

# IBM Storage Defender: Database Protection and Rapid Recovery

Christian Burns

Phil Gerrard


Juan Carlos Jimenez

Julien Sauvanet

Ken Salerno

Jack Tedjai

Christopher Vollmar







IBM Redbooks

**IBM Storage Defender: Database Protection and Rapid Recovery**

January 2025

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

**First Edition (January 2025)**

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

This document was created or updated on January 15, 2025.

© Copyright International Business Machines Corporation 2022. All rights reserved.

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

# Contents

<b>Notices</b> .....	v
Trademarks .....	vi
<b>Preface</b> .....	vii
Authors .....	vii
Now you can become a published author, too! .....	viii
Comments welcome .....	viii
Stay connected to IBM Redbooks .....	ix
<b>Chapter 1. Introduction to IBM Storage Defender</b> .....	1
1.1 Overview of IBM Storage Defender .....	2
1.2 Overview of IBM Defender Data Protect .....	2
1.3 Overview of IBM Defender Data Management Service .....	4
1.4 IBM Defender Data Protect and Database workloads .....	5
1.4.1 DB integration and agents .....	5
1.4.2 Remote Adapter .....	6
1.4.3 Universal Adapter .....	6
1.4.4 SmartTarget .....	6
1.4.5 Logs backed up as part of the policy .....	6
1.4.6 Megafire and Minion .....	7
1.4.7 Full to tape via Storage Protect S3 .....	7
1.4.8 Archive to S3 in the cloud .....	7
<b>Chapter 2. Protecting Microsoft SQL server</b> .....	9
2.1 IBM Data Protect MS SQL Server Protection Overview .....	10
2.1.1 Data Protection for Microsoft SQL CORE Terms .....	10
2.2 Requirements for Microsoft SQL Server Protection .....	10
2.3 Registering a Microsoft SQL Server .....	11
2.4 Recover from Microsoft SQL Server .....	15
2.5 Cloning a Microsoft SQL Server .....	17
2.5.1 Creating the Snapshot of the Microsoft SQL Server Backup .....	17
<b>Chapter 3. Protecting Oracle Databases</b> .....	27
3.1 IBM Data Protect Oracle Server Protection Overview .....	28
3.1.1 Oracle version support .....	28
3.2 Backup using the Oracle Adapter .....	28
3.3 Recovery using the Oracle Adapter .....	35
3.3.1 Recovering a Database via Instant Recovery .....	35
3.4 Backup using the Remote Adapter .....	48
3.5 Recovery using the Remote Adapter .....	58
3.5.1 Restoring a CDB to an alternate host .....	60
<b>Chapter 4. Protecting Microsoft Active Directory</b> .....	69
4.1 Protecting Microsoft Active Directory .....	70
4.2 Protecting the Microsoft Active Directory DB .....	70
4.3 Protect and Recover Microsoft Active Directory .....	72
<b>Chapter 5. Protecting Microsoft Exchange on Premises Data</b> .....	77
5.1 Backing up and restoring Microsoft Exchange Data .....	78

5.2 Register Exchange server as Data Protection Source . . . . .	79
5.3 Restore Exchange data using a recovery database . . . . .	81
<b>Chapter 6. Protecting PostgreSQL Databases.</b> . . . . .	<b>85</b>
6.1 Prerequisites and initial steps . . . . .	86
6.1.1 Versions requirements . . . . .	86
6.1.2 Communication port requirements . . . . .	86
6.1.3 Local user requirements . . . . .	87
6.1.4 Local Command Requirements. . . . .	87
6.1.5 Secure PostgreSQL host to DP cluster communications. . . . .	88
6.1.6 Other Requirements . . . . .	88
6.2 Deployment Overview . . . . .	89
6.3 IBM Storage Defender Data Protect capabilities for PostgreSQL database . . . . .	90
6.3.1 Backup and Recovery Methods . . . . .	90
6.3.2 Special Considerations . . . . .	90
6.3.3 Backup Workflows . . . . .	91
6.3.4 Recovery Workflows . . . . .	93
6.4 practical deployment example. . . . .	98
6.4.1 Download and install the Linux and PostgreSQL Connector agents . . . . .	98
6.4.2 Step 3: Register the PostgreSQL host machine as a source. . . . .	102
6.4.3 Step 4: Create a Protection Group and Protection Policy. . . . .	104
6.4.4 Step 5: Backup and Recovery activities . . . . .	107
6.5 Troubleshooting . . . . .	108

# 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.

AIX®	IBM Cloud®	QRadar®
DB2®	IBM FlashSystem®	Redbooks®
Enterprise Design Thinking®	IBM Security®	Redbooks (logo)  ®
Guardium®	IBM Spectrum®	Storwize®
IBM®	PowerPC®	XIV®

The following terms are trademarks of other companies:

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

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Ceph, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries 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 introduces the new IBM Storage Defender offering for enterprise data management and protection. This IBM Redpaper publication will help you install, tailor and configure this solution for the protection and rapid recovery of databases like Oracle, Oracle (OVM), Oracle RAC, SAP HANA, SAP Oracle, SAP DB2®, SAP MS SQL, SAP Sybase ASE, Sybase IQ & ASE, IBM DB2, MS SQL, Hadoop, IRIS and Cache for EPIC applications.

## Authors

This paper was produced by a team of specialists from around the world working at IBM Redbooks, Poughkeepsie Center.

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

**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.

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

**Julien Sauvanet** has been working in IT for 15+ years, covering many different areas including networking, systems, storage, and for the past 10 years data protection. He continues to share his knowledge and expertise as a contributing author to IBM Redbooks since 2013. Being involved with the ever evolving data protection world, he continues to expand his knowledge beyond the usual focus on data backup. As an SME focused on overall Data Resilience, keeping up to date with various techniques helps him with designing solutions which create additional value beyond just protecting backup data (data reuse, automation and orchestration of recoveries, infrastructure resiliency).

**Kenneth Salerno** is an Open Group Certified Distinguished Technical Specialist working for IBM in the USA. He has 27 years of experience in Information Technology. Prior to joining IBM, he worked 7 years as a Senior Infrastructure Engineer and Architect on Wall Street managing and supporting multiple data centers for mission-critical online financial services.

He holds a degree in Computer Science from CUNY Queens College in New York City. His areas of expertise include operating systems, storage, security, networking, middleware, databases, containers and enterprise data center operations. He has contributed code to various open source projects, currently he is one of the maintainers of the Jenkins CI Docker images, and is also Linux and Cisco certified.

**Jack Tedjai** is an IBM Certified Expert IT Specialist and IBM Systems subject matter expert, working in the Northern Europe Infrastructure Lab Expert Services organization. He joined IBM in 1998, and has more than 25 years of experience in the delivery of Storage, Storage Virtualize, Backup and Cyber Resiliency services for Open Systems. He is mostly involved in architecture and deployments world-wide for IBM Lab Expert Services, with a focus on IBM Storage Protect, IBM Storage Protect Plus and IBM Cloud® Object Storage.

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

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

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

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

[ibm.com/redbooks/residencies.html](https://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](https://ibm.com/redbooks)

- ▶ Send your comments in an email to:

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

- ▶ Mail your comments to:

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

## Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:  
<https://www.linkedin.com/groups/2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:  
<https://www.redbooks.ibm.com/subscribe>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:  
<https://www.redbooks.ibm.com/rss.html>





# Introduction to IBM Storage Defender

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

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

Although the current world of IT may seem like a dangerous place with new and creative attempts to exploit vulnerabilities, careful planning and execution of appropriate data security and data resilience processes can enable organizations to gracefully recover from otherwise dire situations. This Redbooks publication provides guidance on one of IBM's solutions dedicated to these use cases, enabling customers to recover rapidly, and at scale.

In this chapter:

- ▶ 1.1, "Overview of IBM Storage Defender" on page 2
- ▶ 1.2, "Overview of IBM Defender Data Protect" on page 2
- ▶ 1.3, "Overview of IBM Defender Data Management Service" on page 4
- ▶ 1.4, "IBM Defender Data Protect and Database workloads" on page 5

## 1.1 Overview of IBM Storage Defender

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

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

### **Key capabilities of IBM Storage Defender include:**

#### ***Data Resilience and Compliance***

- ▶ Set policies and standards to ensure resilience compliance across the data estate.

#### ***Early Threat detection***

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

#### ***Safe and Fast recovery, at scale***

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

#### ***Flexible licensing***

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

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

Defender is comprised of various components designed to meet customers resilience needs.

## 1.2 Overview of IBM Defender Data Protect

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

**Defender Data Protect supports the following workloads:**

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

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

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

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

***Integrated Cybersecurity:***

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

***Instant Mass Restore:***

- ▶ Enables users to restore a high number of VMs instantly. For example, in testing it has been shown that 200 VMs can be restored in around 10 minutes and 2000 VMs in under an hour. This drastically reduced downtime after an incident like a large-scale malware attack or ransomware attack.

***Global Actionable Search:***

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

***Fast Cloning:***

- ▶ Extremely fast cloning of large databases for devOps, testing, and other development use cases. For example cloning a 2 TB SAP HANA database in under a minute.

***Global Space Efficiency:***

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

***Primary Storage Integration:***

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

## 1.3 Overview of IBM Defender Data Management Service

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

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

**Other noteworthy DMS capabilities include:*****Quorum:***

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

***Security Advisor:***

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

***Simulations:***

- ▶ This functionality offers predictive planning models that can make projections about utilization and storage consumption. This capability is based on historical usage, workloads, and user-defined what-if scenarios. This empowers users to proactively plan for various situations, such as acquiring new nodes, integrating new workloads, optimizing current workloads, and more. Simulations can be created with scenarios using specific clusters and time periods to help better understand and plan environment changes.

***Reports:***

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



## 1.4 IBM Defender Data Protect and Database workloads

### 1.4.1 DB integration and agents

As noted above, there a variety of DB workloads supported leveraging IBM Defender Data Protect. They include: Oracle, Oracle (OVM), Oracle RAC, SAP HANA, SAP Oracle, SAP DB2, SAP MS SQL, SAP Sybase ASE, Sybase IQ & ASE, IBM DB2, MS SQL, Hadoop, IRIS and Cache for EPIC.

There is a variety of options for protecting a DB workload, IBM Storage Defender Native DB integration currently includes the following:

- ▶ Amazon RDS
- ▶ Amazon Aurora
- ▶ Cassandra
- ▶ MongoDB
- ▶ Microsoft SQL Server
- ▶ Oracle Database

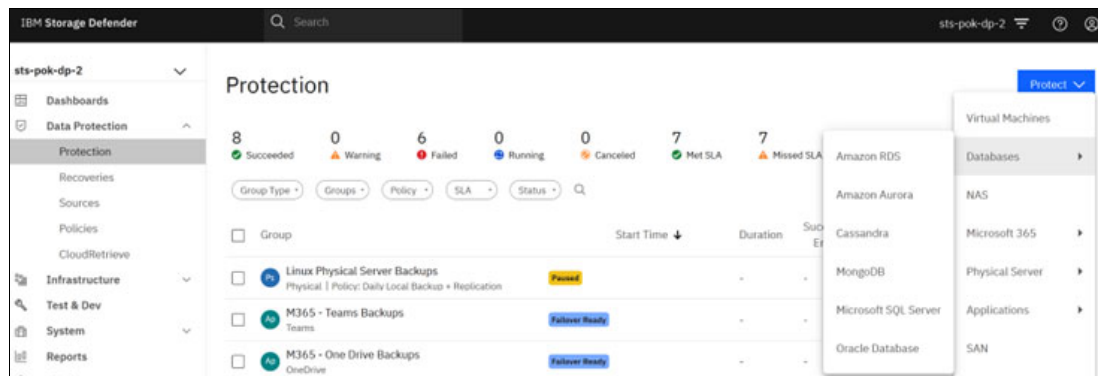


Figure 1-1 IBM Storage Defender Protection dashboard

This integration provides the ability to write directly to the IBM Storage Defender Cluster as a target.

There are additional options for protecting Database workloads which can leverage an agent based approach, leveraging either the Remote Adapter or the Universal adapter.

Figure 1-2 on page 6

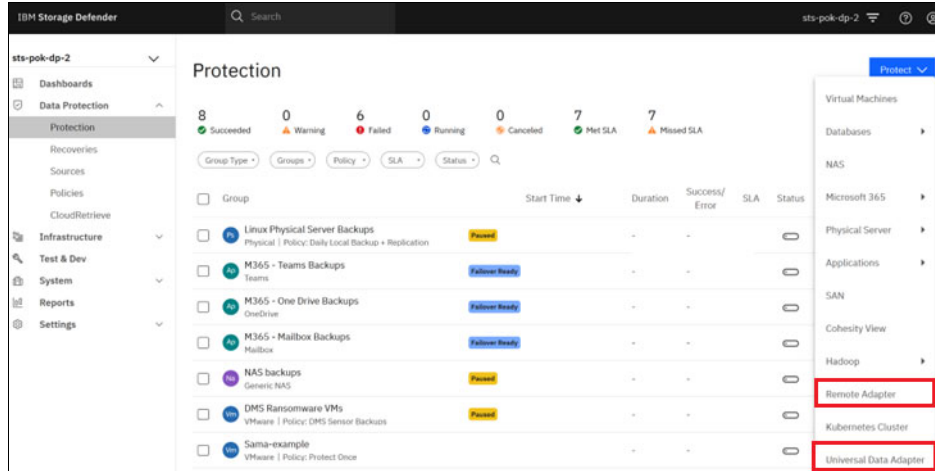


Figure 1-2 IBM Storage Defender Protection Remote Adapter options

### 1.4.2 Remote Adapter

The Remote Adapter offers flexible management of scripts running on remote hosts. This allows the Defender Data Protect cluster to manage data protection processes and schedules as well as provide a consolidated log of all activity.

### 1.4.3 Universal Adapter

The Universal Data Adapter or UDA is a tool that automates and simplifies database protection for a variety of databases. It allows users to register, create backup jobs, add database instances, restore data across multiple database types, and consolidate protection and scheduling across multiple servers and sources

### 1.4.4 SmartTarget

IBM Storage Defender Data Protect also includes the ability to backup to a File System target, such as an NFS or a CIFS/SMB target. This provides the ability to maintain existing scripts and practices that exist in an enterprise and update the backup target to the cluster. This can be leveraged for the DB workloads with native integration as well as those with agent based integration, leveraging the available agents in both scenarios. The advantage to this approach over traditional backups to standard filesystem mounts that then need to be backed up includes, speed to ingest to the backup system and simplified process among other benefits.

### 1.4.5 Logs backed up as part of the policy

When setting the Protection Policy for a DB backup, such as MS SQL, the Log backups can be included. It provides the option to backup the logs to a granularity of hours or minutes to provide additional protection and recovery points.

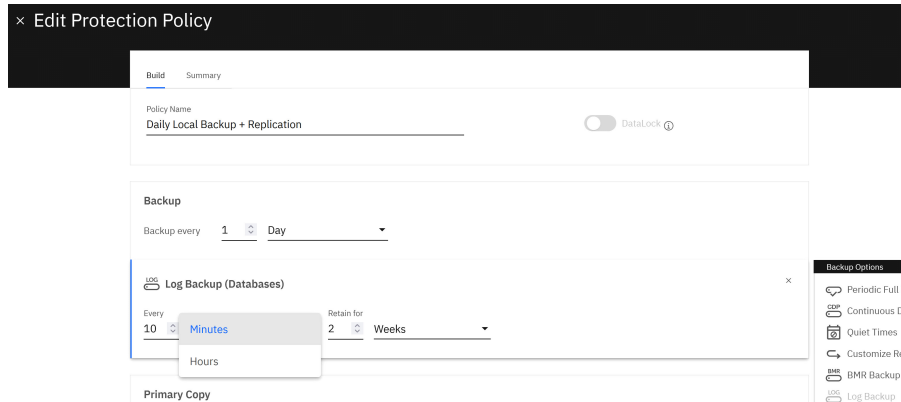


Figure 1-3 Protection Policy log backup configuration panel

## 1.4.6 Megafile and Minion

MegaFile provides quick backups and restores for large, multi-terabyte files (256 GB+). MegaFile breaks large files into smaller chunks and distributes these files across all Data Protect nodes in a defined cluster for parallel backup and recovery. The minimum size of these chunks is determined and optimized to maximize performance.

Minion or MinionBlob provides quick backup for small files. Minion file metadata is consolidated for groups of small files (8MB or smaller) into single logical metadata objects.

## 1.4.7 Full to tape via Storage Protect S3

Defender Data Protect and Storage Protect (SP), formerly Spectrum Protect and TSM, users are able to leverage both solutions to protect data on tape through Storage Protect's S3 container pools. Users are able to leverage this target for tape archive purposes with no additional licensing cost and no additional capacity licensing needed for SP, as long as customer is licensed to use SP. At this moment, only full copies are supported with incremental on the immediate roadmap.

## 1.4.8 Archive to S3 in the cloud

Defender Data Protect users can leverage CloudArchive to back up data on a cluster and then copy it to the cloud for archive purposes. CloudArchive first backs up your data onto a Data Protect cluster and then copies your backups to an external target, examples are AWS, Azure, GCP, and Oracle, as well as S3-compatible storage, or NFS-mounted storage.





# Protecting Microsoft SQL server

This chapter describes the management of Microsoft SQL Server databases with IBM Data Protect Cluster. Microsoft SQL Server is supported as a stand-alone/failover cluster and Always On Availability Groups (AAGs) database.

This chapter includes the following topics:

- ▶ 2.1, “IBM Data Protect MS SQL Server Protection Overview” on page 10
- ▶ 2.2, “Requirements for Microsoft SQL Server Protection” on page 10
- ▶ 2.3, “Registering a Microsoft SQL Server” on page 11
- ▶ 2.4, “Recover from Microsoft SQL Server” on page 15
- ▶ 2.5, “Cloning a Microsoft SQL Server” on page 17

## 2.1 IBM Data Protect MS SQL Server Protection Overview

In this section, we describe the features of IBM Data Protect with Microsoft SQL Server. As of December 2024, the following features are supported:

- ▶ Backup, restore, and recovery of stand-alone/failover MS cluster and Always-On Availability
- ▶ Incremental forever data and log backups, including log truncation
- ▶ Automatic discovery of SQL instance on registered servers
- ▶ Production restore (database is restored by copying data):
  - To original location
  - To alternative location (that is, alternative source path)
- ▶ Restore to alternate instance and / or database name
- ▶ Recover to specific point-in-time (requires log backups enabled)

### 2.1.1 Data Protection for Microsoft SQL CORE Terms

The following Protection feature can be used to protect MS SQL Servers:

**AAG** - AlwaysOn Availability Groups is a database mirroring technique for Microsoft SQL Server that allows administrators to pull together a group of user databases that can fail over together.

**AG Replica** - The term "replica" typically refers to availability replicas. For example, "primary replica" and "secondary replica" always refer to availability replicas.

**FCI** – Failover Cluster Instance is a single instance of SQL Server that is installed across Windows Server Failover Clustering (WSFC) nodes.

**SQL VIP** – The IP address of the SQL Instance that moves from one physical node to another when a node fails.

**SQL System Databases** – A set of four system-level databases (master, model, msdb, tempdb), which are essential for the operation of a server instance.

IBM Data Protect Cluster can protect MS SQL databases on physical servers by utilizing the following IBM Defender Data Protect Agent-based methods:

- ▶ Volume-based backup utilizes VSS to take backup of all MS SQL databases running on a volume (or all volumes) on MS SQL server
- ▶ File-based backup utilizes VSS to take backup of ONLY specified MS SQL databases running on a MS SQL servers
- ▶ VDI-based backup allows the IBM Defender Data Protect to execute SQL server Native backup and restore commands via native VDI API calls

## 2.2 Requirements for Microsoft SQL Server Protection

In this section, we describe the requirements to protect Microsoft SQL Server:

1. Credentials and Privileges for Microsoft SQL Server Protection.

There are three accounts you must consider when installing the IBM Data Protect Agent:

- ▶ **Installation Account:** The account you use to log in to the host and run the installer.
- ▶ **Service Account:** The account under which the IBM Data Protect Agent service runs on the SQL Server host.
- ▶ **SQL Server login account:** The SQL Server account by which the IBM Data Protect Agent has access to the databases. (Configured after installation.)

You can use either the host LOCAL SYSTEM account or an account that meets the requirements to install the IBM Data Protect agent. It is recommended to run the IBM Data Protect Agent service with an Active Directory domain user account that is a member of the local administrator of the SQL Server host. The AD domain user account must be a member of the SQL sysadmin server role. The user account must have log-on rights to the SQL Server host in the local security policy of the SQL Server host.

If you do not use the LOCAL SYSTEM account, ensure the following for the chosen account:

- ▶ The account must be a member of the local Windows Administrators group and local Windows Administrators group on the SQL server.
- ▶ The account must have Log on as a service in the User Rights Assignment on the MS SQL server to install the IBM Data Protect agent.
- ▶ The account must have the sysadmin role in the MS SQL Server instance.

#### 2. Ports Used for Communication:

- ▶ On physical servers or VMs with an ephemeral or installed agent, open the ports 445, 11113, 11117, and 50051.

#### 3. If the Windows Firewall is used:

Inbound rules:

- ▶ Add a rule to accept SQL Server traffic and TCP connections on local port 1433.
- ▶ Set Remote Port to All Ports.

#### 4. Outbound rules (for MS SQL Server 2016 running on Windows 2016):

- ▶ Update the "Block network access for R local user accounts in SQL server instance MSSQLSERVER" rule by going to General tab > Action window > select "Allow the connection".

## 2.3 Registering a Microsoft SQL Server

To register Microsoft SQL Server as data source:

1. Expand to Data Protection
2. click on Sources
3. Databases and select Microsoft SQL Server
4. Fill in the Microsoft SQL Server host dns record.

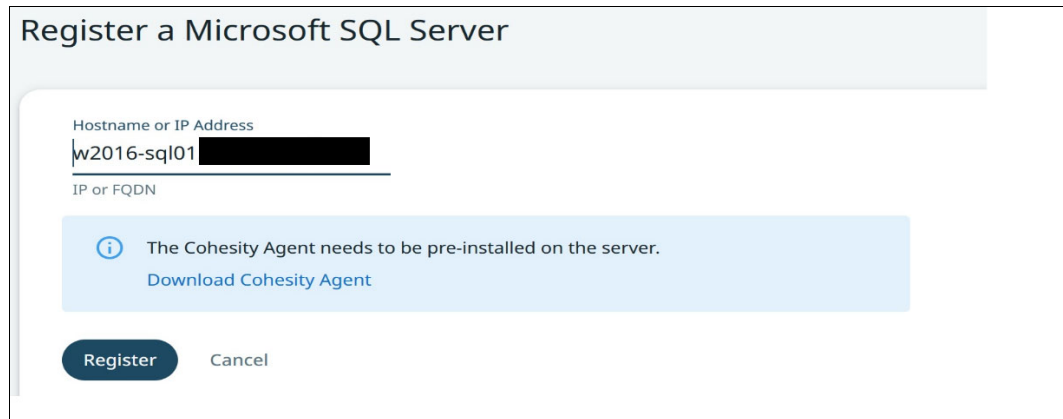


Figure 2-1 Register MSSQL server panel

If the IBM Data Protect (Cohesity) agent is not yet installed click 'Download Cohesity Protection Service':

- ▶ Ensure that the agent has been copied over to the appropriate server.
- ▶ As an AD Domain Admin, run the executable and complete the installation wizard

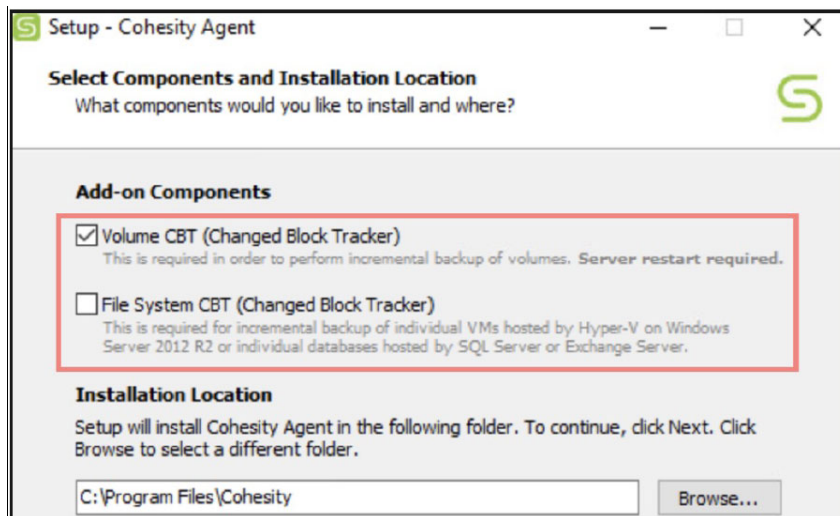


Figure 2-2 Windows Agent Installation panel

Volume CBT (Changed Block Tracker): Install this component for the best incremental backup performance. Installing this component requires a onetime reboot to load the IBM Data Protect Volume CBT driver.

File System CBT (Changed Block Tracker); the reboot is not required but not recommended.

Service Account Credentials: The service can run as the "Local System" account with Exchange admin credentials.

If the SQL requirement are not meet, then you may receive the following message issued:



Source	Protected	Protected Size	Total Size	Last Refreshed
w2016-sql01.STGHDK.LOCAL	No	0 Bytes	99.9 GiB	30 minutes ago

**Health Check Problems**  
Registration was successful, but the following Health Check tests did not pass. You can correct the issues and refresh the Source.

Health Check	Status	Messages
Are privileges sufficient for Microsoft SQL Instance service, SQLWriter service, and Cohesity Agent?	✗	The account NT AUTHORITY\SYSTEM for the service Cohesity Agent Service is not a sysadmin of the sql instance(s) MSSQLSERVER.

Figure 2-3 Required Permissions for Windows Agent

Registering an MS SQL server as a data source steps continued:

5. Right Click [5] and Select Protect Database to Protect the SQL instances
6. Fill in a meaningful Protection Group name [6]
7. Select an existing SLA Policy [7]
8. Select the Logical Storage Domain [8]
9. Select the Backup type File-based, Volume-based or VDI-based

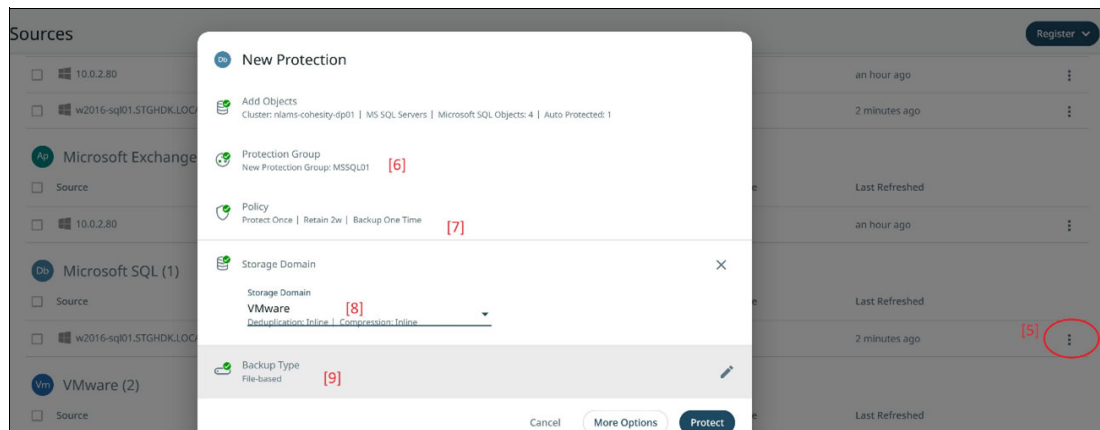


Figure 2-4 Creating a Microsoft SQL Protection group

**Note:** When setting the Protection Policy for a DB backup, such as MS SQL, the Log backups can be included. It provides the option to backup the logs to a granularity of hours or minutes to provide additional protection and recovery points. An example of this configuration is shown in Figure 2-5 on page 14.

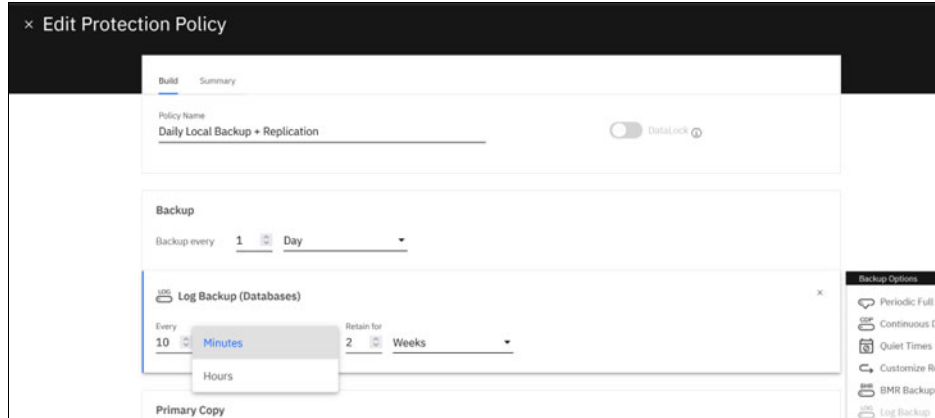


Figure 2-5 Protection policy log backup settings

**Best Practice:** For VDI backups, the backup and restore retention requirements are identical to SQL native dumps. This means that you will need to perform a full backup, incremental backup (equivalent to SQL Server Differential backup), and T-log backups for PIT recoveries.

For example, if you have a retention requirement of 7 days for a SQL Server DB backup, your VDI protection job policy should ensure that the retention period that encompasses the full and incremental retention period is greater than 7 days. Similar, to SQL native dumps, a full and incremental backup is required for restore. Setting the retention period for longer than 7 days to encompass the full, incremental, and even t-log backups will ensure that there is no hole in the recovery when using VDI.

Finish configuring the protection group by clicking the protect button. Once this is complete, monitor the running task to confirm the selected items are successfully protected (Figure 2-6 on page 14).

