

IBM Storage Defender: Data Resiliency Service (DRS)

Ondrej Bláha

Erin Farr

Phillip Gerrard

Meghan Grable

Juan Carlos Jimenez

Alexis Kojic

Ranjith Rajagopalan Nair

Daniel Paulin

Ramakrishna Vadla

Christopher Vollmar



Storage

 Hybrid Cloud



IBM Redbooks

IBM Storage Defender: (DRS)

December 2024

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

First Edition (December 2024)

This edition applies to IBM Storage Defender Data Protect Version 7.1.1 and 7.1.2 and IBM Storage Defender Data Resiliency Service 2.0.9.

This document was created or updated on January 7, 2025.

© Copyright International Business Machines Corporation 2024. All rights reserved.

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

Contents

Notices	v
Trademarks	vi
Preface	vii
Authors	vii
Now you can become a published author, too!	ix
Comments welcome	ix
Stay connected to IBM Redbooks	ix
Chapter 1. Introduction	1
1.1 What is Defender Data Resilience Service (DRS)?	2
1.2 IBM Storage Defender Overview and Vision	3
1.2.1 Why Defender?	3
1.3 Defender components and functions	4
Chapter 2. IBM Storage Defender DRS and Architecture Overview	7
2.1 IBM Storage Defender DRS architecture and elements overview	8
2.1.1 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	12
2.2.3 Sensor control nodes	12
2.2.4 IBM Storage Defender Sensors	13
2.2.5 Recovery Groups	14
2.2.6 IBM Clean Room	18
2.3 Adding Resources In the Connection Manager and Creating Profiles In DRS	20
2.3.1 Adding resources in Connection Manager	20
2.3.2 Creating profiles in DRS	27
2.4 Auto-forward IBM Storage FlashSystem Ransomware Threat Alerts to IBM Storage Defender	33
2.4.1 IBM FlashCore Modules (FCM)	34
2.4.2 Integration between IBM Storage Defender Data Resiliency Service and IBM Storage Insights PRO	34
Chapter 3. IBM Defender Sensors	37
3.1 What do sensors do?	38
3.2 Installing Sensors	41
3.2.1 Installing the sensor control software	41
3.2.2 Adding a sensor control node	42
3.2.3 Removing a sensor control node	42
3.2.4 Installing an IBM Storage Defender sensor by using UI	43
3.2.5 Installing an IBM Storage Defender sensor by using CLI	44
3.2.6 Uninstalling an IBM Storage Defender sensor by using UI	45
3.2.7 Uninstalling an IBM Storage Defender sensor by using CLI	47
3.2.8 Requirements for IBM Storage Defender sensors	47
Chapter 4. Daily Administration, Alerting, Testing and Validation	49
4.1 DRS Dashboard	50

- 4.1.1 Resiliency Monitoring in the Dashboard 50
- 4.1.2 Recovery Group status 52
- 4.1.3 Governance Profile Status 53
- 4.1.4 Recovery Posture 54
- 4.2 User Management Profiles 55
- 4.3 Integrations for Alerting 55
- 4.4 Recovery Testing and Validation 56
- 4.5 Activating the Recovery Plan 60

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.

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

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.

Ansible, 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 Services (DRS). This IBM Redpaper publication will help you set up, tailor and configure this new offering. By doing this, users will be able to leverage new detection mechanisms for their environment to detect threats early, get a full view of the infrastructure by connecting both primary storage arrays like IBM FlashSystem® as well as secondary storage solutions for backup like Defender Data Protect and Storage Protect. Additionally, users can set up governance profiles to ensure their data is meeting internal or regulatory standards

Authors

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

Ondrej Bláha works as a Technology EMEA SME/Architect focusing on IBM® Storage SW with specialization in Data Resilience (Storage Defender strategy for primary and secondary workloads). He has been with IBM for more than 17 years, the last 10 years in several regional roles as SME/CTS or Technical Advisor for key IBM customers. Ondrej is an official IBM instructor for external IBM Software Training organizations and 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, lives in Prague, is married and has two beautiful daughters. He enjoys traveling, experiencing new cultures and loves discovering the culinary qualities of different regions

Erin Farr is a Senior Technical Staff Member in the IBM Storage CTO Office where she explores new technology for future products and shapes strategy in anticipation of industry trends. Her area of focus is Cybersecurity and Cyber Resiliency. She was instrumental in forming the vision for IBM Storage Defender and is passionate about helping customers prevent and recover from cyberattacks. Before joining Storage in 2021, she was the team lead for the IBM Z® Center for Secure Engineering for z/OS®. She also enjoyed product development for the majority of her career, in areas such as z/OS UNIX, analytics, virtualization management, and Open Source.

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

Meghan Grable is a global Growth Product Manager specializing in data management and resilience solutions, both SaaS and software-based, with a strong focus on Product-Led Growth (PLG) strategies. With over five 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 World-Wide Data Resiliency Product Manager based in Dallas, Texas. He is focused on defining roadmap, initiatives, and strategy within the various data resiliency software products that he manages alongside his team. Juan Carlos brings an end-to-end view to cyber resilience leveraging his expertise in both storage and security. Juan Carlos developed 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 based in Canada. With two years of experience in the IT storage and Cyber Resilience field, he holds a degree in Computer Engineering BEng from Toronto Metropolitan University.

Ranjith Rajagopalan Nair is a Software Architect at IBM India. He has worked in IBM for past 20 years, and 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 Masters degree in Computer Science from University of Kerala.

Daniel Paulin is a Storage Software Architect at IBM Croatia. An IT professional since 1997, he has worked as a system 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's storage solutions, particularly the IBM Storage Defender. His work is part of IBM's broader initiative to enhance cyber resiliency and storage security, ensuring 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 a Senior Technical Staff Member (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, World Wide Storage Data Resiliency Architect. Christopher is an IBM Certified IT Specialist (Level 3 Thought Leader) and Storage Architect. He is focused on helping customers design solutions to support Operational and Cyber Resiliency on primary and backup data to complement their Cyber Security practices. He is an author of several IBM Redbooks®, an Enterprise Design Thinking® Co-Creator, and a frequent speaker at events like IBM THINK, and TechXchange.

Special thanks to the following people for their contributions to this project:

Christian Burns
Principal WW Storage Developer / Architect - Data Resilience, IBM

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our 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

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

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

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

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

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

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



Introduction

In this chapter we introduce the new IBM Storage Defender – Data Resiliency Service. Additionally, we share this solution’s overview and vision for the future. Later in this document, we will also cover the different functionalities of the solution, how to set them up, configure them, and run them in order to drive the most value.

In this chapter:

- ▶ 1.1, “What is Defender Data Resilience Service (DRS)?” on page 2
- ▶ 1.2, “IBM Storage Defender Overview and Vision” on page 3
- ▶ 1.3, “Defender components and functions” on page 4

1.1 What is Defender Data Resilience Service (DRS)?

Today, organizations face severe threats to their data as the number of cyberattacks increased and malicious actors continue to become more sophisticated. According to the IBM X-Force® Threat Intelligence Index 2024 report, 43% of all reported incidents involved malware, making it the most common threat, while 20% were attributed to ransomware attacks. In addition to malware, IT organizations are threatened by natural disasters, system failures, human errors, and even sabotage. Events like these may result any number of issues including financial losses and even harm to customer trust if sensitive data is compromised.

IBM Storage Defender Data Resilience Service (DRS) is a purpose-built cloud-based data resilience platform designed to help organizations quickly restart essential business operations in the event of a cyberattack or other unforeseen catastrophic event. DRS provides data resilience and compliance, early threat detection, and safe and fast recovery orchestration for data stored across primary and secondary storage. The DRS software helps to detect and respond to cyber threats, such as malware and ransomware attacks, and allows for rapid recovery of data in case of a security breach or data loss. This allows administrators to 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 data backup, data management, disaster recovery, and data isolation to help organizations protect their data from cyber threats and unexpected disruptions. Additionally, it provides rapid recovery of data and applications in the event of a disaster or data loss, minimizing downtime and ensuring business continuity.

Note: The term 'secondary storage' used above denotes a secondary or backup storage location that provides the ability to leverage and use copies of data in place prior to the data being recovered.

IBM Defender DRS provides the following benefits:

- ▶ **Data Resilience and Compliance**
 - Set resiliency standard to meet compliance across data estate
- ▶ **Early Threat Detection**
 - Near real-time filesystem monitoring, backup anomaly analysis, FlashSystem inline detection, and recovery time scanning.
- ▶ **Safe and Fast Recovery**
 - Air gapped data, immutable snapshots, and clean room recovery enables you to confidently and quickly recover your business operations.
- ▶ **Connect Storage Ops and SecOps**
 - Collects storage and security events to send alerts to support staff and other security tools. Deep integration with IBM QRadar® and Splunk.
- ▶ **FlashSystem Integration**
 - Understand threats down to the FlashSystem volume and virtual machine level speeding up identification and time to initiate remediation. Automatically trigger proactive safeguarded copy snapshots to limit damage and automatically recover to a clean room for testing.
- ▶ **IBM Defender Data Protect Integration**
 - Catalog IBM 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 security operations, storage, and infrastructure tools, DRS provides the capability to monitor end to end data movement and quickly supply critical information allows for teams to make the most intelligent decision on recovery strategies. DRS presents data resilience and recoverability options across primary and secondary 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

In this section we'll cover the vision behind IBM Storage Defender, along with a high-level overview of the functions Defender provides to meet those goals and customer needs.

1.2.1 Why Defender?

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

- ▶ The impact of a natural disaster is regional whereas a cyber attack can be global.
- ▶ Depending on the location of backup data, the expectation may be that typically the data would not be impacted by a natural disaster, but with a cyber attacks 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; 55% in North America¹.
- ▶ Additionally, the probability of a natural disaster is relatively low when compared to that of a cyber attack.

In addition, while an enterprise might practice their DR recovery, our research showed that very few were practicing Cyber Recovery, which includes aspects outside of traditional DR such as:

- ▶ Playbooks to ensure seamless interaction with Incident Responders
- ▶ Antivirus scanning during recovery to avoid re-introduction of dormant malware
- ▶ Practicing identification of good data copies at scale. Cyber attacks aren't as instantaneous as a power outage for example, and the actual points of impact may vary across available recovery points.

The industry has responded to some of these threats with solutions that provided additional protections such as air gap or immutability. Initially many secondary 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 getting added into Storage across the industry. At first, only secondary storage vendors were providing this capability, but IBM saw a need for detection in primary storage as well.

¹ 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

However, as IBM investigated this trend, it quickly became apparent there were a series of concerns which needed to be addressed:

- ▶ If an attack is detected, who is expected to respond? Nobody expects storage admins or data protection teams to suddenly become Incident Responders.
- ▶ Disparate solutions make it difficult to identify and locate the last good copy. Recovery points can be primary storage snapshots or secondary storage backups. Often these are managed not only by different tools, but by different teams. Also, if an incident is actively occurring, a storage admin or incident responder may not have a holistic view across both primary and secondary storage.
- ▶ If an enterprise takes backups once a day, would backup-based detection (only) be fast enough to detect issues?
- ▶ How to determine the scope of damage? Which systems were impacted? The timeline?
- ▶ While storage-based threat detection is important, it is unlikely that someone would swap out their current solution just to get access to threat detection. How can this need for additional features be met?

IBM recognized the need for:

- ▶ A way to provide these cyber recovery features, one that works with existing investments and current storage solutions.
- ▶ The ability to provide a holistic view across both primary and secondary storage, for not only recovery but also threat detection.
- ▶ Features that specifically address cyber recovery. This includes 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 of these needs. Storage Defender DRS is a SaaS management pane, designed and intended to integrate with and sit above an enterprise's existing storage system investments. DRS allows for this holistic view across primary and secondary storage, while providing advanced ransomware detection and recovery features that are needed to address modern threats in storage environments.

1.3 Defender components and functions

IBM Storage Defender Data Resilience Service (DRS) is a multi-faceted offering that has several functions which 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. Additionally, 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.

FlashSystems offers protection through immutable copies of data known as safeguarded copies, which are isolated from production environments and cannot be modified or deleted through user error, malicious actions, or ransomware attacks. Defender Includes IBM Storage hardware integration with FlashSystems and SVC to include the use of safeguarded copy as part of the DRS configuration.

IBM Defender Data Protect offers an immutable secondary 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 an SEIM like Splunk or QRadar, advanced notification aggregation allows for crucial information to be available and used for initiating the next steps between infrastructure and security operations (SecOps) teams. Providing needed information use when deciding whether recovery plans should be implemented immediately or how best to address threats.

Clean room isolation provides the ability to ensure the backups are clean and malware free prior to returning them to a production environment. As a customer managed resource, DRS provides guided testing workflows to recover, test, and isolate backups before pushing to production systems ensuring the ability to confirm 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 to ensure the availability of business operations, this 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 secondary storage. This technology allows users to effectively detect and respond to cyberattacks and other unforeseen threats to storage environments. When put together these features allow Defender DRS to help provide the peace of mind needed to successfully navigate unpredictable events and ensure the continuity of vital business operations and processes.



IBM Storage Defender DRS and Architecture Overview

The following chapter describes 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 as well as alerting.

- ▶ 2.1, “IBM Storage Defender DRS architecture and elements overview” on page 8
- ▶ 2.2, “IBM Storage Defender Connection Manager” on page 10
- ▶ 2.3, “Adding Resources In the Connection Manager and Creating Profiles In DRS” on page 20
- ▶ 2.4, “Auto-forward IBM Storage FlashSystem Ransomware Threat Alerts to IBM Storage Defender” on page 33

2.1 IBM Storage Defender DRS architecture and elements overview

IBM Storage Defender provides end-to-end data resiliency, and it is important to understand Data Resiliency Service (DRS) architecture and the DRS elements as it will help you to properly plan, test, and recover your critical data.

The architecture will show you how DRS fits into IBM Storage Defender and what are the elements that makes 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 IBM Storage Defender Mission

IBM Storage Defender Data Resiliency Service (DRS) is an optional component within IBM Storage Defender solution that provides cyber-resiliency capabilities for the management of primary and secondary data and workloads. DRS introduces concepts to simplify the recovery of complex applications, automated recovery tests, and perform validation of primary and secondary data. In addition, DRS has the ability to send notifications if anomalies are detected and indicate when the trustworthiness of existing primary and secondary data sources have decreased.

DRS is a combination of cloud-based SaaS 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 about your primary and secondary data, and data sources like VMware, while the data itself stays on premises. The telemetry data is communicated to the DRS which helps secure and recover the data that is important to you.

DRS can surface and aggregate the detection of operational threats on your production data. Currently this includes the following system level detections:

- ▶ Detection on the file system level using IBM Storage Defender Sensor technology
- ▶ Detection on the storage block level using IBM Storage FlashSystem - FlashCore Module 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 allows DRS to perform automated test recoveries and to verify automatically if the protection policies setup 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, the Clean Room Profile that defines the environment which 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.

A dashboard is provided with DRS, which presents information relevant to cyber-resiliency in a consolidated and easy to digest view. This dashboard displays the configured Recovery Groups and any potential informational or warning messages related to its cyber-resiliency requirements being met for each of those groups. From the DRS dashboard you can access all of the capabilities and configuration options for the service.

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

- ▶ Supports software protection for multiple operating systems inside a VMware environment.
- ▶ Deploy anomaly based sensor agents on VMware VMs (Defender Sensors).
- ▶ Integrates and aggregates the hardware detection capabilities of IBM Storage FlashSystem to provide the ability to receive alerts from IBM Storage Insights and with Storage Defender. These alerts that can be sent out through integration with QRadar and Splunk SIEM solutions.
- ▶ Provides the capability to recover data from a more recent point in time by creating a SafeGuarded Copy (immutable hardware snapshot) on IBM Storage FlashSystem.
- ▶ Provides the ability to recover IBM Defender Data Protect backups into the clean room for testing as part of the Recovery Group's collection of recovery points.
- ▶ Gives clients a dashboard that can help them to better understand inconsistencies between their primary storage copies and backup copies for the same workload or application.
- ▶ Additional dashboard features include:
 - Ability to 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 virtual machines, data sources, recovery locations, and connection managers
 - A license usage overview highlighting the number of recovery groups and deployed sensors

2.1.2 Architecture overview

IBM Storage Defender Data Resiliency Service (DRS) is a component of IBM Storage Defender which runs in a cloud, using the on-prem IBM Storage Defender Connection Manager to get inventory of available and important resources in a data center.

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

- ▶ Data Sources (IBM FlashSystem, IBM Storage Defender Data Protect, VMware vCenter)
- ▶ 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, following DRS elements can be created:

- ▶ Recovery Groups with resources
- ▶ Profiles which include:
 - Governance profiles
 - Clean Room profiles

The Figure 2-1 on page 10 shows high level overview of DRS.

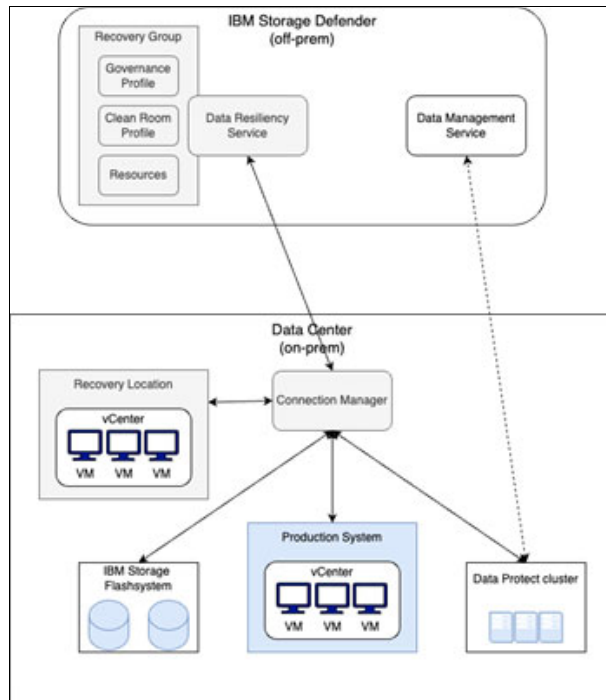


Figure 2-1 IBM Storage Defender DRS - high level overview

2.2 IBM Storage Defender Connection Manager

IBM Storage Defender Connection Manager (Connection Manager) is used to connect to your local environment, and it does inventory operations. It is also used to do test recovery and recovery operations.

