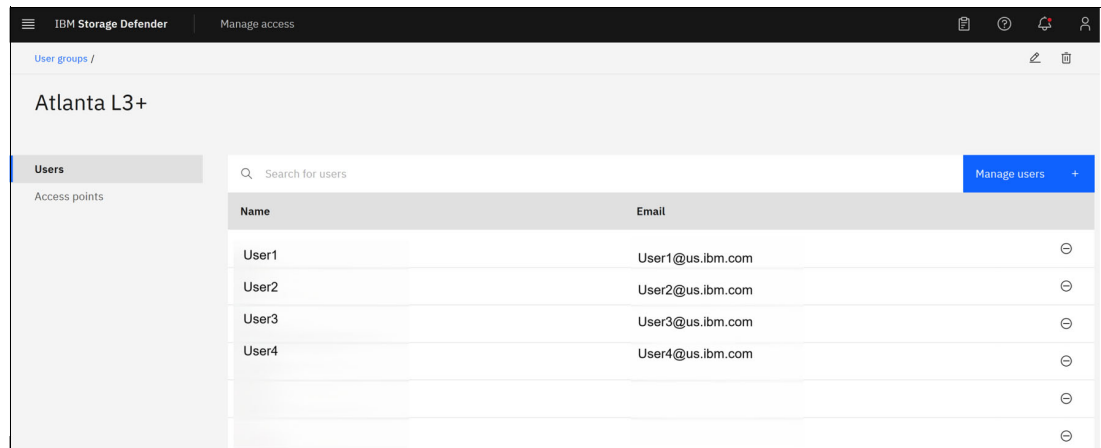


Figure 4-21 User Group member list page

Adding or removing Access points also provides the ability to increase or decrease access for all users in the group at once.

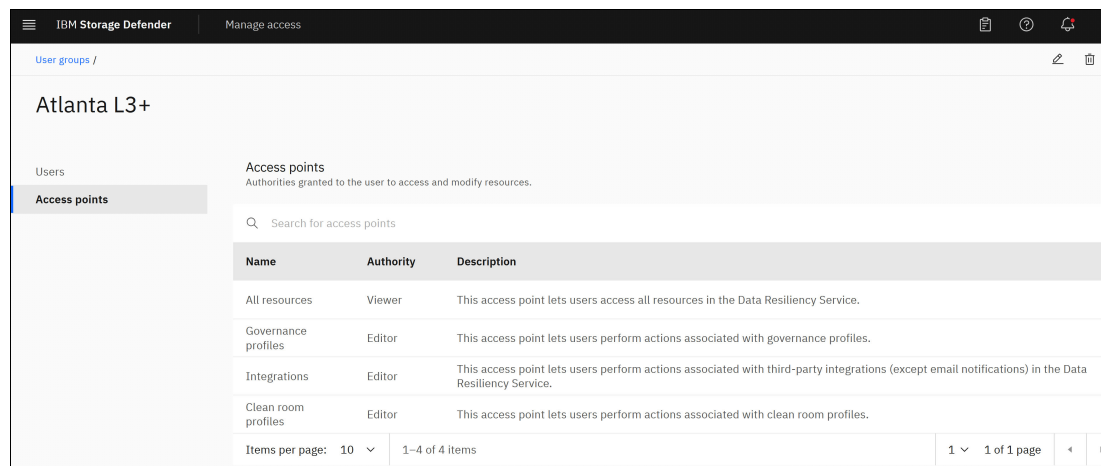


Figure 4-22 Manage access settings panel

## 4.3 Integrations for alerting

You can integrate DRS with the SIEM solutions IBM QRadar and Splunk to improve your security posture while also bridging the storage and security silos that sometimes exist in enterprises storage landscapes.

For more information, see the following resources:

- ▶ [Integrating Data Resiliency to QRadar SIEM](#)
- ▶ [Integrating Data Resiliency to Splunk® SIEM](#)

Figure 4-23 on page 76 shows the Integrations tab in the DRS dashboard.