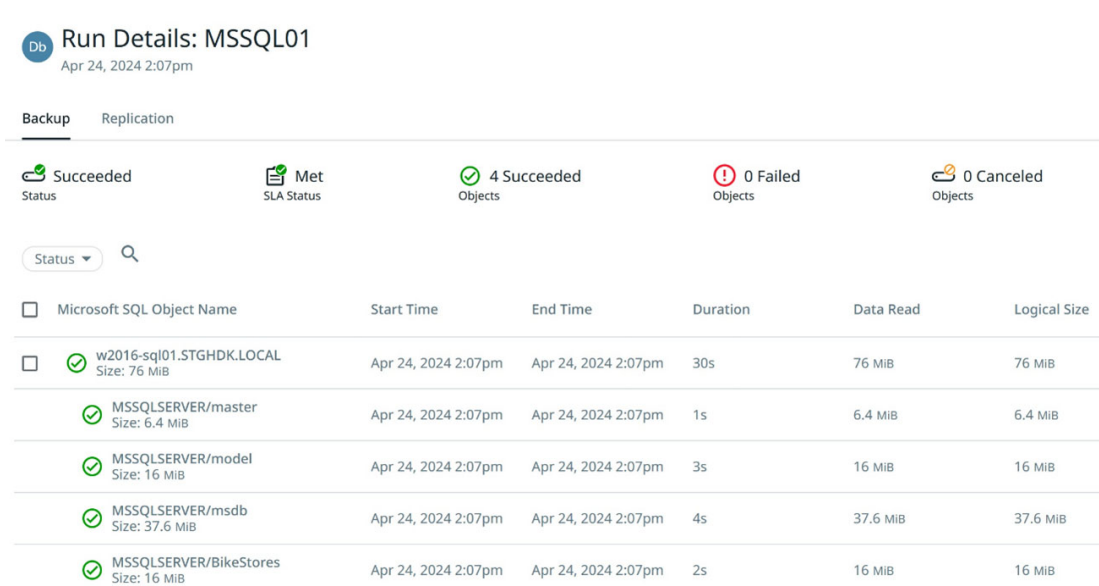


Figure 2-6 Successful Microsoft SQL Protection group backup

## 2.4 Recover from Microsoft SQL Server

To Recover Microsoft SQL Database:

1. Expand to Data Protection and Recoveries
2. On the right panel click Recover, Databases and Microsoft SQL Server
3. Databases and select Microsoft SQL Server
4. On the filter bar search for the Microsoft SQL Databases [4]

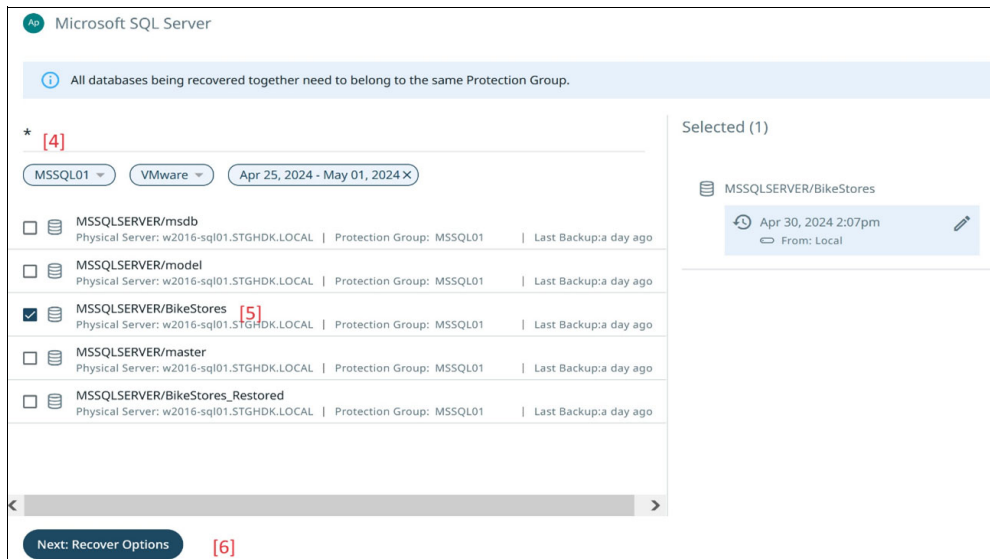


Figure 2-7 Recover Microsoft SQL Database step, server options selection panel

5. Select the Microsoft SQL Database which you like to recover [5]
6. Click Next for more Recover Options [6]
7. Select to Recover as new Database or Overwrite Original Database [7]
8. Select the Microsoft SQL Instance [8]

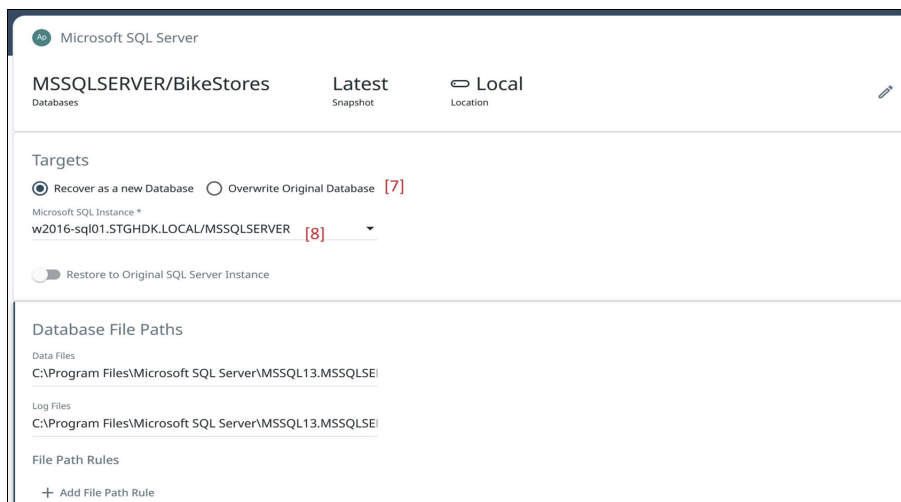


Figure 2-8 Recover Microsoft SQL Database, recovery options selection panel

9. Complete the recover task by clicking the 'Recover' button

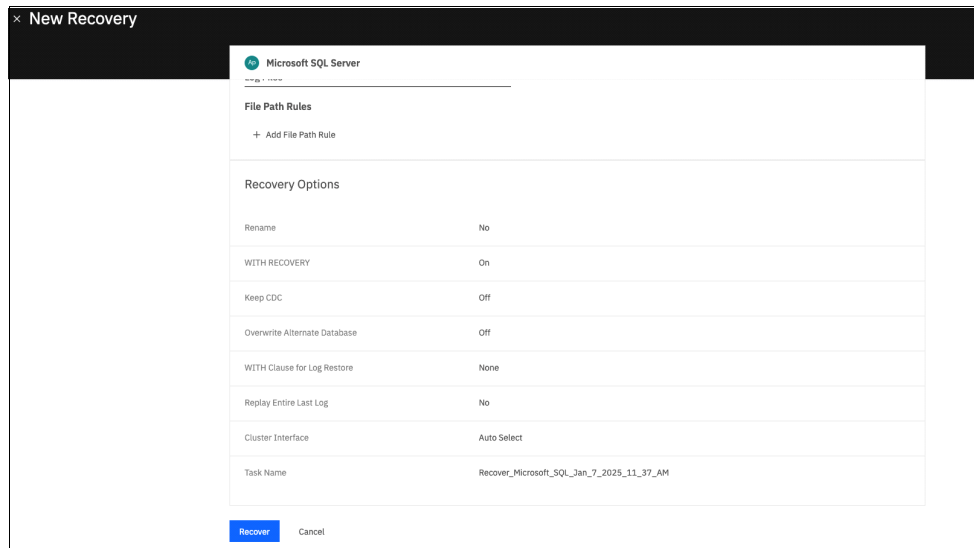


Figure 2-9 Recover Microsoft SQL Database, recovery trigger

When log backups are enabled as part of the Protection Policy recoveries are able to be time adjusted based on desired recovery point. This provides for a recovery that has been pre-set to the designed time stamp.



Figure 2-10 SQL DB recovery point settings panel

After the recover task is finished open Microsoft SQL Server Management Studio and confirm is the database is restored.

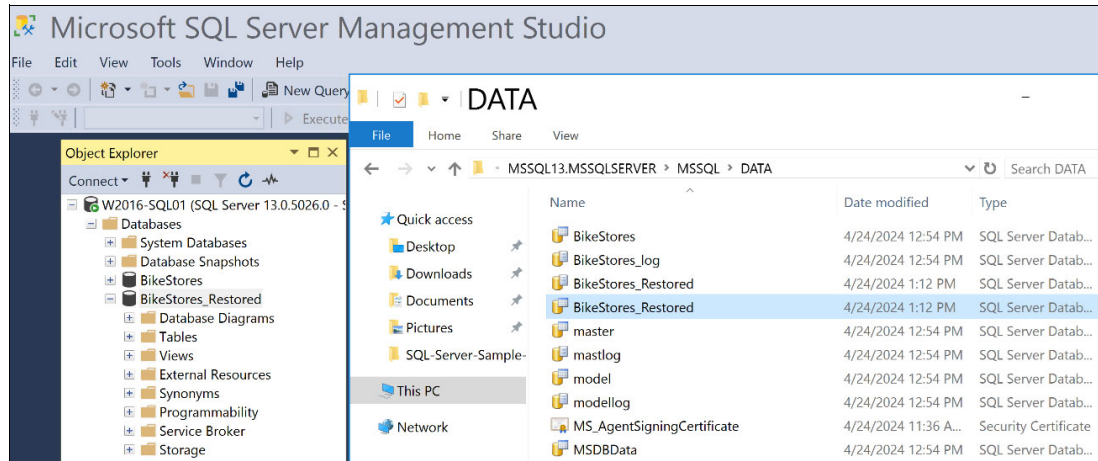


Figure 2-11 Confirming SQL DB restore with the Server Management Studio

## 2.5 Cloning a Microsoft SQL Server

IBM Storage Defender Data Protect also provides the ability to rapidly clone database workloads such as Microsoft SQL Server. It enables the clone to be mounted directly to the database instance but the data remains on the cluster. This is sometimes referred to as a “capacity free clone” as it makes the database instance immediately available, but it does not consume capacity in the primary storage platform. The data stays on the IBM Defender Data Protect cluster using a snapshot of the backup being cloned and placed in a mount that is read/write capable.

### 2.5.1 Creating the Snapshot of the Microsoft SQL Server Backup

There are two potential ways to clone the backup of the database instance, 1) it can be done from the Microsoft SQL specialty page, or 2) it can be done from the Test & Dev section of the left hand navigation area. Both methods are available from IBM Storage Defender DMS.

#### ***Option 1: using the Microsoft SQL tab databases menu:***

From the Microsoft SQL specialty page, view the list of protected database hosts and their database protection status then select the database view to get a list of protected DB's.

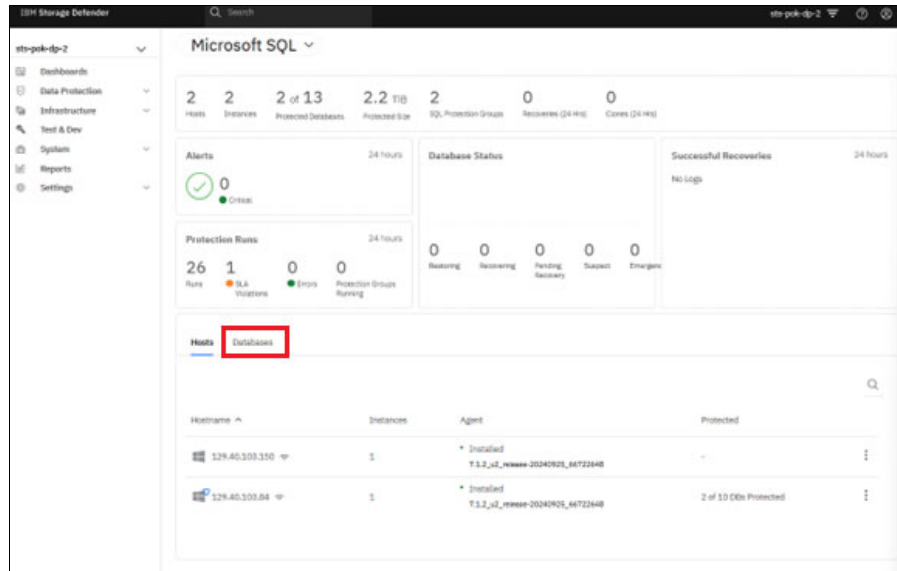


Figure 2-12 Microsoft SQL specialty page

Select the three dots on the desired database and choose the “Clone” option:

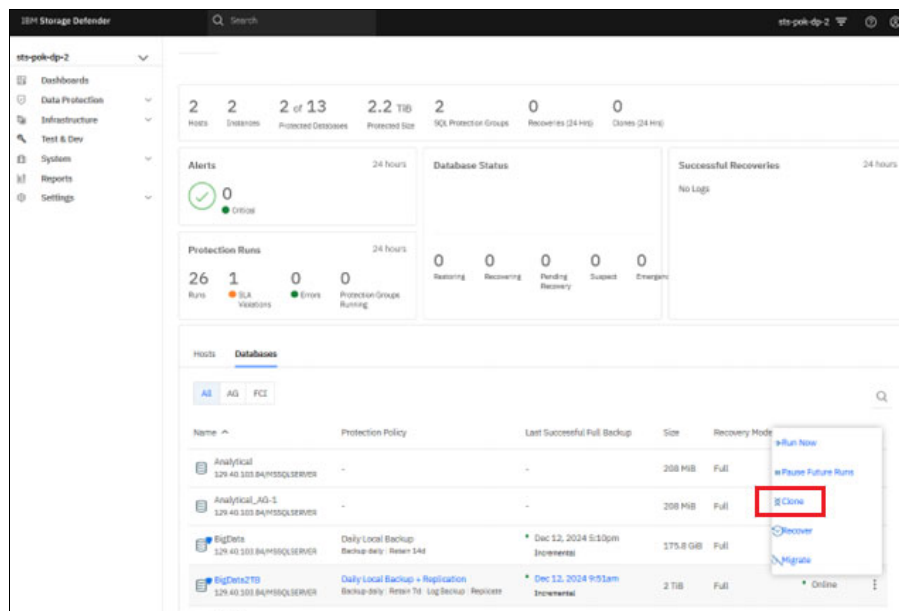


Figure 2-13 Clone a database instance option from the Microsoft SQL specialty page

### Option 2 using the Test & Dev function

Navigate to the side menu and select Test & Dev from the list:

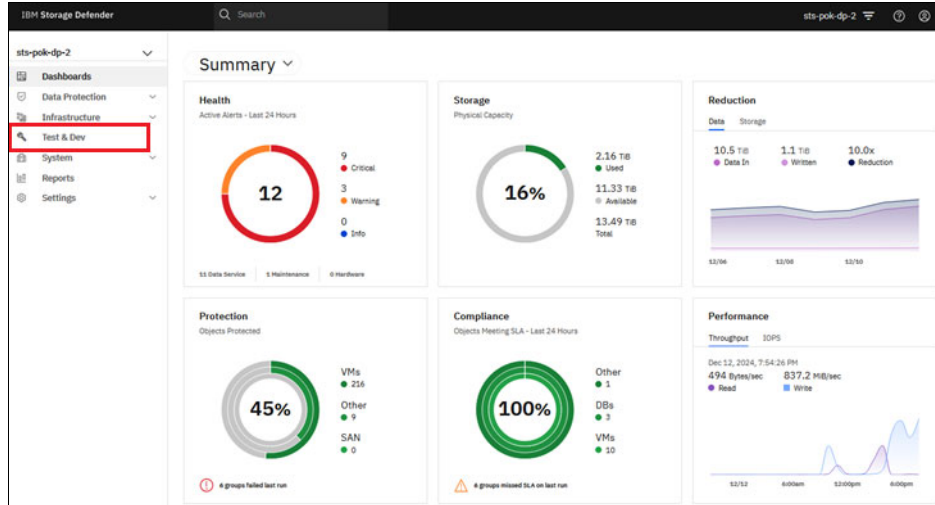


Figure 2-14 Test and Dev summary panel

Select 'Clone' and 'Database' from the top right corner drop down menu:




Figure 2-15 Selecting the clone function

From this page, use the search field to find the desired DB you wish to clone:




Figure 2-16 Find and select the desired DB for cloning

In this case we will select the instance called “**MSSQLSERVER/BigData2TB**”

From there you are able to perform the following actions:

- ▶ Create a Task Name
- ▶ Select a Clone Point (screen shot below)
- ▶ Identify the SQL Host to mount the clone to
- ▶ Identify the SQL Instance
- ▶ Rename the Database if needed
- ▶ Use SCN for the Clone
- ▶ Select the network interface

**Note:** for use in a later step, the name of the IBM Storage Defender Data Protect Cluster in this example is ‘STS-POK-DP-2’

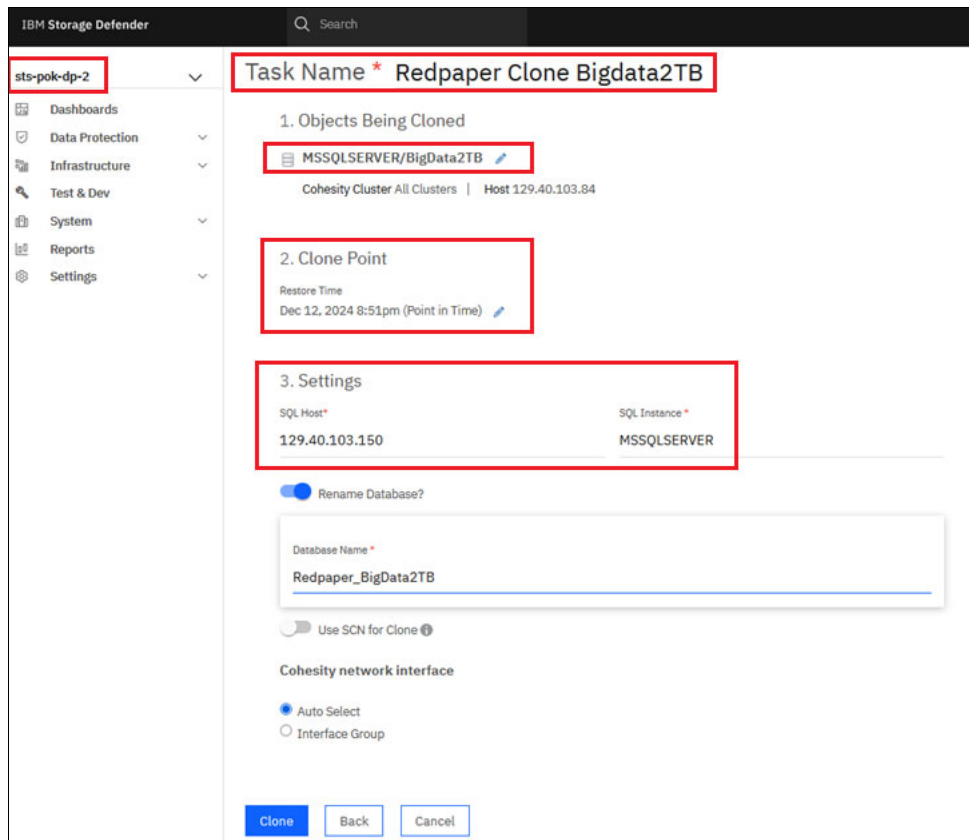


Figure 2-17 DB clone configuration settings

**Clone Point option:** One of the selection for creating the database clone is the ability to select the point in time the have the clone restored to, as mentioned above in the recovery part of this chapter, a further example is included below, on selecting the Clone Point, by being able to select the backup and the time stamp to automatically recover the logs to.



Figure 2-18 Clone Point options panel and point in time selection

Once the Clone Point is selected, and other restore options are set, select “Clone” to generate the clone of the database.



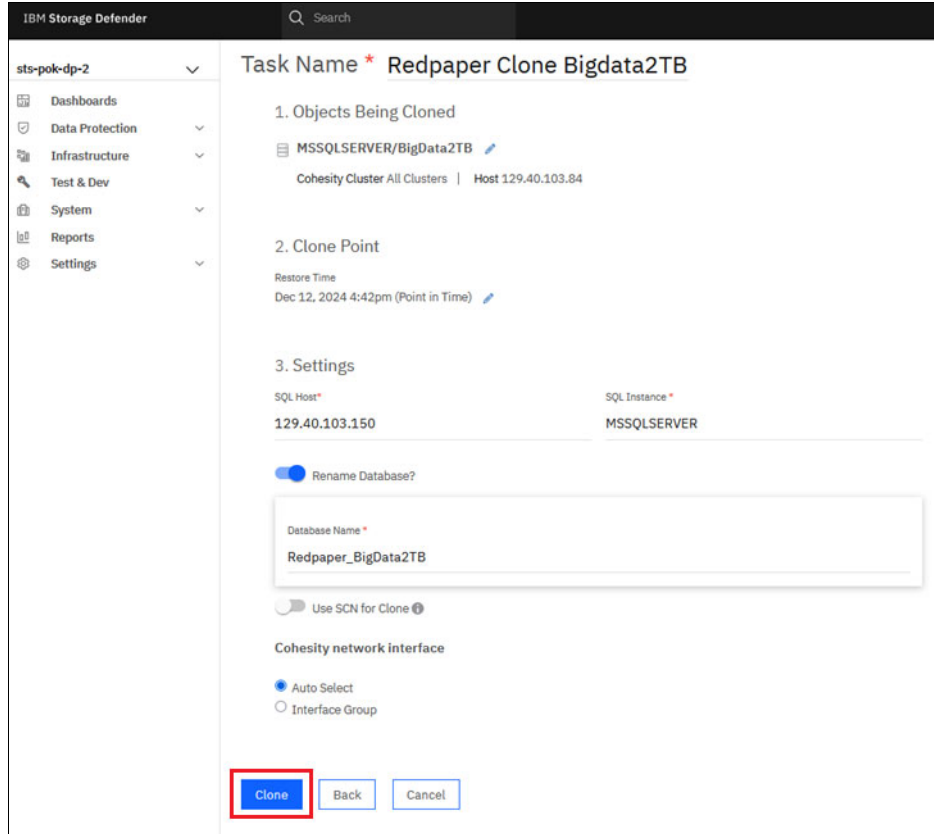


Figure 2-19 Generate clone button

Next, the clone process will begin and the progress and task details page will appear. You are also able to select “Show Subtasks” to check the status of clone process:

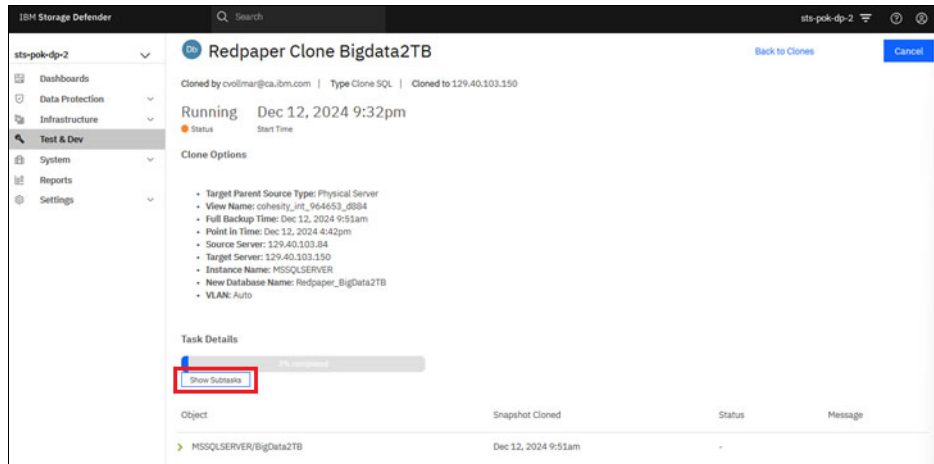


Figure 2-20 Show status of clone process

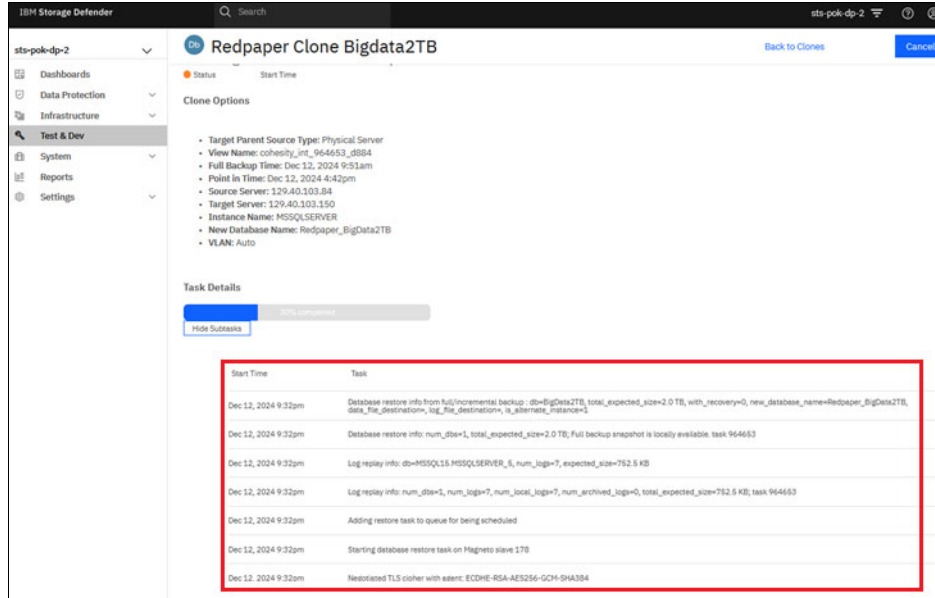


Figure 2-21 DB clone process status details

Once the clone process is completed, the task history will update with current DB information and the status will update to Success.

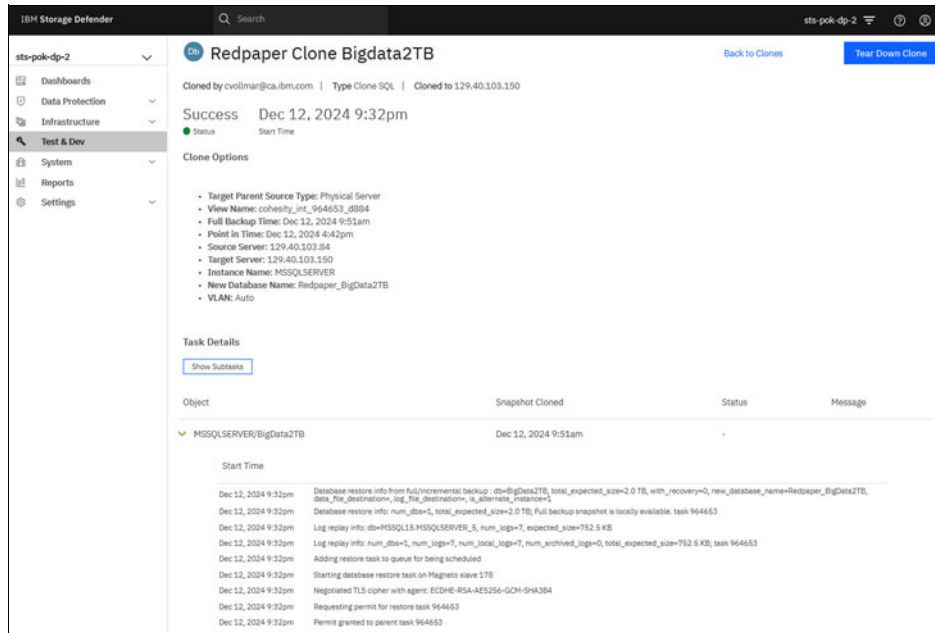


Figure 2-22

Once the status of Success is reported, the DB clone is then accessible via the MS SQL server.

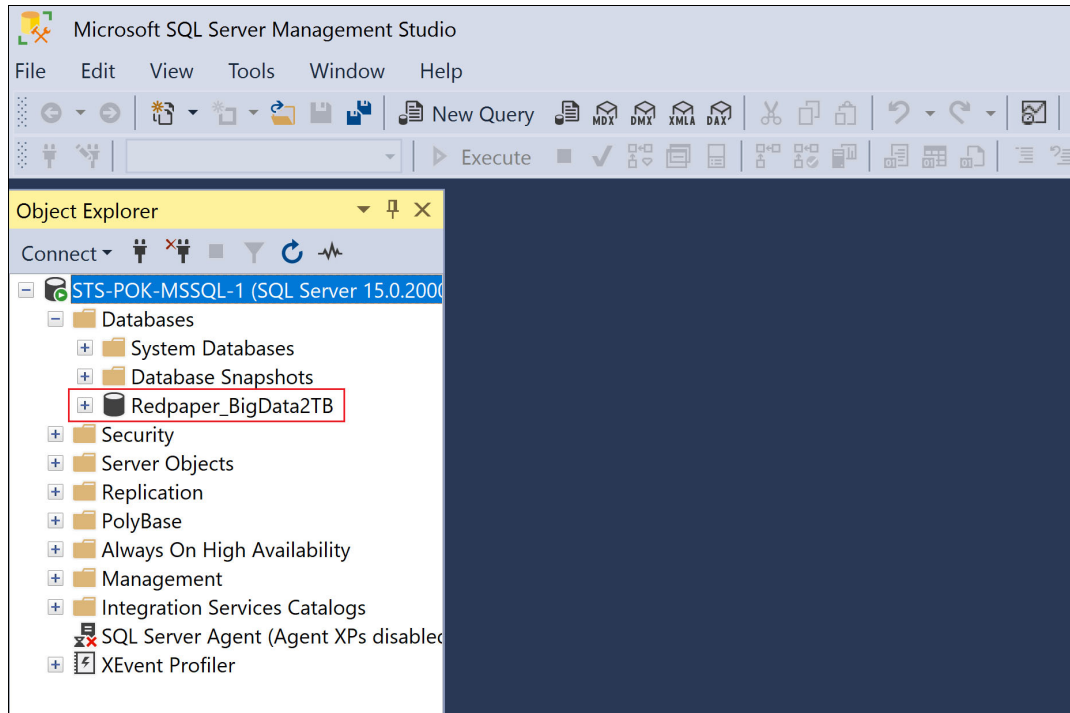


Figure 2-23 Cloned db access via the MS SQL Server Management Studio

The clone of the database is available for testing, or other activities. The database Properties will also provide additional information highlighting where the data resides.

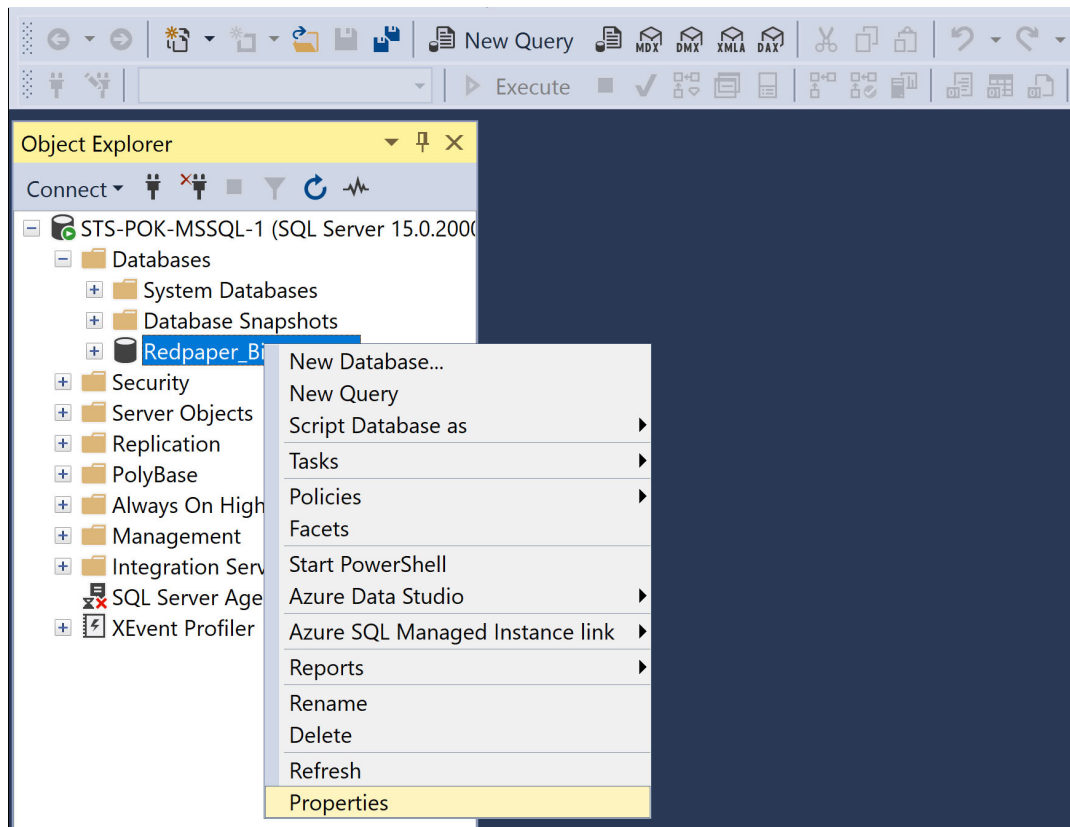


Figure 2-24 Review cloned DB properties

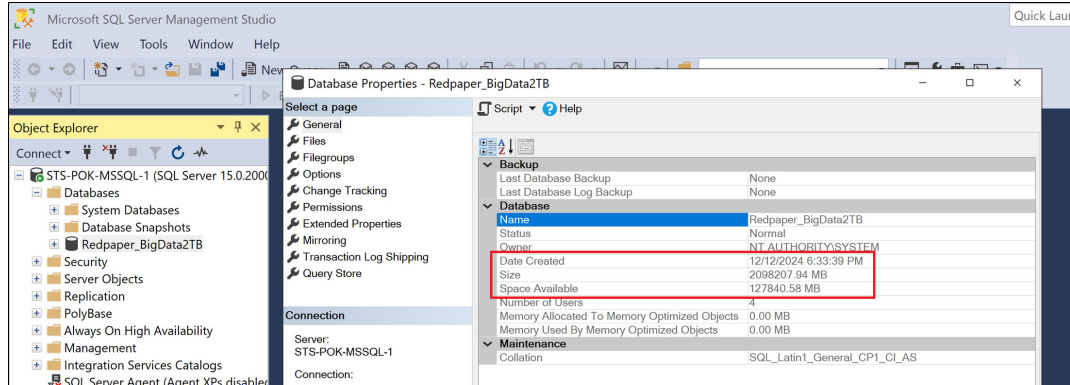


Figure 2-25 Cloned DB details in MS SQL Server Management Studio

It also highlights the mount path which includes the name of the IBM Storage Defender Data Protect Cluster; STS-POK-DP-2'.