The Connection Manager must be installed in an on-prem Data Center or cloud instance and is provided in an OVA format which to be deployed in your local VMware vCenter. Inside the Connection Manager, Red Hat Enterprise Linux is used as the underlying operating system but is part of the installation. The Connection Manager software is built to become active and get connected to your local resources and the IBM Storage Defender DRS that runs in the cloud quickly with less initial configuration prior to initial use.

It can be deployed from OVA or on a bare metal server, you can login to Connection Manager and begin to add Connections.

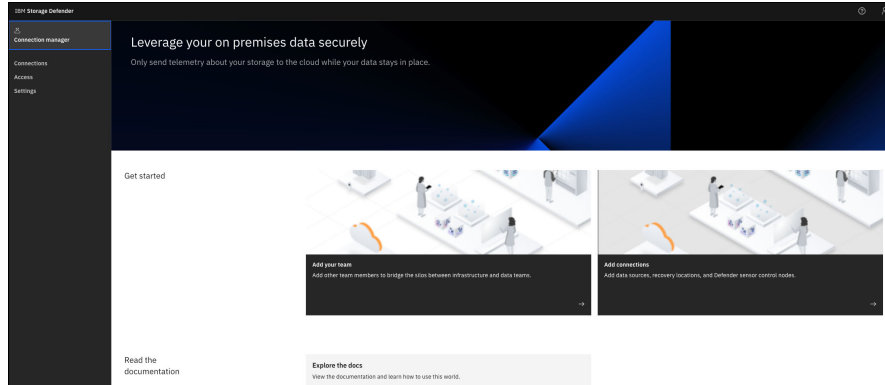


Figure 2-2 IBM Storage Defender Connection Manager

Connections in Connection Manager include Data Sources, Recovery locations and Sensor control nodes. Typically, only one Connection Manager should be deployed at each physical location. Data sources must be registered to the Connection Manager instance located in the same physical location.

Connection Manager also includes Job manager which communicates internally with various workload agents running in Connection Manager, and also catalogs Safeguarded Copies for IBM FlashSystem.

By using built in SIEM agent, Connection Manager integrates with on-premise IBM 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 IBM Storage Defender DRS. Connection Manager supports the following Data Sources:

- ▶ IBM FlashSystem
- ▶ IBM Storage Defender Data Protect
- ▶ VMware vCenter

For IBM FlashSystem, Connection Manager gathers 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 scan results to IBM Storage Defender DRS. It also coordinates recovery of VMs protected by IBM Storage Defender Data Protect.

The following figure shows example of Data Sources in Connection Manager:




Figure 2-3 IBM Storage Defender Connection Manager - Data Sources

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. The Recovery Locations, like hypervisors that you connect to the IBM Storage Defender Connection Manager, are inventoried automatically. The inventory metadata is transferred to the IBM Storage Defender DRS.




Figure 2-4 IBM Storage Defender Connection Manager - Recovery Locations

The Figure 2-5 on page 12 shows an example of the relation between your production environment and the Recovery Location.

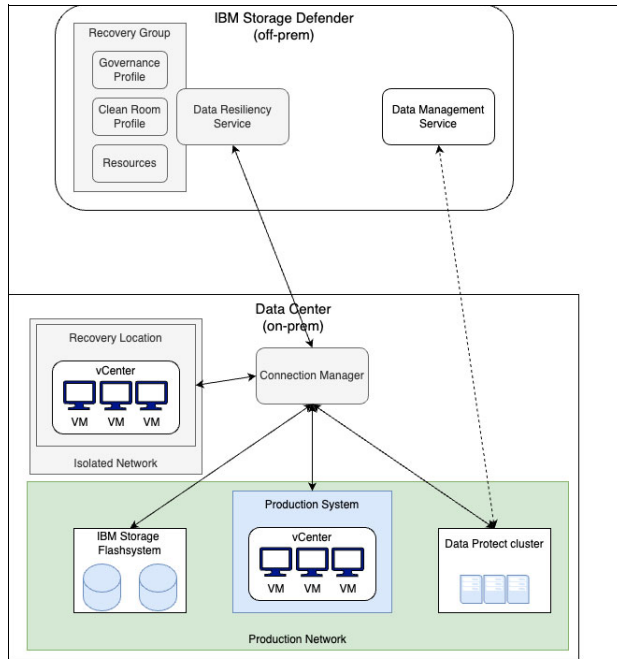


Figure 2-5 Recovery Location example diagram

2.2.3 Sensor control nodes

IBM Storage Defender 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 which get installed on resources like virtual machines.

The sensor control node hosts the sensor software and distributes it to the virtual machines that have sensors installed. These sensors observe the systems 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 IBM Storage Defender DRS.

Connection Manager comes with a built in control node so you can start adding sensors right away. However, if you'd like to use your own control nodes, you can add them through the Connection Manager and use the provided Ansible playbooks to manage the sensors.

Figure 2-6 on page 13 illustrates the sensor control architecture:

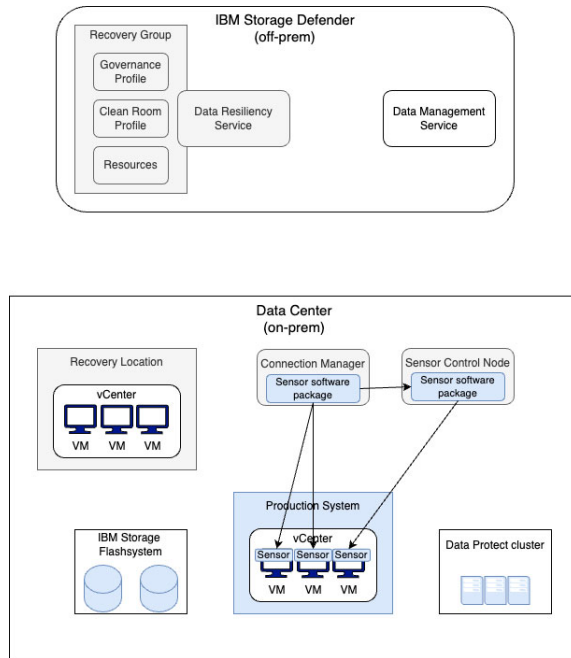


Figure 2-6 Sensor Architecture overview

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 they are installed on. IBM Storage Defender sensors are part of the IBM Storage Defender product, and can be deployed on virtual machines that are part of a recovery group. When the sensors are deployed, the sensors automatically send metadata to the IBM Storage Defender DRS.

A high level example of the workflow and data path for DRS sensors is shown in Figure 2-7 on page 14.

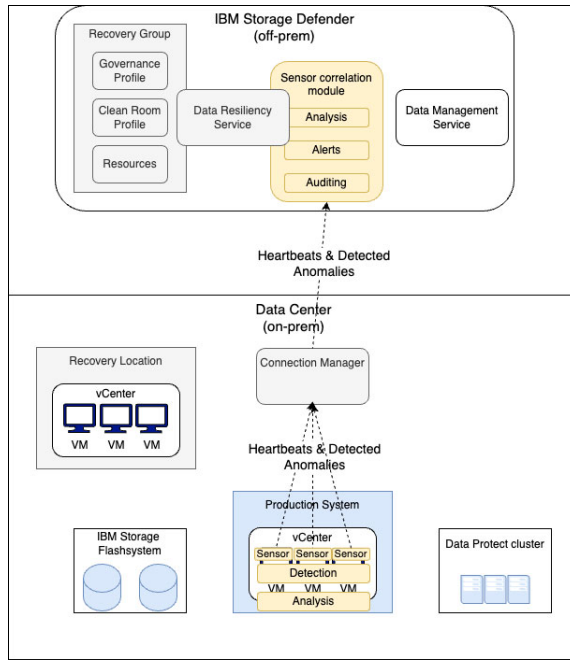


Figure 2-7 Defender Sensor Workflow

IBM Storage Defender sensors operate 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.

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. DRS correlates the information with recovery groups in your tenant.
4. When sensor heartbeat information is missing or when an anomaly is detected for 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.

2.2.5 Recovery Groups

Recovery Groups are a core concept within IBM Storage Defender DRS which 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 with VM1, VM2, VM3, and VM4 then IBM Storage Defender DRS determines whether it can find backup snapshots for these VMs in the secondary data sources within the same location (data center). When IBM Storage Defender DRS correlates the primary and secondary resource data for the assigned VM's, it proceeds to test the recovery group based on the policy and clean room profile settings.

In the IBM Storage Defender DRS dashboard you can find the details about recovery group. The Figure 2-8 shows Recovery Group details.

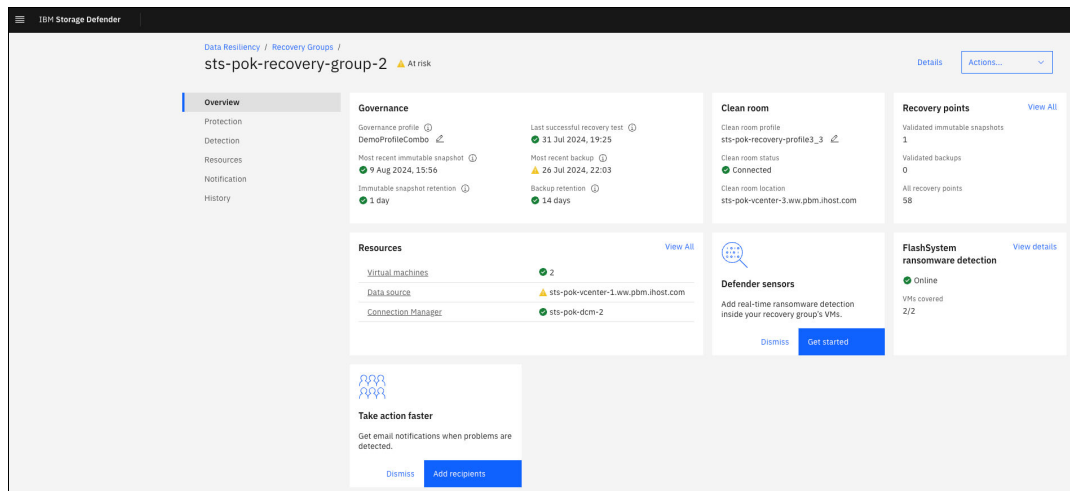


Figure 2-8 DRS Recovery Group Details

Profiles

The governance and clean room profiles are used to define and set recovery objectives of recovery groups and “recovery target environments.

Governance profiles are created and applied to the recovery group and allow for specific recovery objectives to 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 also allows you to specify a threshold time that must be elapsed before the next recovery test will be performed for the recovery group. The separate recovery objectives can be defined for IBM Storage FlashSystem and IBM Storage Defender Data Protect independently.

The governance profile definition allows one of the following three 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 recovery points
- ▶ Observation of both the recovery objectives for IBM Storage FlashSystem recovery points and IBM Storage Defender Data Protect recovery points

The test frequency objective is optional for all use cases. Figure 2-9 shows an overview of recovery objectives configured in the governance profile.

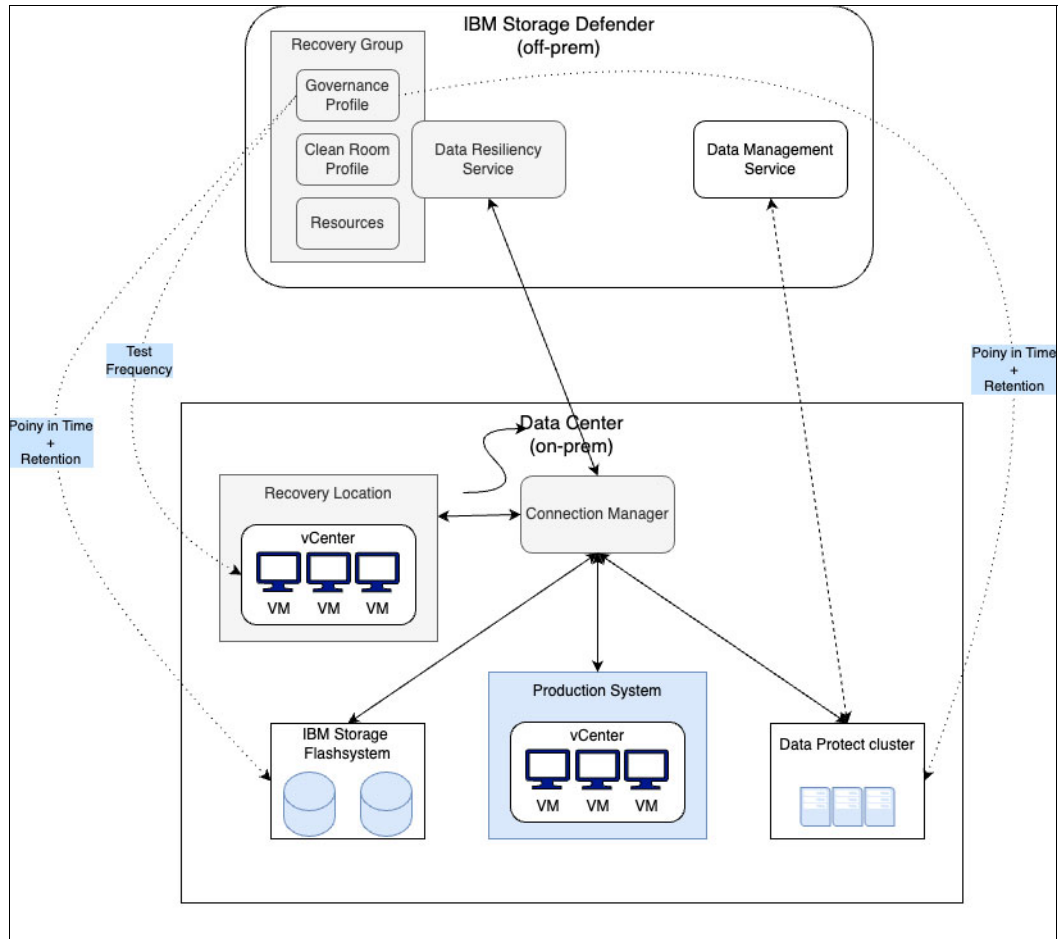


Figure 2-9 Recovery objectives configured in the governance profile

The clean room profiles connect the recovery groups that belong to resources in the production environment with configuration and resources that are defined in IBM Storage Defender 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 2-10 on page 17 illustrates the logical connection between the different components.

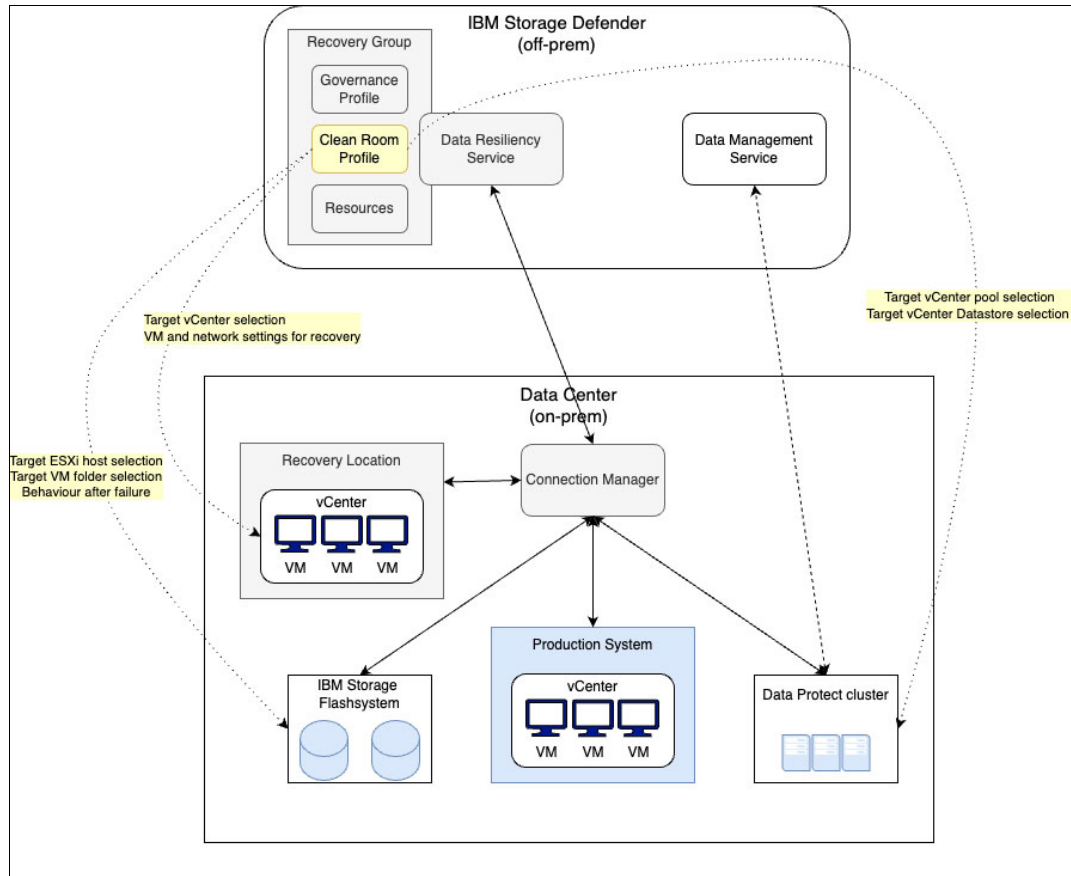


Figure 2-10 Clean Room objectives configured in the Clean Room profile

To ensure the successful recovery of the recovery group that is assigned to the specific clean room profile, configuration requirements must be met. The configuration of a clean room profile allows the usage of the profile for one of the following three different use cases:

1. Recovery from IBM Storage FlashSystem safeguarded snapshots
2. Recovery from IBM Storage Defender Data Protect backup copies
3. Recovery from both IBM Storage FlashSystem safeguarded snapshots and IBM Storage Defender Data Protect backup copies.