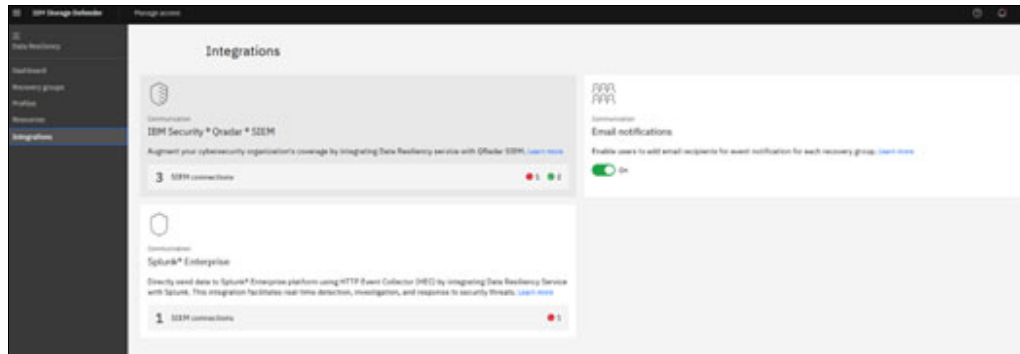


Figure 4-23 Integrations tab in DRS

## 4.4 Recovery testing and validation

You can use DRS to test and validate recovery points for a recovery group. You can test recovery points for a recovery group only when the status for the group is Ready, which means that the recovery group is complete. That is, it has a Governance plan that is assigned and Clean Room that is defined, and it has one or more recovery points.

Figure 4-24 shows the recovery group status of Ready and the details of Governance for the policy.

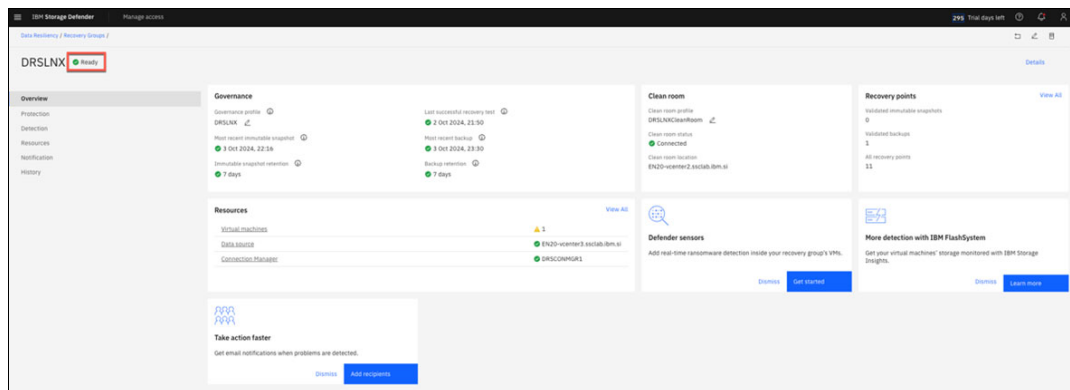


Figure 4-24 Recovery group details

Testing recovery points for a recovery group establishes the recovery plan. This plan is used in response to a cyberevent. From the recovery points of the selected recovery group, you can choose a recovery point that is required for testing. To select a recovery point, go to Recovery group details, and from the **Protection** menu, you see all recovery points (Figure 4-25 on page 77).