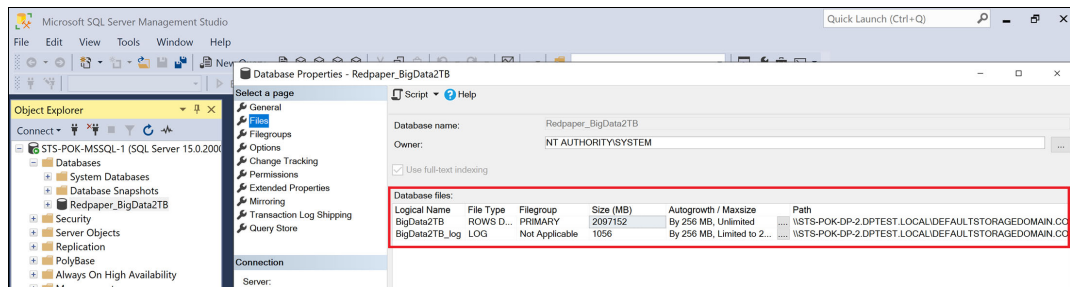


Figure 2-26 Cloned DB mount path

### Removing the Clone

Once testing or other activities have been completed, the clone can simply be removed from the database instance using IBM Storage Defender DMS, by selecting 'Tear Down Clone' from the clone instance information:

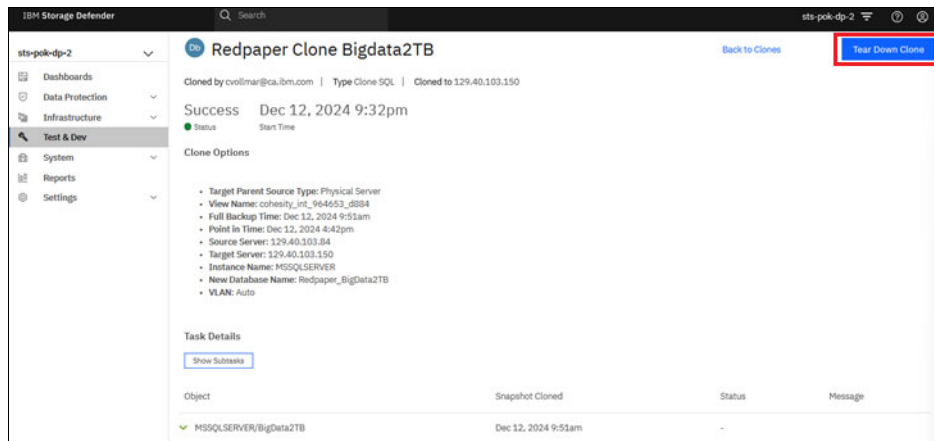


Figure 2-27 Tear Down Clone button

Once the tear down clone button is selected, the confirmation panel will be shown:




Figure 2-28 Activity confirmation dialog

Once the tear down action is confirmed, IBM Storage Defender DMS will confirm the activity is being performed by showing a “Destroying” Status:

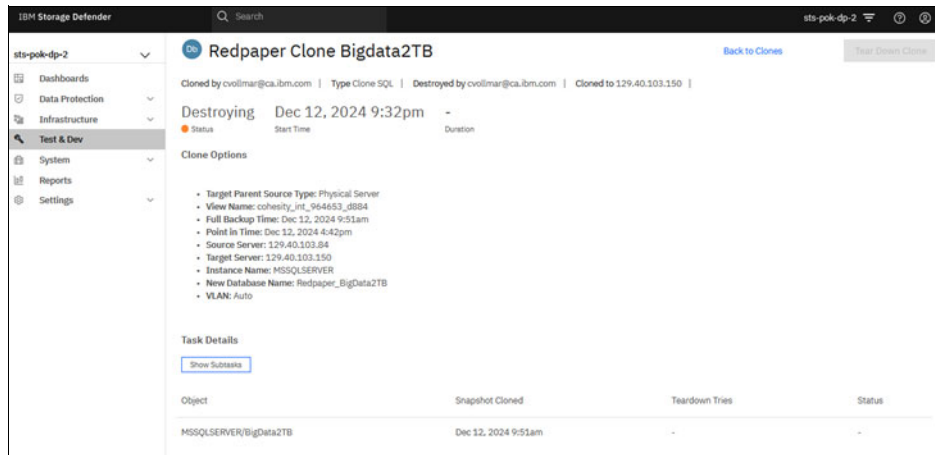


Figure 2-29 Cone tear down in progress

As the tear down task is completed the Defender DMS GUI will update to show a status of “Destroyed” along with details about the task execution time and duration of the task:

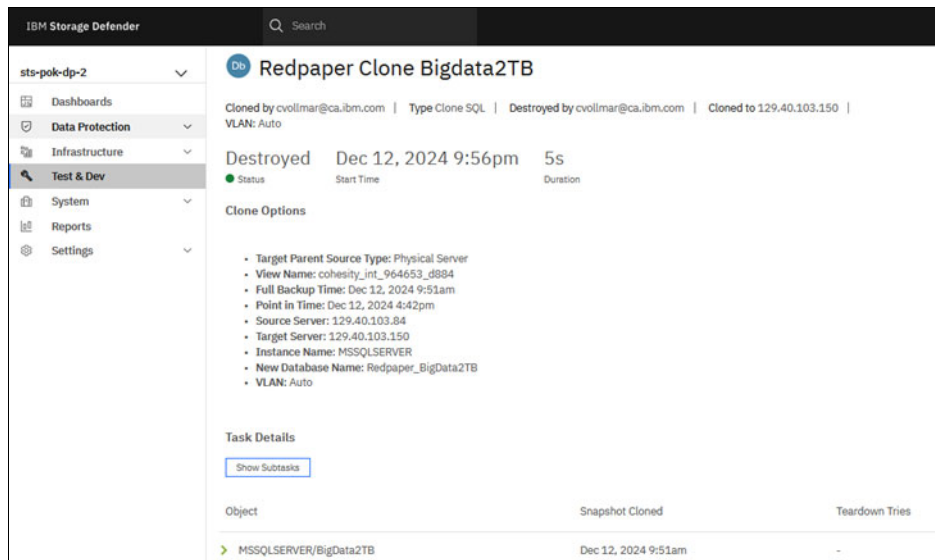


Figure 2-30 Cloned DB tear down completed message

The clone history for DB actions will also be updated and can be reviewed on the Test & Dev page. This will show a history of cloning actions by date as well the status of the cloned DB:

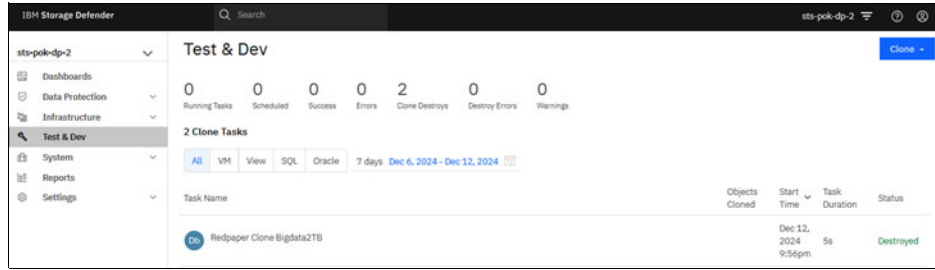


Figure 2-31 Test and Dev activities history and cloned DB status



# Protecting Oracle Databases

In this chapter we discuss the options for protecting an Oracle database with IBM Storage Defender Data Protect. This includes example configurations and the required steps to protect and recover Oracle DBs using both the Oracle Adapter as well as the Remote adapter.

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

- ▶ 3.1, “IBM Data Protect Oracle Server Protection Overview” on page 28
- ▶ 3.2, “Backup using the Oracle Adapter” on page 28
- ▶ 3.3, “Recovery using the Oracle Adapter” on page 35
- ▶ 3.4, “Backup using the Remote Adapter” on page 48
- ▶ 3.5, “Recovery using the Remote Adapter” on page 58

## 3.1 IBM Data Protect Oracle Server Protection Overview

IBM Storage Defender Data Protection allows Oracle databases to be protected by using your choice of either the Oracle Adapter or by using the Remote Adapter.

When selecting the Oracle Adapter, this allows for a simplified backup and recovery process with the use of the Data Protect GUI. This also allows for the use of powerful restore capabilities such as, Instant Recovery that automates the instantiation and recovery of an Oracle database.

The Remote Adapter is available as an alternative to the Oracle Adapter for DBAs who wish to have full control of the backup and recovery of their database environment. By writing or reusing their own RMAN scripts, DBAs can set the Data Protect cluster as the target, which allows Data Protect to not only catalog backups but take advantage of IBM Storage Defender features like immutability and anomaly detection.

### 3.1.1 Oracle version support

Data Protect supports the following versions of Oracle, Oracle Real Application Clusters (RAC) and Oracle Pluggable Databases (PDB):

- ▶ 21c, 19c, 18c, 12cR1, 12cR2 and 11gR2

## 3.2 Backup using the Oracle Adapter

Oracle backups created using the Data Protect Oracle Adapter are immutable, online, incremental forever, block level image copies. By applying RMAN incremental updates to the image copies to the Data Protect cluster over NFS, an immutable snapshot is created following each backup.

For recovery, a restore of the immutable snapshot is presented to RMAN over NFS as the repository for the backup sets. The advantage of this approach results in only one single full backup image taken, eliminating the need for periodic full backups.

Using the Oracle Adapter for backup and recovery requires an Agent to be installed on each database host you intend to backup and each host you intend to restore to.

The following Oracle environments are supported by the Agent:

- ▶ Windows
- ▶ Linux (RPM, Debian, SuSE RPM, PowerPC® RPM, Script installer)
- ▶ AIX (Java agent)
- ▶ Solaris 11
- ▶ HPUX
- ▶ SAP HANA x64 (RPM, Script installer)
- ▶ SAP Oracle (Java agent: RPM, Script installer)
- ▶ SAP HANA PowerPC (Java agent)

To register an Oracle Source host with IBM Storage Defender Data Protect select the following options in the WEB GUI:

- ▶ Data Protection



- ▶ Sources
- ▶ Register
- ▶ Then select Databases
- ▶ Oracle Source

Figure 3-1 on page 29 shows the Register Oracle Source dialog panel to specify the Oracle host address and authentication type.

### Register Oracle Source

Enter Host Address 192.0.2.1 [Browse Registered Source](#)

Authentication Type  OS Authentication  Database Authentication

The Cohesity Agent needs to be pre-installed on the server. [Download Cohesity Agent](#)

[Register](#) [Cancel](#)

Figure 3-1 Register Oracle Source panel

Data Protect detects and supports Oracle Block Change Tracking (BTC) to improve backup performance and reduced backup window size for incremental backups. The use of BTC avoids scanning the datafiles for changed blocks by collecting a record of changed blocks from Oracle via a log file.

**Note:** The Oracle Adapter agent mounts an NFS share from each Data Protect cluster node, then instructs RMAN to allocate channels to each share. For this reason, **by default**, the number of **RMAN channels** are set to be **equal to number of Data Protect nodes**.

Once registered, you can customize the number of RMAN channels from the Protection Group in the source options (Figure 3-2 on page 30).

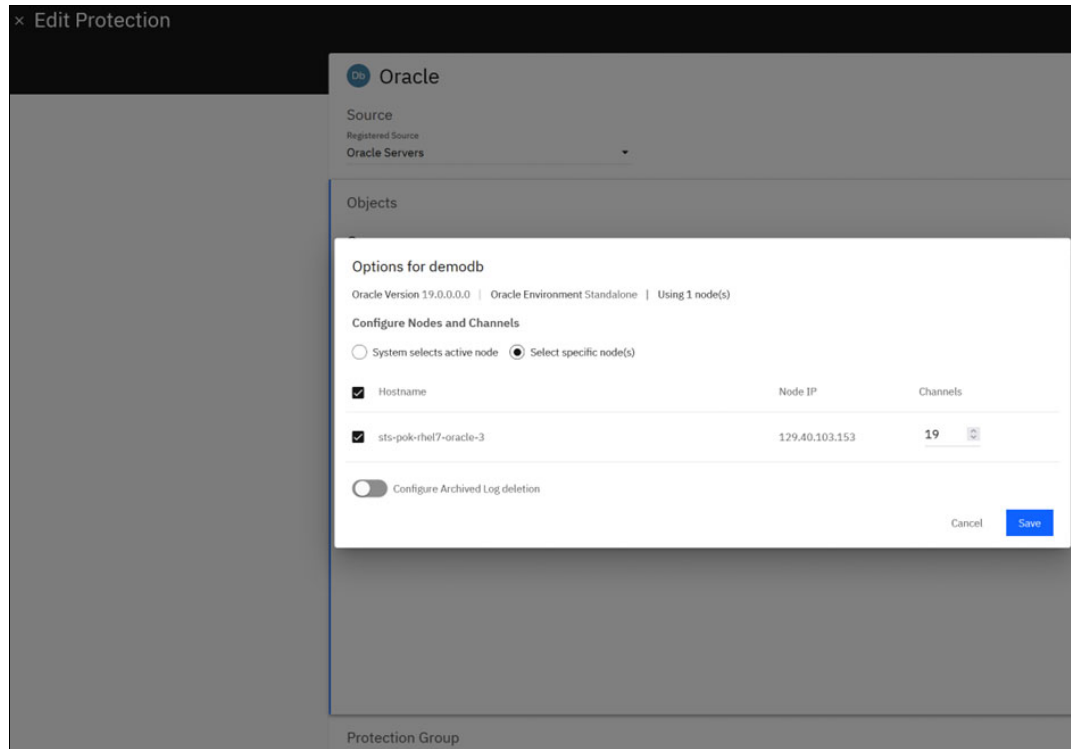


Figure 3-2 Customize number of RMAN channels from Protection Group settings panel

**Note:** The agent sets the **RMAN section size** for datafiles to **200G** to divide large Oracle datafiles for parallel transfer to the Data Protect nodes.

At the start of each backup, the Oracle Adapter runs an RMAN crosscheck and deletes expired backups of the following items:

- ▶ Controlfile
- ▶ SPFile
- ▶ Database

Next, the Oracle Adapter creates an incremental copy of the database files

- ▶ Allocates a channel for each Data Protect cluster node serving the NFS mounts for parallelism
- ▶ Creates incremental datafile copy with section size of 200G to parallelize large individual datafiles
- ▶ Flush current redo log to archived redo logs

Then the Oracle Adapter creates a backupset of the following items:

- ▶ Controlfile
- ▶ SPFile
- ▶ RMAN configuration

The Oracle Adapter updates the level 0 copy of the database files with the incremental updates:

- ▶ recovery copy of database with tag 'cohesity\_nnnnn';

Finally, an immutable snapshot of NFS share is then taken by Data Protect. An example database backup command run by the agent can be found in both agent logs, located under `/var/log/cohesity/oracle_rman_logs/` and from GUI under Protection screen:

*Example 3-1 Backup command issued to Oracle by the agent*

---

```
CONFIGURE CONTROLFILE AUTOBACKUP OFF; CONFIGURE CONTROLFILE AUTOBACKUP ON;
CONFIGURE CONTROLFILE AUTOBACKUP OFF;
run {
allocate channel coh1 device type disk format
"/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_48869_119514670_path0/41379113
56/%U";
allocate channel coh2 device type disk format
"/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_48869_119514670_path1/41379113
56/%U";
allocate channel coh3 device type disk format
"/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_48869_119514670_path2/41379113
56/%U";
allocate channel coh4 device type disk format
"/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_48869_119514670_path3/41379113
56/%U";
allocate channel coh5 device type disk format
"/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_48869_119514670_path4/41379113
56/%U";
allocate channel coh6 device type disk format
"/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_48869_119514670_path5/41379113
56/%U";
allocate channel coh7 device type disk format
"/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_48869_119514670_path6/41379113
56/%U";
allocate channel coh8 device type disk format
"/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_48869_119514670_path7/41379113
56/%U";
backup SECTION SIZE 200G incremental level 1 for recover of copy with tag
'cohesity_48869' database;
backup current controlfile tag 'cohesity_48869';
    sql 'alter system archive log current';
backup spfile tag 'cohesity_48869';
backup current controlfile tag 'cohesity_48869'; RECOVER COPY OF DATABASE WITH TAG
'cohesity_48869';
}
```

---

The separate Oracle Adapter log backup procedure runs on an independent schedule you select to create archived redo log backupsets:

- ▶ SPFile
- ▶ Controlfile
- ▶ backup force tag 'cohesity\_nnnnn' archivelog from time 'DD:MM:YYYY-HH:MM:SS';
- ▶ backup force tag 'cohesity\_nnnnn' archivelog until time 'DD:MM:YYYY-HH\_MM\_SS' not backed up 1 times;
- ▶ Controlfile

Immutable snapshot of NFS share is then taken by Data Protect.

- ▶ Create an Oracle source with Data Protect

1. Create an Oracle source with Data Protect, create a Policy for a daily incremental and periodic archived log backups: <<

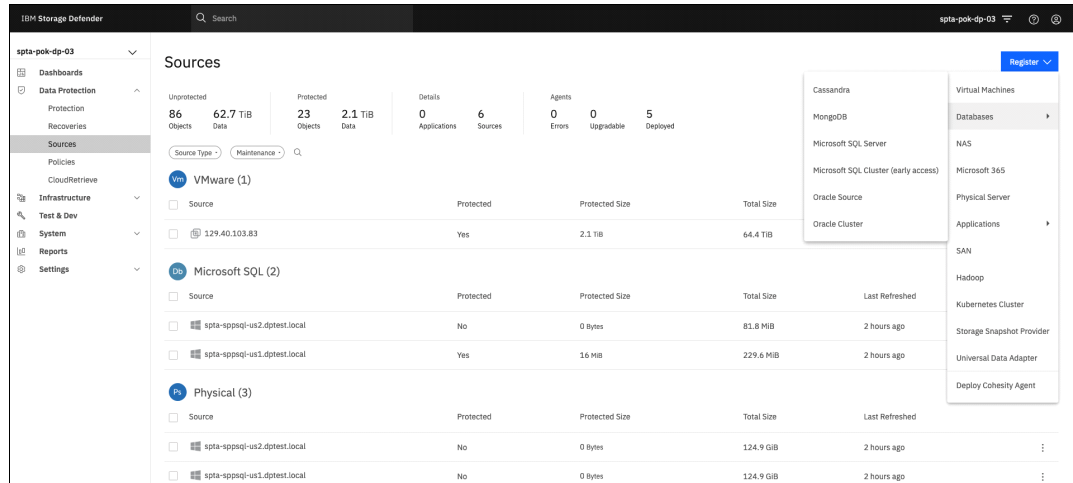


Figure 3-3 Data protect sources panel

**Note:** When configuring this remember, a periodic *full backup* is not needed with the Oracle Adapter

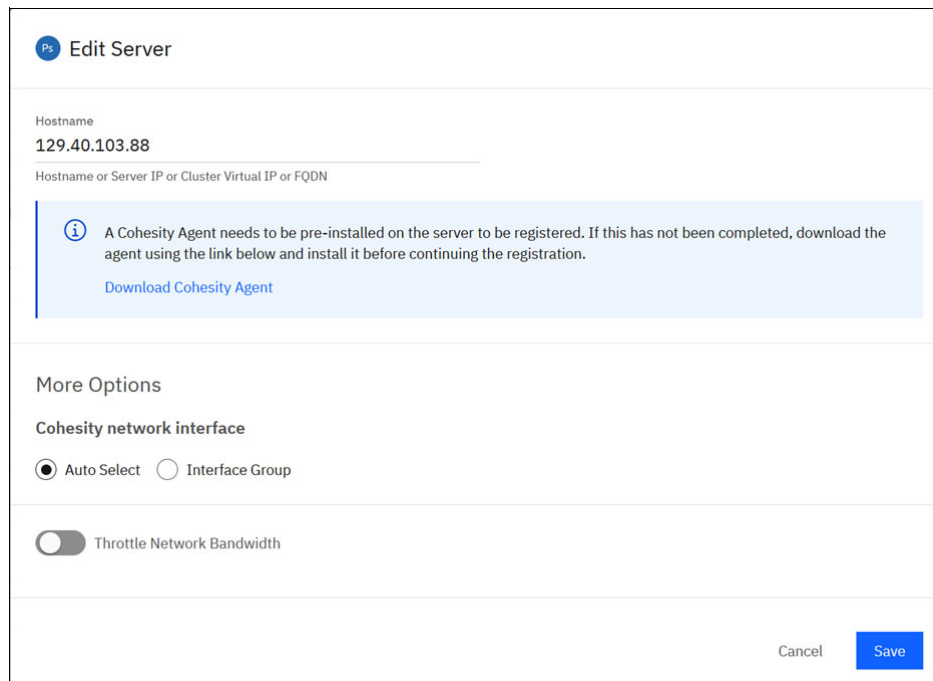


Figure 3-4 Oracle source address panel

2. Next you will need to create a policy to determine the schedule and what type of backup to perform (Full (not required for Oracle Adapter), Incremental, Log) and if additional copies should be replicated and where.

Create a Policy for a daily incremental backup by selecting Data Protection / Policies /

Create Policy. Build a backup policy as shown in Figure 3-4 by providing a Policy name and selecting the desired backup frequency and options:

The screenshot displays the 'Build' tab of a backup policy configuration interface. The 'Policy Name' is 'Daily Local Backup'. The 'Backup' section is configured with 'Backup every 1 Day'. The 'Retry Options' section shows 'Retries' set to 0 and 'Wait (minutes)' set to 1. The 'Log Backup (Databases)' section shows 'Every 1 Hour' and 'Retain for 2 Weeks'. The 'Primary Copy' section shows 'Keep on Local' and 'Retain for 1 Week'. At the bottom, there are buttons for 'Add Replication', 'Add Archive', and 'Add CloudSpin', along with 'Save' and 'Cancel' buttons.

Figure 3-5 Backup Policy configuration for Oracle DB

**Note:** Periodic full backups are not required when performing backups using the Oracle Adapter.

3. Once the policy is created, assign a Protected Group to the new Policy. Protection Groups determine the start time to execute the policy on the selected sources. Create the Protection group by selecting the following:
  - ▶ Data Protection
  - ▶ Protection
  - ▶ Protect
  - ▶ Databases
  - ▶ Oracle Databases

Figure 3-6 on page 34 shows an example of setting up the Source, Policy and start time for the backup:

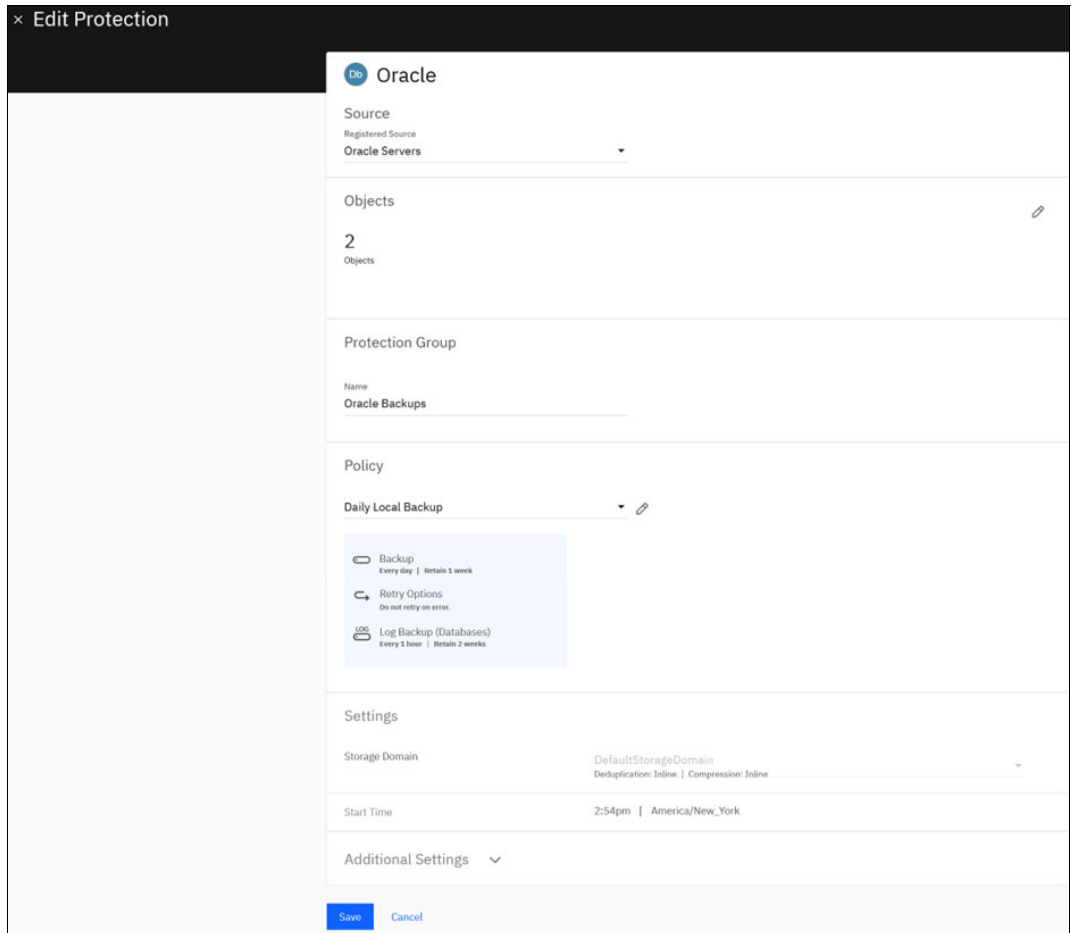


Figure 3-6 Edit protection settings panel

4. Click Save and you can either choose to wait for the scheduled run or select Run now from the Protection screen. Once this is complete the backup will run at the time scheduled.

To view details related to a backup run (Figure 3-7) select the following in the web GUI:

- ▶ Data Protection
- ▶ Protection
- ▶ {Desired Policy Name}

Server Name	Start Time	End Time	Duration	Data Read	Logical Size	Message
129.40.103.153 Size: 38 GiB	Dec 16, 2023 2:54pm	Dec 16, 2023 3:13pm	19m 9s	38 GiB	38 GiB	
demodb Size: 38 GiB	Dec 16, 2023 2:54pm	Dec 16, 2023 3:13pm	18m 54s	38 GiB	38 GiB	
129.40.103.88 Size: 1.1 TiB	Dec 16, 2023 2:54pm	Dec 16, 2023 3:54pm	59m 49s	1.1 TiB	1.1 TiB	
ibmdb Size: 1.1 TiB	Dec 16, 2023 2:54pm	Dec 16, 2023 3:53pm	59m 35s	1.1 TiB	1.1 TiB	

Figure 3-7 Run Details report for Oracle backups

From here you can review the run details of the backup policy including success or failure, run times and size of the backups.

## 3.3 Recovery using the Oracle Adapter

When performing a restore of Oracle data via the adapter there are different recovery options. Choose to recover the database or just recover the archive logs.

### 3.3.1 Recovering a Database via Instant Recovery

To recover the database, select a New Recovery then choose the following:

1. Data Protection
2. Recoveries
3. Recover
4. Databases
5. Oracle

From the Oracle Server panel, select the object you wish to recover and continue to recovery options.

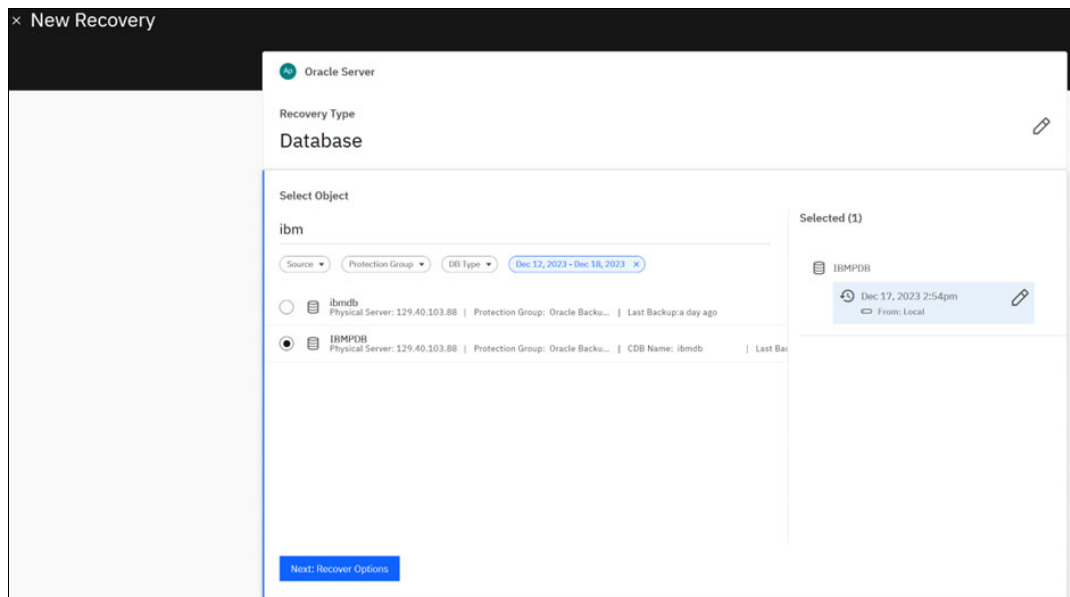


Figure 3-8 Oracle DB recovery via Oracle Adapter panel

Once the object is selected, select the desired recovery point (Figure 3-9 on page 36). Selectable recovery points may be viewed by either list or by timeline view:

Edit recovery point for IBMPDB

Choose a date  
Dec 17, 2023

Timeline

12 AM 6 AM 12 PM 6 PM 12 AM

Time  
02:54:04 PM

Cohesity Incremental

Location:

Cancel Select Recovery Point

Figure 3-9 Recovery database selection panel

Next, choose the location where the data will be restored. An alternative DB or PDB, overwrite the original DB or PDB, are options. It is also possible to perform a rapid recovery which creates an NFS view with DB files, or with Instant Recovery perform a rapid recovery that instantiates the Oracle database in addition to creating an NFS view with DB files. With Instant Recovery, the background migration of datafiles can either be immediate or manually selected later.