Important: If these requirements are not met, the recovery of the virtual machines that belong to the specific recovery group will 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 may have different requirements for their recovery. This is specifically important when recovering from IBM Storage FlashSystem, as the requirements for network infrastructure, mapping of volumes, or SAN zoning may be different. Therefore, it may be beneficial to implement multiple clean room profiles with different configurations to provide you with more flexibility for the recovery scenarios that you want to implement for different recovery groups.

Resources

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

- ▶ Virtual Machines
- ▶ Connection Managers
- ▶ Data Sources
- ▶ Clean Rooms

Resources are added to Recovery Groups during its creation, and are checked during inventories by Connection Manager.

2.2.6 IBM Clean Room

Clean Room plays an important role in the IBM Storage Defender solution by allowing for the recovery of workloads into an isolated environment. This concept introduces the ability to safely restore and investigate resources that might be contaminated with viruses, or other malware without the risk of infecting your production environment.

Protected virtual machines are able to be recovered into the associated Clean Room for verification prior to 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 datastores that are mapped from data protection solutions, a DMZ is implemented to allow access to the isolated portions of a clean room. Figure 2-11 displays the high-level structure of a Clean Room environment.

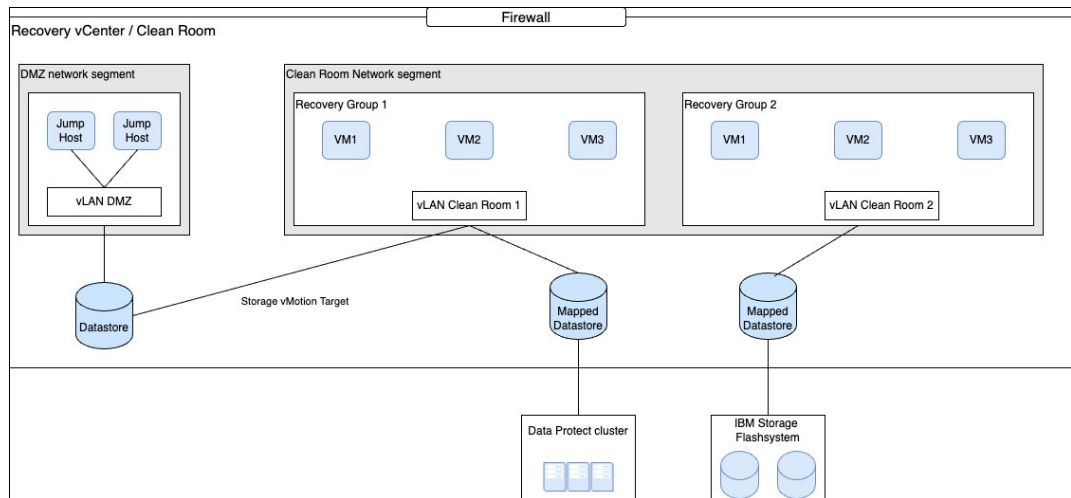


Figure 2-11 Clean Room environment schema

Isolation is an important aspect to consider when implementing clean room functionality. It is important to remember there are multiple dimensions, such as isolation of infrastructure, network, and access management. In addition to the isolation, you must implement monitoring and logging of a clean room environment as well. The following sections describe the different aspects of isolation for a clean room environment.

Infrastructure isolation

The isolation of the infrastructure is an important aspect of a clean room environment. Isolation for physical resources refers to physical separation in the form of a set of computer hardware that is used for a hypervisor independent from any production environment. When a cloud service provider is used, isolation refers to a logical separation configured using different cloud accounts.

Network segmentation and monitoring

Network segmentation comprises multiple aspects:

Logical separation and subnetting: In addition to the recovered virtual machines, the clean room environment contains systems that are used for tools and management. You need to 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, you must use a dedicated VLAN for each recovery group. Apart from the breakout prevention, the positive impact on 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. In addition, enhance security by enforcing rules that are based on source, destination, and port.

Security zones and critical infrastructure protection: Establish security zones, including a De-Militarized Zone (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 ensure visibility and timely detection of security incidents. In addition, 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 to provide an extra layer of security. This can range from using a different set of administrative identities to a total separation of the 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 the user's permissions prevent a user from influencing both environments.

The implementation of auditable logging for all operations in the clean room makes sure that any operation on the recovered data is traceable. This includes the creation and configuration of the clean room, clean room operations such as recovery, data masking, anonymization, or temporary production usage of the data.

Compliance and legal compliance

The bounded usage scope of a clean room environment allows for 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 ensure 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 operations logs allow for

analysis or development and can be used to document events or actions taken, while a comprehensive review of this documentation can be used to help audit the regulatory compliance status of a company and confirm if requirements are being met.

For further information, there is a detailed blueprint for a clean room available here:
<https://www.ibm.com/support/pages/ibm-storage-defender-clean-room-environments>

2.3 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 IBM Storage Defender DRS with these resources.

2.3.1 Adding resources in Connection Manager

After deploying Connection Manager, you can login to the Connection Manager GUI and add resources to be managed by IBM Storage Defender DRS.

To login to Connection Manager enter the following link in a browser `https://<ConnectionManager IP or hostname>/login`. This will bring up the login page where you can enter the desired username and password. Confirm the login by entering confirmation code from your authentication application.

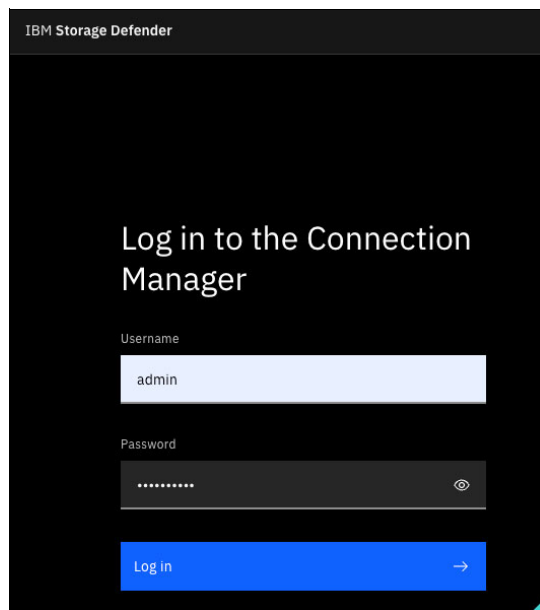


Figure 2-12 Connection Manager Login

Figure 2-13 on page 21 shows the Connection Manager dashboard, from here you can add resources from the Connections menu.

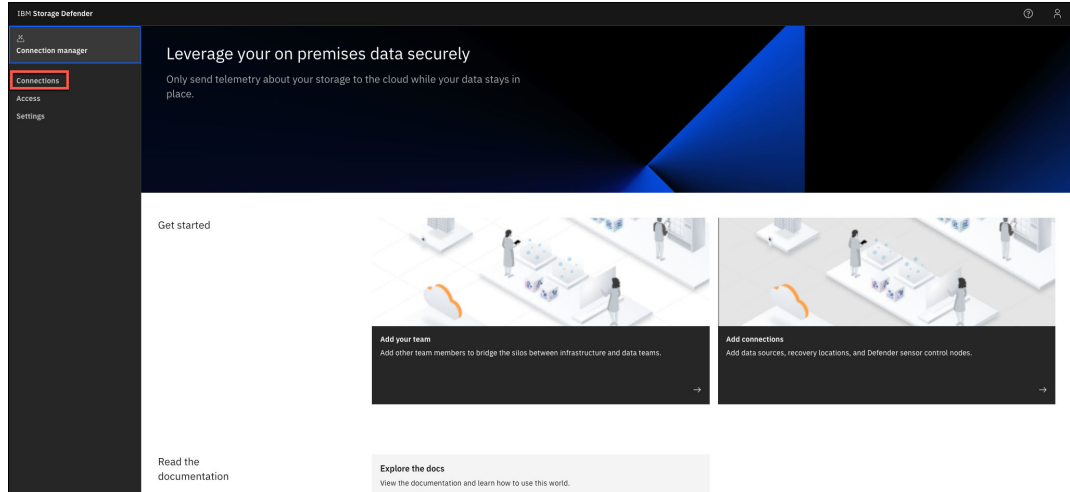


Figure 2-13 Connection Manager - Connections

From the Connections dashboard data sources, recovery locations and sensor control nodes can all be added to the DRS configuration.

Adding data sources

To add data sources in the Connections dashboard select the Data sources tab and click "Add a data source". The wizard will open in the right side of the window to guide you through the process (Figure 2-14). Select the type of data source you would like to add and click Next.

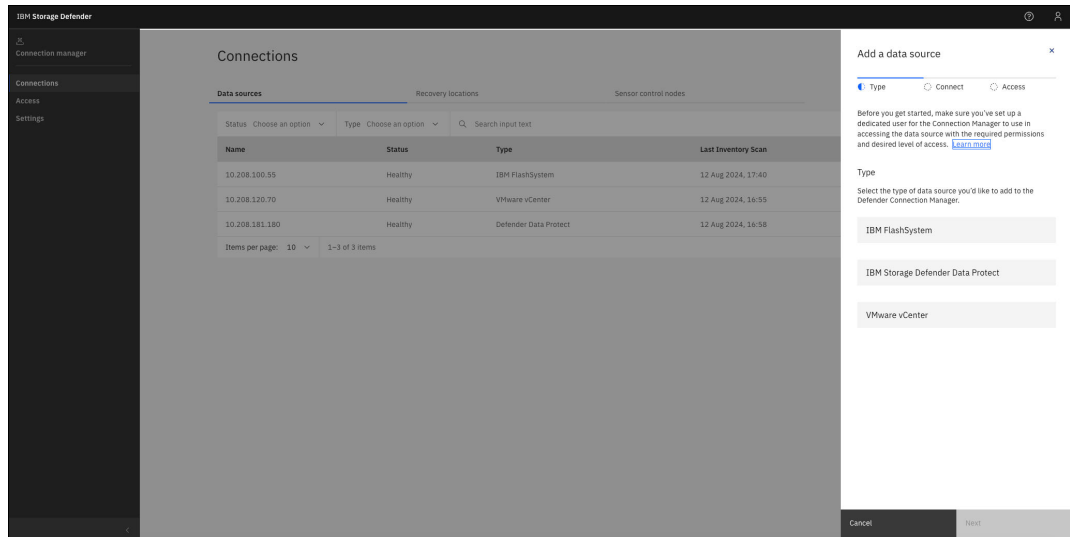


Figure 2-14 Connection Manager - Add a data source panel

Enter the Hostname or IP address of the data source and click Next. (Figure 2-15 on page 22)

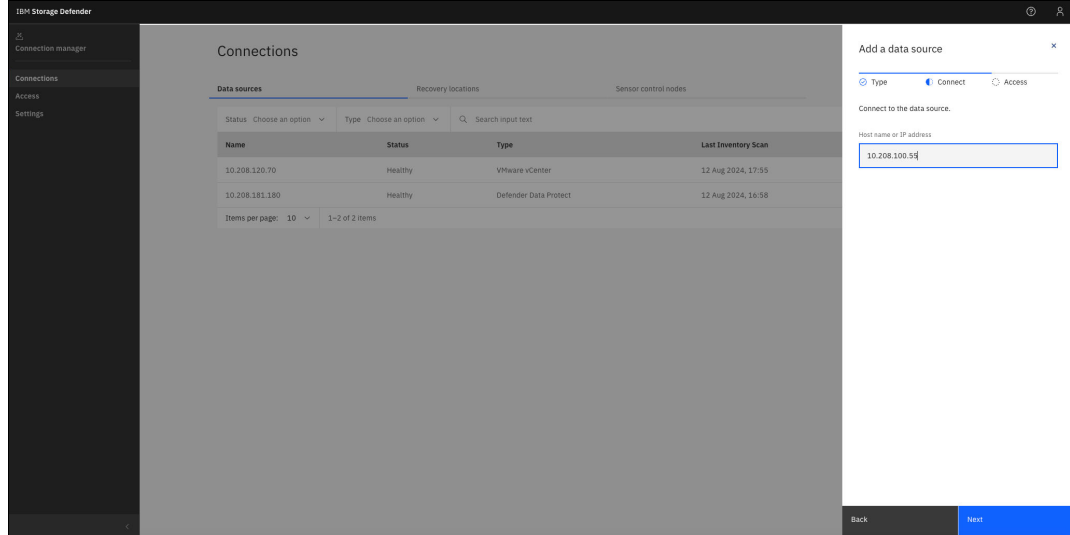


Figure 2-15 Connection Manager - Add a data source details panel

Review the certificate details and click Next. (Figure 2-16 on page 22)

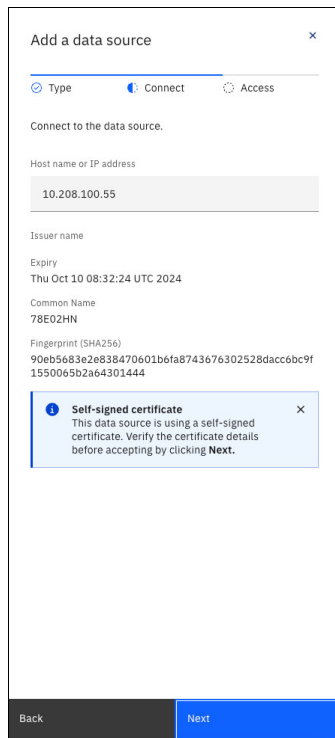


Figure 2-16 Connection Manager - Add a data source details panel

In the following window enter the credentials that will be used by Connection Manager to access this data source. (Figure 2-17 on page 23)

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 2-17 Connection Manager - Add a data source credentials panel

Click Add. Once complete the new data source will be added to Connection Manager as shown in Figure 2-18.

Name	Status	Type	Last Inventory Scan
10.208.100.55	Healthy	IBM FlashSystem	12 Aug 2024, 18:04
10.208.120.70	Healthy	VMware vCenter	12 Aug 2024, 17:55
10.208.181.180	Healthy	Defender Data Protect	12 Aug 2024, 17:58

Figure 2-18 Connection Manager - new data source

Adding Recovery Locations

To add recovery locations in the Connections dashboard, select the Recovery locations tab and click "Add recovery location". The wizard will open in the right side of the window. Enter the hostname or IP address of VMware vCenter that you would like to add and click Next. (Figure 2-19 on page 24)

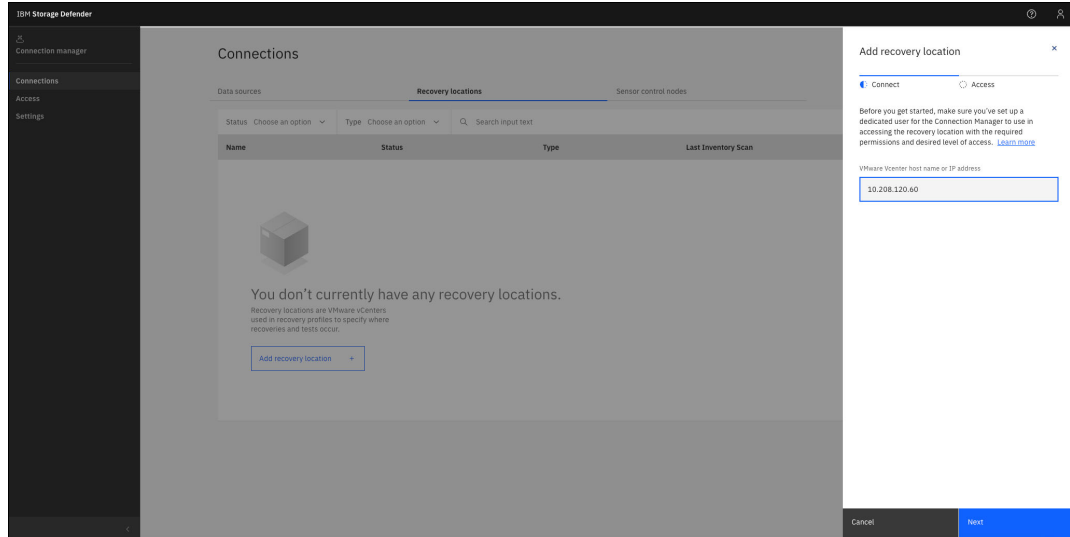


Figure 2-19 Recovery location - adding recovery location panel

Review the certificate details and click Next. (Figure 2-20 on page 24)

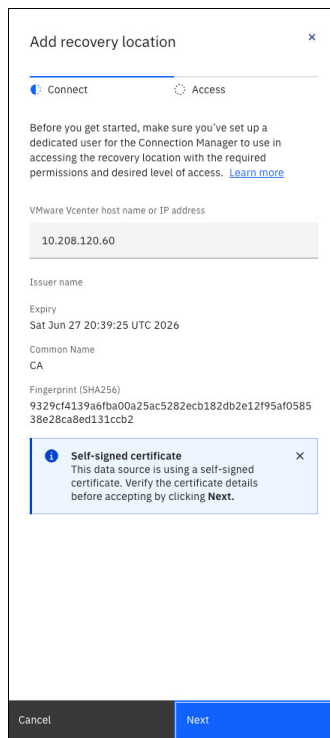


Figure 2-20 Recovery location - certificate details panel

Enter the dedicated credentials with the required permissions and desired level of access to the environment for this recovery location and click Add. (Figure 2-21 on page 25)

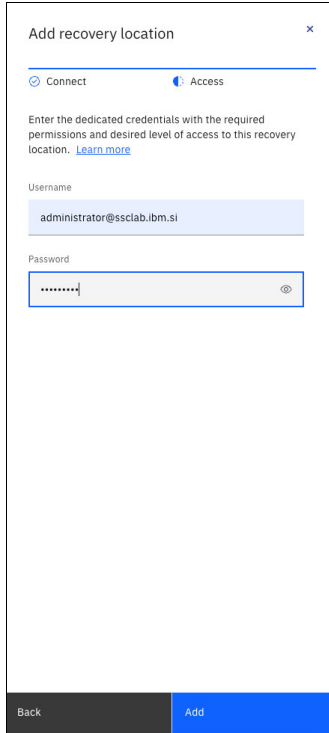


Figure 2-21 Recovery location - Add location panel

Figure 2-22 show the new recovery location is successfully added to the Connections list.