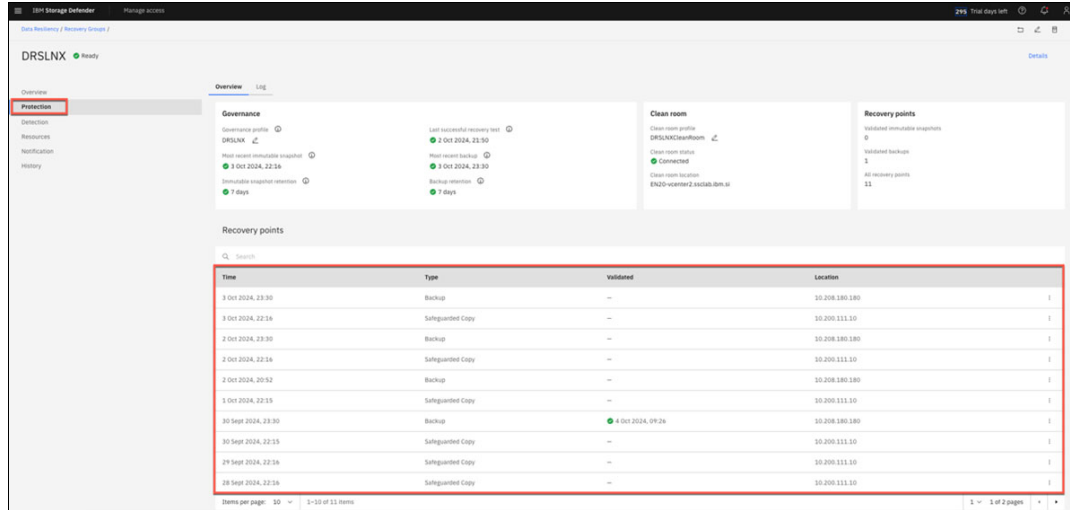


Figure 4-25 Recovery points details

You can use these recovery points to test or activate a recovery plan. Figure 4-26 shows the options that you can select for each recovery point.

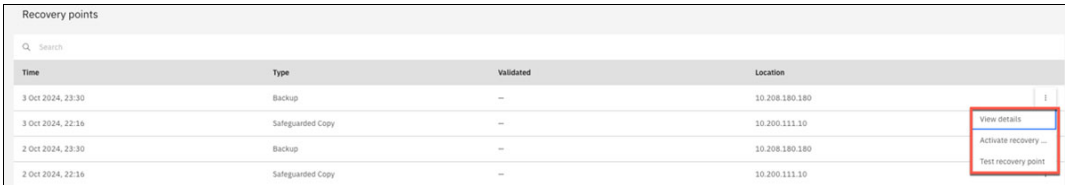


Figure 4-26 Recovery point details

Click **Test recovery point** to test a recovery of the virtual machines (VMs) that belong to the recovery group. These VMs are recovered by using the information that is stored in the Clean Room profile that is associated with the recovery group. Depending on the configuration of the Clean Room profile, the VMs either start and connect to the defined network or they do not. When the test recovery finishes successfully, the status of the recovery point is updated from “Recovery in progress” to “Awaiting validation”, as shown in Figure 4-27.

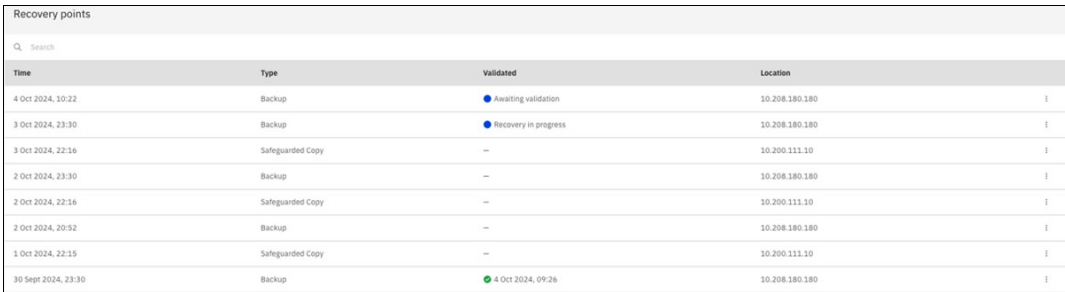


Figure 4-27 Recovery point status window

After the recovery point is recovered to the Clean Room and ready for validation, a blue box appears across the top of the page with a link to confirm that the validation passed or failed testing, as shown in Figure 4-28.