**Note:** A rapid recovery with an NFS view will instantly mount a snapshot of the DB files and start the instance with the option to copy the DB files to the host in the background.

Finally, customize the parameters for the recovery DB (Figure 3-10 on page 37):



Figure 3-10 Recovery container database options with Oracle Adapter

The Recovery wizard will automatically generate a PFILE based on the source databases PFILE.

When the target host for the restore has different resource characteristics, there are some important parameters to customize for the target host. These can be found in the generated PFILE. The following settings should be reviewed and adjusted as needed on the target:

- ▶ SGA\_TARGET
- ▶ DB\_RECOVERY\_FILE\_DEST\_SIZE
- ▶ DB\_CREATE\_ONLINE\_LOG\_DEST\_1
- ▶ DB\_RECOVERY\_FILE\_DEST
- ▶ DB\_CREATE\_FILE\_DEST
- ▶ CONTROL\_FILES
- ▶ DB\_WRITER\_PROCESSES
- ▶ MAX\_DUMP\_FILE\_SIZE
- ▶ PGA\_AGGREGATE\_TARGET

Below (Example 3-2 on page 38) is a sample of a customized PFILE generated for an Instant Recovery of a large, 10 TiB Oracle database. This database is from a host with many processors and a large amount of RAM, and being restored to a smaller host with modest resources:

*Example 3-2 Customized PFILE example*


---

```

PROCESSES=1920
LOG_ARCHIVE_FORMAT='%t_%s_%r.arc'
NLS_LANGUAGE='AMERICAN'
SGA_TARGET=6g
ENABLE_PLUGGABLE_DATABASE=true
AUDIT_TRAIL='db'
IBMDB.__INMEMORY_EXT_RWAREA=0
FILESYSTEMIO_OPTIONS='SETALL'
RECYCLEBIN='OFF'
FAL_SERVER=''
REMOTE_LOGIN_PASSWORDFILE='EXCLUSIVE'
NLS_TERRITORY='AMERICA'
DB_UNIQUE_NAME=KEN1
AUDIT_FILE_DEST='/u01/app/oracle/admin/KEN1/adump'
DIAGNOSTIC_DEST='/u01/app/oracle'
DB_RECOVERY_FILE_DEST_SIZE=8000g
DB_CREATE_ONLINE_LOG_DEST_1=/pocdb/orafra
OPEN_CURSORS=500
DB_CREATE_ONLINE_LOG_DEST_2=''
DB_RECOVERY_FILE_DEST='/pocdb/orafra/fast_recovery_area/KEN1'
DB_CREATE_FILE_DEST=/pocdb/oradata
CLUSTER_DATABASE=FALSE
DB_FILES=1024
UNDO_TABLESPACE='UNDOTBS1'
CONTROL_FILES='/u01/app/oracle/oradata/KEN1/control01.ct1'
COMPATIBLE='19.0.0'
FAL_CLIENT=''
DB_WRITER_PROCESSES=10
MAX_DUMP_FILE_SIZE='2G'
LOG_FILE_NAME_CONVERT='/ibmpoc/orafra/IBMDB/online1og','/pocdb/oradata'
DB_BLOCK_SIZE=8192
CLUSTER_INTERCONNECTS=''
PGA_AGGREGATE_TARGET=3g

```

---

Figure 3-11 shows a GUI panel with options for instant recovery of a DB using the Oracle Adapter and the ability to edit the generated PFILE:

Oracle Server

Recovery Type  
Database

Select Object  
ibmdb CDB    Latest Snapshot    Local Location    All (1) Selected PDBs

Targets  
 Alternate CDB     Overwrite Original CDB     Create Cohesity View with DB Files     Instant Recovery

*Instant Recovery creates a clone of the database with datafiles on a Cohesity view. After the clone is created, the datafiles can be migrated to your production storage while your database remains open for transaction.*

Use Case  
 Disaster Recovery

Oracle Hosts and RACs  
 129.40.103.89  
 OS Type: Linux

[Configure Channels](#)

Datafile Migration Method  
 Instant Migration     Manual Migration

*Once recovered successfully, you can manually start the migration from the 'Recovery job' summary page.*

Figure 3-11 Instant Recovery panel with Oracle Adapter

You can select to restore to a different host than the source by selecting the drop-down list of Oracle hosts. When performing a Disaster Recovery, it might be desirable to recover to the source host, but most cases you would want to recover an entire database to a different target host, whether for testing or data reuse purposes in addition to surgical restores.

**Note:** For a host to appear in the drop-down list, it must be registered with the Oracle Adapter.

Further down the form you can add Shell Environment variables to pass to the recovery process. One useful variable is `SKIP_NID_STEP`, which when set to 1 (TRUE), will not run the Oracle new ID utility (NID). The purpose of running the NID utility is to assign a new DBID to the instance (useful if you intend to permanently keep the recovered instance and need RMAN to catalog both this new instance and the source instance it was recovered from simultaneously to have a unique DBID).

For temporary or isolated recovered database instances however, this NID step is unnecessary and can cost a lot of time for large database instance recoveries with a large amount (>1,000) datafiles.

The screenshot shows the 'Oracle Server' interface for 'Recovery Options'. The configuration includes:

- Restore Database Files to:** /pocdb/oradata (Newly created database files will reside in this path)
- Oracle Home:** t/19.0.0/dbhome\_1 (ORACLE\_HOME value where the database is restored)
- Base Directory:** /u01/app/oracle (Directory for the database)
- Target Database Name:** KEN1
- Pfile:** Cohesity Generated Pfile (with edit icon)
- Leave database in Recovery mode:** Off
- Shell Environment:** A section titled 'Use this to configure shell environment variables for your Recover workflow'. It contains a variable `SKIP_NID_STEP` set to `1`. A '+ Add Environment Variable' button is visible below the input field.
- Cluster Interface:** Auto Select
- Task Name:** Instant\_Recover\_Oracle\_Nov\_22\_2024\_9\_55\_AM

Figure 3-12 Shell Environment variables in the Instant Recovery panel with Oracle Adapter

**Note:** For large databases with many datafiles, you can save time on the recovery by skipping the NID utility step that would have reassigned a new DBID that could be unnecessary depending on your intentions for the recovered instance.

Once the instant recovery is complete, the results of the Instant Recovery job for this 10 TiB Oracle database to a new host can be reviewed in the job log. Figure 3-13 below shows an example log file:



Figure 3-13 Instant Recovery job log using Oracle Adapter example

Example 3-3 shows the `df` output listing the NFS mounts for the snapshot DB files. These mounts will automatically be created on the target host as part of the recovery process when running a Instant Recovery job:

*Example 3-3 listing mounts with the `df` command*

```

[oracle@oracle2 ~]$ df -Th
FilesystemType      Size UsedAvail Use% Mounted on
devtmpfsdevtmpfs   32G  32G  0% /dev tmpfstmpfs32G  32G  0% /dev/shm tmpfs
tmpfs 32G27M        32G    1% /run
tmpfstmpfs32G032G0% /sys/fs/cgroup
/dev/mapper/rhel-rootxfs 36G16G20G 45% /
/dev/sda1xfs1014M 183M832M 19% /boot
/dev/mapper/oradata-lv1xfs 16T 39M 16T 1% /pocdb/oradata
/dev/mapper/orafra-lv1 xfs 8.0T 36M 8.0T 1% /pocdb/orafra
tmpfs tmpfs 6.3G 44K 6.3G 1% /run/user/0
tmpfs tmpfs 6.3G 0 6.3G 0% /run/user/1001
tmpfs tmpfs 6.3G 12K 6.3G 1% /run/user/42
129.40.103.129:/DefaultStorageDomain/cohesity_int_437150_31b26/fs nfs 9.3T
4.7T 4.7T51%
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path0
129.40.103.130:/DefaultStorageDomain/cohesity_int_437150_31b26/fs nfs 9.3T 4.7T
4.7T 51% /opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path1
129.40.103.131:/DefaultStorageDomain/cohesity_int_437150_31b26/fs nfs 9.3T
4.7T 4.7T 51%
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path2
129.40.103.132:/DefaultStorageDomain/cohesity_int_437150_31b26/fs nfs 9.3T
4.7T 4.7T 51% /opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path3
129.40.103.133:/DefaultStorageDomain/cohesity_int_437150_31b26/fs nfs 9.3T
4.7T 4.7T 51% /opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path4
129.40.103.134:/DefaultStorageDomain/cohesity_int_437150_31b26/fs nfs 9.3T
4.7T 4.7T 51% /opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path5
129.40.103.135:/DefaultStorageDomain/cohesity_int_437150_31b26/fs nfs 9.3T
4.7T 4.7T 51% /opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path6
129.40.103.136:/DefaultStorageDomain/cohesity_int_437150_31b26/fs nfs 9.3T 4.7T
  
```

```
4.7T 51%
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path7
```

---

Example 3-4 shows the location of the online Redo logs, changetracking file, temp tablespace datafiles and FRA are written to local storage locations specified in PFILE:

*Example 3-4 displaying file location for .log files related to recovery process*

---

```
[oracle@oracle2 ~]$ find /pocdb -type f
/pocdb/orafra/fast_recovery_area/KEN1/KEN1/autobackup/2024_05_11/o1_mf_s_116869853
2_m3zg747o_.bkp
/pocdb/orafra/fast_recovery_area/KEN1/KEN1/autobackup/2024_05_11/o1_mf_s_116869978
2_m3zhg6o5_.bkp
/pocdb/orafra/fast_recovery_area/KEN1/KEN1/autobackup/2024_05_11/o1_mf_s_116870033
8_m3zhm45_.bkp
/pocdb/orafra/KEN1/onlinelog/o1_mf_1_m3zgjpod_.log
/pocdb/orafra/KEN1/onlinelog/o1_mf_2_m3zgovrf_.log
/pocdb/orafra/KEN1/onlinelog/o1_mf_3_m3zgv1z2_.log
/pocdb/orafra/KEN1/onlinelog/o1_mf_4_m3zh0mp5_.log
/pocdb/orafra/KEN1/onlinelog/o1_mf_5_m3zh6hx9_.log
/pocdb/oradata/KEN1/changetracking/o1_mf_m3zg75df_.chg
/pocdb/oradata/KEN1/datafile/o1_mf_temp_m3zhk0mg_.tmp
/pocdb/oradata/KEN1/datafile/o1_mf_temp_m3zhk2hg_.tmp
/pocdb/oradata/KEN1/datafile/o1_mf_temp_m3zhk2n6_.tmp
```

---

Example 3-5 Confirm the ORACLE\_SID that was specified in PFILE is running and open:

*Example 3-5 Oracle commands to confirm instance creation and running status*

---

```
[oracle@oracle2 ~]$ lsnrctl status
```

```
LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 11-MAY-2024 15:11:29
```

```
Copyright (c) 1991, 2019, Oracle. All rights reserved.
```

```
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=oracle2)(PORT=1521)))
STATUS of the LISTENER
```

```
-----
```

```
Alias                LISTENER
Version              TNSLSNR for Linux: Version 19.0.0.0.0 - Production
Start Date           09-FEB-2024 11:57:24
Uptime               92 days 2 hr. 14 min. 5 sec
Trace Level          off
Security              ON: Local OS Authentication
SNMP                 OFF
```

```
Listener Parameter File
```

```
/u01/app/oracle/product/19.0.0/dbhome_1/network/admin/listener.ora
```

```
Listener Log File
```

```
/u01/app/oracle/diag/tnslsnr/oracle2/listener/alert/log.xml
```

```
Listening Endpoints Summary...
```

```
(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=oracle2)(PORT=1521)))
```

```
(DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC1521)))
```

```
(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=oracle2)(PORT=5500))(Security=(my_walle
```

```
t_directory=/u01/app/oracle/admin/KEN1/xdb_wallet))(Presentation=HTTP)(Session=RAW
))
Services Summary...
Service "KEN1" has 1 instance(s).
  Instance "KEN1", status READY, has 1 handler(s) for this service...
Service "ff0b60e27b816bb9e05358672881609d" has 1 instance(s).
  Instance "KEN1", status READY, has 1 handler(s) for this service...
Service "ibmpdb" has 1 instance(s).
  Instance "KEN1", status READY, has 1 handler(s) for this service...
The command completed successfully
```

```
[oracle@oracle2 ~]$ sqlplus system/manager@//localhost:1521/KEN1
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Sat May 11 15:12:47 2024
Version 19.3.0.0.0
```

```
Copyright (c) 1982, 2019, Oracle. All rights reserved.
```

```
Last Successful login time: Fri May 10 2024 05:30:02 -04:00
```

```
Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0
```

```
SQL> select instance_name, status, database_status from v$instance;
```

INSTANCE_NAME	STATUS	DATABASE_STATUS
KEN1	OPEN	ACTIVE

```
SQL> connect /as sysdba
```

```
Connected.
```

```
SQL> show pdbs
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	IBMPDB	READ WRITE	NO

```
SQL> alter session set container=ibmpdb;
```

```
Session altered.
```

```
SQL> select file_name from dba_data_files;
```

```
FILE_NAME
-----
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path4/BKP_6_436557_
data_D-IBMDB_I-2755005093_TS-SYSTEM_FNO-9_ee2q58kh

/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path6/BKP_1_436557_
data_D-IBMDB_I-2755005093_TS-SYSAUX_FNO-10_eb2q58iu

/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path6/BKP_7_436557_
data_D-IBMDB_I-2755005093_TS-UNDOTBS1_FNO-11_e82q5890
```

```
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path1/BKP_2_436557_
data_D-IBMDB_I-2755005093_TS-USERS_FNO-12_ej2q5812
```

```
FILE_NAME
-----
```

```
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path3/BKP_7_436557_
data_D-IBMDB_I-2755005093_TS-IBMPOCTAB01_FNO-333_a12q4mrk
```

```
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path3/BKP_0_436557_
data_D-IBMDB_I-2755005093_TS-IBMPOCTAB02_FNO-334_am2q4mrk
```

```
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path6/BKP_1_436557_
data_D-IBMDB_I-2755005093_TS-IBMPOCTAB03_FNO-335_an2q4mrk
```

```
/opt/cohesity/mount_paths/nfs_oracle_mounts/oracle_437150_23_path0/BKP_2_436557_
.....
```

```
SQL> select sum(bytes)/1024/1024 as MiB from dba_segments where owner='IBMPOC';
```

```
          MIB
-----
9975036.75
```

```
SQL> Disconnected from Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 -
Production
Version 19.3.0.0.0
```

```
[oracle@oracle2 ~]$ sqlplus ibmpoc/ibmpoc@//localhost:1521/ibmpdb
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Sat May 11 15:15:29 2024
Version 19.3.0.0.0
```

```
Copyright (c) 1982, 2019, Oracle. All rights reserved.
```

```
Last Successful login time: Sat May 11 2024 15:15:13 -04:00
```

```
Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.3.0.0.0
```

```
SQL> select count(1) from tab;
```

```
          COUNT(1)
-----
          90
```

```
SQL> select count(1) from ibmpoctest01;
```

```
          COUNT(1)
-----
        13469880
```

```
SQL> Disconnected from Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 -
Production
```



Version 19.3.0.0.0

Once the instant restore is initiated, until “migrate” is selected or if you selected Instant Migration, the datafiles continue to reside on the NFS mounts. Figure 3-14 show an example of the migration options available for mounted Recoveries.

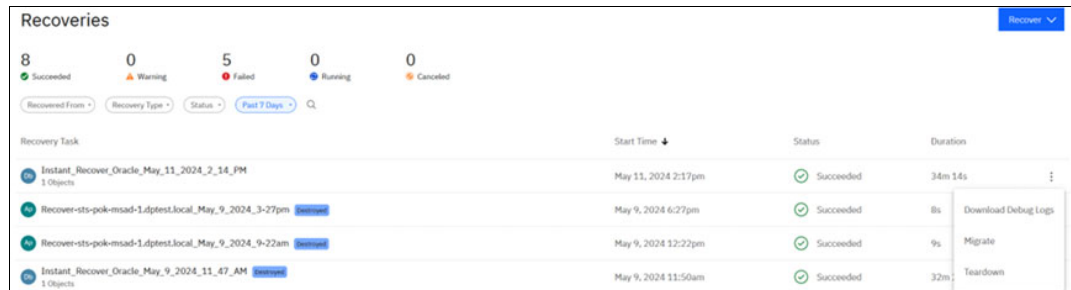


Figure 3-14 Instant Recovery - Migrate option with Oracle Adapter

When finished with the Instant Recovery database, first select the Teardown option on the Recovery, then cleanup the admin/diag and fast\_recovery\_area of your target host

Once Teardown is selected, the recoveries page as show in Figure 3-15 will update to show the NFS paths have been unmounted and the database is destroyed.

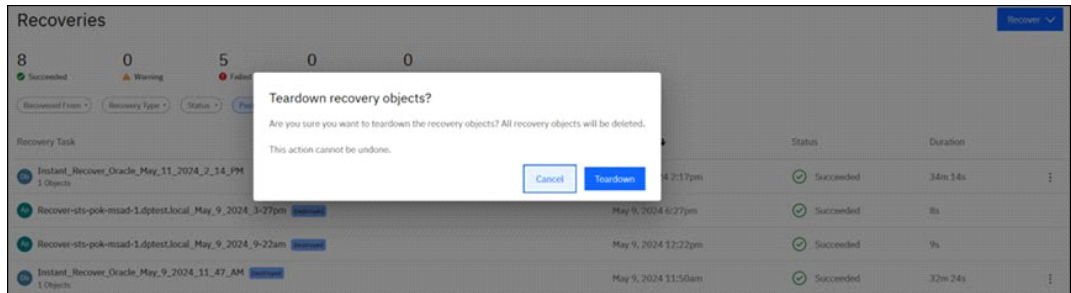


Figure 3-15 Teardown instant recovery objects confirmation dialog

Example 3-6 df output showing removal of temporary mounts completed after teardown

```
[oracle@oracle2 ~]$ df -Th
Filesystem                Type      Size   Used Avail  Use% Mounted on
devtmpfs                  devtmpfs 32G    0    32G   0%   /dev
tmpfs                     tmpfs     32G    0    32G   0%   /dev/shm
tmpfs                     tmpfs     32G   474M  31G    2%   /run
tmpfs                     tmpfs     32G    0    32G   0%   /sys/fs/cgroup
/dev/mapper/rhel-root     xfs       36G   16G   20G   44%   /
/dev/sda1                 xfs      1014M  183M  832M  19%   /boot
/dev/mapper/orafra-lv1    xfs       8.0T   36M   8.0T   1%   /pocdb/orafra
```

Example 3-7 admin/diag and fast\_recovery\_area cleanup

```
[oracle@oracle2 ~]$ rm -fr $ORACLE_BASE/admin/KEN1
[oracle@oracle2 ~]$ rm -fr $ORACLE_BASE/diag/rdbms/ken1
[oracle@oracle2 ~]$ rm -fr /pocdb/orafra/fast_recovery_area
```

Verify the database instance is no longer running and the local datafiles are gone:

*Example 3-8 Confirm local DB files are removed*

---

```
[oracle@oracle2 ~]$ ps -fu oracle
UID      PID  PPID  C  STIME TTY      TIME    CMD
oracle  5490  1     0  Feb09 ?        00:01:37 /u01/app/oracle/product/19.0.0/d
oracle  6063 6059  0   15:28 ?        00:00:00 sshd: oracle@pts/1
oracle  6070 6063  0   15:28 pts/1    00:00:00 -bash
oracle  7450 6070  0   15:45 pts/1    00:00:00 ps -fu oracle

[oracle@oracle2 ~]$ find $ORACLE_BASE -name \*KEN\*

(no results should be displayed)
```

---

Figure 3-16 shows an example of the PDB recovery options in Database recovery panel.

Oracle Server

Recovery Type  
Database

Select Object  
IBMPDB PDB    Latest Snapshot    Local Location

Targets  
 Alternate PDB     Overwrite Original PDB     Create Cohesity View with DB Files  
Oracle Hosts and RACs

Recovery Options

Restore Database Files to  
Newly created database files will reside in this path    /example/path

Oracle Home  
ORACLE\_HOME value where the database is restored    /u01/app/oracle/product/11.2.0.3/db\_1

Base Directory  
Directory for the database    /u01/app/oracle

Target Database Instance Name  
CDB

Use NOFILENAMECHECK    No

Rename PDB    -

Customize Destination Paths with SET NEWNAME    None

Shell Environment    0 environment variables configured.

Cluster Interface    Auto Select

Task Name    Recover\_Oracle\_Dec\_18\_2023\_12\_47\_PM

Recover    Cancel

Figure 3-16 Recovery of pluggable database (PDB) options with Oracle Adapter

When choosing to restore data to an alternate DB or PDB, select the target server from the drop-down menu:

Targets

Alternate PDB     Overwrite Original PDB     Create Cohesity View with DB Files

Search

None

129.40.103.154  
OS Type: Linux

129.40.103.153  
OS Type: Linux

129.40.103.89

/path

Figure 3-17 Target selection for Recovery database location

To restore only the database archive logs, select log sequence to restore by selecting:

- ▶ Data Protection
- ▶ Recoveries
- ▶ Recover
- ▶ Databases
- ▶ Oracle
- ▶ Archive Logs

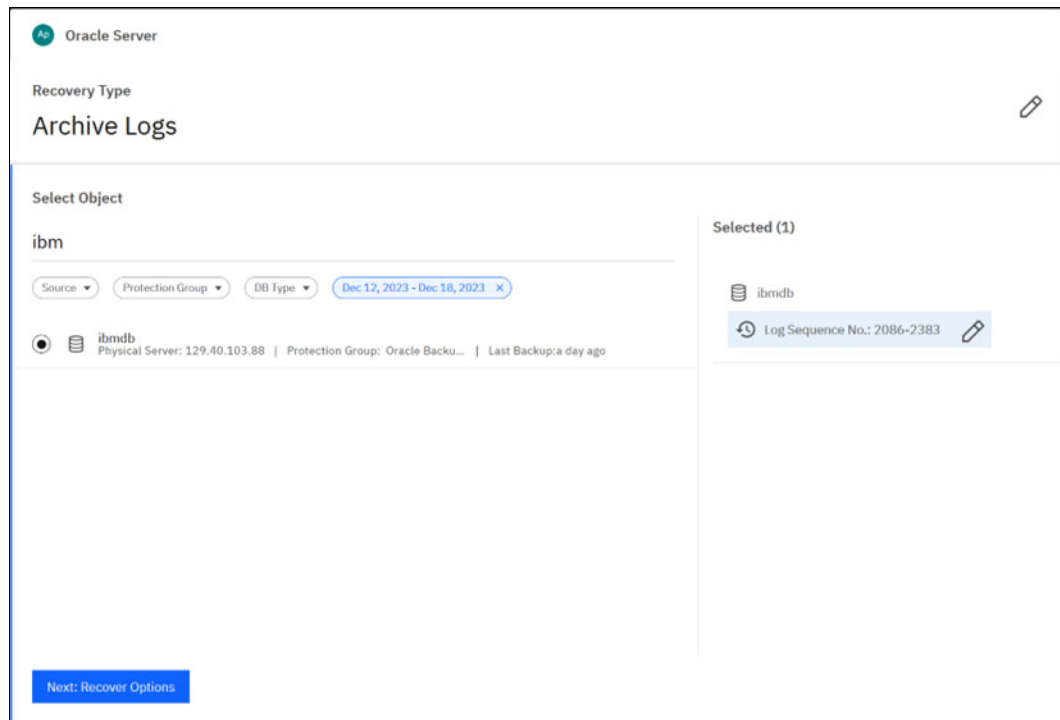


Figure 3-18 Recovery type archive logs selection

As demonstrated above, the Oracle Adapter is a great choice to automate the backup and recovery of Oracle databases without the need to maintain custom RMAN scripts or run manual Oracle commands to restore.

### 3.4 Backup using the Remote Adapter

For DBAs who prefer complete control over their own RMAN scripts that they have already written and maintain for backup and recovery, the Remote Adapter for Oracle is their best option.

When choosing to perform restores via the Remote Adapter, first create a Policy that matches your RMAN script requirements. In this example, we want to retain one week of backups, however When running differential or cumulative incremental backups that require a periodic full backup, be sure to keep your full backup an extra week for restoring the prior week's incrementals as well as keeping an extra week of differential incrementals (if not using cumulative incrementals, because differential rman restores must be applied in sequence from the last full backup that preceded them). Also select archive log backups in your Policy to create the script input field in the Protection Group.

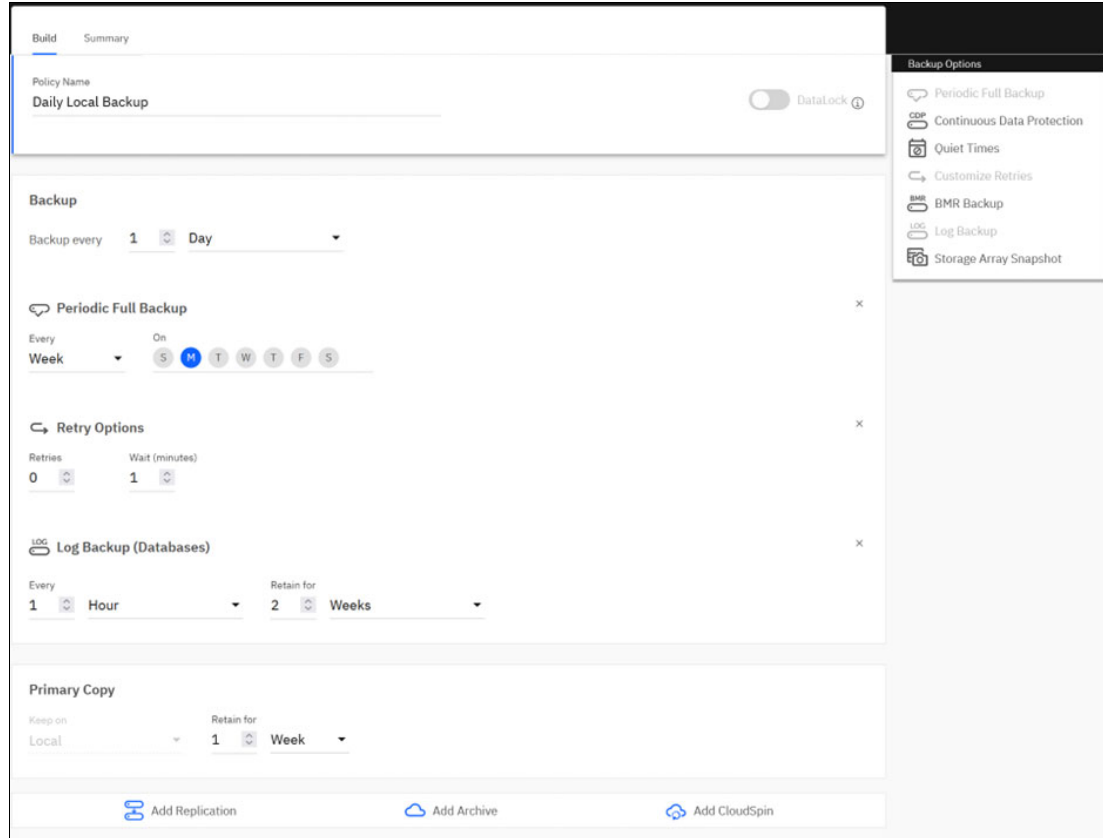
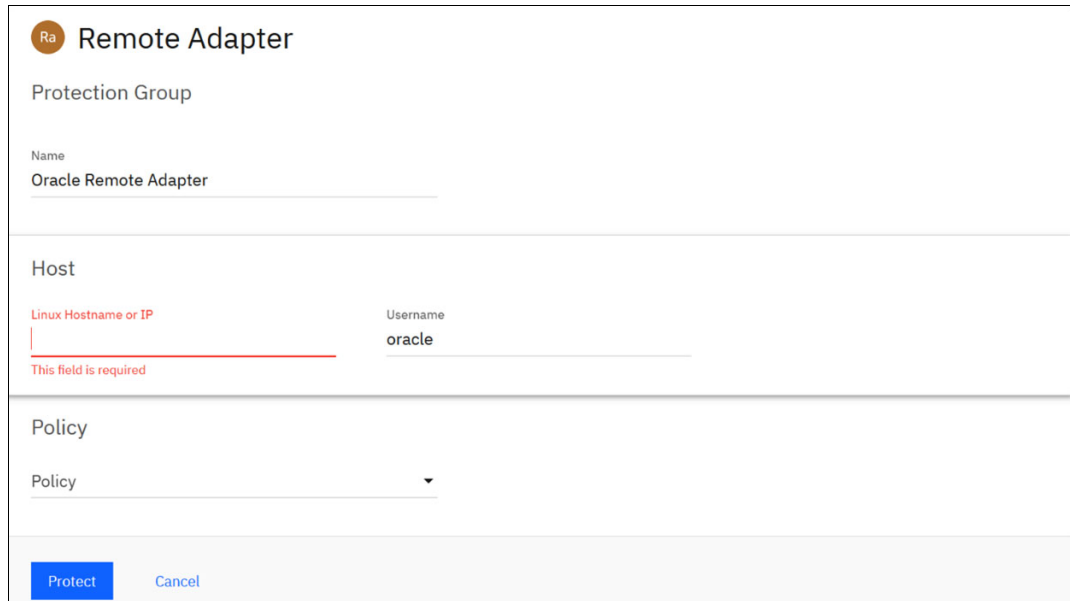


Figure 3-19 Remote Adapter based Protection Policy example

Next, create a Protection group for the Oracle Remote Adapter based backups (Figure 3-20 on page 50) by selecting:

- ▶ Data Protection
- ▶ Protection
- ▶ Protect
- ▶ then select Remote Adapter



The screenshot displays the 'Remote Adapter' configuration window. At the top left, there is a circular icon with 'Ra' and the title 'Remote Adapter'. Below the title, the 'Protection Group' section contains a 'Name' field with the text 'Oracle Remote Adapter'. The 'Host' section has two fields: 'Linux Hostname or IP' (which is empty and has a red error message 'This field is required' below it) and 'Username' (which contains the text 'oracle'). The 'Policy' section has a dropdown menu labeled 'Policy'. At the bottom, there are two buttons: 'Protect' (in blue) and 'Cancel'.

Figure 3-20 Protection group creation example for Remote Adapter backup

Enter the IP Address or Hostname of the Oracle host to generate an SSH Public Key. This key will need to be copied to your host to allow the connection between the Data Protect and the Oracle host.

Ra

## Remote Adapter

Protection Group

**Name**  
Oracle Remote Adapter

---

**Host**

Linux Hostname or IP <b>129.40.103.153</b>	Username <b>oracle</b>
---	---------------------------

Cluster SSH Public Key

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDHYNhF31S58LGyZrUHqGEjt1EjyWbnrHLkuOd/FKvH/4LF3lMepQv3190mm0W6mc07h1T
SF5ChN57IlazivpzrfksoG2ck5tu1nnkYQN6tuCXLcCsiSRl7pcr5Yj202ZpNJYHpkEQelxQLjFXnV7nYhv6e9j38y/iPYyu3321BgpK113S8Csgn1V
zCMMqUk1cY+bi254Pae/ROopzIQUroEqOTMbJ8HGHEOGdLFkF7xEyJuwPYj7C63Kp4Ybp6XR0K95MsXnrXqYr820New6BD5juN1uZZKB
seSiu2DahHbTb14IPuETeipRP8Sw87yK0wYUj/Nt8UNqdP2SFn cohesity@ve-005056b10363-esx
```

i To allow Cohesity Cluster to run the scripts remotely on Linux system, copy SSH Public Key to Clipboard, login to the Linux System with the username specified earlier and set up the permission.