Figure 2-22 Add Recovery location

Adding Sensor control nodes

Connection Manager comes with a built in control node. If you would like to use your own control nodes, they can be added through the Connection Manager GUI and the provided Ansible playbooks can be used to manage the sensors.

To add a control node in the Connections dashboard, select the Sensor control nodes tab and click "Add control node"(Figure 2-23 on page 26). The wizard will open in the right side of the window. Enter the Ansible control node hostname and click Next.

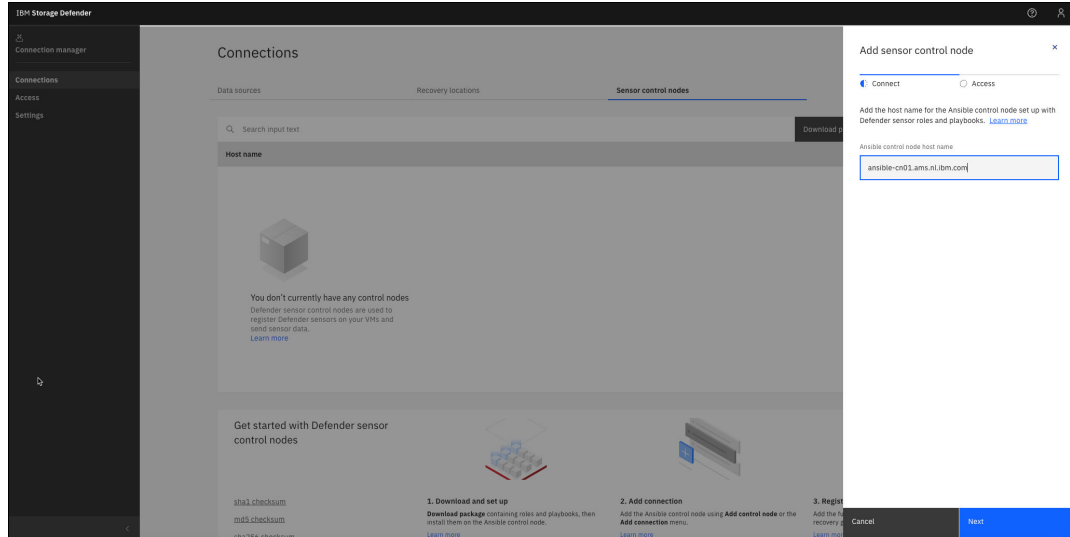


Figure 2-23 Add sensor control node panel

In the following step (Figure 2-24 on page 26) enter credentials created on the Ansible control node during the Defender sensor setup and click Add.

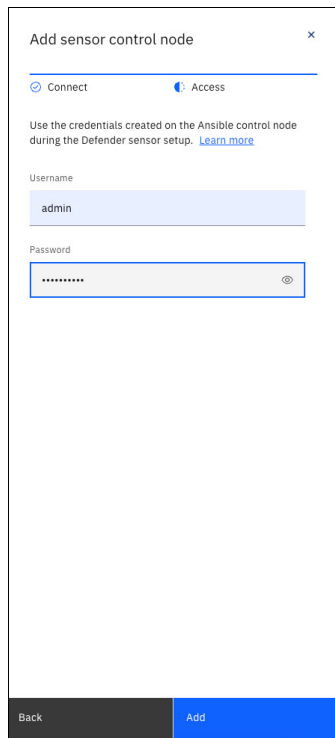


Figure 2-24 Add sensor control node panel

New sensor control node will be added to Connection Manager. (Figure 2-25 on page 27)

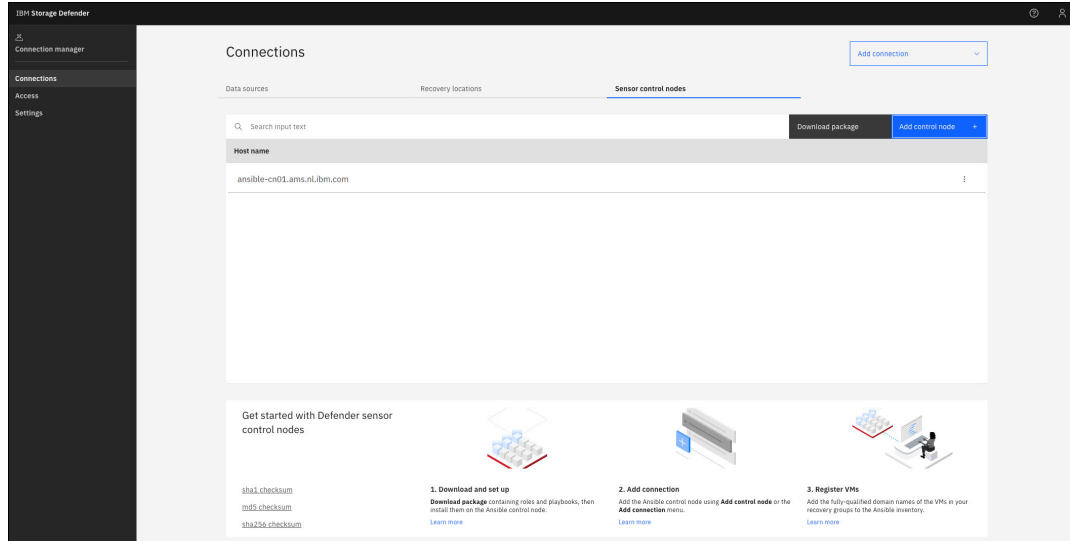


Figure 2-25 Sensor control node connections list

2.3.2 Creating profiles in DRS

Profiles used in IBM Storage Defender 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 needed to recover data. To create profiles in IBM Storage Defender DRS, select Profiles from the menu to bring up the Profiles dashboard in the GUI.

To create a Governance profile, select the Governance tab and click "Create profile" and the Create governance profile window will open (Figure 2-26). Under the Details tab, enter the name for a governance profile and description. Click Next.

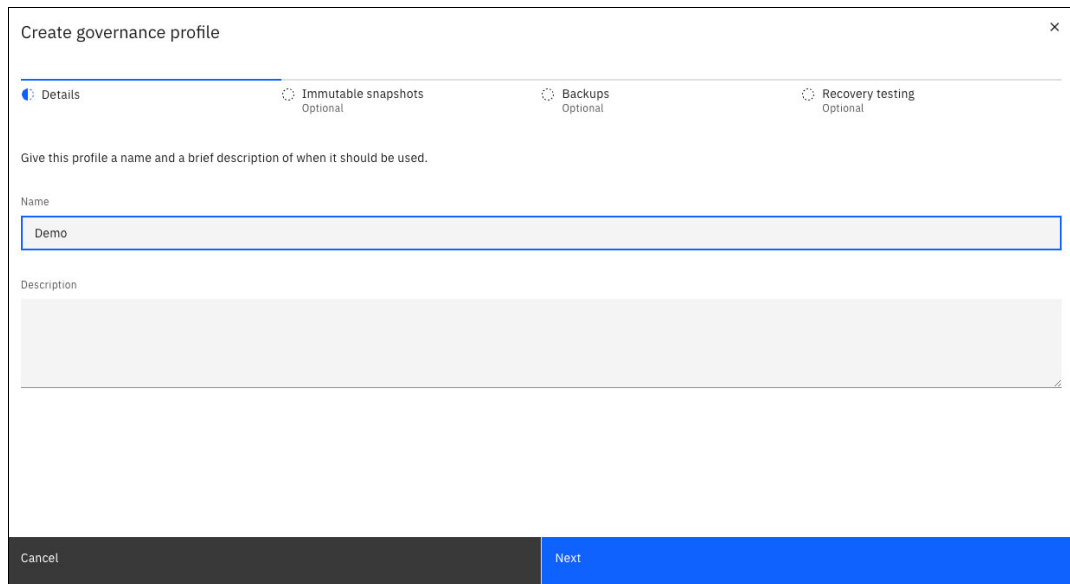


Figure 2-26 Creating a Governance profile panel

When creating a Governance profile, the Immutable snapshots tab allows you to select thresholds for immutable snapshot recovery points available from the IBM FlashSystem (Figure 2-27). You can select the check box to enable the point in time verification and the retention time verification for specified time interval. Click Next.

The screenshot shows a 'Create governance profile' dialog with four tabs: 'Details', 'Immutable snapshots Optional', 'Backups Optional', and 'Recovery testing Optional'. The 'Immutable snapshots' tab is active. Below the tabs, the text reads 'Select the thresholds for immutable snapshot recovery point objectives.' There are two checked checkboxes: 'At least one immutable snapshot within the last' and 'Immutable snapshots kept for at least'. Each checkbox is followed by a numeric input field containing '1' and a 'Days' dropdown menu. At the bottom, there are 'Back' and 'Next' buttons.

Figure 2-27 Governance profile - Immutable snapshots verification panel

Under the Backups tab, you can select thresholds for backup copy recovery points available from IBM Storage Defender Data Protect (Figure 2-28). Select the check box to enable the point in time verification and the retention time verification for specified time interval. Click Next.

The screenshot shows the same 'Create governance profile' dialog, but with the 'Backups' tab active. The text reads 'Select the thresholds for backup recovery point objectives.' There are two checked checkboxes: 'At least one backup within the last' and 'Backups kept for at least'. Each checkbox is followed by a numeric input field containing '1' and a 'Days' dropdown menu. At the bottom, there are 'Back' and 'Next' buttons.

Figure 2-28 Governance profile - Backups verification panel

In the Recovery testing tab select the thresholds for successful recovery testing. Select the check box to enable the test frequency verification and specify time interval (Figure 2-29 on page 29).

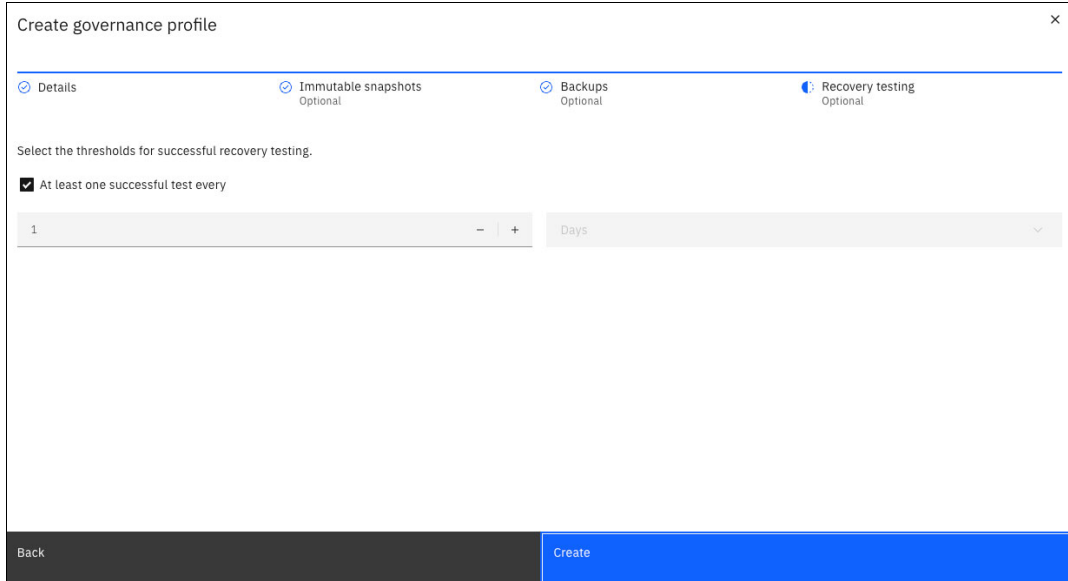


Figure 2-29 Governance profile - Recovery Testing panel

Click Create and new Governance profile will be created.

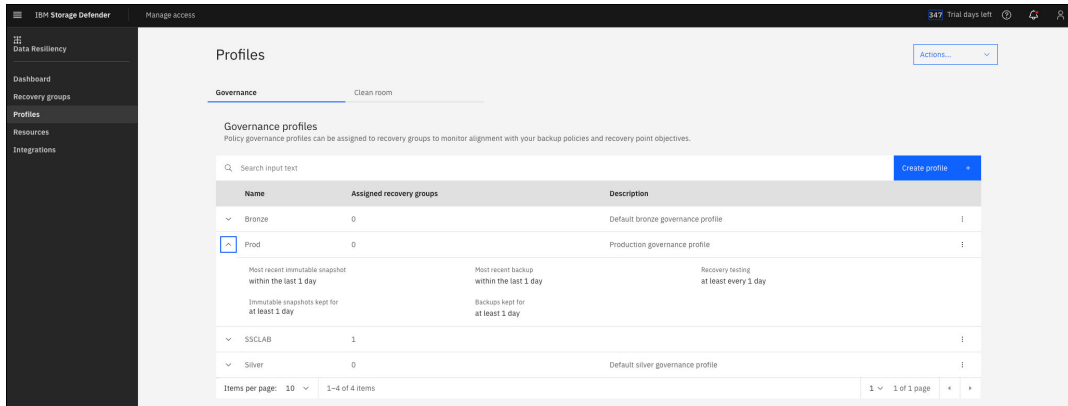


Figure 2-30 New Governance profile created in the profiles list

To create Clean Room profile, select Clean Room tab and click Create profile (Figure 2-31 on page 30).

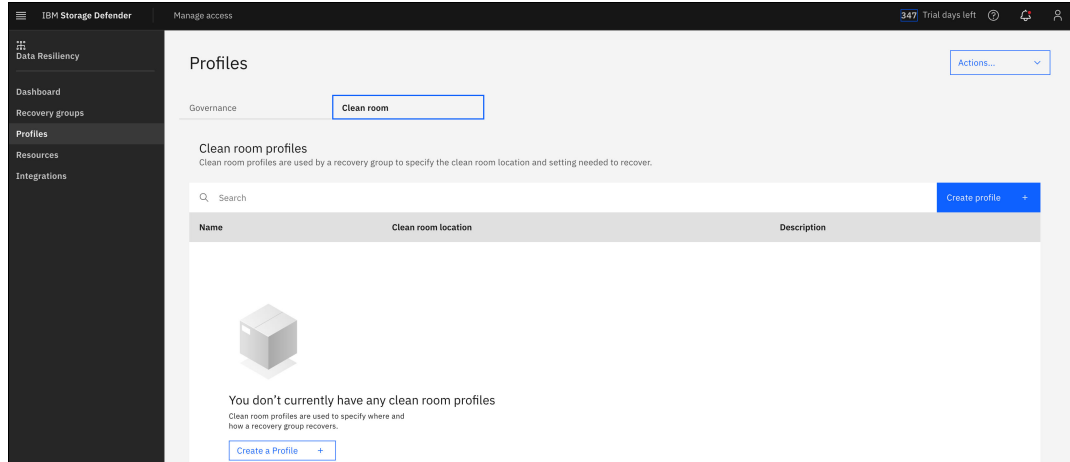


Figure 2-31 Creating a Clean Room profile

On the Create clean room profile window, under the Details tab, specify the name for a clean room profile and you can also provide description of the clean room profile (Figure 2-32). Click Next.

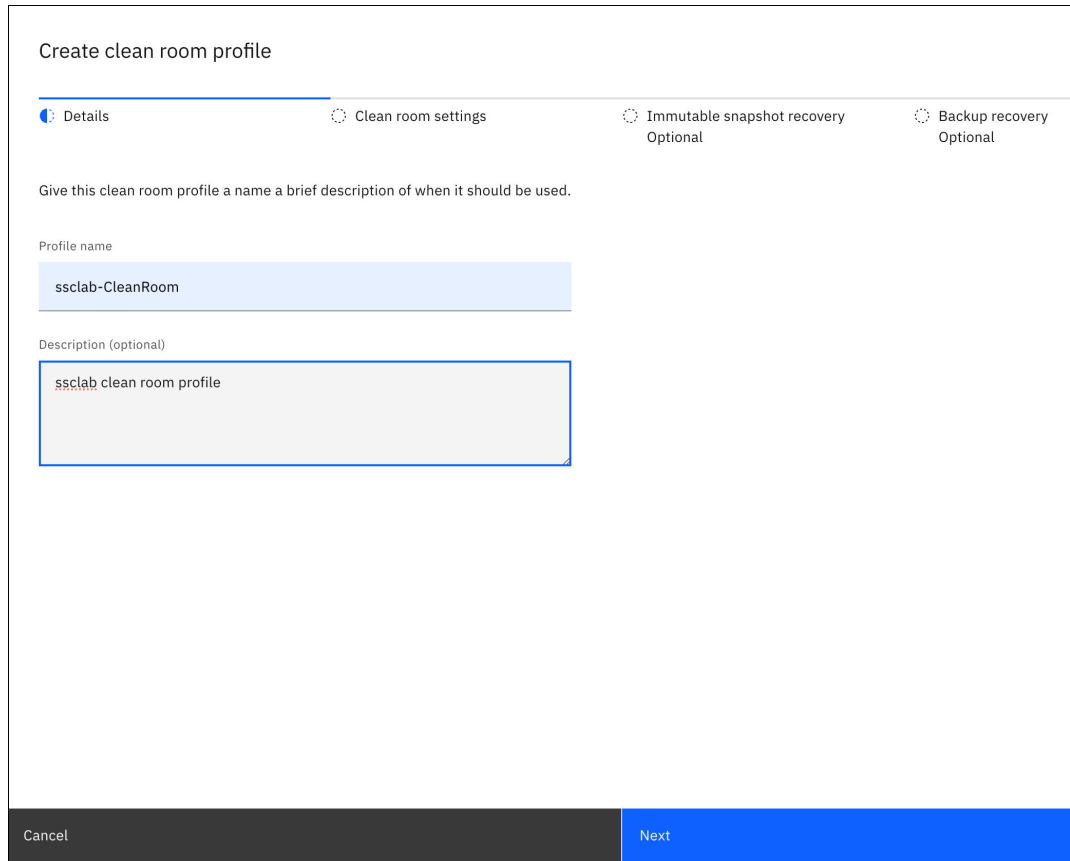


Figure 2-32 Clean Room profile details panel

Under the Clean room settings tab, you can enter the clean room location and recovery preferences (Figure 2-33 on page 31). 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.