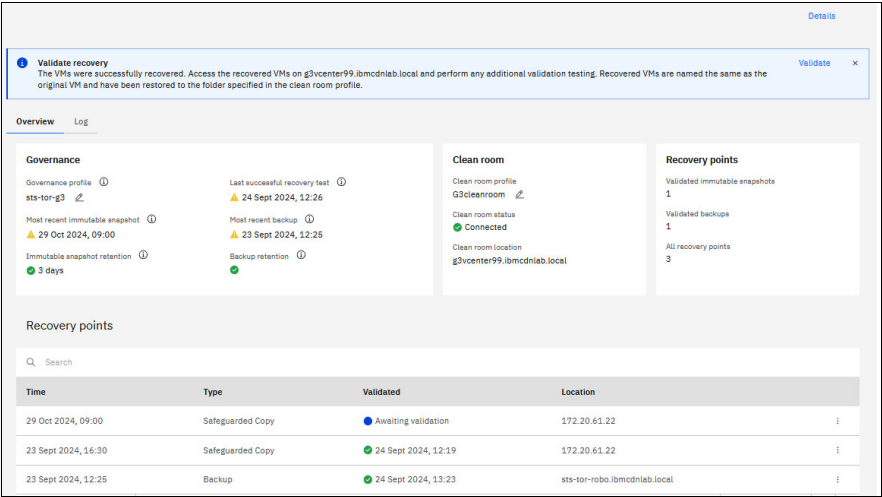


Figure 4-28 Test-only confirmation dialog

Figure 4-29 highlights the ability to validate the recovery point after the restoration of the recovery group to the Clean Room. You can identify the use case of starting the recovery and defining the status of the action as “Test Only” or whether the activity was part of a Recovery Plan resulting from a cyberincident. Then, you can mark it as valid or not.

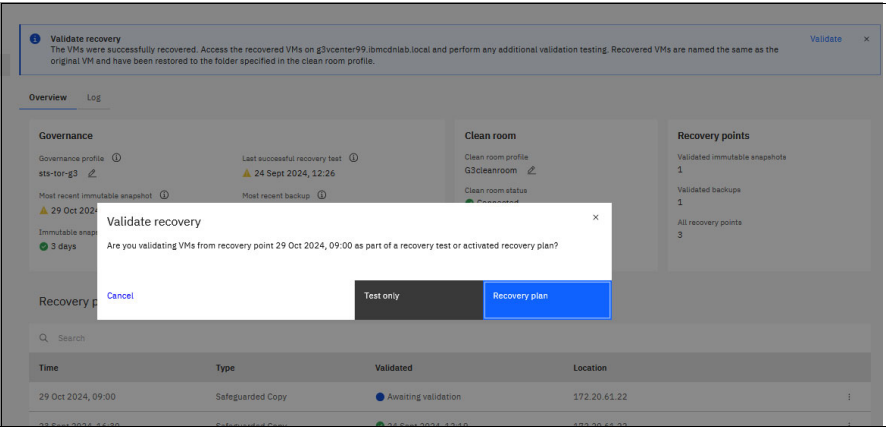


Figure 4-29 Validate recovery dialog

Figure 4-30 on page 79 confirms the results of the recovery to the Clean Room and confirms the results.