**Policy**

Daily Local Backup oracle ▼ ✎

📁 Backup  
Every day | Retain 2 weeks

☁ Archive to Ceph - No Lock - SPTA-lab  
Every run | Retain 2 weeks

🔄 Periodic Full Backup  
Every week on Monday | Retain 2 weeks |  
DataLock 2 weeks

↺ Retry Options  
Do not retry on error.

📄 Log Backup (Databases)  
Every 1 hour | Retain 2 weeks

Figure 3-21 Protection group settings for Remote Adapter

Once generated, copy the Cluster SSH Public Key to the file `/home/oracle/.ssh/authorized_keys` on the Oracle host and make file readable only by the file Owner:

*Example 3-9 Updating SSH public key file access*

```
[oracle@sts-pok-rhel7-oracle-3 ~]$ chmod 600 .ssh/authorized_keys
[oracle@sts-pok-rhel7-oracle-3 ~]$ ls -l .ssh/authorized_keys
-rw-----. 1 oracle oinstall 410 Aug 31 18:28 .ssh/authorized_keys
```

In the Protection Group settings for the Remote Adapter (Figure 3-22 on page 52), fill in fields for the location and parameters of your incremental and full RMAN backup and your archive log backup scripts that you wrote and placed on your database host as shown in Figure 3-22 on page 52.

Ra

## Remote Adapter

NFS View

oracle ▼

NFS Mount Path

sts-pok-dp-3.www.pbm.ihost.com:/oracle

i In order for this Protection Group to capture the results of the script, the associated View must be mounted on your system and the script must write to a directory on the mounted View.

Script Information

### Script Information for Incremental Schedule

Script (with Full Path)

`/home/oracle/testtables-oracle/utills/oracle_rmanincr.sh`

---

Parameters

`14 3 /demodb/orafra /mnt/oracle-3-1/ora /mnt/oracle-3-2/ora /mnt/oracle-3-3/ora /mnt/oracle-3-4/ora`

### Script Information for Full Schedule

Script (with Full Path)

`/home/oracle/testtables-oracle/utills/oracle_rmanfull.sh`

---

Parameters

`14 3 /demodb/orafra /mnt/oracle-3-1/ora /mnt/oracle-3-2/ora /mnt/oracle-3-3/ora /mnt/oracle-3-4/ora`

### Script Information for Log Schedule

Script (with Full Path)

`/home/oracle/testtables-oracle/utills/oracle_rmanlogs.sh`

---

Parameters

`14 3 /demodb/orafra /mnt/oracle-3-1/ora /mnt/oracle-3-2/ora /mnt/oracle-3-3/ora /mnt/oracle-3-4/ora`

Figure 3-22 Protection group creation example for Remote Adapter backup

It is likely that the View permissions (Figure 3-21 on page 40) will need to be customized to allow an NFS mount to be created on the host. If using an older version of Oracle which requires an older Linux version, you may also need to set the protocol to NFS version 3 rather than 4.1. To edit the View that was created for the Remote Adapter, navigate to your Defender URL and append **/platform/views**: e.g. <https://usea-prod.storage-defender.ibm.com/platform/views>



### Edit View

View Name  
**oracle**

Category  
 File Shares    Backup Target    Object Services ⓘ

Storage Domain  
DefaultStorageDomain

Read/Write Protocol  
NFS v3

Less Options ^

---

Protection      Protection Group: Oracle Remote Adapter | Policy: Daily Local Backup

---

Audit Logs      Off

---

Case Sensitive File or Folder Names      Off (Cannot be edited once the View is created)

---

Performance      TestAndDev High | Pin View to SSD: Off

---

Security

IP Allowlist  
 Override Global IP Allowlist    Extend Global IP Allowlist

**Subnet Allowlist**  
Add the subnets (in IP ranges) that have permission for all Views. Add

🔍

Subnet	NFS Permissions	NFS Squash	
129.40.103.0/24	Read/Write	None	✎ 🗑

**Root Squash** ⓘ  
User ID (UID)      Group ID (GID)

Figure 3-23 Editing the View settings for Oracle backup target settings

In order to use your View, you must Mount NFS View on backup source:

*Example 3-10 creating mount locations for data protect view on host*

---

```
mkdir /mnt/ora-1-1 /mnt/ora-1-2 /mnt/ora-1-3 /mnt/ora-1-4 /mnt/ora-1-5
/mnt/ora-1-6 /mnt/ora-1-7 /mnt/ora-1-8
```

---

Next, you will need to know the Virtual IP Addresses of your Data Protect cluster nodes for the NFS mounts on your host:

- ▶ Settings
- ▶ Networking
- ▶ VIPs

Figure 3-24 on page 54 shows an example of the networking page to collect this information.

The screenshot shows the 'Networking' configuration page in IBM Storage Defender. The left sidebar lists various settings like Dashboards, Data Protection, Infrastructure, Test & Dev, System, Reports, and Settings. The main content area is titled 'Networking' and includes sections for 'Interface Group', 'Inbound DNS (Optional)', and a table of interface configurations. The 'Interface Group' section shows 'intf\_group1' with a Subnet of '129.40.103.0/24' and an FQDN of 'st5-pok-dp-3.www.pbm.ihost.com'. Below this, there are 'Add' and 'Update' buttons. The 'Inbound DNS (Optional)' section has a '+ Add' button and an 'Update' button. At the bottom, a table lists interface configurations:

Interface Group	VIP Address	FQDN	Zones
intf_group1	129.40.103.129	st5-pok-dp-3.www.pbm.ihost.com	-
intf_group1	129.40.103.130	st5-pok-dp-3.www.pbm.ihost.com	-
intf_group1	129.40.103.131	st5-pok-dp-3.www.pbm.ihost.com	-
intf_group1	129.40.103.132	st5-pok-dp-3.www.pbm.ihost.com	-
intf_group1	129.40.103.133	st5-pok-dp-3.www.pbm.ihost.com	-
intf_group1	129.40.103.134	st5-pok-dp-3.www.pbm.ihost.com	-
intf_group1	129.40.103.135	st5-pok-dp-3.www.pbm.ihost.com	-
intf_group1	129.40.103.136	st5-pok-dp-3.www.pbm.ihost.com	-

Figure 3-24 Data Protect cluster node VIPs for NFS mounts

Add the following to `/etc/fstab` to automatically mount on reboot. Must specify NFS option `_netdev` in `fstab` to avoid a panic on boot if the NFS server is not available:

**Example 3-11** *fstab entries*

```
129.40.103.129:/ora /mnt/ora-1-1 nfs defaults,_netdev,noatime 0 0
129.40.103.130:/ora /mnt/ora-1-2 nfs defaults,_netdev,noatime 0 0
129.40.103.131:/ora /mnt/ora-1-3 nfs defaults,_netdev,noatime 0 0
129.40.103.132:/ora /mnt/ora-1-4 nfs defaults,_netdev,noatime 0 0
129.40.103.133:/ora /mnt/ora-1-5 nfs defaults,_netdev,noatime 0 0
129.40.103.134:/ora /mnt/ora-1-6 nfs defaults,_netdev,noatime 0 0
129.40.103.135:/ora /mnt/ora-1-7 nfs defaults,_netdev,noatime 0 0
129.40.103.146:/ora /mnt/ora-1-8 nfs defaults,_netdev,noatime 0 0
```

Mount NFS View using the `'mount -a'` command. Then verify NFS View is mounted correctly:

**Example 3-12** *Using df output to confirm NFS mount*

```
df -Th
Filesystem                Type      Size  Used Avail Use% Mounted on
devtmpfs                  devtmpfs  7.7G   0    7.7G   0% /dev
tmpfs                     tmpfs     7.7G  16K   7.7G   1% /dev/shm
tmpfs                     tmpfs     7.7G  34M   7.7G   1% /run
tmpfs                     tmpfs     7.7G   0    7.7G   0% /sys/fs/cgroup
/dev/mapper/rhel-root     xfs       61G   12G   50G   19% /
/dev/mapper/rhel-home     xfs       30G  332M   30G   2% /home
/dev/sda2                 xfs      1014M 259M   756M  26% /boot
/dev/sda1                 vfat     599M   5.8M  594M   1% /boot/efi
/dev/mapper/orafra-lv1    xfs      8.0T   3.8T   4.8T  48% /ibmpoc/orafra
/dev/mapper/oradata-lv1  xfs       16T    13T   3.5T  79% /ibmpoc/oradata
tmpfs                     tmpfs     1.6G   0    1.6G   0% /run/user/1000
129.40.103.129:/ora      nfs4     15T   1.5T   13T  11% /mnt/ora-1-12
129.40.103.131:/ora      nfs4     15T   1.5T   13T  11% /mnt/ora-1-3
129.40.103.132:/ora      nfs4     15T   1.5T   13T  11% /mnt/ora-1-4
129.40.103.133:/ora      nfs4     15T   1.5T   13T  11% /mnt/ora-1-5
```

129.40.103.134:/ora	nfs4	15T	1.5T	13T	11%	/mnt/ora-1-6
129.40.103.135:/ora	nfs4	15T	1.5T	13T	11%	/mnt/ora-1-7
129.40.103.136:/ora	nfs4	15T	1.5T	13T	11%	/mnt/ora-1-8
tmpfs	tmpfs	1.6G	36K	1.6G	1%	/run/user/0

Create subfolder with appropriate permissions for backups on NFS View so that the oracle user has permission to write backup files to the folder:

*Example 3-13 Creating mount subfolder for NFS*

```
mkdir /mnt/ora-1-1/ora
chown oracle:dba /mnt/ora-1-1/ora
```

Now you are ready to start writing backups to your target View location.

For reference, the following are the RMAN full, incremental and archived redo log backup scripts used in this Remote Adapter example:

*Example 3-14 RMAN full backup script*

```
#!/bin/sh
#
# User must belong to group backupdba
#

ORACLE_SID=DEMO DB
ORACLE_HOME=/u01/app/oracle/product/19.3.0.0.0/dbhome_1
PATH=$PATH:$ORACLE_HOME/bin
export ORACLE_HOME ORACLE_SID

test $# -lt 4 && { echo "usage: $0 [RETENTION_DAYS] [SERVER_LOG_RETENTION_DAYS]
[LOGDEST] [/PATH...]"; exit 1; }
DAYS=$1
LOGDAYS=$2
LOGDEST="'$3/%'"
TARGETPATH=$4

c=1
while [ ! -z "$4" ]; do
    CHANNELS="$CHANNELS allocate channel c$c device type disk
        format '$4/%U';
    "
    shift
    c=$((c + 1))
done

rman target 'system/manager as sysbackup' <<RMAN
CONFIGURE CONTROLFILE AUTOBACKUP ON;
run {
    SET CONTROLFILE AUTOBACKUP FORMAT
        FOR DEVICE TYPE DISK TO "$TARGETPATH/cf%F";
    $CHANNELS
    SQL 'ALTER SYSTEM ARCHIVE LOG CURRENT';
    BACKUP
```

```

SECTION SIZE 500M
INCREMENTAL LEVEL 0
TAG 'weekly_full'
KEEP UNTIL TIME "SYSDATE + $DAYS"
DATABASE;
BACKUP
ARCHIVELOG
FROM TIME "SYSDATE - 1"
TAG 'alog_backup'
KEEP UNTIL TIME "SYSDATE + $DAYS";
DELETE NOPROMPT ARCHIVELOG UNTIL TIME "SYSDATE - $LOGDAYS"
LIKE $LOGDEST;
DELETE NOPROMPT OBSOLETE;
BACKUP CURRENT CONTROLFILE TAG 'ctl_backup';
}
RMAN

```

---

*Example 3-15 RMAN incremental backup script*

---

```

#!/bin/sh
#
# User must belong to group backupdba
#

ORACLE_SID=DEMOB
ORACLE_HOME=/u01/app/oracle/product/19.3.0.0.0/dbhome_1
PATH=$PATH:$ORACLE_HOME/bin
export ORACLE_HOME ORACLE_SID

test $# -lt 4 && { echo "usage: $0 [RETENTION_DAYS] [SERVER_LOG_RETENTION_DAYS]
[LOGDEST] [/PATH...]"; exit 1; }
DAYS=$1
LOGDAYS=$2
LOGDEST="'$3/%'"
TARGETPATH=$4

c=1
while [ ! -z "$4" ]; do
    CHANNELS="$CHANNELS allocate channel c$c device type disk
        format '$4/%U';
    "
    shift
    c=$((c + 1))
done

rman target "system/manager as sysbackup" <<RMAN
CONFIGURE CONTROLFILE AUTOBACKUP ON;
run {
    SET CONTROLFILE AUTOBACKUP FORMAT
        FOR DEVICE TYPE DISK TO "$TARGETPATH/cf%F";
    $CHANNELS
    SQL 'ALTER SYSTEM ARCHIVE LOG CURRENT';
    BACKUP

```

```

        SECTION SIZE 500M
        INCREMENTAL LEVEL 1
        TAG 'daily_incr'
        KEEP UNTIL TIME "SYSDATE + $DAYS"
        DATABASE;
    BACKUP
        ARCHIVELOG
        FROM TIME "SYSDATE - 1"
        TAG 'alog_backup'
        KEEP UNTIL TIME "SYSDATE + $DAYS";
    DELETE NOPROMPT ARCHIVELOG UNTIL TIME "SYSDATE - $LOGDAYS"
        LIKE $LOGDEST;
    DELETE NOPROMPT OBSOLETE;
    BACKUP CURRENT CONTROLFILE TAG 'ctl_backup';
}
RMAN

```

---

*Example 3-16 RMAN log backup script*

---

```

#!/bin/sh
#
# User must belong to group backupdba
#

ORACLE_SID=DEMO
ORACLE_HOME=/u01/app/oracle/product/19.3.0.0.0/dbhome_1
PATH=$PATH:$ORACLE_HOME/bin
export ORACLE_HOME ORACLE_SID

test $# -lt 4 && { echo "usage: $0 [RETENTION_DAYS] [SERVER_LOG_RETENTION_DAYS]
[LOGDEST] [/PATH...]"; exit 1; }
DAYS=$1
LOGDAYS=$2
LOGDEST="'$3/%'"
TARGETPATH=$4

c=1
while [ ! -z "$4" ]; do
    CHANNELS="$CHANNELS allocate channel c$c device type disk
        format '$4/%U';
    "
    shift
    c=$((c + 1))
done

rman target "system/manager as sysbackup" <<RMAN
CONFIGURE CONTROLFILE AUTOBACKUP ON;
run {
    SET CONTROLFILE AUTOBACKUP FORMAT
        FOR DEVICE TYPE DISK TO "$TARGETPATH/cf%f";
    $CHANNELS
    SQL 'ALTER SYSTEM ARCHIVE LOG CURRENT';
    BACKUP
        ARCHIVELOG

```

```

FROM TIME "SYSDATE - 1"
TAG 'alog_backup'
KEEP UNTIL TIME "SYSDATE + $DAYS";
DELETE NOPROMPT ARCHIVELOG UNTIL TIME "SYSDATE - $LOGDAYS"
LIKE $LOGDEST;
DELETE NOPROMPT OBSOLETE;
BACKUP CURRENT CONTROLFILE TAG 'ctl_backup';
}
RMAN

```

## 3.5 Recovery using the Remote Adapter

After protecting an Oracle DB using the remote adapter, the first step to recovering the protected Oracle database and its archive redo logs is to create a Clone View. A Clone View is a copy of an immutable snapshot to ensure backup integrity. This is preferred over mounting the same View the source has mounted in the event it was compromised and keeps the recovery operation in an isolated namespace.

You can also create a Clone View from the Test & Dev screen: Figure 3-25

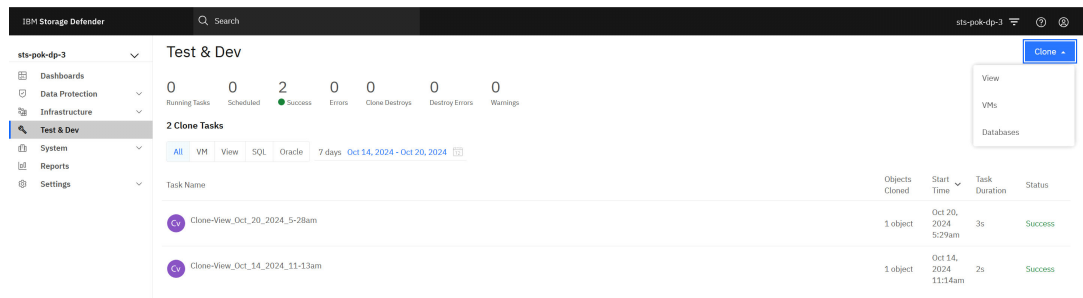


Figure 3-25 Test & Dev screen to create Clone View

To create the clone view from the Recovery menu screen, select the following options in the Data Protect GUI:

1. Select Data Protection
2. Recoveries
3. Cohesity view
4. select Clone View
5. Finally, search for the name of the View you want to clone (Figure 3-26)




Figure 3-26 Clone View search and filter panel

Once the specific View to clone is selected, this will bring up the Clone View panel (Figure 3-27) allowing options to be customized for the Clone View:

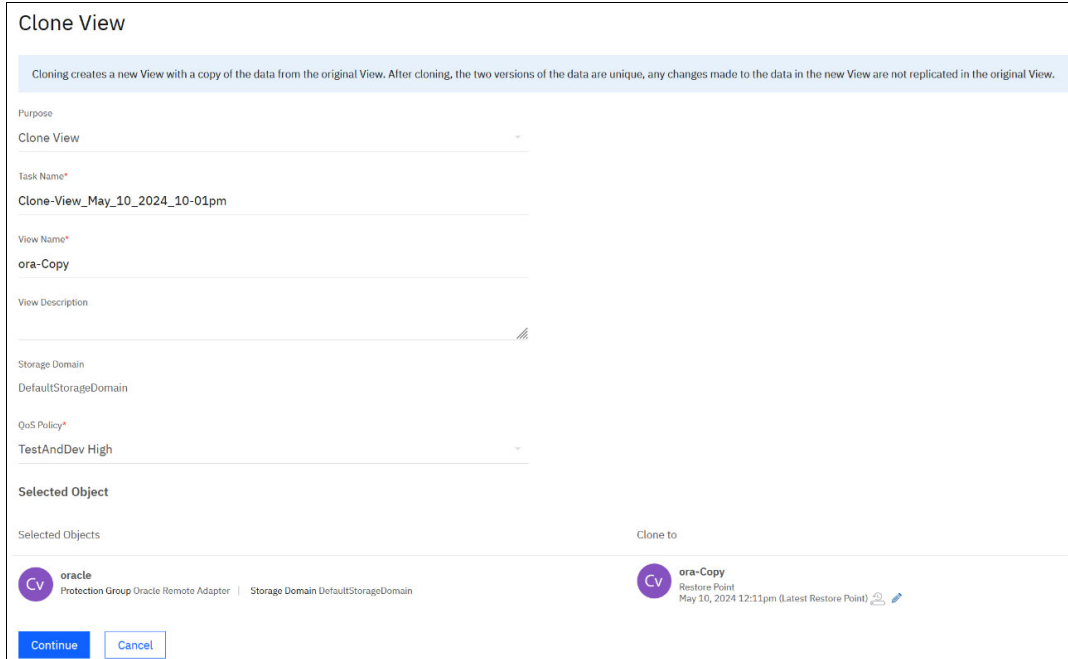


Figure 3-27 Clone view options panel

As shown in Figure 3-28 on page 59, select from the options presented to customize the desired point-in-time to create the Clone View from:

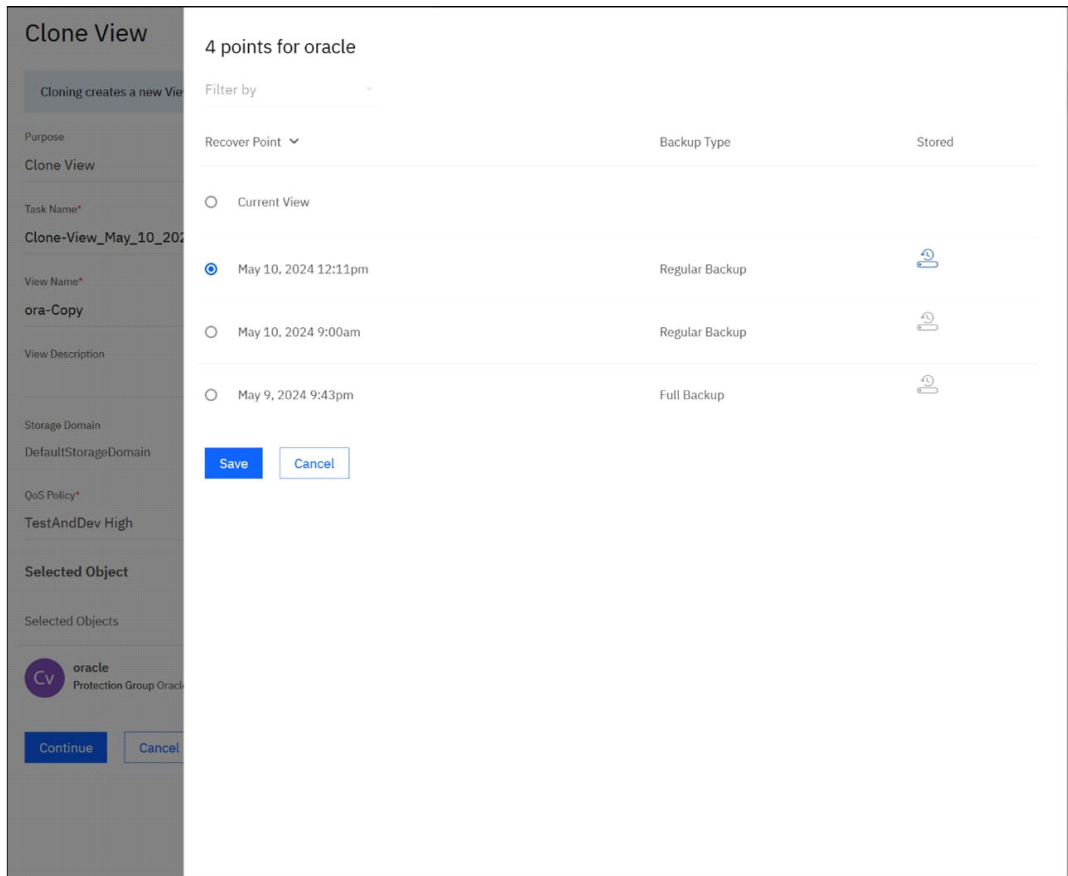


Figure 3-28 Point in time selection for clone view creation

After selecting the desired point in time, click Save and the Clone View is created. Next, manually mount the NFS view on the target host and proceed with the DB recovery using RMAN or any custom scripts you have written.

Example 3-17 shows manual creation of mount points for Clone View on target host:

*Example 3-17 /etc/fstab entries*

---

```
129.40.103.129:/ora-Copy /mnt/oracle-3-1 nfs defaults,_netdev,noatime 0 0
129.40.103.130:/ora-Copy /mnt/oracle-3-2 nfs defaults,_netdev,noatime 0 0
129.40.103.131:/ora-Copy /mnt/oracle-3-3 nfs defaults,_netdev,noatime 0 0
129.40.103.132:/ora-Copy /mnt/oracle-3-4 nfs defaults,_netdev,noatime 0 0
129.40.103.133:/ora-Copy /mnt/oracle-3-5 nfs defaults,_netdev,noatime 0 0
129.40.103.134:/ora-Copy /mnt/oracle-3-6 nfs defaults,_netdev,noatime 0 0
129.40.103.135:/ora-Copy /mnt/oracle-3-7 nfs defaults,_netdev,noatime 0 0
129.40.103.136:/ora-Copy /mnt/oracle-3-8 nfs defaults,_netdev,noatime 0 0
```

```
mount -a
```

---

Confirm NFS mount points of Clone View are attached Example 3-18:

*Example 3-18 df -Th output*

---

```
[oracle@sts-pok-rhel7-oracle-4 ~]$ df -Th
Filesystem                Type      Size  Used Avail Use% Mounted on
devtmpfs                  devtmpfs  7.8G   0  7.8G   0% /dev
tmpfs                     tmpfs     7.8G 560M  7.3G   8% /dev/shm
tmpfs                     tmpfs     7.8G 170M  7.6G   3% /run
tmpfs                     tmpfs     7.8G   0  7.8G   0% /sys/fs/cgroup
/dev/mapper/rhel-root     xfs       36G   23G   13G  64% /
/dev/sda1                 xfs      1014M 183M  832M  19% /boot
/dev/mapper/oradata-lv1  xfs      320G 265G   56G  83% /demodb/oradata
/dev/mapper/orafra-lv1   xfs      320G  41G  280G  13% /demodb/orafra
tmpfs                     tmpfs     1.6G   0  1.6G   0% /run/user/1001
tmpfs                     tmpfs     1.6G  12K  1.6G   1% /run/user/42
129.40.103.129:/ora-Copy nfs        9.3T  4.6T  4.7T  50% /mnt/oracle-3-1
129.40.103.130:/ora-Copy nfs        9.3T  4.6T  4.7T  50% /mnt/oracle-3-2
129.40.103.131:/ora-Copy nfs        9.3T  4.6T  4.7T  50% /mnt/oracle-3-3
129.40.103.132:/ora-Copy nfs        9.3T  4.6T  4.7T  50% /mnt/oracle-3-4
129.40.103.133:/ora-Copy nfs        9.3T  4.6T  4.7T  50% /mnt/oracle-3-5
129.40.103.134:/ora-Copy nfs        9.3T  4.6T  4.7T  50% /mnt/oracle-3-6
129.40.103.135:/ora-Copy nfs        9.3T  4.6T  4.7T  50% /mnt/oracle-3-7
129.40.103.136:/ora-Copy nfs        9.3T  4.6T  4.7T  50% /mnt/oracle-3-8
```

---

### 3.5.1 Restoring a CDB to an alternate host

The following is a working example of restoring a CDB to a different host where we have mounted a Clone View called ora-Copy

**Step 1:** Capture the DBID of the original database we are restoring to the new host:

This can be done by either finding the DBID in the job log, located in the messages for each backup of the source database as shown in Figure 3-29





Figure 3-29 Oracle Remote Adapter job log showing DBID

Or, connect to the database via rman and connect to the source database which was backed up as shown in Example 3-19

*Example 3-19 Gather DBID via rman*

```

rman
connect target /

connected to target database: DEMODB (DBID=4137911356)
quit

```

**Note:** Depending on the RMAN settings, a control file backup file name may contain the DBID as well (e.g. `cfbc-4137911356-20240513-09`).

**Step 2:** Create a database parameter file (PFile) on the target host with your desired instance name, set as both `DB_UNIQUE_NAME` in your PFile and your `ORACLE_SID` environment variable, and you must set `DB_NAME` and `DBID` from the database you want to restore:

Here we will perform the following steps:

1. Set variables `DBNAME`, `DATADIR`, `LOGDIR`, `ORACLE_SID`
2. Variables `PFILE`, `SGA`, `PGA` and `THREADS` are calculated for you in the script below
3. Generate the custom PFILE contents based on variables above
4. Note: match `DB_NAME` to source database name

This serves as an example that is customized for our target host set for a restore of `DBNAME=DEMO` with `DBID 4137911356` where we have chosen an `ORACLE_SID` instance name of `KEN2`.

Your customizations for your target host may vary depending on the exact environment you are attempting to restore to. You could also simply save a copy of your PFile as part of your RMAN backup script or copy directly from the source host to edit by hand.

*Example 3-20 Customizing the KEN2 Oracle DEMODB restore parameter file config for the target host*

```

DBNAME=DEMO
DATADIR=/demodb/oradata
LOGDIR=/demodb/orafra
ORACLE_SID=KEN2
export ORACLE_SID

PFILE=$ORACLE_HOME/dbs/init$ORACLE_SID.ora
SGA=$(free | head -2 | tail -1 | awk '{

```

```

    printf("%dG", $2/1024/1024/2/1.5)
}')
PGA=$(free | head -2 | tail -1 | awk '{
    printf("%dG", $2/1024/1024/2/1.5/2)
}')
THREADS=$(grep -c processor /proc/cpuinfo)

echo "db_name='$DBNAME'
memory_target=0
processes = 1000
parallel_max_servers=$((THREADS * 20))
db_block_size=8192
db_domain=''
db_recovery_file_dest='$LOGDIR/fast_recovery_area'
db_recovery_file_dest_size=200G
diagnostic_dest='$ORACLE_BASE'
dispatchers='(PROTOCOL=TCP) (SERVICE=${ORACLE_SID}XDB)'
open_cursors=500
remote_login_passwordfile='EXCLUSIVE'
undo_tablespace='UNDOTBS1'
# You may want to ensure that control files are created on separate physical
# devices
control_files = ($DATADIR/$ORACLE_SID/controlfile/${ORACLE_SID}_control1,
$LOGDIR/$ORACLE_SID/controlfile/${ORACLE_SID}_control2)
compatible = '19.0.0'
db_create_file_dest='$DATADIR'
db_create_online_log_dest_1='$LOGDIR'
enable_pluggable_database=TRUE
db_unique_name=$ORACLE_SID
filesystemio_options=setall
db_writer_processes=$THREADS
db_files=1024
max_dump_file_size=2G
recyclebin=off
sga_target=$SGA
pga_aggregate_target=$PGA" >$PFILE

```

---

This (Example 3-21) is the result of this example `initKEN2.ora` pfile we generated above in Example 3-20:

*Example 3-21 example contents for pfile*

---

```

db_name='DEMODB'
memory_target=0
processes = 1000
parallel_max_servers=320
db_block_size=8192
db_domain=''
db_recovery_file_dest='/demodb/orafra/fast_recovery_area'
db_recovery_file_dest_size=200G
diagnostic_dest='/u01/app/oracle'
dispatchers='(PROTOCOL=TCP) (SERVICE=KEN2XDB)'
open_cursors=500
remote_login_passwordfile='EXCLUSIVE'
undo_tablespace='UNDOTBS1'
# You may want to ensure that control files are created on separate physical

```

```
# devices
control_files=(/demodb/oradata/KEN2/controlfile/KEN2_control1,
demodb/orafra/KEN2/controlfile/KEN2_control2)
compatible='19.0.0'
db_create_file_dest='/demodb/oradata'
db_create_online_log_dest_1='/demodb/orafra'
enable_pluggable_database=TRUE
db_unique_name=KEN2
filesystemio_options=setall
db_writer_processes=16
db_files=1024
max_dump_file_size=2G
recyclebin=off
sga_target=5G
pga_aggregate_target=2G
```

---

Create the required adump, BCT (for both the new SID and temporarily for the source DBNAME) and FRA directories for your new instance in advance of the restore attempt to avoid RMAN failing to open the new database instance:

*Example 3-22 Creating DB directories for recovery instance*

---

```
mkdir -p $ORACLE_BASE/admin/$ORACLE_SID/adump \
$DATADIR/$DBNAME/changetracking \
$DATADIR/$ORACLE_SID/changetracking \
$LOGDIR/fast_recovery_area
```

---

**Step 3:** Restore the protected DB data to the alternate Demo DB

1. Connect to the newly configured SID
2. Set DBID to match source from Step 1
3. Startup nomount
4. Create an spfile from your customized pfile
5. Restore the controlfile from the known View Clone location based on RMAN backup settings:
  - a. Mount the database
  - b. Run restore of the database
  - c. Recover the DB archived redo logs
  - d. Open
  - e. database and reset redo logs
6. Toggle block change tracking to relocate datafile to new location, remove old directory

*Example 3-23 DB Recovery process example*

---

```
export ORACLE_SID=KEN2

rman <<RMAN connect target /
set DBID 4137911356;
startup nomount;
create spfile from pfile;
set controlfile autobackup format for device type disk to
'/mnt/oracle-3-1/ora/cf%F';
```

```

restore controlfile from autobackup;
alter database mount;
list incarnation; run {
    allocate channel c1 device type disk format '/mnt/oracle-3-1/ora/%U';
    allocate channel c2 device type disk format '/mnt/oracle-3-2/ora/%U';
    allocate channel c3 device type disk format '/mnt/oracle-3-3/ora/%U';
    allocate channel c4 device type disk format '/mnt/oracle-3-4/ora/%U';
    allocate channel c5 device type disk format '/mnt/oracle-3-5/ora/%U';
    allocate channel c6 device type disk format '/mnt/oracle-3-6/ora/%U';
    allocate channel c7 device type disk format '/mnt/oracle-3-7/ora/%U';
    allocate channel c8 device type disk format '/mnt/oracle-3-8/ora/%U';
    restore database;
    recover database until sequence 12297;
}
alter database open resetlogs;
alter database disable block change tracking;
alter database enable block change tracking;
quit
RMAN
rmdir $DATADIR/$DBNAME/changetracking

```

---

The following Example 3-24 shows the restored DEMODB database on our target host with instance KEN2: non-OMF datafile names need to be renamed manually.