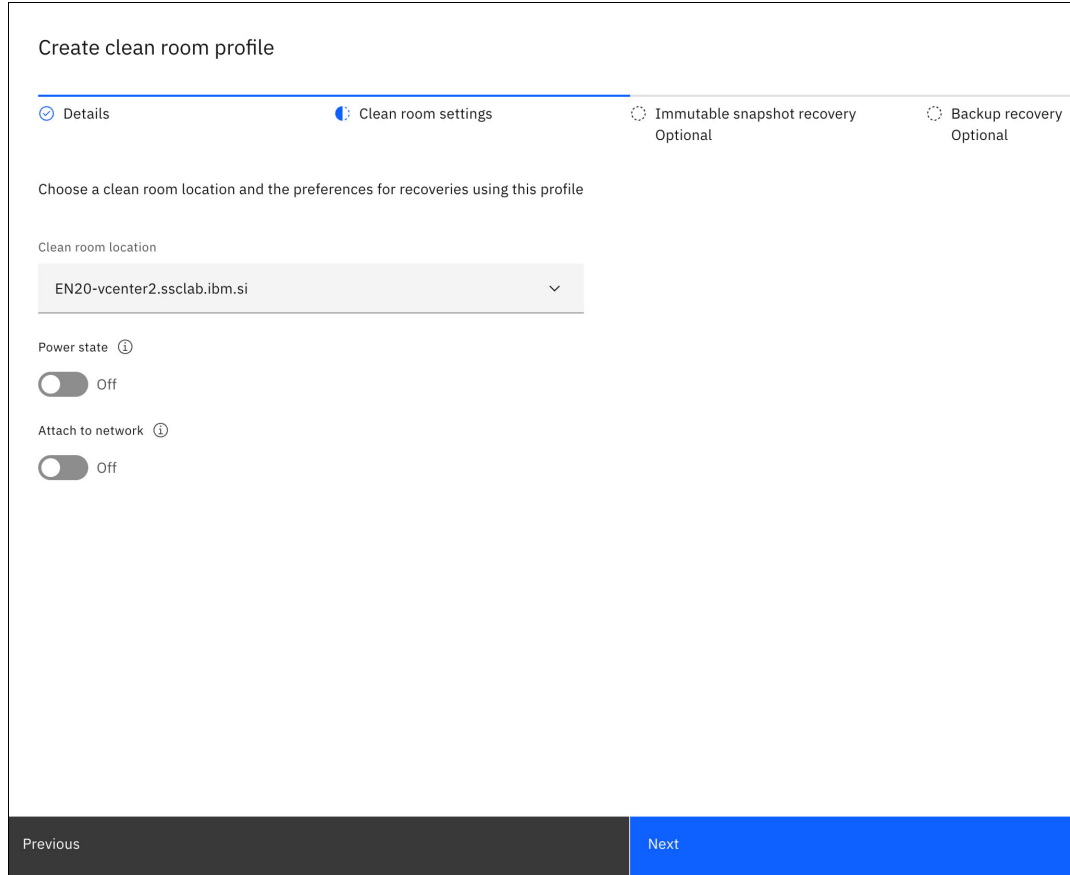


Figure 2-33 Clean Room settings panel

Under the Immutable snapshot recovery tab, you can enter recovery preferences when recovering from immutable snapshots with IBM Storage FlashSystem (Figure 2-34 on page 32).

The screenshot shows a web interface for creating a clean room profile. At the top, the title is 'Create clean room profile'. Below it are four tabs: 'Details', 'Clean room settings', 'Immutable snapshot recovery' (which is active and highlighted in blue), and 'Backup recovery'. Under the active tab, there are sub-labels 'Optional' and 'Optional'. A message states: 'Choose the settings to be used when recovering using an immutable snapshot with IBM FlashSystem.' Below this, there are three sections: 'Recovery settings' with a toggle switch set to 'On'; 'ESXi host' with a dropdown menu showing '10.200.120.61'; and 'vCenter folder' with a dropdown menu showing 'CleanRoom'. At the bottom, there is a 'Clean up on failure' toggle switch set to 'Off'. At the very bottom of the panel are two buttons: 'Previous' and 'Next'.

Figure 2-34 Clean Room - Immutable snapshots settings panel

Under the Backup recovery tab, specify the recovery preferences when recovering from IBM Storage Defender Data Protect (Figure 2-35 on page 33). If you plan to recover from a backup, select the vSphere resource pool from the drop-down list that can be used for recovery. The default resource pool on each vCenter is the pool that is called Resources. In addition, you can have other resource pools that you have created in your vCenter. All available resource pools can be selected for the recovery. In addition you can select the desired vCenter datastore from the drop-down list that you wish to be used for recovery with this policy. Click Create to create a clean room profile with the specified values.

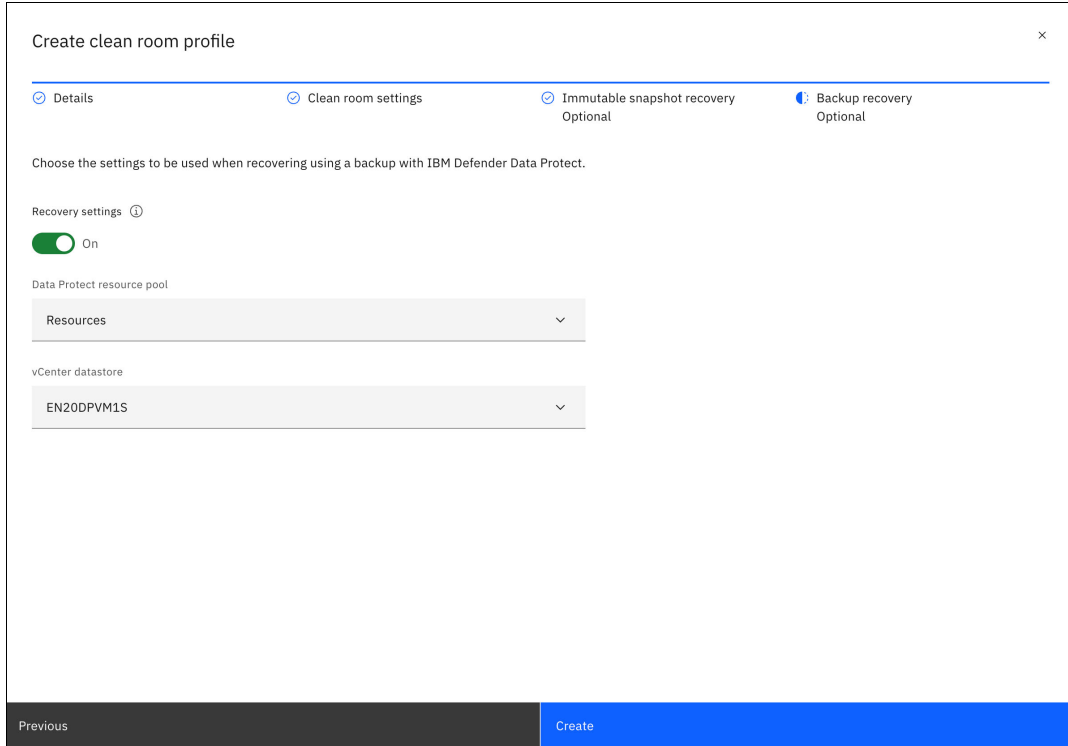


Figure 2-35 Clean Room - Backup recovery settings panel

The clean room profile is created under the Clean room tab (Figure 2-36).

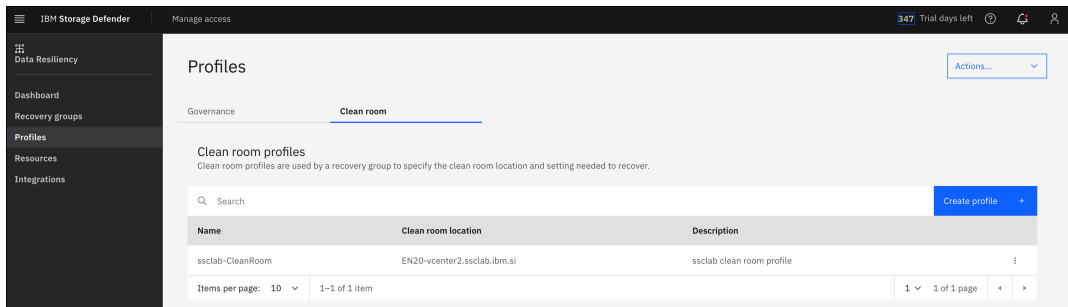


Figure 2-36 Clean Room profile list

By creating the governance and clean room profiles, recovery objectives of recovery groups and recovery target environments are configured.

2.4 Auto-forward IBM Storage FlashSystem Ransomware Threat Alerts to IBM Storage Defender

IBM Storage Defender Data Resiliency Service integrates with IBM Storage Insights 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 4th generation of FlashCore Modules (FCM4) and designed to continuously monitor statistics 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 is able to be executed.

These ransomware alerts 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 a customers subscribing to IBM Storage Insights Pro and IBM Storage Defender, this will enable enhanced protection from ransomware attacks with simply and easy integration.

2.4.1 IBM FlashCore Modules (FCM)

IBM has been delivering high-performance, highly-reliable customized flash modules for many years. With FCM, the control path and the data path within the module have been separated, ensuring that data can be accessed and transferred without any performance degradation caused by the control path. To enhance endurance and reliability, FCM modules have implemented endurance features and RAID within the modules themselves. Numerous additional technologies and benefits were implemented. For a detailed outline, please visit the IBM Redbooks site for ones of the publications on the IBM FlashSystem family such as ['IBM Storage FlashSystem 7300 Product Guide: Updated for IBM Storage Virtualize 8.7'](#)

How Flashcore modules detect and report ransomware threats.

In 2024, IBM introduced FCM4. This brought another industry leading breakthrough called Ransomware Threat Detection, which is a process that identifies and responds to security threats before they can damage data or systems. The FCM4 collects detailed statistics on every I/O operation (IOP) for each virtual disk (VDisk). This data is then intelligently summarized for efficient processing. The FCM4 transmits this summary to 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. Upon detection, an immediate alert is sent to IBM Storage Insights Pro, allowing for swift action. Additionally, the information can be shared with IBM Storage Defender if available, further strengthening your security posture.

With IBM Storage Virtualize software 870 and FCM's with FCM firmware 4.1, the ransomware threat detection is further improved by the following process:

- ▶ IBM FlashCore modules collect and analyze detailed ransomware statistics from every I/O with no performance impact.
- ▶ IBM Storage Virtualize runs an AI engine on every FlashSystem that is fed Machine Language (ML) models developed by IBM Research® trained on real-world ransomware.
- ▶ The AI engine learns what's normal for the system and detects threats using data from FCM.
- ▶ IBM Storage Insights Pro collects threat information from connected FlashSystems, alerts trigger SIEM/SOAR software to initiate a response.
- ▶ Statistics are fed back to IBM to improve ML models.

2.4.2 Integration between IBM Storage Defender Data Resiliency Service and IBM Storage Insights PRO

IBM Storage Insights PRO.

IBM Storage Insights Pro is a subscription-based Software as a Service (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 brings a lot of AIOps capabilities which will help the customers to plan for the future, and manage their infrastructure more efficiently.

For IBM FlashSystems running firmware version 8.6.3 and later, FlashCore modules (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 detected on all monitored IBM FlashSystems and generates alerts accordingly. These alerts are sent to the storage administrator via email and are also displayed in the IBM Storage Insights Pro user interface. Additionally, IBM Storage Insights Pro identifies affected volumes, marking them as having detected ransomware threats.

Enable Data Resiliency Service integration in Storage Insights Pro.

Storage administrators who subscribe to both IBM Storage Insights Pro and IBM Storage Defender Data Resiliency Service have the option to direct ransomware alerts generated in IBM Storage Insights Pro to IBM Storage Defender Data Resiliency Service. When an IBM FlashSystem is onboarded to both IBM Storage Insights Pro and IBM Storage Defender, Storage Defender sends an integration request to IBM Storage Insights Pro. The IBM Storage Insights Pro administrator receives this request via the user interface and decides whether to forward ransomware alerts to Defender. Once the administrator approves the request, the system is enabled to send ransomware alerts to Defender. When a ransomware threat is detected on any volume or volume group of the monitored IBM FlashSystem, the alert is forwarded to the Defender webhook, including details such as storage system information, volume specifics, and the ransomware timestamp. Upon receiving and acknowledging the alert, IBM Storage Insights Pro notifies the user that Defender has acknowledged the alert and is actively addressing it. From there the alert is also available to be sent via IBM Storage Defender DRS to any connected SIEM systems to also notify the Security Operations team, currently IBM QRadar and Splunk are supported.

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

- ▶ IBM FlashCore® Module version 4 technology is built in the IBM Storage FlashSystem that is used.
- ▶ 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 and ransomware threats to your IBM Storage Insights Pro tenant.
- ▶ The IBM Storage FlashSystem needs to be registered in IBM Storage Defender Data Resiliency Service. This registration is done in the user interface of Connection Manager.
- ▶ The IBM Storage Insights Pro communicates with Data Resiliency Service. The health status of your IBM Storage FlashSystem is sent to the Data Resiliency Service so that if 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 FlashCore Module 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. This could include alerts being sent to a connected SIEM.

Viewing ransomware threats in IBM Storage Defender DRS.

IBM Storage Insights Pro reports the ransomware threats at the volume or volume groups for an IBM Flash system and the alert is sent and shown in IBM Storage Defender Resiliency Service (Figure 2-37 on page 36). IBM Storage Defender Data Resiliency Service uses recovery groups to group the related virtual machines. When a ransomware

alert is received, IBM Storage Defender Data Resiliency Service correlates the volume in the alert to the datastore, and opens a case for the recovery group where the VM using the datastore resides. The newly opened cases can be viewed on the IBM Storage Defender DRS, and a recovery plan can be activated to recover to the last copy, or last best copy available.

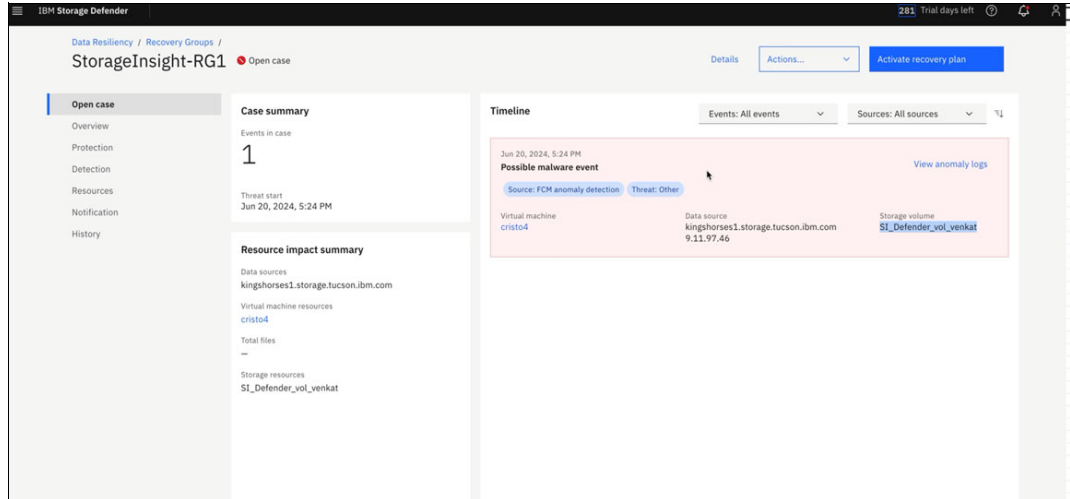


Figure 2-37 DRS showing malware event notification message



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 provides, describes, discusses, or contains the following:

- ▶ 3.1, “What do sensors do?” on page 38
- ▶ 3.2, “Installing Sensors” on page 41

3.1 What do sensors do?

IBM Storage Defender sensors are small, lightweight pieces of software that get installed into each VM 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 actual ransomware variants tend to use. Sensors also use a pre-built Machine Learning model, this model is trained on known ransomware variant patterns which the sensors use to help identify similar patterns on the host where they are installed. Finally, if malicious activity is suspected based on those factors, a third check using file introspection is performed to determine if suspected victim files are encrypted.