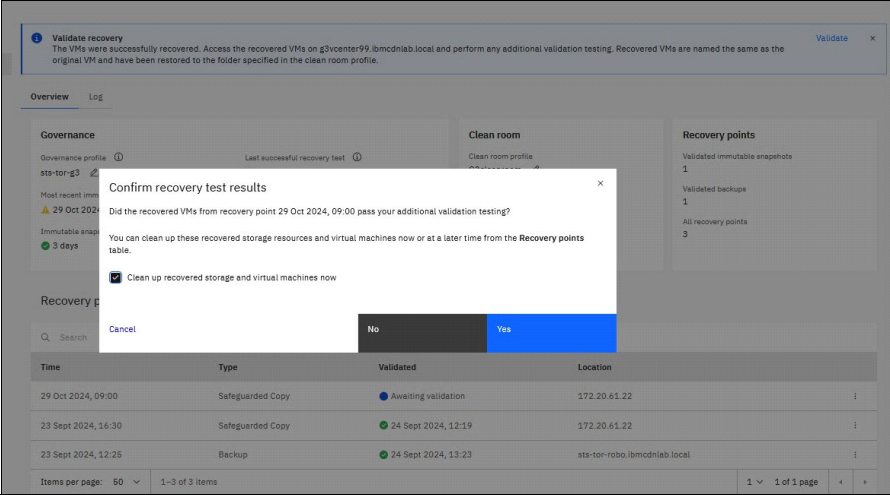


Figure 4-30 Confirm recovery results dialog

After you determine whether the recovery point is valid, you can mark it as **Valid** or **Not Valid**. As part of the validation process, the recovery points are kept in the history of the recovery group until their policies expire them from the inventory of their supporting services.

Depending on the decision that you make, the status of the recovery point is updated from “Awaiting validation” to “Validated” or “Not valid”.

Figure 4-31 shows the different categorizations of a recovery point.

Recovery points				
Search				
Time	Type	Validated	Location	
4 Oct 2024, 10:22	Backup	Not valid	10.208.180.180	
3 Oct 2024, 23:30	Backup	Not valid	10.208.180.180	
3 Oct 2024, 22:16	Safeguarded Copy	—	10.200.111.10	
2 Oct 2024, 23:30	Backup	Awaiting validation	10.208.180.180	
2 Oct 2024, 22:16	Safeguarded Copy	—	10.200.111.10	
2 Oct 2024, 20:52	Backup	—	10.208.180.180	
1 Oct 2024, 22:15	Safeguarded Copy	—	10.200.111.10	
30 Sept 2024, 23:30	Backup	4 Oct 2024, 09:28	10.208.180.180	

Figure 4-31 Validation status window with invalid recovery points

After the recovery test data is verified, if the cleanup option is selected in the test results window (Figure 4-32), the data is confirmed as validated. The system cleans up the VM data that was restored as part of the validation test. If the cleanup option is not selected, the VMs remain in the Clean Room and can be manually removed later.

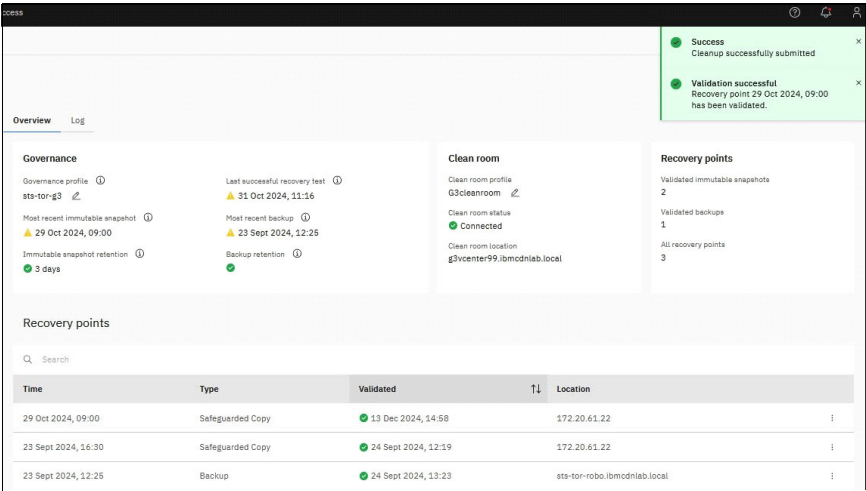


Figure 4-32 Validation and cleanup notification messages

## 4.5 Activating the recovery plan

The Recovery group option **Activate recovery plan** describes the recovery of resources that are associated with a recovery group. This option uses an existing and valid recovery point to recover your application after a cyberattack or disaster.