*Example 3-24 Showing KEN2 details from target host*

---

```

[oracle@sts-pok-rhel7-oracle-4 ~]$ find /demodb -type f
/demodb/oradata/KEN2/changetracking/o1_mf_lbm4wx5g_.chg
/demodb/oradata/KEN2/controlfile/KEN2_control1
/demodb/oradata/KEN2/datafile/o1_mf_undotbs1_m3zgwkkx_.dbf
/demodb/oradata/KEN2/datafile/o1_mf_users_m3zhkk9j_.dbf
/demodb/oradata/KEN2/datafile/o1_mf_system_m3zhkljn_.dbf
/demodb/oradata/KEN2/datafile/o1_mf_sysaux_m3zhkpx0_.dbf
/demodb/oradata/KEN2/datafile/o1_mf_temp_m3zhnbdd_.tmp
/demodb/oradata/KEN2/datafile/temp012023-07-08_11-18-32-277-AM.dbf
/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_sysaux_m3zg
w10f_.dbf
/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_undotbs1_m3
zhkk14_.dbf
/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_undotbs1_m3
zgp1m4_.dbf
/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_undotbs1_m3
zsdbf8s_.dbf
/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_undotbs1_m3
zis9skm_.dbf
/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_system_m3zh
kp1s_.dbf
/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_sysaux_m3zg
w1gw_.dbf
/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_undotbs1_m3
zhk126_.dbf
/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_users_m3zhk
kdf_.dbf
/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_system_m3zh
klny_.dbf

```

```

/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_temp_m3zhnf
wb_.tmp
/demodb/oradata/DEMODB/datafile/ibmpoc01.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc02.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc03.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc04.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc05.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc06.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc07.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc08.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc09.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc10.dbf
/demodb/orafra/KEN2/controlfile/KEN2_control2
/demodb/orafra/KEN2/onlinelog/o1_mf_1_m3zh1c6c_.log
/demodb/orafra/KEN2/onlinelog/o1_mf_2_m3zh1c76_.log
/demodb/orafra/KEN2/onlinelog/o1_mf_3_m3zh1c85_.log
/demodb/orafra/KEN2/onlinelog/o1_mf_4_m3zh1c91_.log
/demodb/orafra/KEN2/onlinelog/o1_mf_5_m3zh1c9w_.log
[oracle@sts-pok-rhel7-oracle-4 ~]$ export ORACLE_SID=KEN2
[oracle@sts-pok-rhel7-oracle-4 ~]$ sqlplus /nolog
SQL*Plus: Release 19.0.0.0.0 - Production on Sat May 11 15:00:01 2024
Version 19.3.0.0.0
Copyright (c) 1982, 2019, Oracle. All rights reserved.

```

```

SQL> connect /as sysdbaConnected.
SQL> show pdbs

```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	DEMOPDB	MOUNTED	

```

SQL> alter session set container=DEMOPDB;

```

Session altered.

```

SQL> startup;
Pluggable Database opened.
SQL> alter pluggable database demopdb save state;

```

Pluggable database altered.

```

SQL> show con_name

```

```

CON_NAME
-----
DEMOPDB

```

```

SQL> select instance_name, status, database_status from v$instance; INSTANCE_NAME
STATUSDATABASE_STATUS
KEN2OPENACTIVE

```

```

SQL> select sum(bytes)/1024/1024 as MiB from dba_segments where owner='IBMPOC';
MIB
195176.063

```

```
SQL> select file_name from dba_data_files;
FILE_NAME
/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_system_m3
zhklny_.dbf
/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_sysaux_m3
zgwlgw_.dbf
/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_undotbs1_
m3zhk126_.dbf

/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_undotbs1_
m3 zgplm4_.dbf

/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_undotbs1_
m3 zsdbf8s_.dbf

/demodb/oradata/KEN2/FFFC432AA58638E4E055025056B152C9/datafile/o1_mf_undotbs1_
m3 zis9skm_.dbf
/demodb/oradata/KEN2/FFFC6A9752B245FBE055025056B152C9/datafile/o1_mf_users_m3z
hkkdf_.dbf
```

```
FILE_NAME
```

```
/demodb/oradata/DEMODB/datafile/ibmpoc01.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc02.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc03.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc04.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc05.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc06.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc07.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc08.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc09.dbf
/demodb/oradata/DEMODB/datafile/ibmpoc10.dbf
14 rows selected.
```

```
SQL>
[oracle@sts-pok-rhel7-oracle-4 ~]$ lsnrctl status
```

```
LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 21-OCT-2024 10:25:41
```

```
Copyright (c) 1991, 2019, Oracle. All rights reserved.
```

```
Connecting to
(DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=sts-pok-rhel7-oracle-4)(PORT=1521)))
STATUS of the LISTENER
```

```
-----
Alias                LISTENER
Version              TNSLSNR for Linux: Version 19.0.0.0.0 - Production
Start Date           16-OCT-2024 08:58:41
Uptime                5 days 1 hr. 27 min. 0 sec
Trace Level          off
Security              ON: Local OS Authentication
SNMP                  OFF
Listener Parameter File
/u01/app/oracle/product/19.0.0/dbhome_1/network/admin/listener.ora
```

```

Listener Log File
/u01/app/oracle/diag/tnslsnr/sts-pok-rhel7-oracle-4/listener/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=sts-pok-rhel7-oracle-4)(PORT=1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC1521)))

  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=sts-pok-rhel7-oracle-4)(PORT=5500))(Security=(my_wallet_directory=/u01/app/oracle/admin/KEN2/xdw_wallet))(Presentation=HTTP)(Session=RAW))
Services Summary...
Service "KEN2" has 1 instance(s).
  Instance "KEN2", status READY, has 1 handler(s) for this service...
Service "KEN2XDB" has 1 instance(s).
  Instance "KEN2", status READY, has 1 handler(s) for this service...
Service "demopdb" has 1 instance(s).
  Instance "KEN2", status READY, has 1 handler(s) for this service...
Service "fffc6a9752b245fbc055025056b152c9" has 1 instance(s).
  Instance "KEN2", status READY, has 1 handler(s) for this service...

The command completed successfully

```

```

[oracle@sts-pok-rhel7-oracle-4 ~]$ sqlplus ibmpoc/ibmpoc@//localhost:1521/demopdb
SQL*Plus: Release 19.0.0.0.0 - Production on Sat May 11 15:05:35 2024
Version 19.3.0.0.0

```

```

Copyright (c) 1982, 2019, Oracle. All rights reserved.
Last Successful login time: Thu May 09 2024 16:56:58 -04:00 Connected to:
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production Version
19.3.0.0.0

```

```

SQL> select count(1) from tab; COUNT(1)
11

```

```

SQL> select count(1) from ibmpoctest01; COUNT(1)
2295270

```

```

SQL> Disconnected from Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 -
Production
Version
19.3.0.0.0

```

---

The above example illustrates the steps for a typical standalone Oracle database restore and recovery from backup to a new host. The Remote Adapter gives complete flexibility how you choose to backup and restore your particular environment and is a good choice for experienced DBAs who require this level of control.

In contrast a backup and restore with the Oracle Adapter, as we saw in the previous section 3.3 where no detailed knowledge of RMAN was required and Oracle commands to create and recover the database is automated, is a suitable choice for most environments compared to the Remote Adapter.







# 4

## Protecting Microsoft Active Directory

In this chapter we discuss the options for protecting Microsoft Active Directory with IBM Storage Defender Data Protect. This includes example configurations and the required steps to protect and recover Active Directory as well as individual Active Directory elements.

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

- ▶ 4.1, “Protecting Microsoft Active Directory” on page 70
- ▶ 4.2, “Protecting the Microsoft Active Directory DB” on page 70
- ▶ 4.3, “Protect and Recover Microsoft Active Directory” on page 72

## 4.1 Protecting Microsoft Active Directory

In today's large organizations and enterprises, Active Directory is an increasingly critical component of a growing data infrastructure. Given its central and growing role in managing and protecting an organization's information and resources, it is critical to protect Active Directory from any number of today's data threats:

1. Accidental data deletion
2. Insider attack
3. Security breaches
4. Administrative error
5. Ransomware and malware attacks

The IBM Data Protect Cluster solution for Active Directory (AD) includes many features that make your backups much more valuable, including:

### **Granular Object Restore:**

Once your Active Directory is protected, it gives you the flexibility to restore everything from an entire snapshot to a whole Microsoft Organizational Unit (OU).

In addition, IBM Data Protect Cluster granular object restore uses a comparison of your live Active Directory with a mounted backup, which allows you to identify the differences between live data and protected data quickly. You can easily spot the difference and then restore just the objects and attributes you need.

### **Flexibility:**

IBM Data Protect Cluster gives you the ability to browse and search across all your snapshots, and the desired data to different locations on different servers.

## 4.2 Protecting the Microsoft Active Directory DB

To protect your Active Directory database, the Windows Agent must first be installed on your Active Directory (AD) server. The Windows agent is designed to work specifically with the Windows operating system and is compatible with Windows versions 2012R2 and above. If there are multiple Active Directory hosts, you will need to install the agent on each host that needs to be protected.

### ***To protect a Microsoft Active Directory DB perform the following steps:***

1. Install the Window Agent and Register Microsoft Active Directory by navigating to:
  - ▶ Sources
  - ▶ Register
  - ▶ Applications
  - ▶ Active Directory

The agent is lightweight and has a small memory footprint. This agent carries out the tasks that are defined in the IBM Data Protect Cluster Protection Group. The agent ties together technologies and capabilities already in Windows, like Windows VSS with new technologies, like Changed Block Tracker (CBT), which allows the system to tackle data management more efficiently (Figure 4-1 on page 71).

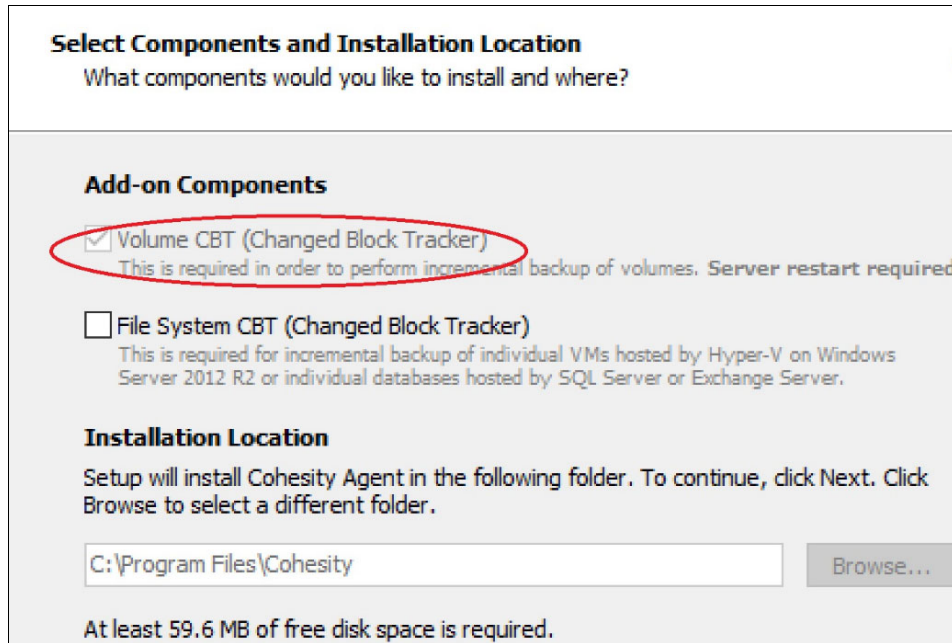


Figure 4-1 Installing the Windows Agent

The Volume CBT (Changed Block Tracker) component is required to perform incremental backups and requires a reboot. Until you reboot, you can only perform volume-based *full* backups.

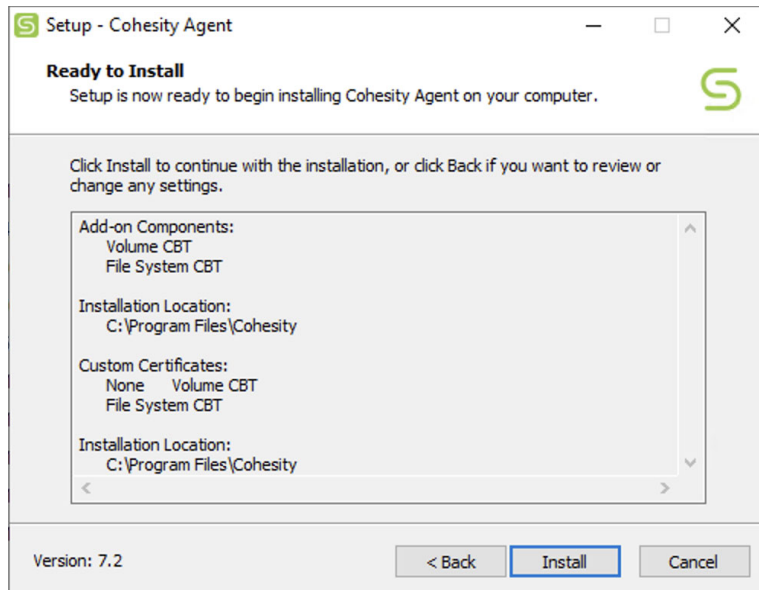


Figure 4-2 Windows Agent installation changes confirmation dialog

2. Register Active Directory as a Data source with IBM Defender Data Protect.

From the Data Protection web GUI Select the following from the left hand side menu

- ▶ Sources
- ▶ Register

- ▶ Applications
- ▶ Active Directory

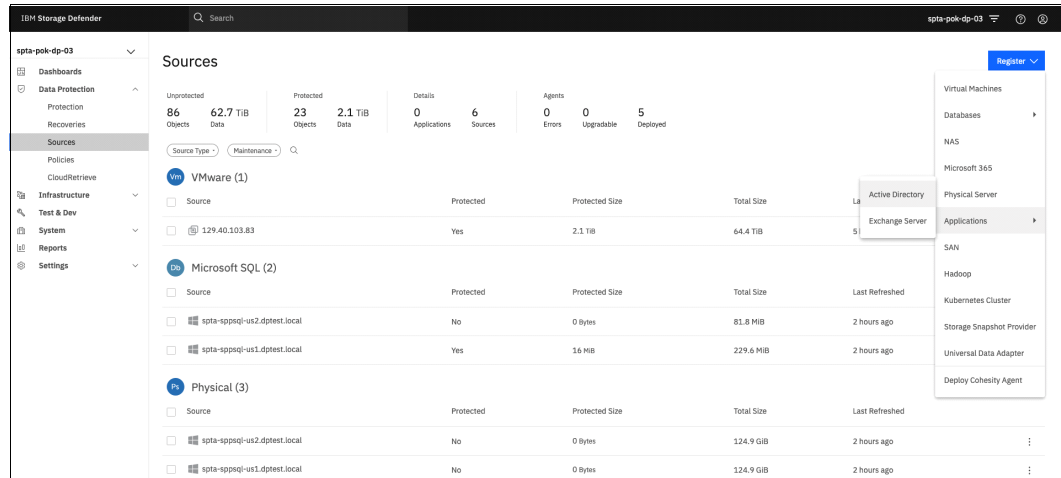


Figure 4-3 Registering Active Directory

3. Fill in the Hostname/IP Address of the Active Directory host

## Register Active Directory

Figure 4-4 Register Active Directory Hostname/IP field

4. Select the desired Active Directory source to complete the registration process

Active Directory (1)				
<input type="checkbox"/>	Source	Protected	Protected Size	Total Size
<input type="checkbox"/>	10.0.2.1	Yes	58 MiB	58 MiB

Figure 4-5 Select desired active directory source

## 4.3 Protect and Recover Microsoft Active Directory

Once the Windows Adapter is installed and registered with Data Protect, use the following steps to protect the Active Directory DB:

1. Expand to Data Protection
2. Select Protection
3. Click Protect and Select Application Active Directory

Once presented with the Active Directory panel, configure the desired options for the backup as shown in Figure 4-6 on page 73.

4. Search and Select the Active Directory Object to restore [4]
5. Click on Protection Group and provide a meaningful protection name [5]

6. Click on Policy and select an existing SLA policy [6]
7. Click on Storage Domain [7]

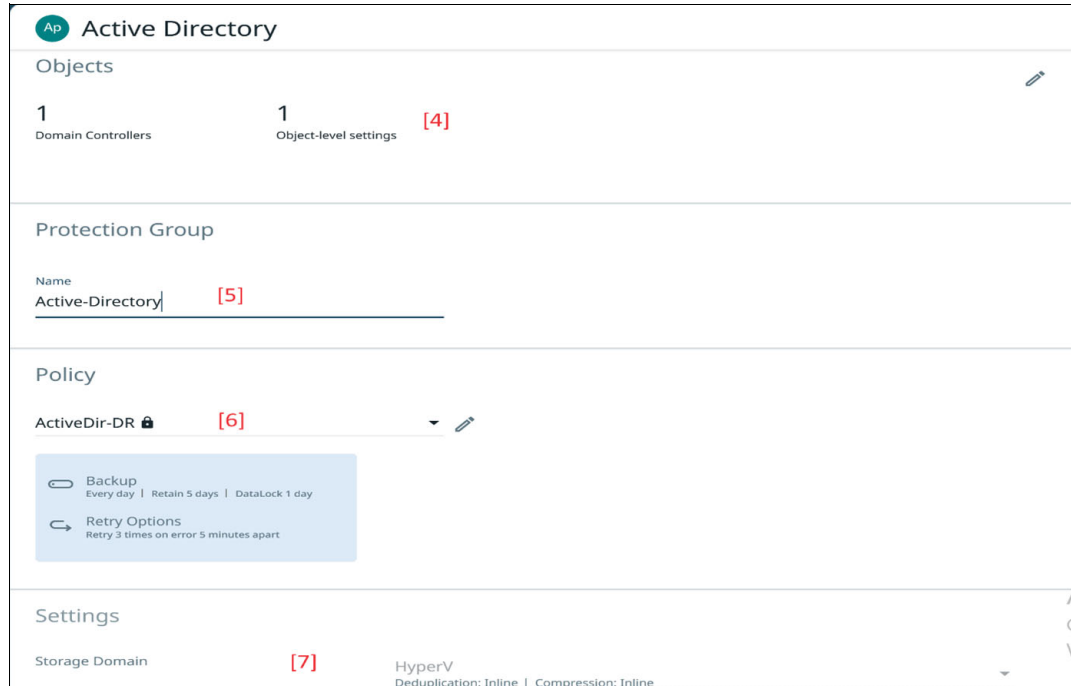


Figure 4-6 Active directory recovery panel

With the Agent installed, the option exists to restore specific objects from the Active Directory DB rather than just performing a full restore. This is referred to as a Granular recovery.

To Granular recovery Active Directory object (Figure 4-7 on page 74):

1. Expand to Data Protection
2. Select Recover
3. Click Recover and Select Application Active Directory
4. Select the Active Directory Server and fill in the Recovery options [4]
  - a. AD Administrative Username:
    - i. Enter the username to use for the recovery. It must be one of the following:
      - ii. Domain Admin or Enterprise Admin
      - iii. Domain user with delegated permissions to the OU where objects will be recovered and reanimate permission on the AD Recycle Bin, if enabled.
  - b. Password: Enter the account password
  - c. Port: Enter the port number to use for the recovery. This port and the 3 consecutive port numbers after it must be currently unused and opened for inbound traffic in the Windows Firewall.

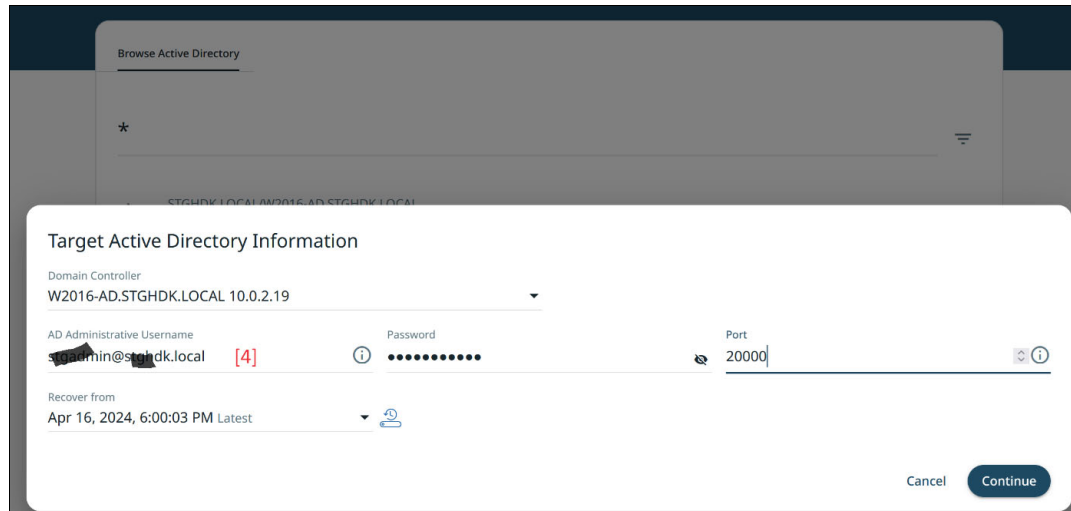


Figure 4-7 Target Active Directory server information panel

5. After successfully mount of the snapshot:
  - a. you can browse the snapshot to review differences between the backup and live data.
  - b. Select Data Protection > Recoveries.
  - c. Click the task name.
  - d. Click Browse Snapshot (Figure 4-8) [5]

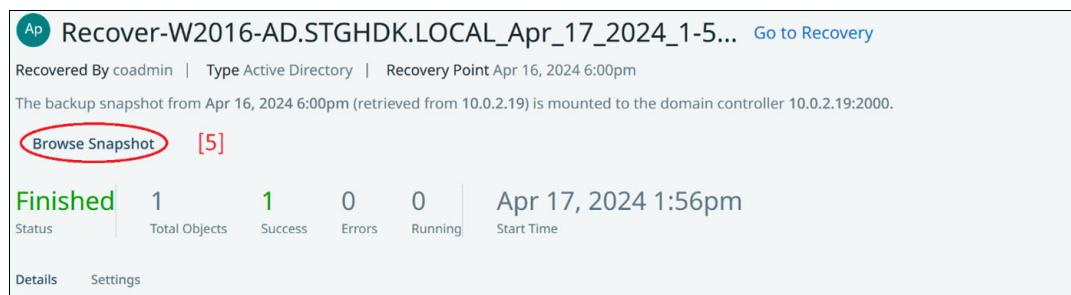


Figure 4-8 Active Directory snapshot version selection panel

6. The Browse Snapshot feature (Recover AD) will mark the differences between the backup set objects (snapshot) and the live Active Directory objects (Figure 4-9) [6]
7. You can use the search bar for text searches [7]. In this example the account 'co-operator' is missing [8] and we can click Recover [9]



Figure 4-9 Recover AD search panel example

**Best Practice:** Search queries are executed against the currently selected entity hierarchy level. To search the entire hierarchy, ensure the top level is selected before running a search.

8. Assign the Recover Object a temporary password (Figure 4-10).

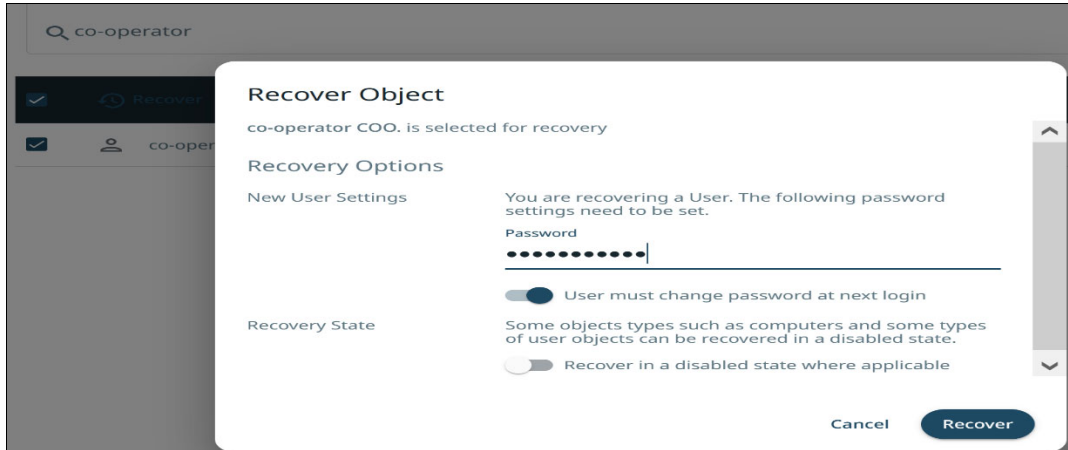


Figure 4-10 Recovery object panel temporary password example

Once the password is set, select the recovery button to begin the 'Recover' process.

9. After a recovery, the status of the AD object will display as "Recovered" if the process was successful.

Q co-operator				
<input type="checkbox"/>	Name	Type	Description	Difference
<input type="checkbox"/>	co-operator COO.	Recovered	User	Yes

Figure 4-11 Recovered Object status







# Protecting Microsoft Exchange on Premises Data

In this chapter we discuss the options for protecting Microsoft Exchange Server data with IBM Storage Defender Data Protect. This includes a configuration example and the required steps to protect and recover Exchange Databases, and prepare the recovered database data for use with Exchange recovery tools.

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

- ▶ 5.1, “Backing up and restoring Microsoft Exchange Data” on page 78
- ▶ 5.2, “Register Exchange server as Data Protection Source” on page 79
- ▶ 5.3, “Restore Exchange data using a recovery database” on page 81

## 5.1 Backing up and restoring Microsoft Exchange Data

Microsoft Exchange is a widely used mailing solution embraced by businesses both big and small as such, being able to protect and recover this data is extremely important. IBM Data Protect Cluster provides the ability to protect Microsoft Exchange data, and restore or recover at the database or single item (mail, contact, or calendar entry) level.

With the removal of the client Access server role in Exchange 2016, only the Mailbox server role is supported with Data Protect. The Mailbox server role hosts the on-premises recipient mailboxes and communicates with the Exchange Online organization by proxy via the on-premises Client Access server. By default, a dedicated Send connector is configured on the Mailbox server role to support secure hybrid mail transport.

### MS Exchange Server Requirements:

The IBM Data Protect agent service logon should be running as a specific AD account (not Local System or local computer account such as local administrator) which has sufficient privileges to run Exchange Management PowerShell and query AD for Exchange objects. Exchange Management PowerShell is required for executing the following sets of cmdlets to get the Exchange server, DAG, and database topologies:

- ▶ Get-ExchangeServer
- ▶ Get-Mailboxdatabase
- ▶ Get-Mailbox
- ▶ Get-DatabaseAvailabilityGroup

### Ensure the following requirements are met to register and backup Exchange Servers:

- ▶ Service account permissions:
  - Is a member of the Backup Operators group on AD Domain
  - Is a member of the Exchange Servers and Organization Management groups under the Microsoft Exchange Security Groups Organizational Unit
  - Is a member of the Local Administrators group on the Exchange Server
  - The Exchange server must have joined the same AD domain as the IBM Data Protect cluster for SMB authentication.
- ▶ Software Prerequisites:
  - IBM Data Protect Agent on the Microsoft Exchange Servers
  - Exchange mailbox recovery tooling
  - A Windows server with 32-bit Microsoft Outlook installed

**Best Practice:** IBM recommended to install the Exchange recovery tooling on a remote management server and not on the Exchange server its self.

### Microsoft Exchange also offers the following built-in data loss prevention options:

- ▶ Deleted item retention:

Whenever a user permanently deletes items in their mailbox database, these items are not purged immediately. Depending on the deleted item retention of the Mailbox Database

(default 14 days) this deleted item is kept in the Mailbox Database and available for self-service restores.

► Deleted User retention:

Comparable to the deleted item retention, user mailboxes that are deleted from a Mailbox Databases are still kept for a specific number of days in this Mailbox Database (default 30 days).

► Database availability groups:

Database availability groups are a great feature to avoid service interruption if a Mailbox Server needs a downtime, is corrupted, or even lost. In this case, the Mailbox database is activated on another copy and the users can access their mailboxes without any interruption.

IBM Data Protect Cluster adds data protection capabilities that can be used whenever the built-in solutions are not satisfying or in case of a disaster.

## 5.2 Register Exchange server as Data Protection Source

To register Exchange Server as data source:

1. Expand Data Protection
2. Click on Sources
3. Choose Applications and then select Exchanger Server
4. Fill in the Exchange server host DNS record (Figure 5-1)

Register Microsoft Exchange Server

Exchange server host or DAG endpoint or DAG host  
10.0.1.80 (4)

Fully Qualified Domain Name (FQDN) is recommended.

**i** The Cohesity Protection Service needs to be installed on the server for registration and backups to work.  
[Download Cohesity Protection Service](#)

Note: Exchange Server configuration discovery may take a few minutes.

Cancel Register

Figure 5-1 Microsoft Exchange server registration panel

**Note:** If the IBM Data Protect (Cohesity) agent is not installed yet please click the 'Download Cohesity Protection Service' link to download and install the agent. If required, ensure that the agent has been copied over to the appropriate server. As an AD Domain Admin, run the executable and complete the installation wizard.

Figure 5-2 shows the add-on component options when installing the IBM Storage Defender (Cohesity) Agent.