If each of these criteria are met, an event is raised to Defender's Data Resiliency Service indicating a possible malware event and a "case" is opened, as shown in Figure 3-1

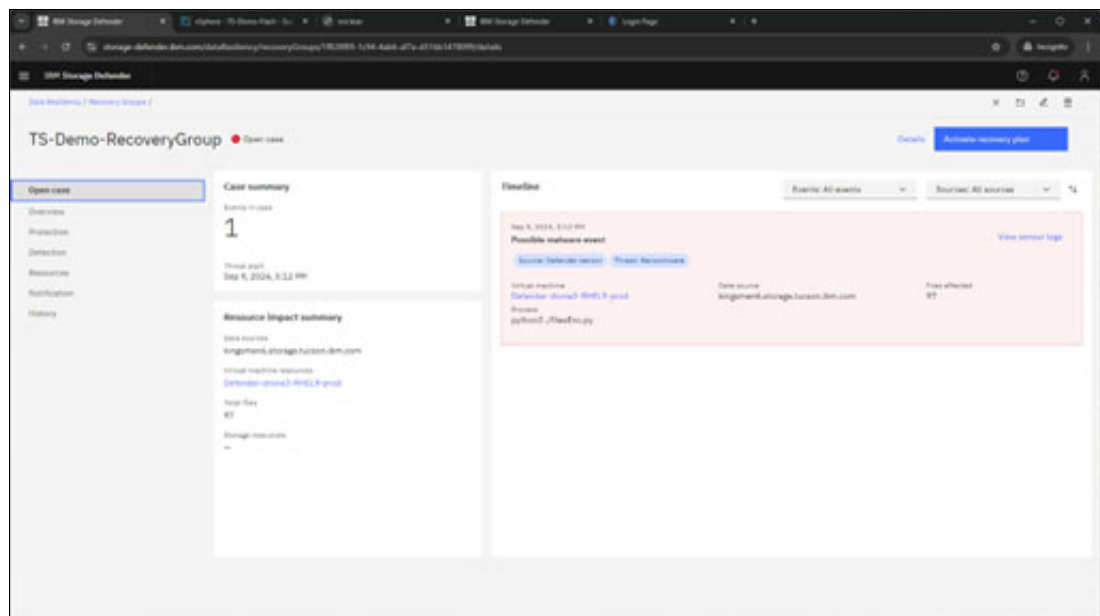


Figure 3-1 IBM Storage Defender Timeline sensor event

In the details for the specific event (Figure 3-2 on page 39) both informational and actionable information is provided at a glance, this includes (in our example):

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



Figure 3-2 Event details panel

You can also drill down further and review detailed information for the event by choosing View sensor logs. A sample sensor log is shown in Figure 3-3.



Figure 3-3 Defender sensor log details

This shows additional actionable information such as the hostname (FQDN) of the VM and the Process ID (pid) which can be used to identify and kill the suspicious process. It also shows the user ID (uid) which would allow an admin with appropriate rights, to lock out that user. Detailed logs can also be useful information for incident responders as the absolute pathnames of every impacted file are also shown.

A summary of the impact is provided at the end of the log, including the total number of files (Figure 3-4 on page 40 shows suspected malicious accesses event details). Regardless of the number of files impacted, recovery would happen at a volume level and all files can then be recovered to an earlier, unimpacted state.

Note: Currently, the encryption detection identifies encryption only on files > 4K, so it is likely that if specific files are identified as impacted, it's probable smaller files in these locations are too.

```
Log details
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata162.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata145.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata185.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata128.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/vm_plants_onGrid.pdf.aes: [72.6866], uid = 1001, size
= 207719, pids = [977784], encStatus = Detected
File: /home/defensor/RWdemo/TestFiles/Solaris_10_Summary.pdf.aes: [72.6866], uid = 1001,
size = 1076271, pids = [977784], encStatus = Detected
File: /home/defensor/RWdemo/TestFiles/xdata116.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/16.pdf.aes: [72.6866], uid = 1001, size = 1122695,
pids = [977784], encStatus = Detected
File: /home/defensor/RWdemo/TestFiles/xdata118.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata111.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata152.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata171.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata157.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata114.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata12.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata131.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
File: /home/defensor/RWdemo/TestFiles/xdata134.txt.aes: [72.6866], uid = 1001, size = 391,
pids = [977784], encStatus = NotChecked
----- END Files Involved in Malicious accesses -----
Summary: totalMaliciousAccesses: 97
Summary: totalMaliciousEncryptions: 16
Summary: isAlert: True
End analysis: 2024-09-09_14:12:29
***** END Log Entry *****
```

Figure 3-4 Sensor log details cont.

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



Figure 3-5 Sensor Heartbeat warning message

For any of these events, a case is opened so actions can be reviewed and easily communicated between team members and teams. For example, once an event is analyzed by an admin or responder, information from the event is reviewed and the cause can be addressed or confirmed. Once appropriate remediation is taken to resolve the issue, the case can be closed. When a case is closed the corresponding event messages as shown in Figure 3-5 are cleared.

However, past events can still be viewed from the Detection panel as shown in Figure 3-6 on page 41.

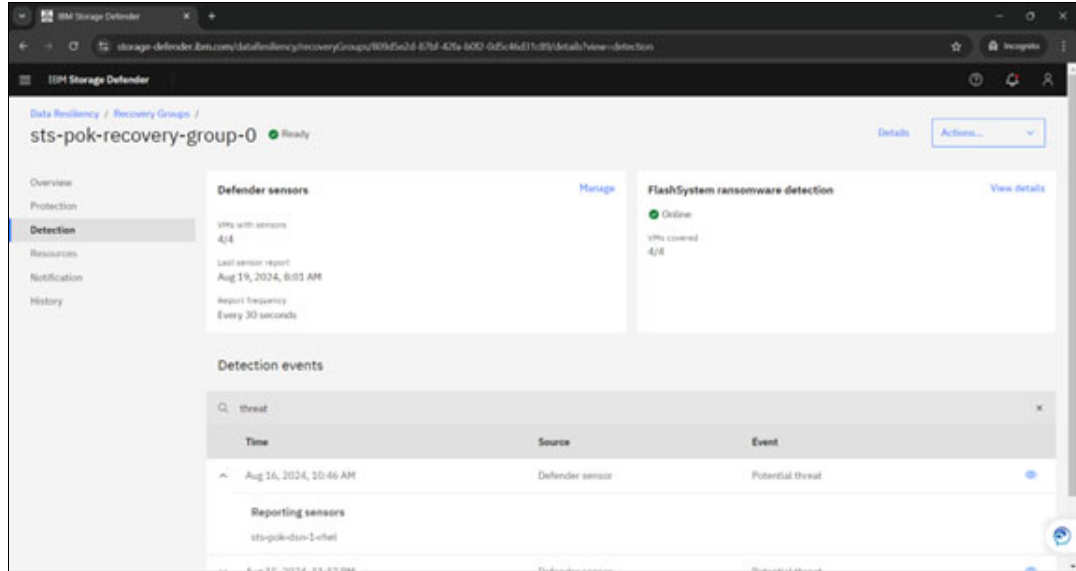


Figure 3-6 Recovery group detection panel

After closing the case, the DRS dashboard will continue to allow access to historical events and includes the ability to search on previous threat events, as well as drill down to review details of those events.

3.2 Installing Sensors

This section will explain how to install an IBM Storage Defender sensor on one or more systems using either the UI or CLI. These sensors monitor the systems on which they are installed, enabling real-time detection of cyber threats such as ransomware attacks. Controlling the deployment of sensors can be done 1 of 2 ways; 1) automatically from the Connection Manager with the built-in facilities, or 2) via 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 the desired 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 .tar 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:

Hostname:	FQDN of the Connection Manager.
Username:	Define a username that is to be used to register IBM Storage Defender sensors that are installed on virtual machines for the sensor control node.
Password:	Define a password that is related to the username.
Vault password:	The username and password that is defined before is 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 desired 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 needs to be added through the steps listed above. If you attempt to add more than one control node using the same username, the following error will occur in the UI:

Error getting source native ID: User name already in use. Please select a different user name.

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

Once this is completed you will see the registered sensor control node in the UI.

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 desired Connection Manager instance.
2. On the home page of the Connection Manager, click Connections > Sensor control nodes.

3. In the table that lists all the sensor control nodes, scroll to the sensor control node of interest.
4. In the row of the sensor control node, click the overflow menu (shown below), 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 UI

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

Before you begin:

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

The process to install an IBM Storage Defender sensor on one or more systems requires executing the following these steps:

1. Log in to IBM Storage Defender.
2. Click the hamburger menu (three horizontal lines) in the upper left corner of the page.
3. Navigate to Data Resiliency > Recovery Groups.
4. From the list of recovery groups, select the row for the recovery group you wish to install the sensor on.
5. In the Overview panel, locate the Defender sensors tile and click Get started.

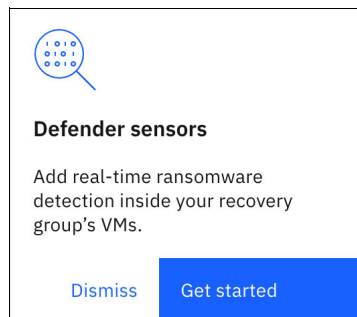


Figure 3-8 Defender Sensors Tile

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

6. In the Manage sensors pop-up window, select one or more virtual machines (VMs) by checking the corresponding boxes.
7. Click Add sensor + in the title bar.
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.
10. The selected VMs will 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.

Once installation is complete, the sensor will automatically begin monitoring file access activity on the system. If it detects any unusual access patterns, such as those associated with ransomware attacks, the sensor will generate an alert. This alert is sent to the on-premises Connection Manager, which securely forwards it to the IBM Storage Defender Data Resiliency Service (DRS). The sensor also periodically sends heartbeat messages through the Connection Manager to confirm it is operating normally.

3.2.5 Installing an IBM Storage Defender sensor by using CLI

You can install an IBM Storage Defender sensor on one or multiple systems by using the UI 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, follow these steps:

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

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

3. Create an inventory file:
 - ▶ Create an inventory file containing the Fully Qualified Domain Names (FQDN) of all systems you want to install the sensor on.
4. Edit the Ansible hosts file:
 - ▶ 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 you choose to do this, use the ``-i /your-directory/your-file`` argument in the next step.

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 using 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:

- ▶ Execute the following Ansible playbook command to begin installation

Example 3-2 Ansible playbook install command

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

7. Enter the Ansible vault password:

- ▶ When prompted, enter the vault password 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` and `--ask-become-pass` to provide the SSH and sudo passwords during playbook execution.

After installation, the sensor will automatically monitor file access activities on the system. If any unusual access patterns resembling ransomware attacks are detected, the sensor will send alert messages to the on-premises Connection Manager, which forwards these alerts securely to the IBM Storage Defender Data Resiliency Service (DRS). The sensor also sends periodic heartbeat messages to the DRS via the Connection Manager, indicating normal operation.

3.2.6 Uninstalling an IBM Storage Defender sensor by using UI

To uninstall an IBM Storage Defender sensor from one or more systems using the UI, follow these steps:

1. Log in to IBM Storage Defender and access the IBM Storage Defender dashboard.

2. Navigate to the recovery group:
 - ▶ Click the hamburger menu (three horizontal lines) in the upper left corner of the page.
 - ▶ Select Data Resiliency > Recovery Groups.
 - ▶ From the list of recovery groups, click the row corresponding to the recovery group where you want to uninstall sensors.
3. Manage sensors:
 - ▶ In the recovery group's Overview dashboard, locate the Defender Sensors tile and click Manage.
 - ▶ 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 using "localhost" as the FQDN
- ▶ VMs with duplicate FQDNs

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

4. Uninstall the sensor:
 - ▶ Click the Remove Sensor button in the title bar.
 - ▶ Enter either the username and password, or the SSH key for the virtual machines.
 - ▶ Click Remove Sensor to submit the uninstallation request.

Note: All selected VMs must share the same login credentials.

5. Monitor the process:
 - ▶ The status of the selected VMs will change to Uninstalling.
 - ▶ You can monitor the Notification menu for updates on 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 associated with a Connection Manager that has been destroyed or improperly backed up and restored during a Connection Manager OVA upgrade, the uninstallation will fail. For troubleshooting, refer to the guide on "Resolving an IBM Storage Defender sensor uninstallation failure."

Once uninstallation is complete, the IBM Storage Defender sensor service will be removed from the selected VMs.

3.2.7 Uninstalling an IBM Storage Defender sensor by using CLI

You can uninstall an IBM Storage Defender sensor from one or more systems using either the UI or CLI. To proceed with the CLI method, follow these 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:
 - ▶ Create an inventory file listing the Fully Qualified Domain Names (FQDN) or IP addresses of all systems from which you want to uninstall the sensor.
3. Edit the Ansible hosts file:
 - ▶ 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 you choose to do this, use the `-i /your-directory/your-file`` argument in the next step.

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 following Ansible playbook command:

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.

Once the playbook is executed the sensor will be 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 version 5.9 or later
 - libgomp
 - python3
- ▶ SUSE Linux Enterprise Server 15 SP5
 - Required packages:
 - bash
 - kernel version 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 version 5.9 or later
 - -python3

Supported File Systems

- ▶ XFS
- ▶ EXT4



Daily Administration, Alerting, Testing and Validation

This chapter provide an overview on how to bring together all of the elements for daily administration and how to test your recovery points.

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

- ▶ 4.1, “DRS Dashboard” on page 50
- ▶ 4.2, “User Management Profiles” on page 55
- ▶ 4.3, “Integrations for Alerting” on page 55
- ▶ 4.4, “Recovery Testing and Validation” on page 56
- ▶ 4.5, “Activating the Recovery Plan” on page 60

4.1 DRS Dashboard

For daily administration and status 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, all of the Connection Managers, open cases, Recovery Group status and additional information. Figure 4-1 includes a view of the DRS Dashboard.

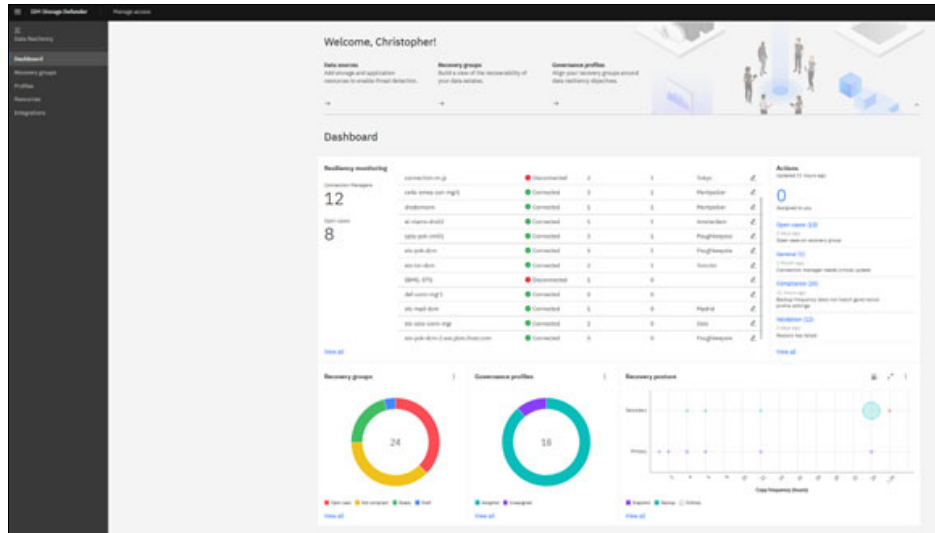


Figure 4-1 DRS Dashboard overview page

This dashboard contains several elements that enable users to leap out to additional information and context. These capabilities include;

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

4.1.1 Resiliency Monitoring in the Dashboard

The DRS Dashboard provides an at-a-glance of the Resiliency Monitoring which highlights the status of the connection managers and any open cases.

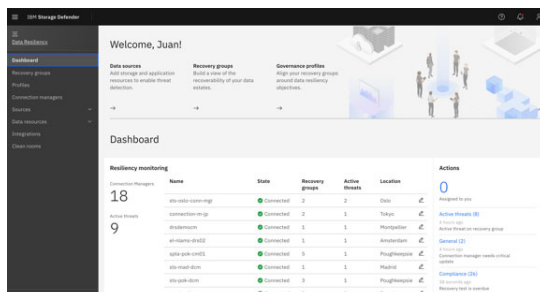


Figure 4-2 Resiliency Monitoring Dashboard panel

This provides the ability to denote locations, see the state/status and drill in deeper on the managed connection managers. The ‘view all’ link, enables the user to see the Connection Managers tab of the Resources page. This highlights the Connection Managers, their state, their type, Host name, version and if an update is required as depicted in Figure 4-3.

Name	State	Type	Host name	Version
connection-mgr	Disconnected	VH	connection-mgr.us.oracle.com	2.0.8.273479145
DRS@-070	Disconnected	VH	localhost	2.0.8.273479145
oracle-ams.com-mgr1	Connected	VH	defender-connection-manager.us.oracle.com	2.0.8.273479145
def-conn-mgr1	Connected	VH	def-conn-mgr1.us.oracle.com	2.0.8.273479145
defconnmgr	Connected	VH	defconnmgr.us.oracle.com	2.0.8.273479145
defconnmgr2	Connected	VH	defconnmgr2.us.oracle.com	2.0.8.273479145
lga-pdb-ams	Connected	VH	lga-pdb-ams.us.oracle.com	2.0.8.273479145
lga-ams-us	Connected	VH	lga-ams-us.us.oracle.com	2.0.8.273479145
lga-pdb-ams-2.us.oracle.com	Connected	VH	lga-pdb-ams-2.us.oracle.com	2.0.8.273479145
lga-pdb-ams	Connected	VH	lga-pdb-ams.us.oracle.com	2.0.8.273479145
lga-pdb-ams-2.us.oracle.com	Connected	VH	lga-pdb-ams-2.us.oracle.com	2.0.8.273479145
lga-pdb-ams	Connected	VH	lga-pdb-ams.us.oracle.com	2.0.8.273479145

Figure 4-3 DRS Monitoring Resources overview panel

Figure 4-4 shows the Actions which can be taken. This includes open cases, assigned actions, required updates and other issues

Actions
Updated 11 hours ago

0
Assigned to you

Open cases (10)
2 days ago
Open case on recovery group

General (1)
1 Month ago
Connection manager needs critical update

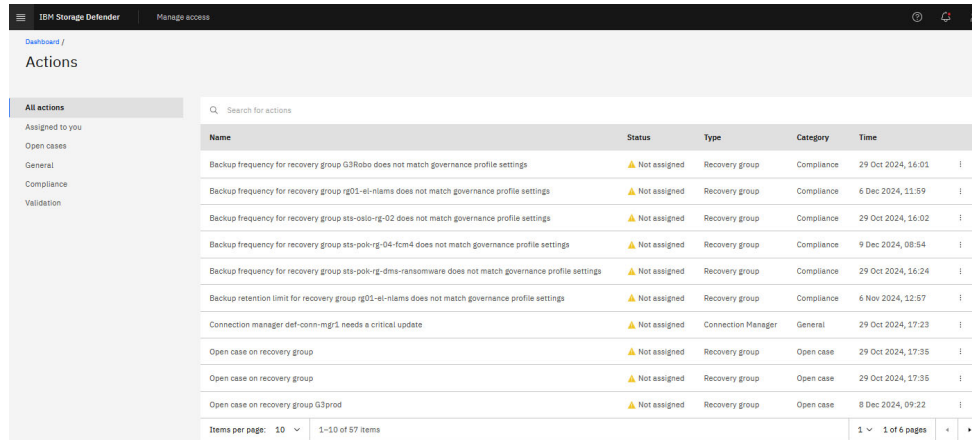
Compliance (26)
11 hours ago
Backup frequency does not match governance profile settings

Validation (12)
2 days ago
Restore has failed

[View all](#)

Figure 4-4 Actions summary panel

Figure 4-5 on page 52 is the view available when clicking on the actions panel. This will lead to a deeper view of the actions in order to review recommendations and resolve issues, see pending actions or view the history.