In contrast to the manual recovery test, the activate recovery plan process provides the flexibility to specify a new Clean Room profile for the recovery point. With this option, you can use a dedicated recovery environment to test the recovery point again and prepare a recovery point for a downstream promotion into your production environment.

Figure 4-33 shows the Activate recovery plan options where you select the required recovery plan.

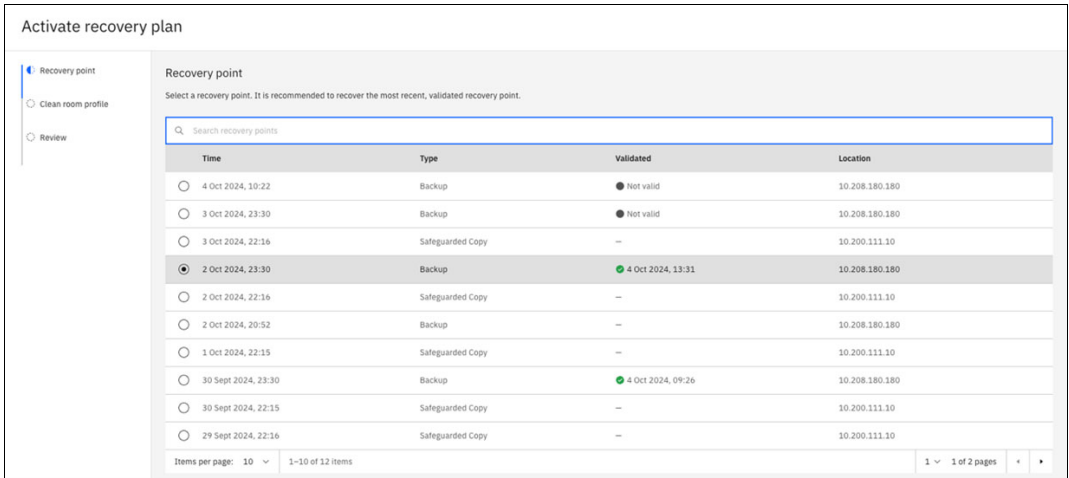


Figure 4-33 Activate recovery plan window

Now, specify a Clean Room profile (Figure 4-34). After you review the profile settings (Figure 4-35), click **Done** and wait for the recovery to complete.