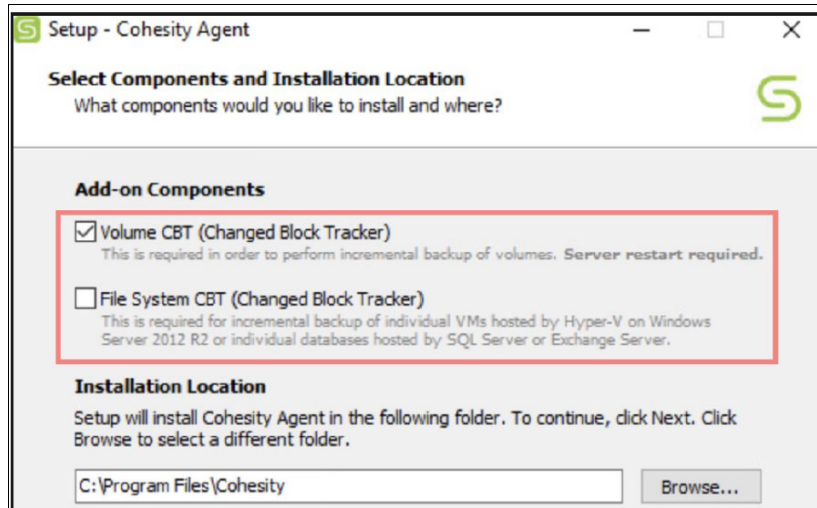


Figure 5-2 Windows Agent installation options

**Volume CBT (Changed Block Tracker):** Install this component for the best incremental backup performance. Installing this component requires a onetime reboot to load the IBM Data Protect Volume CBT driver.

**File System CBT (Changed Block Tracker):** the reboot is not required, but is recommended.

**Service Account Credentials:** The service can run as the “Local System” account with Exchange admin credentials.

Once the agent is installed, select the following options to configure protection for the Exchange Server Databases:

1. Expand to Data Protection
2. Select Protection
3. Click Protect
4. Click on Add Object to select the already register Exchange Server [4]
5. Click on Protection Group and provide a meaningful protection name [5]
6. Click on Policy and select an existing SLA policy [6]
7. Click on Storage Domain [7]

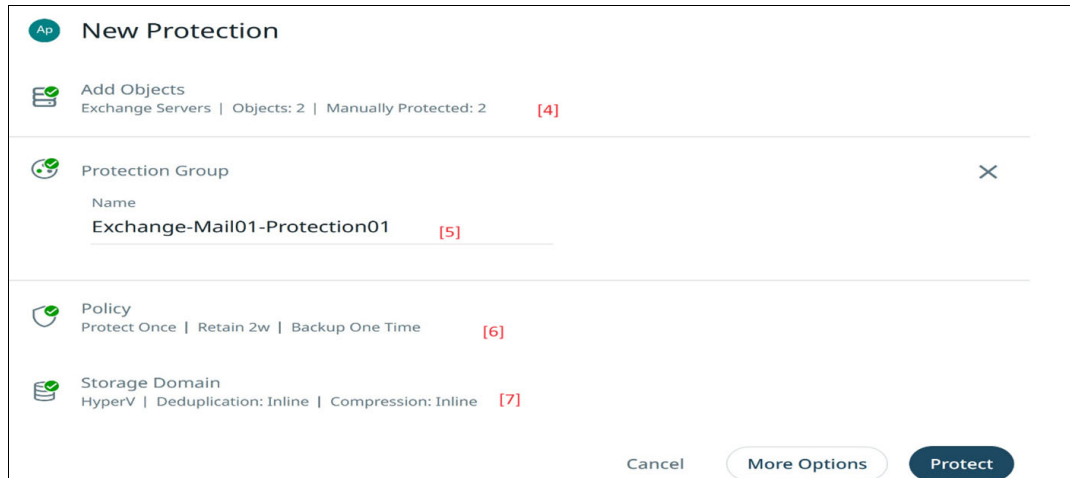


Figure 5-3 Configure a new Protection Group panel

8. Click the 'Protect' button on the bottom right corner to finish and trigger the protection job

## 5.3 Restore Exchange data using a recovery database

Microsoft Exchange Server supports the ability to restore data directly to a recovery database (RDB). Mounting the recovered data as a recovery database allows the administrator to restore individual mailboxes or individual items in a mailbox.

A recovery database (RDB) is a special kind of mailbox database that allows for the temporary mount of a restored mailbox database to extract data from the restored database as part of a recovery operation. You can use the 'New-MailboxRestoreRequest' cmdlet to extract data from an RDB. After extraction, the data can be exported to a folder or merged into an existing mailbox. The use of an RDB enables the recovery of data from a backup or copy of a database, without disturbing user access to current data.

To recover an Exchange Server Databases:

1. Expand to Data Protection
2. Select Recoveries
3. Select Application and Exchange recover tab
4. In the Exchange Recover field filter by Protection Group [4]
5. In the Exchange Recover field type \* to query all protected database [5]



Figure 5-4 Exchange Server recovery options panel

6. Select the desired Exchange Mail Database to recover followed by the next button. For more recovery options, fill in the DNS record field with the address of the server where the recovery tool is running [6].

**Note:** Recovery is performed using third party tools such as Ontrack

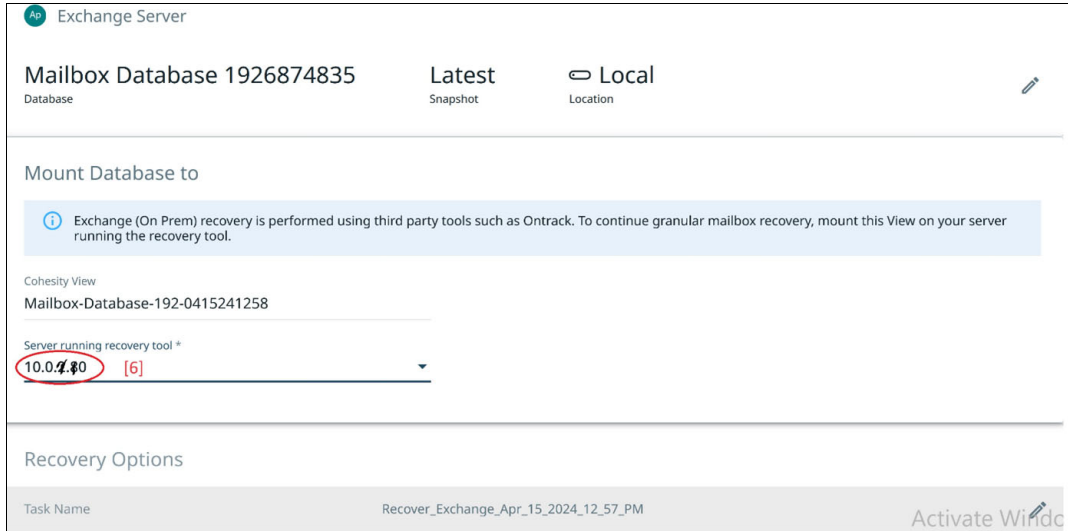


Figure 5-5 Mounting data to Exchange Server for recovery panel

7. Finally click to Recover and a recover task will be created
8. The recovered DB data is available and can be found in the SMB view on the server where the recovery tools are installed. PowerShell or other Recovery tooling can be used to restore single items from the recovered DB.

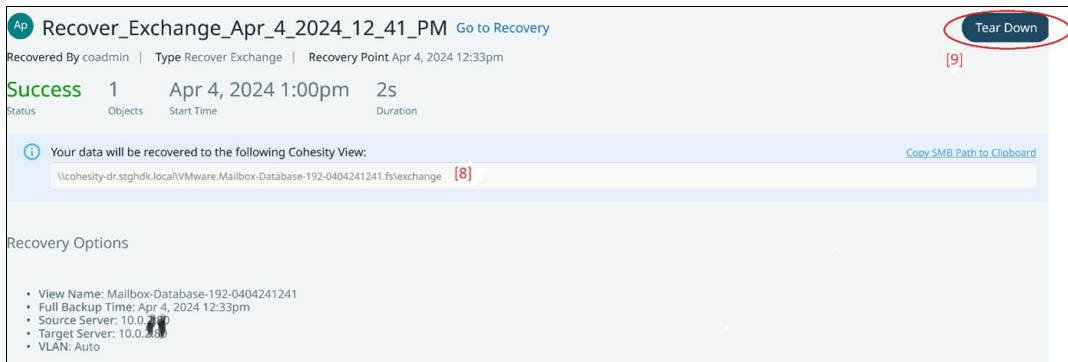


Figure 5-6 Exchange Recovery task details panel

9. After the recovery is no longer needed, select the Tear Down button on the Recovery view panel to remove the view.
10. Using the mlink command, create a hardlink to the mounted exchange data location that can then be manage with Powershell or Recovery tooling.

*Example 5-1 mlink command example*

```

mklink EX-RDB
"\\cohesity-dr.stghdk.local\VMware.Mailbox-Database-192-0404241241.fs\exchange"

```

```

[PS] C:\Windows\system32>Get-MailboxDatabase
Name                               Server           Recovery         ReplicationType
----                               -
Mailbox Database 1926874835        W2016-CSM01     False           None
SMTP                               W2016-CSM01     False           None

```

```

#New-MailboxDatabase -Recovery -Name RDB01 -Server W2016-CSM01 -EdbFilePath
"C:\Users\stgadmin\EX-RDB\Mailbox Database 1926874835.edb" -LogFolderPath
"C:\Users\stgadmin\EX-RDB"

```

```

[PS] C:\Windows\system32>Get-MailboxDatabase
Name                               Server           Recovery         ReplicationType
----                               -
Mailbox Database 1926874835        W2016-CSM01     False           None
SMTP                               W2016-CSM01     False           None
RDB01                              W2016-CSM01     True            None

```

```

[PS] C:\Windows\system32>Get-Mailbox stgadmin@stghdk.local |select
Name,ExchangeGuid
Creating a new session for implicit remoting of "Get-Mailbox" command...
Name                               ExchangeGuid
----                               -
stgadmin STG. stgadmin 8a3e8ed6-9253-4403-9a4e-7bf978543f4a

```

11. Confirm the required data was successfully mounted and is accessible via the Exchange recovery tools.

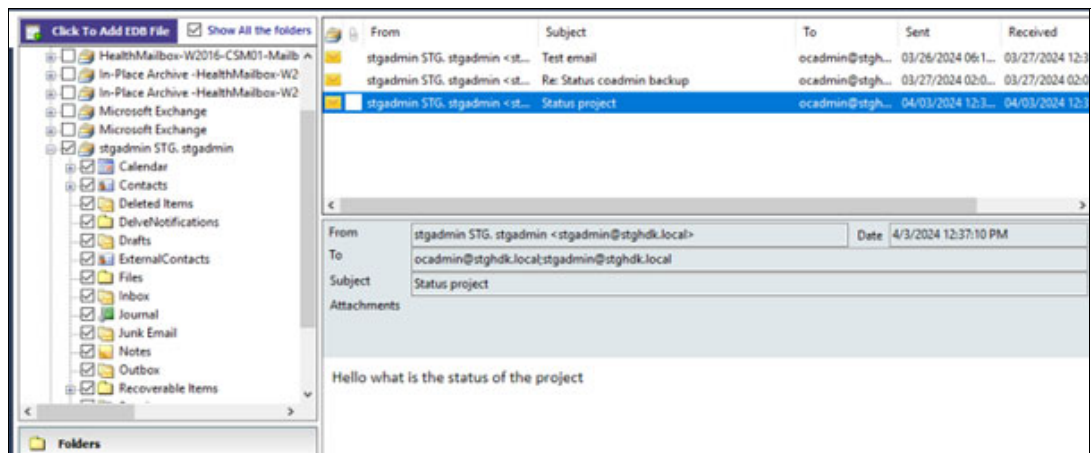


Figure 5-7 Recovered mail information shown with Exchange recovery tools

12. Once the data is confirmed as being accessible, continue to use the Exchange recovery tools to access and restore any individual Exchange objects.







# Protecting PostgreSQL Databases

PostgreSQL has existed since 1986, and started to be more widely adopted as a database engine in the 2000's. Much appreciated for its robustness and reliability in large and complex environments PostgreSQL is often used as the primary data store or data warehouse for many web, mobile, geospatial, and analytics applications.

PostgreSQL comes with many features aimed to help developers build applications, protect data integrity, build fault-tolerant environments and help administrators manage data no matter how big or small the dataset.

This chapter will explain what features IBM Storage Defender Data Protect brings to secure PostgreSQL database on x86-64 platforms.

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

- ▶ 6.1, “Prerequisites and initial steps” on page 86
- ▶ 6.2, “Deployment Overview” on page 89
- ▶ 6.3, “IBM Storage Defender Data Protect capabilities for PostgreSQL database” on page 90
- ▶ 6.4, “practical deployment example” on page 98
- ▶ 6.5, “Troubleshooting” on page 108

## 6.1 Prerequisites and initial steps

Consider the following requirements before starting any implementation of the IBM Storage Defender Data Protect to protect your PostgreSQL database.

### 6.1.1 Versions requirements

The following versions of the PostgreSQL database can be protected by IBM Storage Defender Data Protect.

Table 6-1 Supported PostgreSQL versions

PostgreSQL Versions		Platform	Operating System
PostgreSQL	Releases: 11.x, 12.x, 13.x, 14.x, 15.x	x86-64	Red Hat Enterprise Linux (RHEL) 7.x, and 8.x.  CentOS Linux 7.x and 8.x.  SUSE Linux Enterprise Server (SLES) 12 and 15.
EDB Postgres Releases	Releases: 11.x, 12.x, 14.x		Red Hat Enterprise Linux (RHEL) 7.x, and 8.x.

### 6.1.2 Communication port requirements

There are bidirectional communications between the host where the PostgreSQL database is running and the IBM Storage Defender Data Protect cluster where the data is being backed up. The table below summarizes the list of ports and data flow that needs to be opened in an environment where a firewall is filtering network traffic.

Table 6-2 Agent Communication port usage information

Port	Use	Source	Target	Direction	Network Protocol
50051	Local Agent - Required for backup and recovery	PostgreSQL	Each Data Protect Cluster node	Bidirectional	Tcp/ip
59999	Local Agent - Required for local to local communication for self-monitoring and debugging	PostgreSQL	Each Data Protect Cluster node	Bidirectional	Tcp/ip
11113	Local Agent - Required for backup and recovery operations	PostgreSQL	Each Data Protect Cluster node	Unidirectional	Tcp/ip

**Note:** Port 59999 is required when your PostgreSQL deployment is comprised of multiple nodes (such as a high availability configuration).

### 6.1.3 Local user requirements

As part of the setup, a local agent must be deployed on the system hosting the PostgreSQL database. This installation can be done either with root privileges or with a specific user. When using a non-root user, grant appropriate privileges through the sudo configuration to that user so that the user can perform the required actions.

When planning to use a non-root user to perform the agent and connector installation, add the below line to the `/etc/sudoers` configuration file:

*Example 6-1 Sudo privileges for non-root user required to install local PostgreSQL connector*

---

```
cohesityagent ALL=NOPASSWD:SETENV: /bin/chmod, /bin/chown, /bin/mkdir, /bin/rm,
/bin/psql, /usr/bin/ps, /usr/sbin/runuser, /bin/java, /usr/bin/netstat
```

---

### 6.1.4 Local Command Requirements

Before installing the local Linux agent, be sure that any dependencies are met to allow the following commands/utilities to be available on the host. These commands are used as part of either backup or recovery operations which are triggered by Data Protect, through the local PostgreSQL connector:

*Table 6-3 List of required packages and commands, used by the Linux Agent*

RHEL	SUSE	CentOS	Ubuntu	Debian
rsync	rsync	rsync	rsync	rsync
mount	mount	mount	mount	mount
lsuf	lsuf	lsuf	lsuf	lsuf
umount	umount	umount	umount	umount
cp	cp	cp	cp	cp
chown	chown	chown	chown	chown
chmod	chmod	chmod	chmod	chmod
mkdir	mkdir	mkdir	mkdir	mkdir
rm	rm	rm	rm	rm
tee	tee	tee	tee	tee
hostname	hostname	hostname	hostname	hostname
stat	stat	stat	stat	stat
blkid	blkid	blkid	blkid	blkid
ls	ls	ls	ls	ls
losetup	losetup	losetup	losetup	losetup
dmsetup	dmsetup	dmsetup	dmsetup	dmsetup
timeout	timeout	timeout	timeout	timeout

RHEL	SUSE	CentOS	Ubuntu	Debian
lvs	lvs	lvs	lvs	lvs
vgs	vgs	vgs	vgs	vgs
lvcreate	lvcreate	lvcreate	lvcreate	lvcreate
lvremove	lvremove	lvremove	lvremove	lvremove
lvchange	lvchange	lvchange	lvchange	lvchange
null	libpcap-progs	null	null	null
nfs-utils	nfs-utils	nfs-utils	nfs-utils	nfs-utils
wget	wget	wget	wget	wget

**Note:** nfs-utils is Required for Instant Volume Mount, file-folder recovery from block-based backup and VMware backup.

### 6.1.5 Secure PostgreSQL host to DP cluster communications

To securely connect to the cluster while reading or writing data, use secure gRPC. To enable secure gRPC, use the following steps:

1. Generate the certificate config file on the node using the following command:

*Example 6-2 Certificate config file generation command*

```
~/<release-version>/crux/bin/client_tls_cert_generator_exec -cert_name
<clustername> -output_file /home/cohesity/<clustername>.cfg
```

2. Copy /home/cohesity/<clustername>.cfg from the node to all PostgreSQL nodes at the same location. For example, /opt/certs/<clustername>.cfg.

**Note:** Ensure that the root or non-root user used to install the Linux agent and PostgreSQL Connector has read access to the generated certificate config file.

3. While registering the PostgreSQL source, provide the path of the certificate config file in the SSL Settings field.

### 6.1.6 Other Requirements

The local agent will execute scripts to facilitate the registration of the PostgreSQL database to the Data Protect cluster.

As part of the initial registration, one of the scripts relies on the jq command to parse the output of some local commands. JQ might not be installed by default on your Linux platform. Refer to your operating system documentation to install jq package.

### ***PostgreSQL listening service***

By default PostgreSQL service is listening only connection from localhost. This will prevent any external connection to the database, including the required communication for backup and restore activity triggered by Data Protect.

To change this, you must modify the default settings in the PostgreSQL configuration file, generally located into the database directory, for example `/var/lib/pgsql/16/data/postgresql.conf`. Find the line containing the parameter `listen_addresses` and change the value to allow non local connection, as shown below:

#### *Example 6-3 Listen address option example*

---

```
listen_addresses = '*'           # what IP address(es) to listen on;
```

---

For more security you can configure the local Linux firewall to allow very specific IP or set of IPs addresses allowed to connect remotely to the PostgreSQL database.

Once the value changed, restart the PostgreSQL service using the operating system command `systemctl`.

#### *Example 6-4 systemctl restart command*

---

```
systemctl restart postgresql-16
```

---

## **6.2 Deployment Overview**

The following is an overview of the components involved in the PostgreSQL database protection, as well as the various actions that need to be completed to properly configure the PostgreSQL database protection when using IBM Storage Defender Data Protect.

Configuration will require multiple steps which are summarized in the figure below.

The left side of the figure shows the configuration steps to be executed within the host where the database is running.

The right side of the figure shows the configuration steps to be performed using the IBM Storage Data Protect graphical user interface and its configuration wizards.

Later in this chapter we will go into the details of each step, numbered from 1 to 5 in the figure below (Figure 6-1).

- ▶ Step 1: Download and install the Linux Agent
- ▶ Step 2: Download and install the PostgreSQL Connector
- ▶ Step 3: Register the PostgreSQL host machine as a source of data in Data Protect
- ▶ Step 4: Create and Protection Group and Protection Policy
- ▶ Step 5: Perform backup and recovery activities as required

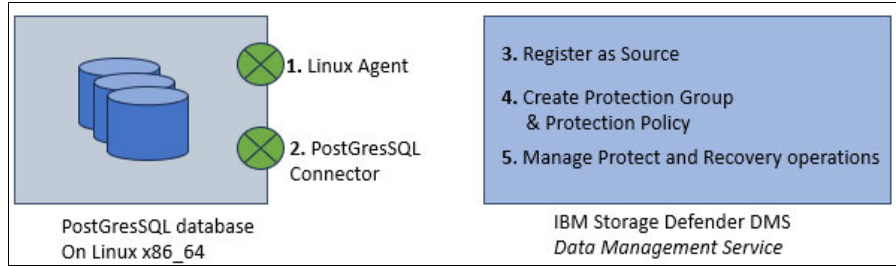


Figure 6-1 PostgreSQL component installation overview

## 6.3 IBM Storage Defender Data Protect capabilities for PostgreSQL database

Data Protect provides a PostgreSQL connector for backup and restore of PostgreSQL databases. This PostgreSQL connector uses the JDBC interface to connect the PostgreSQL Server. It allows a tight integration between the PostgreSQL database and the Defender Data Protect cluster, creating the ability to have online consistent backup capabilities as well as fast recoveries.

### 6.3.1 Backup and Recovery Methods

Backup operations can be scheduled via Protection policy or executed on demand, by manually triggering the protection policy. The following backup methods are available when using Data Protect:

- ▶ Full backup
- ▶ Incremental backup
- ▶ Log backup

The following recovery methods are available, either on the same host or to an alternate host:

- ▶ Full recovery (Regular and Instant Restore methods are available)
- ▶ Point in time recovery

The recommended backup method for PostgreSQL database is to do a FULL database backup on a regular basis, and a INCREMENTAL backup on a daily basis. To further enhance the protection level, you can complete this backup schema with an hourly LOG backup. This backup strategy will allow you to recover to a specific point in time by rolling forward through the available LOGS.

This recommended backup schema will be illustrated in 6.4, “practical deployment example” on page 98, showing how it translates into a Defender DataProtect Protect Policy.

### 6.3.2 Special Considerations

The following is the list of things to be aware of when protecting PostgreSQL databases with IBM Storage Defender Data Protect:

- ▶ Backup or restore is performed at the cluster level. Object-level backup or restore is not supported.

- ▶ In the case of the HA cluster, the backup is executed from the active node only.
- ▶ In case of a failover in the HA cluster, the next backup run will be a full backup.
- ▶ Restore across different PostgreSQL versions is not supported.
- ▶ For optimum log backups performance, it is recommended to set the WAL size to 1GB or higher.
- ▶ When the incremental backup chain is broken, a point-in-time recovery (PITR) is possible only after a subsequent successful full or incremental backup is completed.
- ▶ Data Protect supports backup and restore of PostgreSQL databases running on dual-stack (IPv4 and IPv6) mode or single-stack (only IPv6) mode.

### 6.3.3 Backup Workflows

There are three different types of backups available via Data Protect. Full database backup, Incremental database backup (both covering the datafiles), and Log backup, which is backing up the WAL files of the PostgreSQL database.

Note that the backup schema FULL + Incremental + Logs is the recommended strategy.

If the logs are being backed up, they will be taken after the FULL or INCREMENTAL backup operations, automatically being triggered by Defender DataProtect. LOGS are also backed up as per the LOG schedule and not necessarily after a database backup, as defined in the Defender DataProtect protection policy.

#### **Full backup**

First backup will always be a FULL backup. Beyond the first backup, it is important to make regular FULL backup as PostgreSQL rely on a FULL backup to be able to recover any database.

When a backup is triggered, some checks are being done before actually transferring the data. Data Protect PostgreSQL connector ensures that the database is in correct state to perform the backup. It checks whether there is a recovery on going, where the logs are, where the data files are in the local host. Here under is the list of pre backup queries that ensure the database is in appropriate state for backup, as well as allows the PostgreSQL connector to gather required information to properly configure the backup command.

#### *Example 6-5*

---

```

QUERY : select pg_is_in_recovery()
QUERY : checkpoint
QUERY : show log_directory
QUERY : show data_directory
Command : /usr/sbin/runuser -l postgres -c /usr/pgsql-16/bin/pg_ctl -VExitcode = 0
QUERY : SELECT system_identifier FROM pg_control_system()
QUERY : select substring(pg_walfile_name(pg_current_wal_lsn()), 1, 8) as timeline
QUERY : show archive_mode
QUERY : show archive_command

```

---

The database backup is an actual file copy, from the host where the database is located up to a specific location (in the SpanFS structure) onto the Data Protect cluster. File transfer is happening using the gRPC protocol. All files being identified by the initial backup process are being transferred over to the Data Protect cluster local storage, on a specific and dedicated view that is then being snapshotted at the end of the data transfer, therefore creating a specific point in time copy of this database.

Data Protect Backup process leverages the `select pg_backup_start('Timestamp',true)` PostgreSQL statements to inform the database that an online backup is about to happen.

This instruction prepares the server to begin an on-line backup. The only required parameter is an arbitrary user-defined label for the backup. In the case of Data Protect, it is a backup start timestamp. The second parameter given is `true`, it specifies executing `pg_backup_start` as quickly as possible. This forces an immediate checkpoint which will cause a spike in I/O operations, slowing any concurrently executing queries.

At the end of the backup, the `pg_backup_stop(True)` statement is used. It is used to inform that the system can do the different tasks required to finish an on-line backup. Specifying the "True" argument in the `pg_backup_stop` call implies we will wait for WAL to be archived when archiving is enabled.

### **Incremental Backup**

Incremental backup is using the same workflow as the FULL database backup when it starts. Getting the information on the database paths used to locate data files.

Then the postgresql adapter is comparing the inventory of the current actual data files with the last backup.

Once identified, the files that has changed (length and/or last updated timestamp) are being transferred over gRPC to Data Protect local storage, making the incremental backup point in time copy using the snapshot feature of the SpanFS.

As for the FULL backup, the PostgreSQL is aware of the backup and take appropriate action when the PostgreSQL connector uses the `pg_backup_start()` and `pg_backup_stop procedure()` call.

### **Log backup**

In a PostgreSQL database system, the actual database 'writes' to an addition file called write-ahead log (WAL) which is located on disk storage. These logs contains a record of the write actions which were made in the database. In case of crash, these log files can be used to repaired/recover the database. Protection and maintaining access to these files is important, as it allows for point in time recovery of the database.

**Note:** Postgres manages its log backups outside of Data Protect. It is optional to schedule Log backups for Postgres in Data Protect. It is important only if you want to copy these log files outside of the production system, to your backup environment.

To check if the archive log is enabled on the PostgreSQL database, use the below command, logged as PostgreSQL user.

#### *Example 6-6 Check Archive log status*

```
[postgres@jsa-rhel-01 ~]$psql
psql (16.1)
Type "help" for help.

postgres=# SHOW archive_mode;
 archive_mode
-----
on
```



(1 row)

**Note:** The archivelog mode will be automatically set to **ON** when you enable the Log backup as part of the Protection Policy on IBM Storage Defender Data Protect configuration wizard.

**The PostgreSQL configuration update can be seen into the first FULL backup log, available in /var/log/Cohesity/uda/full-backup.xxx.STDOUT:**

```
2024-06-19 08:27:30.535:Updating Postgres archive command.
2024-06-19 08:27:30.535:QUERY : alter system set archive_command
= '/opt/cohesity/postgres/scripts/stream-log.sh %p %f
/opt/cohesity/postgres/scripts/archive_config/199374'
2024-06-19 08:27:30.541:QUERY : select pg_reload_conf()
2024-06-19 08:27:31.542:QUERY : show archive_command

2024-06-19 08:27:31.545:Postgres archive command is set to :
/opt/cohesity/postgres/scripts/stream-log.sh %p %f
/opt/cohesity/postgres/scripts/archive_config/199374
```

Additionally, to confirm what script and configuration is being used locally to transfer the log onto the Data Protect Cluster, the command below provides the details of what is being executed each time a log backup is initiated:

*Example 6-7 Show archive command output*

```
postgres=# show archive_command;
              archive_command
-----
/opt/cohesity/postgres/scripts/stream-log.sh %p %f
/opt/cohesity/postgres/scripts/archive_config/65964
(1 row)
```

The log backup is executing a script named “/opt/cohesity/postgres/scripts/stream-log.sh” which is configured and use parameter which are specific to the environment from where it runs, that is the PostgreSQL database server.

This script queries the LOG file location and current LSN, to determine what LOG files must be transferred from the last log backup.

The file transfer is happening between the local server and the Data Protect cluster using the gRPC protocol. As for the FULL backup, a dedicated view is being used to store and snapshotted to record the point in time LOG backup.

### 6.3.4 Recovery Workflows

There are two types of recovery methods:

- ▶ **Regular recovery**, meaning that the data files are copied over from the Data Protect local storage to the PostgreSQL host.
- ▶ **Instant Recovery**: meaning that the data are mounted from the Data Protect local backup repository and exposed directly to the host, and immediately accessible in read write mode.

Database backups, differentials, and logs depend on a FULL backup to perform a database restore.

Postgres databases require that you start with a FULL backup recovery before applying any transaction logs. This means your backup retention policy must keep a FULL backup along with its LOG backups to successfully restore a database.

- ▶ It is recommended to retain two sets of FULL backups with their DIFFERENTIAL.
- ▶ Recovering a PostgreSQL database consist in sequentially adding the captured changes to the database: FULL+DIFF+Log1+Log2+... + LogN = Restored database.
- ▶ Independently of the two methods Regular or Instant, the same recovery schema applies

**Note:** When restoring a database, an empty target database must be created first

### **Full Recovery – Regular**

Triggering a database recovery, from a full, an incremental or a specific point in time using log, is done through the Data Management Service interface, selecting the required Data Protect Cluster from the welcome screen, and then navigating to the *Data Protection > Recoveries* menu.

From there select the *Recover* drop down button and select *Universal Data Adapter*.