Name	Status	Type	Category	Time
Backup frequency for recovery group G3Robo does not match governance profile settings	▲ Not assigned	Recovery group	Compliance	29 Oct 2024, 16:01
Backup frequency for recovery group rg01-el-nlams does not match governance profile settings	▲ Not assigned	Recovery group	Compliance	6 Dec 2024, 11:59
Backup frequency for recovery group sts-oso-rig-02 does not match governance profile settings	▲ Not assigned	Recovery group	Compliance	29 Oct 2024, 16:02
Backup frequency for recovery group sts-pok-rig-04-fcm4 does not match governance profile settings	▲ Not assigned	Recovery group	Compliance	9 Dec 2024, 08:54
Backup frequency for recovery group sts-pok-rig-dms-ransomware does not match governance profile settings	▲ Not assigned	Recovery group	Compliance	29 Oct 2024, 16:24
Backup retention limit for recovery group rg01-el-nlams does not match governance profile settings	▲ Not assigned	Recovery group	Compliance	6 Nov 2024, 12:57
Connection manager def-conn-mgr1 needs a critical update	▲ Not assigned	Connection Manager	General	29 Oct 2024, 17:23
Open case on recovery group	▲ Not assigned	Recovery group	Open case	29 Oct 2024, 17:35
Open case on recovery group	▲ Not assigned	Recovery group	Open case	29 Oct 2024, 17:35
Open case on recovery group G3prod	▲ Not assigned	Recovery group	Open case	8 Dec 2024, 09:22

Figure 4-5 Actions panel history details

4.1.2 Recovery Group status

The Recovery Group pie chart (Figure 4-6) shows a summarized view of this tenant's recovery groups, indicating the percentage of recovery groups that are ready, in draft, not compliant, or have an open case / threat recorded on them that needs to be addressed.




Figure 4-6 Recovery Group pie chart on DRS Dashboard

If users select recovery groups on the left of the dashboard, they will be brought to the recovery group list (Figure 4-7), where users can drill down on any recovery group previously created or create a new one.

Name	Status	Governance profile	Connection Manager	Last updated
CE45-EMEA-classroom	Ready	CE45-EMEA-Classroom-Profile	ce45-emea-con-mig1	19 Nov 2024, 06:17
CE45-EMEA-demo-RG	Ready	CE45-EMEA-demo-Profile	ce45-emea-con-mig1	14 Nov 2024, 05:35
CE45-EMEA_RecoveryGroup00	Open case	CE45-EMEA-sensor-only	ce45-emea-con-mig1	18 Nov 2024, 05:12
CE45APAC	Not compliant	Bronze		
G3prof	Open case	sto-tor-g3	sto-tor-dcm	
G3Rdb	Not compliant	sto-tor-g3	sto-tor-dcm	29 Oct 2024, 09:39
IBMExp-Cluster	Open case	xi-name-profile	xi-name-dm02	16 Aug 2024, 09:00
IBMGL-GMP	Not compliant	IBMGL-every Day for 7 days	IBMGL-ST0	10 Nov 2024, 08:52
redbook	Not compliant	Bronze	sto-pak-dcm	
rg01-xi-nams	Open case	xi-name-profile		7 Nov 2024, 05:09
RG01-MAD	Not compliant	GovProfile-MAD	sto-mad-dcm	19 Nov 2024, 11:49
SPTADemo-dj01-AD	Ready	SPTA-Demo-Profile-Default	sp1a-pak-cm01	27 Nov 2024, 08:52
SPTADemo-dj01-MVC	Open case	SPTA-Demo-Profile-MVC	sp1a-pak-cm01	20 Nov 2024, 10:28
SPTADemo-dj01-RG02	Ready	SPTA-Demo-Profile-MVC	sp1a-pak-cm01	19 Nov 2024, 09:43
sto-dach-recovery-group-1	Open case	sto-dach-governance	sto-dach-dcm	17 Sept 2024, 06:42
sto-gp-test03	Not compliant	Silver	connection-mg-gp	26 Oct 2024, 02:52

Figure 4-7 Recovery Group list

4.1.3 Governance Profile Status

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




Figure 4-8 Governance profile pie chart on DRS Dashboard

Users can create and modify their existing governance and clean room profiles within the Profiles tab.

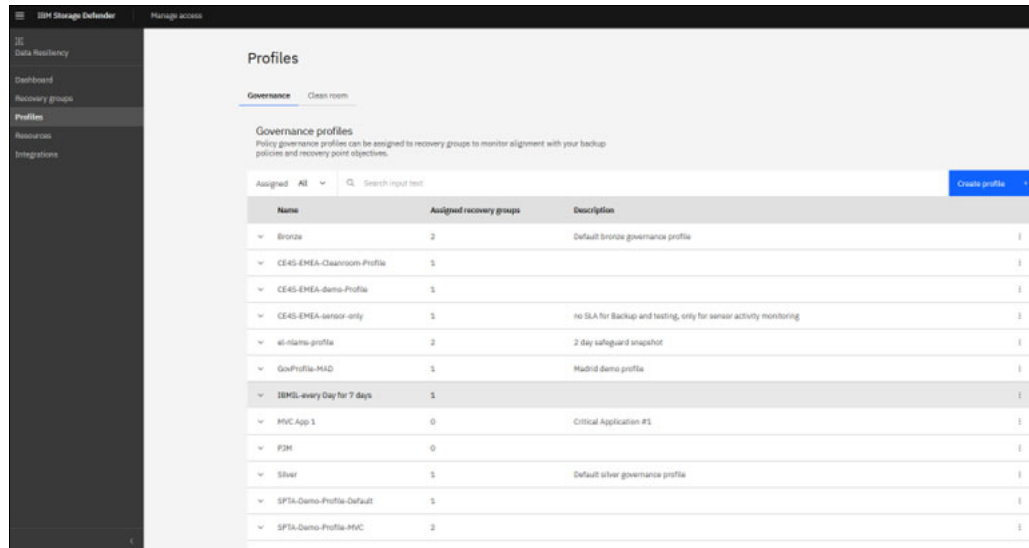


Figure 4-9 Profiles tab on DRS

4.1.4 Recovery Posture

The recovery posture graphic (Figure 4-10) helps users quickly get an understanding of their recovery posture. On the Y axis we can see Secondary and Primary, this refers to Secondary storage for example backups in Defender Data Protect, and Primary storage for example IBM FlashSystem. On the X axis we observe copy frequency, which tells how often we create copies. By combining these two, we can quickly view what are frequency policies are for our environment for both Primary and Secondary copies.

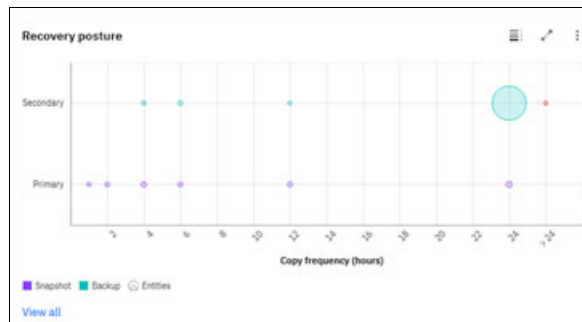


Figure 4-10 Recovery Posture graphic in DRS Dashboard

Users can navigate to the Resources tab (Figure 4-11) by clicking on “Resources” in the left side column of the GUI. Within this view, users can gather more information on their available resources, available copies, connections, and connection managers.

Name	Latest copy date	Type	Frequency	Data source	Policy	Copies
217 VMware vCenter Server 7	16 Nov 2024, 20:48	VM	1 day	vs-for-nsib-ibmchfab.local	G3 Production	32
217 VMware vCenter Server 7	16 Nov 2024, 20:48	VM	1 day	vs-for-dpcluster-ibmchfab.local	G3 Production	32
ChroVM1	17 Nov 2024, 14:48	VM	1 day	vs-for-nsib-ibmchfab.local	ChroPDemo	7
ChroVM1	16 Nov 2024, 14:48	VM	1 day	vs-for-dpcluster-ibmchfab.local	ChroPDemo	14
ChroVM2	17 Nov 2024, 14:48	VM	1 day	vs-for-nsib-ibmchfab.local	ChroPDemo	7
ChroVM2	16 Nov 2024, 14:48	VM	1 day	vs-for-dpcluster-ibmchfab.local	ChroPDemo	14
DDPTest1	28 Nov 2024, 16:49	Volume	4 hours	172.20.41.22	DDPTest1	750
DEFENDER DRS SENSORS	28 Nov 2024, 14:05	Volume	4 hours	97200_appdemo.tam	DEFENDER DRS4h	26
DHL-Datastore	27 Nov 2024, 07:04	Volume	2 hours	192.168.130.130	Every 2hours for 7 days	85
DPtargetW2013	4 Nov 2024, 03:25	VM	—	192.168.81.30		4

Figure 4-11 Resources tab on DRS

4.2 User Management Profiles

Administrators may navigate to the All Resources tab in DRS (Figure 4-12) where they can view the list of authorized users for the solution. Additionally, admins may manage user access in order to add new users and assign authority permissions or modify/remove existing users.

Role	Description
Administrator	As an administrator, you can perform all actions in the service.
Recovery group operator	As a recovery group operator, you can create, activate, and perform all actions associated with recovery groups.
Resource operator	As a resource operator, you can download the Connector Manager package and API keys. You can also create and manage recovery profiles.
Service operator	As a service operator, you can manage Defender services for all recovery groups.
Viewer	As a viewer, you can view the service, but you can't modify it.

Figure 4-12 User Management in DRS

4.3 Integrations for Alerting

Users can integrate their Defender Data Resiliency Service with two SIEM solutions, QRadar and Splunk in order to improve their security posture while also bridging the storage and security silos that sometimes exist in enterprises storage landscapes.

Step by step instructions and more information on this can be found here:

QRadar -

<https://www.ibm.com/docs/en/storage-defender/base?topic=integrations-integrating-d-ata-resiliency-qradar-siem>

Splunk -

<https://www.ibm.com/docs/en/storage-defender/base?topic=integrations-integrating-d-ata-resiliency-splunk-siem>

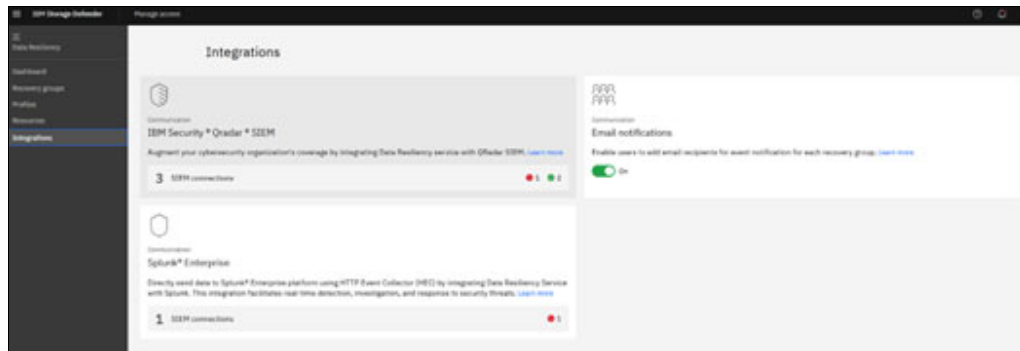


Figure 4-13 Integrations tab in DRS

4.4 Recovery Testing and Validation

The IBM Storage Defender Data Resiliency Service (DRS) provides the tools and ability to test and validate recovery points for a recovery group. Testing of recovery points for a recovery group can be done only when the status for the group is 'Ready', which means the recovery group is complete, with a governance plan assigned and clean room defined, with one or more recovery points.

Figure 4-14 shows recovery group status of Ready and the details of governance for the policy.

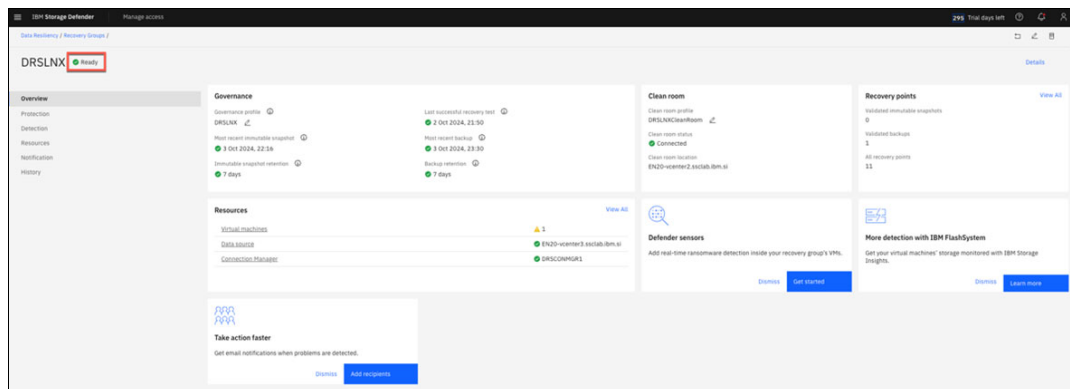


Figure 4-14 Recovery group details

The testing of recovery points for a recovery group establishes the recovery plan. This plan is then used if needed, in response to the occurrence of a cyber event. From the recovery points of selected recovery group, you can choose a recovery point which is required for testing. To select recovery point go to recovery group details, and from the Protection menu you will see all recovery points (Figure 4-15).

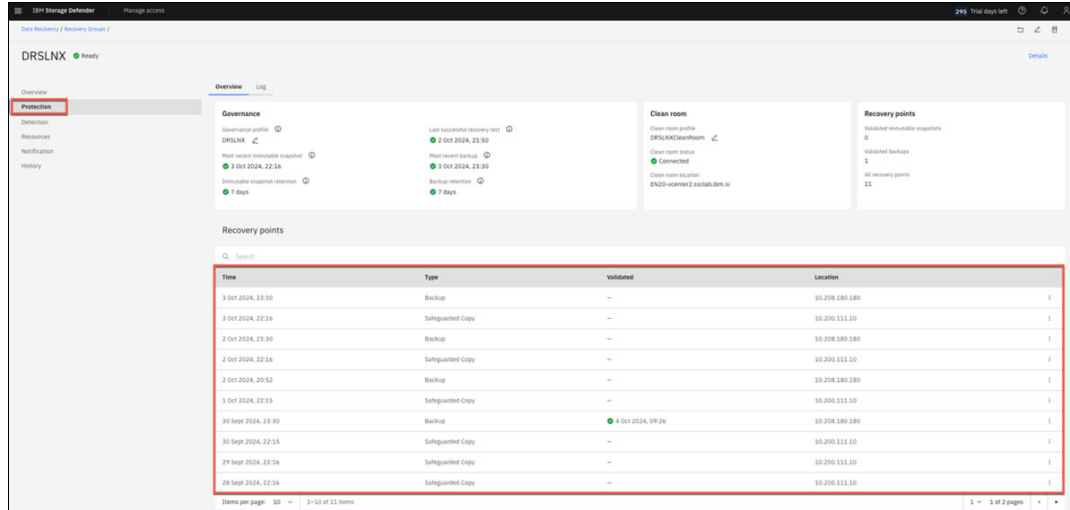


Figure 4-15 Recovery points details

These recovery points can be used to test or activate recovery plan. Figure 4-16 on page 57 shows the options that can be selected for each recovery point.




Figure 4-16 Recovery point details

Use “Test recovery point” to test a recovery of the the virtual machines that belong to the recovery group. These VM’s are then recovered by using the information that is stored in the clean room profile associated with the recovery group. Depending on the configuration of the clean room profile, the virtual machines are started and connected to the defined network or not. When the test recovery is finished successfully the status of the recovery point will be updated from Recovery in progress to Awaiting validation as shown in Figure 4-17.

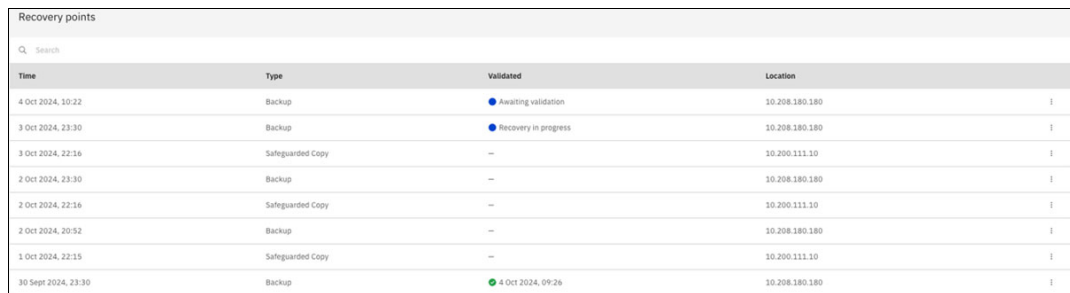


Figure 4-17 Recovery point status panel

Once the Recovery Point has been recovered to the clean room and is ready for validation a blue box will appear across the top of the page with a link to confirm that validation pass or failed testing Figure 4-18 on page 58.