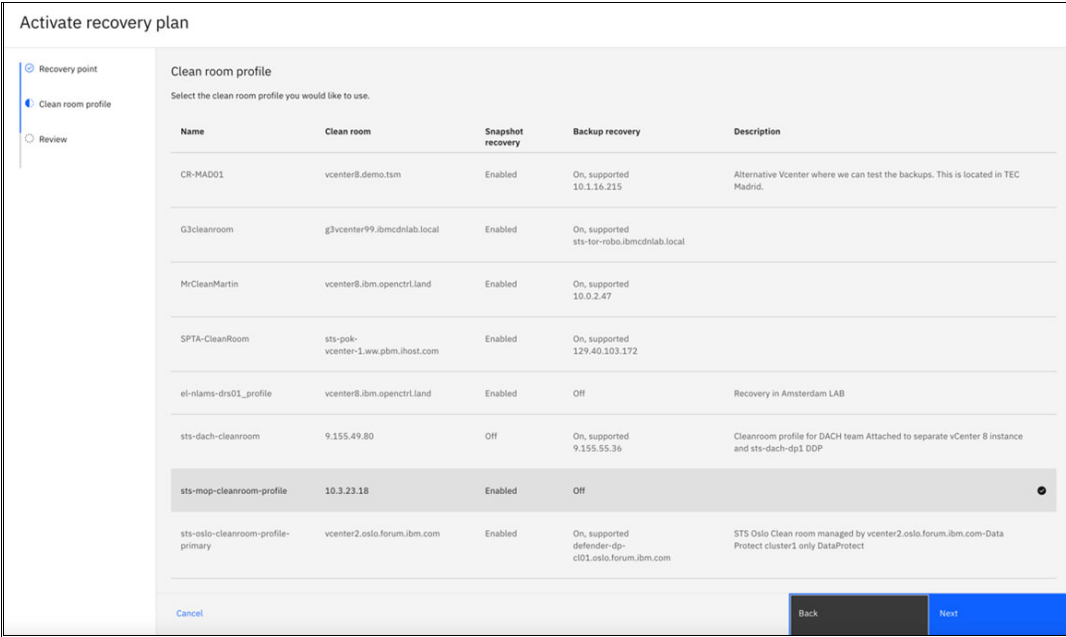


Figure 4-34 Activate recovery plan: Clean Room profile

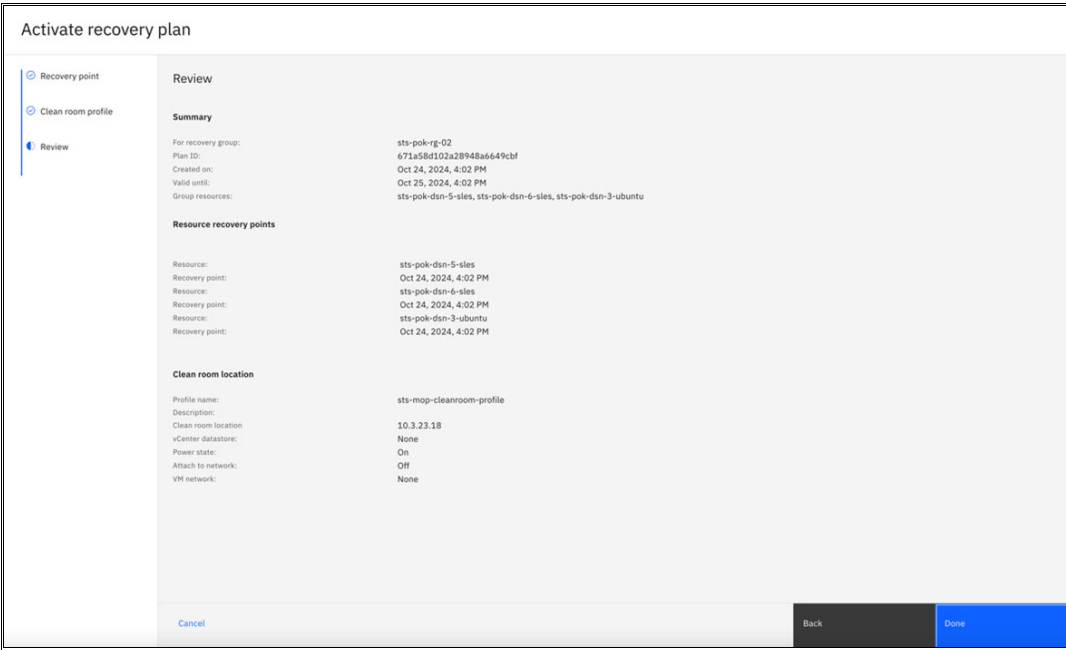


Figure 4-35 Activate recovery plan

After it is confirmed, the Recovery in Progress window in the Recovery group's Overview window shows the progress, as shown in Figure 4-36 and Figure 4-37.