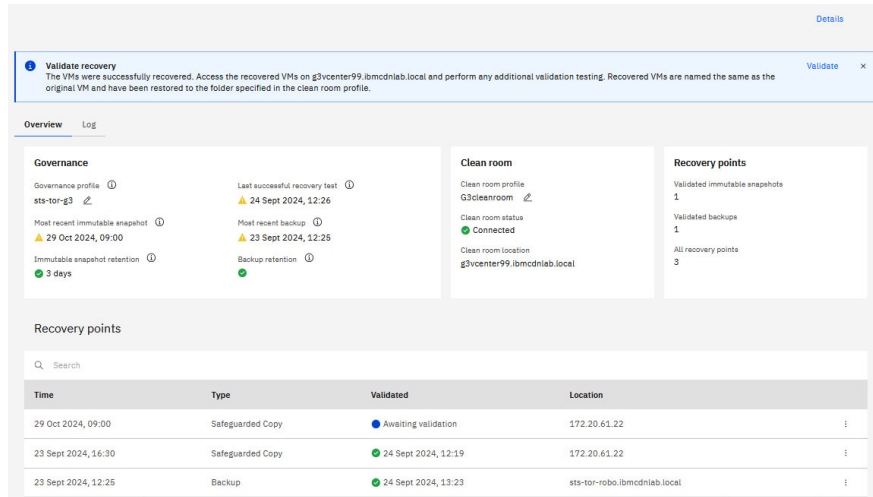


Figure 4-18 Test only confirmation dialog

Figure 4-19 highlights the ability to validate the recovery point following the restoration of the Recovery Group to the clean room. It allows user to identify the use case of bringing up the recovery and to define the status of the action as 'Test Only' or if the activity was part of a 'Recovery Plan' resulting from a Cyber incident. Then to be able to mark it as valid or not.

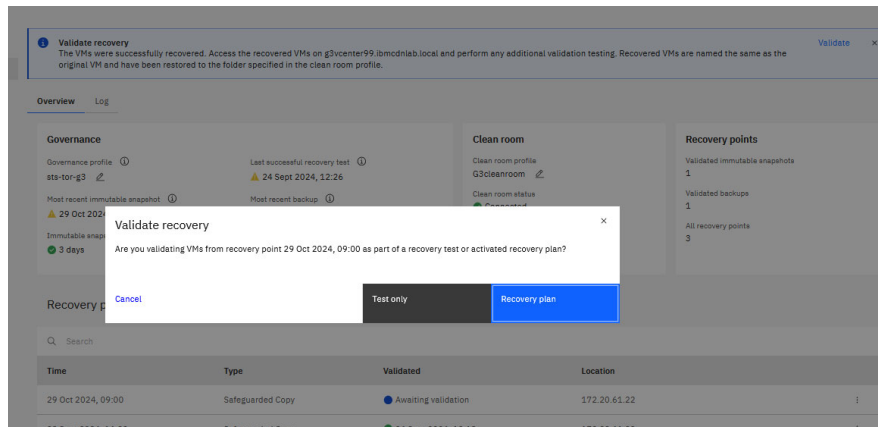


Figure 4-19 Validate recovery dialog

Figure 4-20 on page 59 allows the user to confirm the results of the recovery to the clean room and confirm the results.

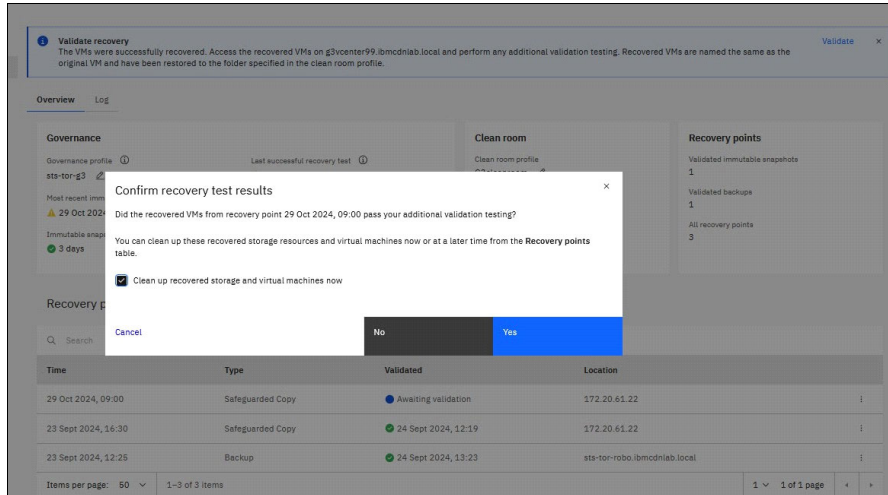


Figure 4-20 Confirm recovery results dialog

Once it has been determined if the recovery point is valid, the user is able to mark that 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 will be updated from Awaiting validation to Validated or Not valid. Figure 4-21 shows the different categorizations of a recovery point.

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, 22:36	Safeguarded Copy	—	10.200.111.10
2 Oct 2024, 23:30	Backup	● Awaiting validation	10.208.180.180
2 Oct 2024, 22:36	Safeguarded Copy	—	10.200.111.10
2 Oct 2024, 20:52	Backup	—	10.208.180.180
1 Oct 2024, 22:35	Safeguarded Copy	—	10.200.111.10
30 Sept 2024, 23:30	Backup	● 4 Oct 2024, 09:26	10.208.180.180

Figure 4-21 Validation status panel with invalid recovery points

Once the recovery test data is verified, if the cleanup option is selected in the test results dialog panel as shown in Figure 4-22 on page 60 the data will be confirmed as validated and the system will then cleanup the VM data that was restored as part of the validation test. If cleanup is not selected, the VMs will remain in the clean room and can be removed manually at a later time.

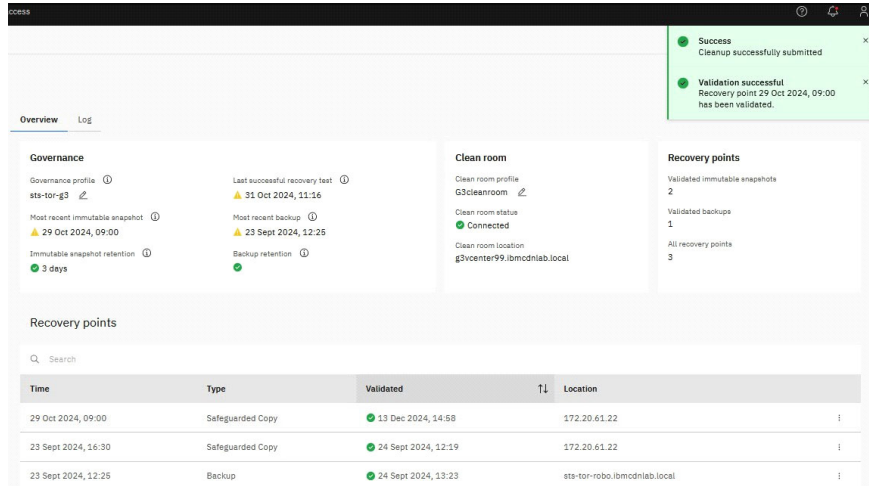


Figure 4-22 Validation and cleanup notification messages

4.5 Activating the Recovery Plan

The recovery group option ‘Activate recovery plan’ describes the actual recovery of resources that are associated with a recovery group. This Activate recovery plan option uses an existing and valid recovery point to recover your application after a cyber-attack 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 given 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 4-23 on page 60 shows the Activate recovery plan options where you select required recovery plan.

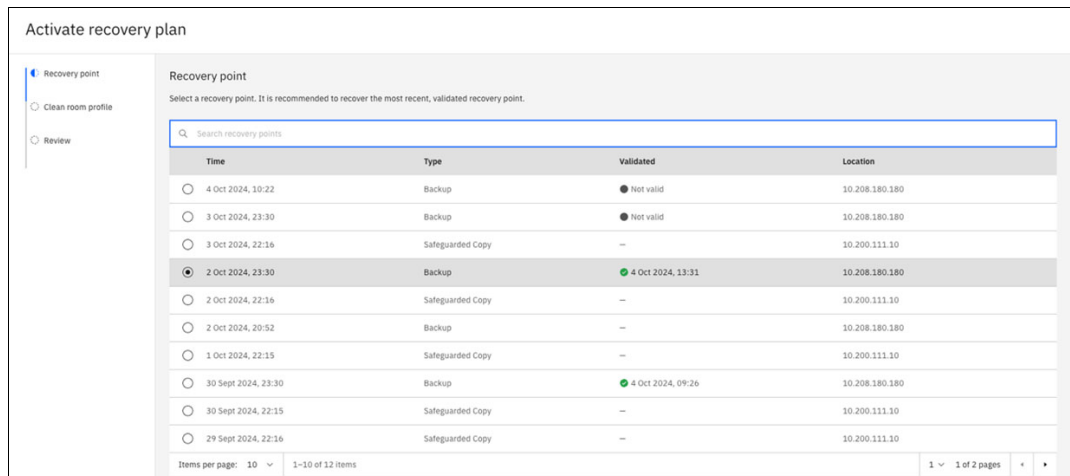


Figure 4-23 Activate recovery plan panel

In the next step you specify a clean room profile (Figure 4-24). After you review the profile settings (Figure 4-25 on page 61) click Done and wait for the recovery to complete.

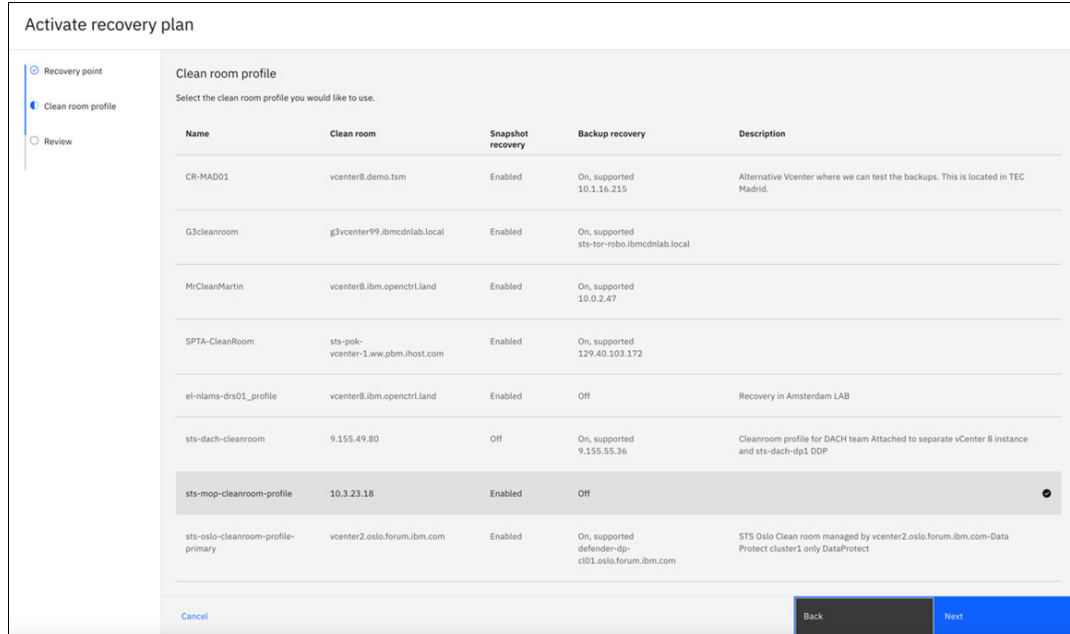


Figure 4-24 Activate recovery plan – Clean room profile

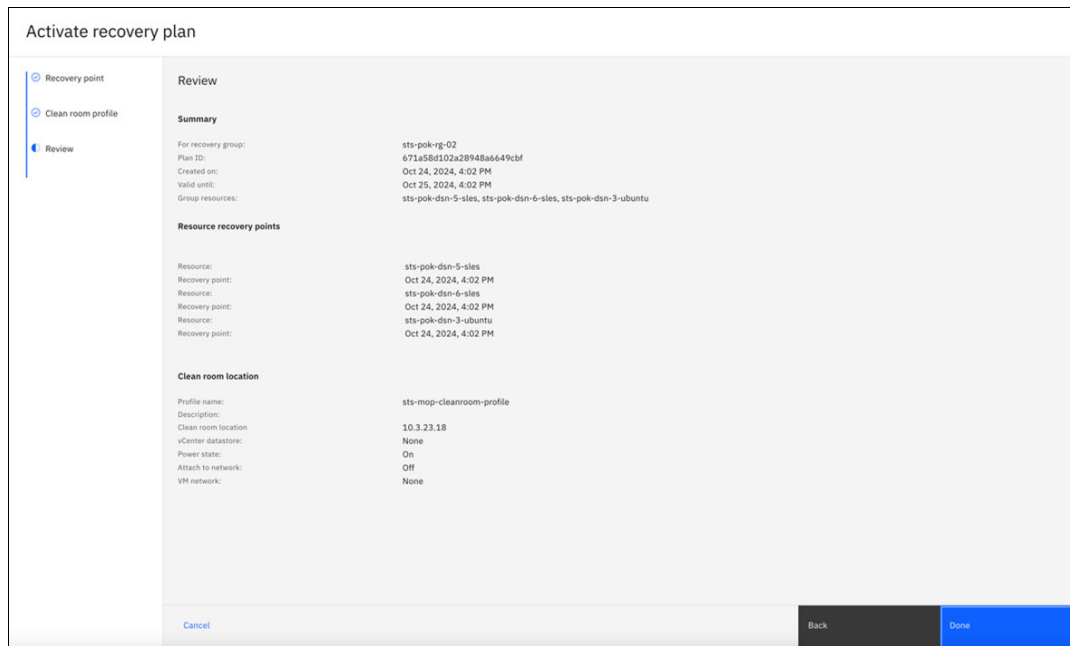


Figure 4-25 Activate recovery plan

Once confirmed, the Recovery in Progress panel in the Recovery Group's Overview panel will show the progress as seen in Figure 4-26 on page 62 and Figure 4-27 on page 62.

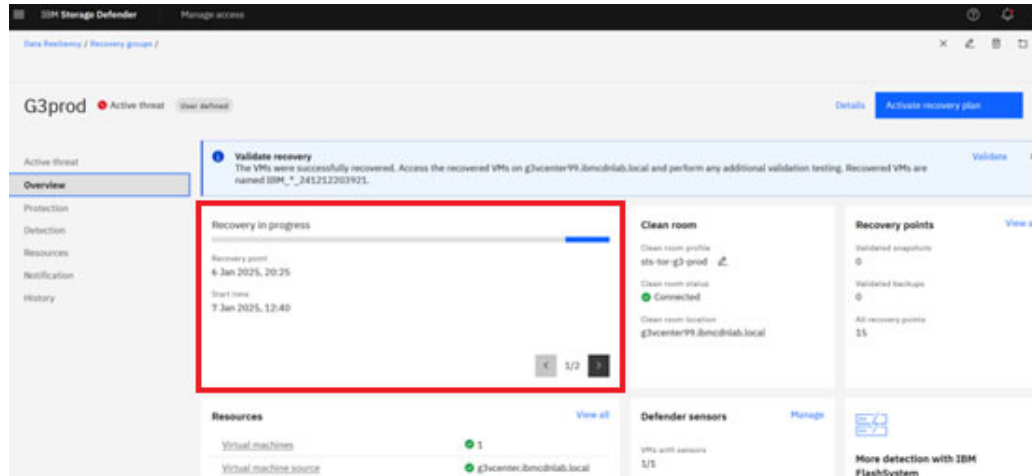


Figure 4-26 Recovery progress information in Recovery Group Overview panel example 1

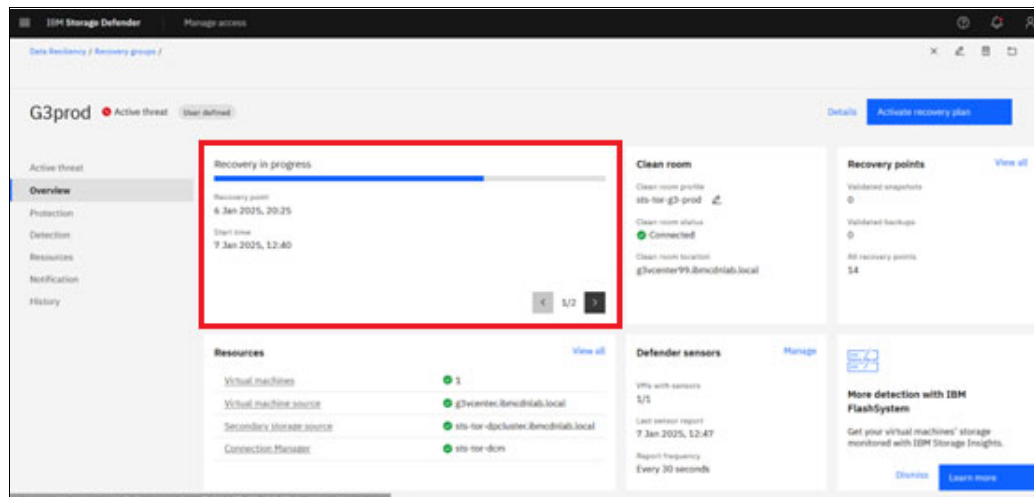


Figure 4-27 Recovery progress information in Recovery Group Overview panel example 2

From here, once the recovery process is completed, the user is able to access the Virtual Machines that were recovered to the clean room environment and return them to production as needed.



REDP-5744-00

ISBN

Printed in U.S.A.

Get connected