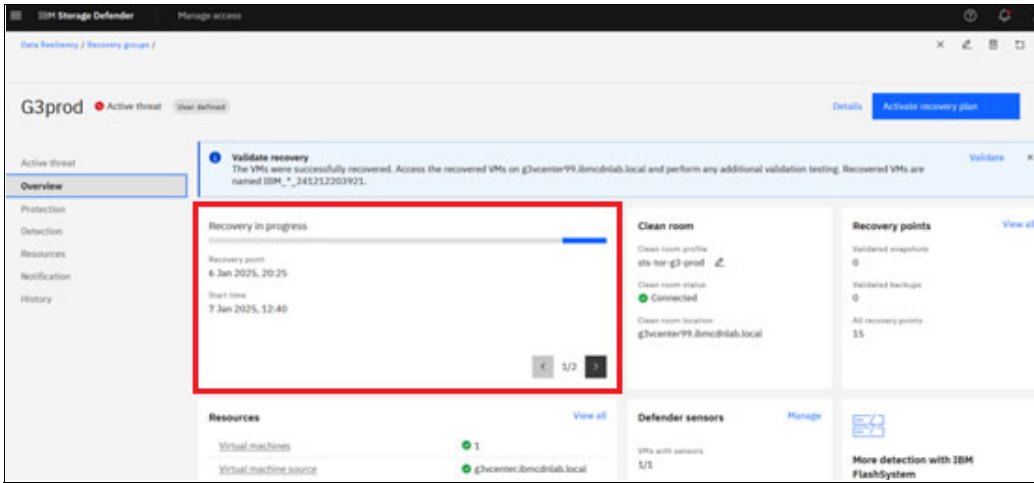


Figure 4-36 Recovery progress information in the Recovery group Overview window: Example 1

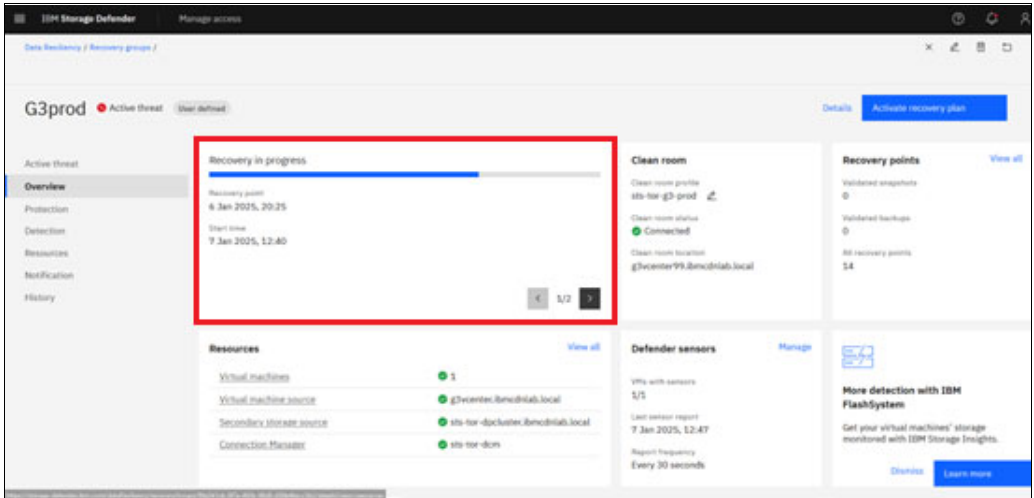


Figure 4-37 Recovery progress information in the Recovery group Overview window: Example 2

From here, after the recovery process is completed, you can access the VMs that were recovered to the Clean Room environment and return them to production as needed.

# Abbreviations and acronyms

<b>ACL</b>	access control list
<b>CLI</b>	command-line interface
<b>DR</b>	disaster recovery
<b>DRS</b>	Data Resiliency Service
<b>FCM</b>	FlashCore Module
<b>FQDN</b>	Fully Qualified Domain Name
<b>IOP</b>	I/O operation
<b>ML</b>	machine language
<b>PLG</b>	Product-Led Growth
<b>SID</b>	single system ID
<b>SIEM</b>	security information and event management
<b>SME</b>	subject matter expert
<b>SP4VE</b>	Storage Protect for Virtual Environments
<b>SaaS</b>	Software as a Service
<b>SecOps</b>	security operations
<b>VDisk</b>	virtual disk
<b>VM</b>	virtual machine







REDP-5744-01

ISBN 0738462225

Printed in U.S.A.

Get connected