Search for your PostgreSQL Data Protection Group and choose the appropriate date for your recovery as shown in Figure 6-2:

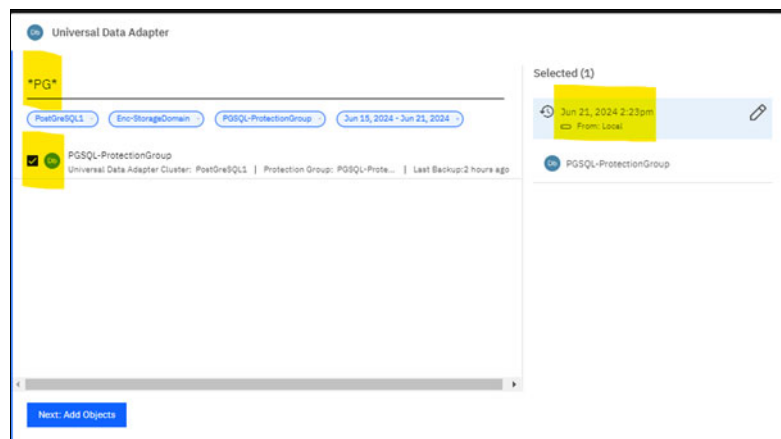


Figure 6-2 Data Protect Universal Data Adapter Recovery wizard – Select resource to recover

Click Next, and specify the other recovery option, such as:

- ▶ Host where you would like to recover. If different from the original, the PostgreSQL source must be registered and prepared beforehand.
- ▶ The date directory location where you would like Data Protect to copy data back on the database host.
- ▶ Request Data Protect to start the PostgreSQL instance after the recovery completed. Note that the instance will be started using the data path as specified in the previous option (*Data directory for restore*)
- ▶ Number of streams can be tuned depending on your environment

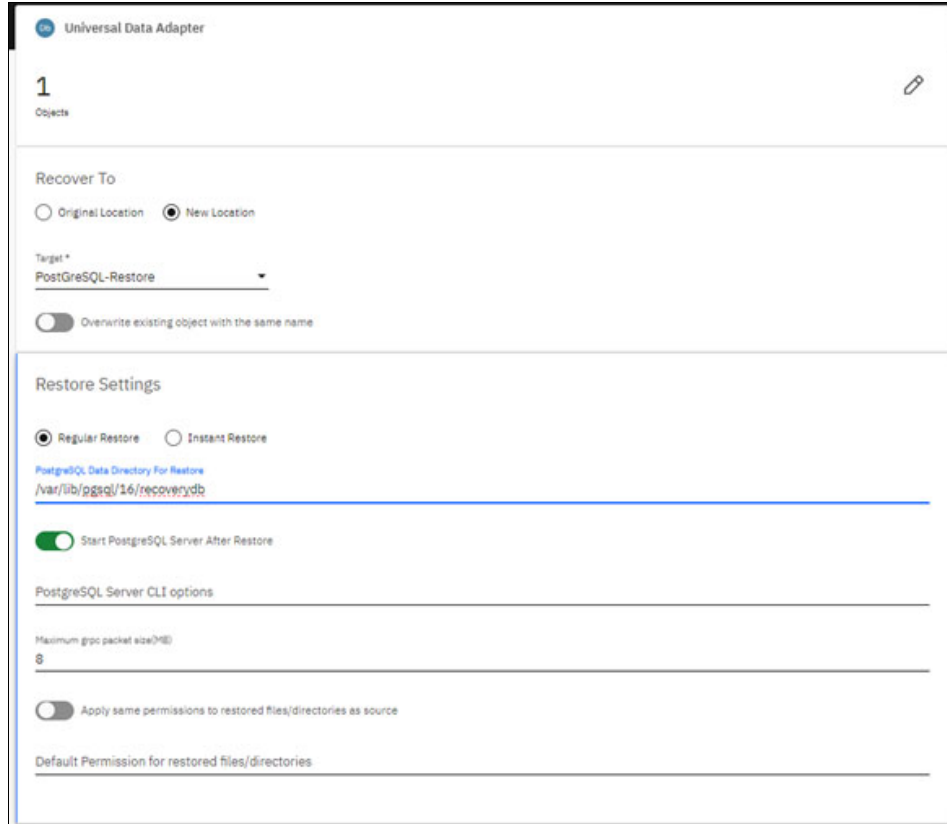


Figure 6-3 Data Protect Universal Data Adapter Recovery wizard - restore settings panel

**Note:** For Full recovery to be successful, the PostgreSQL database service must be stopped. Port 5432 (default) must be available otherwise recovery fails. Reason being that as part of the recovery process, Data Protect restart the PostgreSQL instance using the data path where you recovered the data, as specified in the recovery wizard.

The recovery process is performing the following tasks:

- ▶ Creating clone of the existing snapshot in the Data Protect local storage ( the one representing the point in time backup image for data)
- ▶ Creating clone of the existing snapshot in the Data Protect local storage ( the one representing the point in time backup image for WAL logs)
- ▶ Copy process is then initiated through the gRPC protocol from the Data Protect cluster to the PostgreSQL host.
- ▶ Then Data Protect PostgreSQL connector is configuring the PostgreSQL recovery tasks to apply the specific logs to reach the specific date and time as configured in the recovery wizard (calling the local `/opt/cohesity/postgres/scripts/pitr.sh` script ). This script is setting instruction like `recovery_target_time = '2024-06-20 00:33:29'` and `recovery_target_action = 'promote'`
- ▶ The PostgreSQL instance is then started, and as instructed previously the database engine is doing the rollforward steps to the specific point in time.
- ▶ Once database is up and running on the host, the Data Protect clones used for recovery are being deleted.

### Full Recovery - Instant

It is important to understand that Instant Recovery is not actually an end to end recovery operation managed by Data Protect, meaning that the data are not automatically copied from the backup infrastructure to the target host.

Instant recovery gives the database administrator, instantaneously access to the data from the backup repository.

Data Protect is mounting two mount points, through NFS protocol, between the Data Protection cluster local storage and the database host. The data are then accessible for any operations, including writes. Local copy commands can then be used to copy the data from the backup repository another storage local to the database host.

Here are the steps done by Data Protect, during an instant recovery operation for PostgreSQL database:

- ▶ Creates clone of the backup corresponding to the specified dates, for both the data files and the log files
- ▶ Creates a view and exposes this view as a NFS mount point to the PostgreSQL host
- ▶ Assigns proper privileges to the mounted NFS resources (`chown -R postgres` and `chmod 700` commands)
- ▶ Start the PostgreSQL database on the target host using the mounted NFS resources as data files & log files location for the database

The Instant Recovery procedure stops here, and the database is available for use, in read write mode, from the PostgreSQL host.

The mounted resources can be used for testing, or copy, or any other scenarios that require access to the database.

Figure 6-4 and Figure 6-5 show an example PostgreSQL host when the instant recovery is running. You see the mounted resources and the PostgreSQL server running on these mounted resources.

```
[root@restore-pgsql_uda]# df
Filesystem            1K-blocks    Used Available Use% Mounted on
devtmpfs              8176124      0  8176124   0% /dev
tmpfs                 8204392     1224  8203168   1% /dev/shm
tmpfs                 8204392    59228  8154164   1% /run
tmpfs                 8204392      0  8204392   0% /sys/fs/cgroup
/dev/sda3             163291628 14386744 89906876 14% /
tmpfs                 8204392      0  8204392   0% /tmp
/dev/sda2             479680    221888  248000  33% /boot
/dev/sda1             339532    18776  320756   5% /boot/efi
tmpfs                 1640876      4  1640872   0% /run/user/1000
129.40.183.172:/5219051150900661_200670 83866080 561814784 277846616 67% /opt/cohesity/mount_paths/nfs_uda_mounts/uda_5219051150900661_16908879
61521_200670_129975443_129.40.183.172
tmpfs                 1640876      0  1640872   0% /run/user/26
```

Figure 6-4 PostgreSQL instant recovery mounted resources

```
postgres 4000 24798 0 02:27 ? 00:00:00 /usr/libexec/gvfd-fuse /run/user/26/gvfs -f -o big_writes
postgres 24888 1 0 02:27 ? 00:00:00 /usr/pgsql-16/bin/postgres -D /opt/cohesity/mount_paths/nfs_uda_mounts/uda_5219051150900661_16908879
129975443_129.40.183.172/5219051150900661
```

Figure 6-5 PostgreSQL instant recovery service running on mounted resources

When database administrator has completed his operations, the Instant Recovery must be dismounted from the host. When dismounting, all modifications being done on the mounted resources will be lost.

To dismount and clean up the Instant Recovery, use the Cancel button, from the Storage Defender Data Management Service interface, selecting the Data Protect cluster where the instant recovery is running, and then navigating the *DataProtection > Recoveries* menu, locating your recovery task, finally use the three dots menu on the same job line to select the *Cancel* option as shown in Figure .




Figure 6-6 Data Protect recovery menu - Cancel a PostGreSQL database instant recovery

A `Cancel recovery` popup will appear, as a confirmation. This operation will dismount the volumes from the PostGreSQL host, and delete the cloned backup from the Data Protect cluster storage repository.

**Note:** For instant recovery to be successful, the PostGreSQL database service must be stopped. Port 5432 (default) must be available otherwise Instant recovery fails. Reason being that as part of the recovery process, Data Protect restart the PostGreSQL instance using the mounted resources as data path for the recovered database.

### Point in Time Recovery

Point in Time Recovery means that you can recover to a very specific time using the combined recovery of FULL + Incremental + Logs, in that specific order, until the database reflects the very specific time (as close as the second), as specified in the recovery wizard.

To achieve this, go to the `Data Protection > Recoveries` menu and click the `Recover` button.

Then Select Universal Data Adapter, enter the name of the PostGreSQL protection group , then select the protection group and use the pen icon next to the backup date so you can access the recovery point wizard as shown in Figure 6-7. Be sure to select the **Timeline** view so you can navigate and select a very specific date and time.

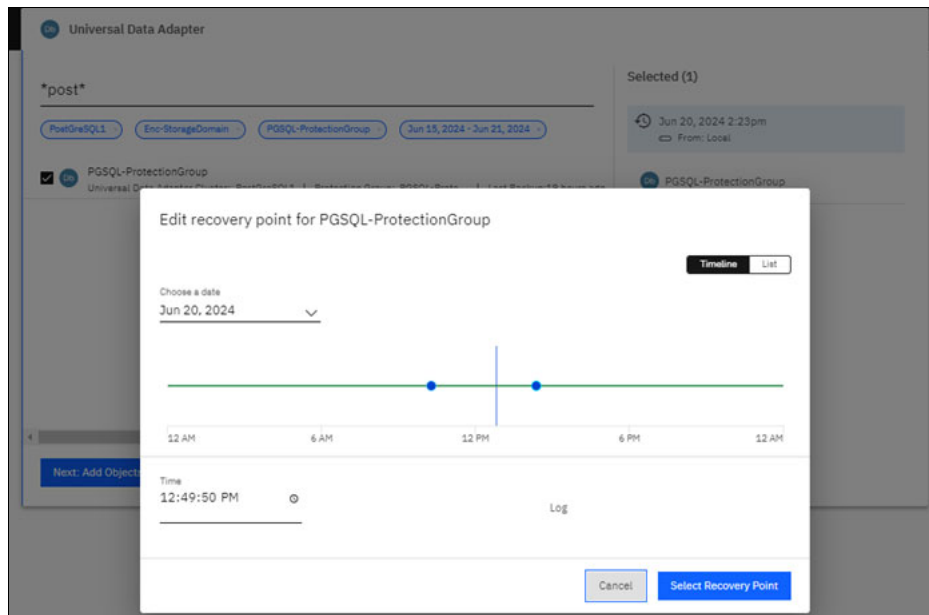


Figure 6-7 Data Protect PostGreSQL Point in Time recovery Timeline

**Note:**

- ▶ The blue dots represent Full or Incremental database backup points.
- ▶ The green line represents possible point in time selections to restore, down to specific second granularity. This is made possible via the use of database logs that have been protected as part of the backup strategy.

## 6.4 practical deployment example

For the deployment example provided in this chapter, the environment being used for deployment is Linux RHEL 8.4, where we installed a PostgreSQL database version 16.

On that environment a small database has been created, containing a table filled with data using pgbench utility. Pgbench is calling a simple set of sql instruction, executed every 30min, to add, remove and update entries in this database to simulate workload, this allows us to generate few log files to better illustrate the LOG backup process.

### 6.4.1 Download and install the Linux and PostgreSQL Connector agents

Perform the following steps on the host where the PostgreSQL database is running.

The Linux agent is available for different installer packages, providing support for multiple Linux distributions. Depending on selection in the download page, you will find RPM (for RHEL and its derivative), Suse RPM or Script installer (All supported Linux operating systems).

At the time of writing this publication, the agent binaries are not available through the IBM Defender Data Management Service portal. You must connect to the local UI interface to Download the Linux Agent.

Once connected to the local User Interface, navigate through the menu:

1. Data Protection
2. Sources
3. Click the Register button at the top right of the screen
4. From the drop down menu that appears, select Universal Data Adapter menu
5. Click the link Download Agent, as show in the figure Figure 6-8 below

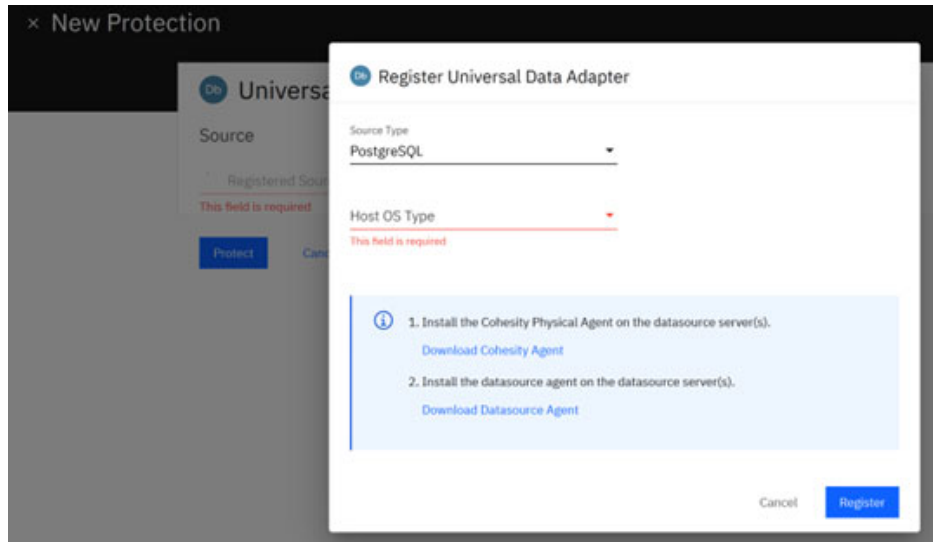


Figure 6-8 Linux Agent Download Screen from the local UI.

**Note:** The agents are also, always available from the local IBM Storage Defender Data Protect Cluster UI.

There are two packages uploaded into the /home/spectrum folder of this example machine.

*Example 6-8 file list for agent files*

---

```
[root@jsa-rhel-01 spectrum]#ls -l *.rpm
-rw-rw-r--. 1 spectrum spectrum 59098606 May 14 03:26
cohesity-postgres-connector-7.1-1.x86_64.rpm
-rw-rw-r--. 1 spectrum spectrum 117704223 May 14 03:26
e1-cohesity-agent-7.1-1.x86_64.rpm
```

---

Next, check whether the required packages are installed (as documented in chapter 6.1.4, “Local Command Requirements” on page 87).

*Example 6-9 Check that all needed system packages are present prior to agent install*

---

```
[root@jsa-rhel-01 spectrum]#for c in rsync mount lsof umount cp chown chmod mkdir
rm tee hostname stat blkid ls losetup dmsetup timeout lvs vgs lvcreate lvremove
lvchange wget; do which $c ; done
/usr/bin/rsync
/usr/bin/mount
/usr/bin/lsof
/usr/bin/umount
alias cp='cp -i'
      /usr/bin/cp
/usr/bin/chown
/usr/bin/chmod
/usr/bin/mkdir
alias rm='rm -i'
      /usr/bin/rm
/usr/bin/tee
/usr/bin/hostname
/usr/bin/stat
```

```

/usr/sbin/blkid
alias ls='ls --color=auto'
    /usr/bin/ls
/usr/sbin/losetup
/usr/sbin/dmsetup
/usr/bin/timeout
/usr/sbin/lvs
/usr/sbin/vgs
/usr/sbin/lvcreate
/usr/sbin/lvremove
/usr/sbin/lvchange
/usr/bin/wget

```

```

[root@jsa-rhel-01 spectrum]#for r in libpcap-progs nfs-utils; do rpm -qa | grep $r
> /dev/null || echo $r missing ; done
libpcap-progs missing

```

```

[root@jsa-rhel-01 spectrum]#

```

---

Once it is confirmed that the required packages installed or present on the host and the required commands are available for the RHEL v8 environment (**libpcap-progs** not required for RHEL 8). proceed with the Agent and PostgreSQL connector installation, using the 2 rpm packages.

**Note:** Using root to install the packages in this example, the local agent and PostGreSQL connector will be running as root user. If you would like to use non-root user for installation and service owner, create a dedicated user account for this on the host and grant appropriate sudo privileges to allow this non-root user to run the required commands.

For PostGreSQL connector, the sudo configuration shown in Example 6-10 is required, assuming the non-root user that was created is named **cohesityagent** (please note that the path to this command might differ in your environment. If needed, use the **which** command as shown in Example 6-9 on page 99 to get the right path to commands).

*Example 6-10 sudo configuration for non root Agent*

```

cohesityagent ALL=NOPASSWD:SETENV: /bin/chmod, /bin/chown,
/bin/mkdir, /bin/rm, /bin/psql, /usr/bin/ps, /usr/sbin/runuser,
/bin/java, /usr/bin/netstat

```

---

At this point the Agent is ready for install either as root or the desired service user id.

*Example 6-11 Linux Agent installation*

```

[root@jsa-rhel-01 spectrum]#rpm -ivh el-cohesity-agent-7.1-1.x86_64.rpm
Verifying... ##### [100%]
Preparing... ##### [100%]
Environment variable COHESITYUSER not defined..will Use root account
Service has already been stopped
Updating / installing...
 1:cohesity-agent-7.1-1 ##### [100%]
Environment variable COHESITYUSER not defined..Using root to run the service!
LVM: Using lvs from path: /usr/sbin/lvs
LVM: 2.03.11(2)-RHEL8(2021-01-28)
Using resultant set_env file as: /opt/cohesity/agent/software/crux/bin/set_env.sh

```



```

Writing env of this upgrade to: /opt/cohesity/agent/software/crux/bin/set_env.sh
Adding systemd service
Synchronizing state of cohesity-agent.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable cohesity-agent
Register cohesity-agent service to systemd succesful.

```

---

Once installed, check the status of the agent service to confirm that the installation went as expected and the agent service was successfully started. At this time the Agent should be ready and listening on the host:

*Example 6-12 Checking Agent status*

```

[root@jsa-rhel-01 spectrum]#systemctl status cohesity-agent
cohesity-agent.service - Linux Agent
   Loaded: loaded (/usr/lib/systemd/system/cohesity-agent.service; enabled; vendor preset: disabled)
   Active: active (running) since Tue 2024-05-14 10:27:14 EDT; 3min 6s ago
   Main PID: 3150441 (linux_agent_exe)
     Tasks: 373 (limit: 102200)
    Memory: 35.3M
   CGroup: /cohesity.slice/cohesity-agent.service
           ..3150441 /opt/cohesity/agent/software/crux/bin/linux_agent_exec
           --log_dir=/var/log/--max_log_size=30 --stop_logging_if_full_disk=true --logbufl>
           ..3150443 /opt/cohesity/agent/software/crux/bin/linux_agent_exec
           --log_dir=/var/log/--max_log_size=30 --stop_logging_if_full_disk=true --logbufl>

May 14 10:27:12 jsa-rhel-01 linux_agent.sh[3150000]: root
May 14 10:27:12 jsa-rhel-01 linux_agent.sh[3150000]: uid=0(root) gid=0(root)
groups=0(root) context=system_u:system_r:unconfined_service_t:s0
May 14 10:27:12 jsa-rhel-01 linux_agent.sh[3150000]: Starting linux_agent_exec...
May 14 10:27:12 jsa-rhel-01 linux_agent.sh[3150000]: GNU coreutils version = 8.30
May 14 10:27:12 jsa-rhel-01 linux_agent.sh[3150000]: Timeout command:
/usr/bin/timeout does support --kill-after option
May 14 10:27:12 jsa-rhel-01 linux_agent.sh[3150000]: USER : root
May 14 10:27:12 jsa-rhel-01 linux_agent.sh[3150000]: Init system is systemd, will
run linux_agent_exec in background
May 14 10:27:14 jsa-rhel-01 linux_agent.sh[3150000]: Parent process pid=3150441
May 14 10:27:14 jsa-rhel-01 linux_agent.sh[3150000]: 3150443 3150441
May 14 10:27:14 jsa-rhel-01 systemd[1]: Started Linux Agent.

```

---

Next, install the connector for the PostgreSQL DB:

*Example 6-13 Installing the PostgreSQL connector*

```

[root@jsa-rhel-01 spectrum]#rpm -ivh cohesity-postgres-connector-7.1-1.x86_64.rpm
Verifying...                               ##### [100%]
Preparing...                               ##### [100%]
OS Version : Red Hat Enterprise Linux release 8

PostgreSQL connector installation is happening at /opt/cohesity/postgres
Updating / installing...
 1:cohesity-postgres-connector-7.1-1##### [100%]
Installation successful for PostgreSQL connector
Scripts Directory   : /opt/cohesity/postgres/scripts

```

---

At this point all components should be installed and ready to be configured for use by Data Protect.

### 6.4.2 Step 3: Register the PostgreSQL host machine as a source

On the host side, to check whether the registration process succeed and that the PostgreSQL connector is working correctly, you check under the `/var/log/cohesity/uda` directory. first action being triggered is the verify-source script, hence, look for the **verify-source.\*.STDOUT** file to see if exit code is 0. If the verification ends with exit code 0, it means that all went fine, and the database will show up in the Defender Management Service portal, under the appropriate Defender Cluster.

To register the PostgreSQL database as a source, use the left-hand side menu, under the `DataProtection > Sources` menu.

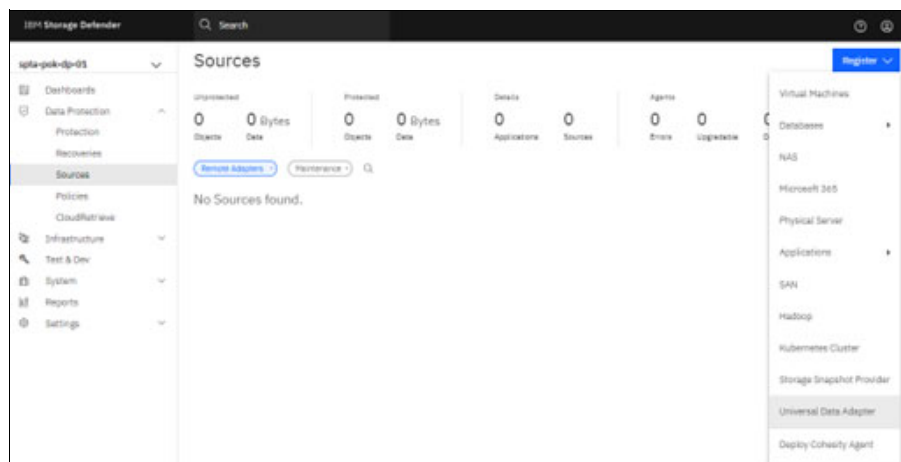


Figure 6-9 Data Protect Register PostgreSQL host as a source #1

Select the Source Type as PostgreSQL from the drop down list and select the appropriate host type (Linux in our example).

Specify the IP address and the Datasource agent installation path, which by default is pointing to `/opt/Cohesity/postgres/scripts`.

Then specify with which user PostgreSQL related commands will be executed. In our example shown in Figure 6-10 on page 103 a dedicated user named “postgres” has been created to interact with PostgreSQL database. This user creation is not covered in this document, you can find information about this user in the PostgreSQL database installation documentation.

Register Universal Data Adapter

Source Type  
PostgreSQL

Host OS Type  
Linux

Hostnames/IP Addresses  
129.40.103.210 Hostnames/IP Addresses  
One or more comma separated hostnames/IP addresses

Datasource Agent Installation Path  
/opt/cohesity/postgres/scripts

Authentication Settings

Password  Kerberos

Username  
postgres

Password  
.....

Cancel Register

Figure 6-10 Data Protect Register PostGreSQL host as a source #2

Finally in the **Source Settings** section of the source registration wizard, give the source a meaningful name, *PostGreSQL1* in our example, specify the IP address of the PostGreSQL controlling node and the port that is used for listening to the external connections.

Specify the path where the PostGreSQL binaries are located. These binaries are used by the PostGreSQL connector to perform database and logs backups as well as recoveries. Figure 6-11 on page 104

Register Universal Data Adapter

PostgreSQL Client SSL Settings

Cohesity SSL Settings

Source Settings

PostgreSQL Datasource Name  
PostGreSQL1

PostgreSQL Server hostname/IP  
129.40.103.210

PostgreSQL Port  
5432

Check Database Connection

Directory Path For PostgreSQL Binaries  
/usr/pgsql-16/bin/

Cancel Update

Figure 6-11 Data Protect Register PostGreSQL host as a source #3

**Note:** It is possible that by default the PostGreSQL database listener is accepting only local connection. For the database backup and recovery operations with Data Protect, it is mandatory that non-local connections are allowed by the listener.

To do this, you need to update the postgresql.conf file (located under the installation directory, for example /var/lib/pgsql/16/data/postgresql.conf) and allow specific or all IPs. See 6.1.6, “Other Requirements” on page 88 earlier on this document for detailed explanation.

### 6.4.3 Step 4: Create a Protection Group and Protection Policy

As explained previously in this chapter, the recommended backup strategy is to configure a regular FULL database backup, a daily INCREMENTAL backup and regular LOG backups. This backup schema translates into the Protection Policy into Defender Data Protect.

Accessing the Protect Policy creation via the left hand side bar, under the Data Protection > Policies menu. The Protect Policy is illustrated in the figure below.

In the Protection Policy named “PGSQL”, we have configured a regular FULL database backup, every Week on Saturday, a Daily backup (INCREMENTAL), and a LOG backup every

hour. All of this being kept for 2 weeks on the local Data Protect cluster storage (Primary Copy=Local) Figure 6-12 on page 105

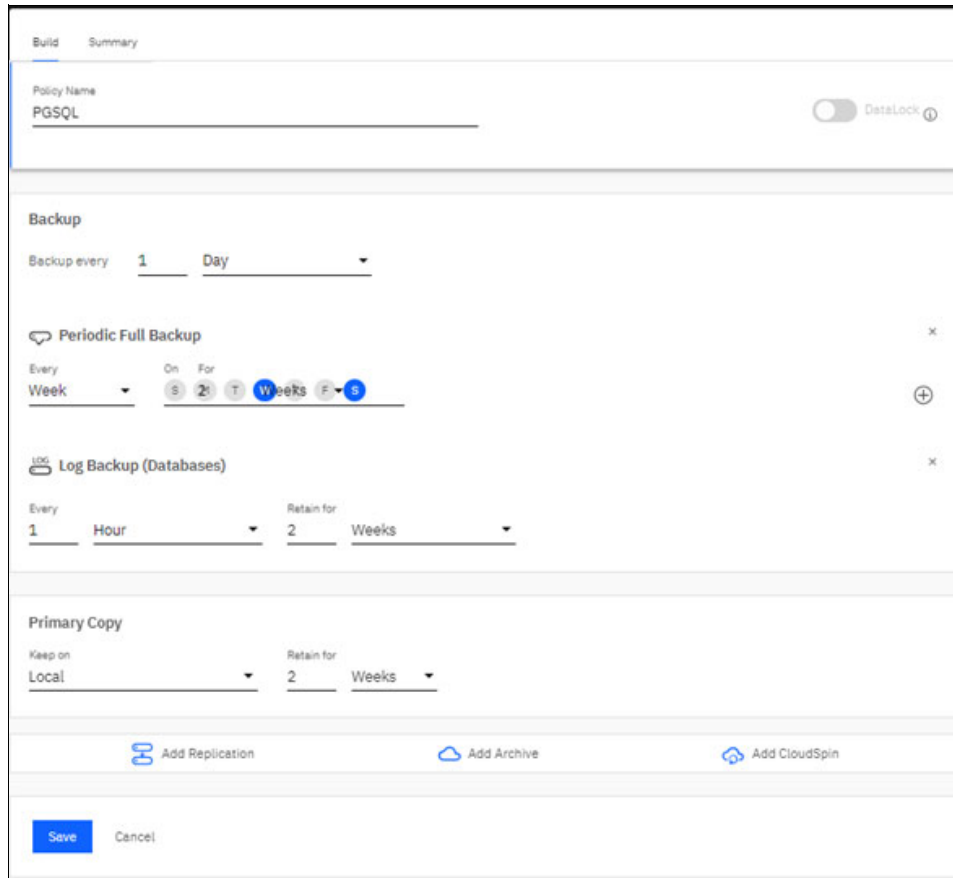


Figure 6-12 Defender Data Protect Protection Policy for PostgreSQL database

Once the protection Policy created, Protection group can be configured to associate the defined PostgreSQL source with the newly configured Protection Policy.

One way of doing it is to use the Protection menu from the *DataProtection > Protection* panel.

From there, select *Universal Data Adapter*. From drop down list that is appearing, select the Registered Source corresponding to your PostgreSQL environment, as show in figure below, PostgreSQL1 in our example. Figure 6-13

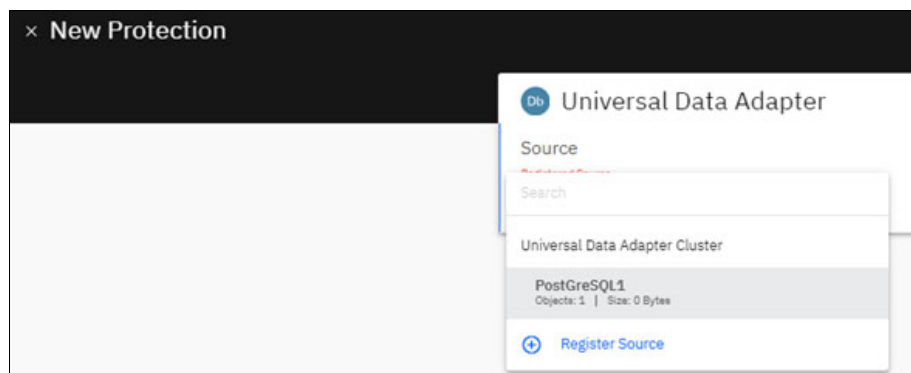


Figure 6-13 Defender Data Protection New Protection Group for PostgreSQL

Select the appropriate source (**PostgreSQL1** in our example) and specify a meaningful object name (**PSQL-DB1** in our example). Then Select the appropriate Protection Policy (the one created just before, **PGSQL** in our example).

Other options can be left with default values.

The screenshot shows the 'Universal Data Adapter' configuration page. The 'Source' dropdown is set to 'PostgreSQL1'. The 'Object Name' field contains 'PSQL-DB1'. The 'Protection Group' name is 'PGSQL-ProtectionGroup'. The 'Policy' dropdown is set to 'PGSQL'. Below the policy dropdown, a list of backup options is visible: 'Backup' (Every day | Retain 2 weeks), 'Periodic Full Backup' (Every week on Wed, Sat | Retain 2 weeks), 'Retry Options' (Retry 3 times on error 5 minutes apart), and 'Log Backup (Databases)' (Every 3 hour | Retain 2 weeks).

Figure 6-14 Data Protection PostgreSQL New Protection Group

Once you defined this Protection Group and its associated policy, Data Protection will trigger the first backup, which will be a FULL database backup, immediately followed by a LOG backup (If you enabled the LOG backup)

This can be seen from the interface as shown in Figure 6-15.

The screenshot shows the 'Group Details: PGSQL-ProtectionGroup' page. It displays a table of backup activities. The table has columns for Start Time, Duration, Backup Type, Data Read, Data Written, Success/Error, S.A, and Status. Two backup activities are listed: a LOG backup on Jun 14, 2024 at 2:28pm and a Full backup on Jun 14, 2024 at 2:27pm.

Start Time	Duration	Backup Type	Data Read	Data Written	Success/Error	S.A	Status
Jun 14, 2024 2:28pm	2m	LOG	0 Bytes	0 Bytes	0/0 objects		
Jun 14, 2024 2:27pm	40s	Full	34.2 MB	0 Bytes	1/0 objects		

Figure 6-15 Data Protection PostgreSQL Data Protection activity

From this DMS view, you can access all the details and logs. Execution logs are also available on the system where the database is running, under the `/var/log/cohesity/uda` folder.

## 6.4.4 Step 5: Backup and Recovery activities

As soon as the protection group and protection policy has been assigned to the PostgreSQL source, all is being managed by Data Protect. It will orchestrate the FULL, INCREMENTAL, LOGs backup as per the schedule as well as controlling the data expiration as per the specified data retention.

All the backup and recovery activity are being managed from the Data Management Service portal.

Using the `Data Protection > Protection` menu, there is a list of all protection activities, including the PostgreSQL backup we just configured.

By selecting the PGSQL-ProtectionGroup Protection Group we created all details about backup activities can be accessed, as well consumption information, as shown in screenshot below, by switching to the Consumption tab.

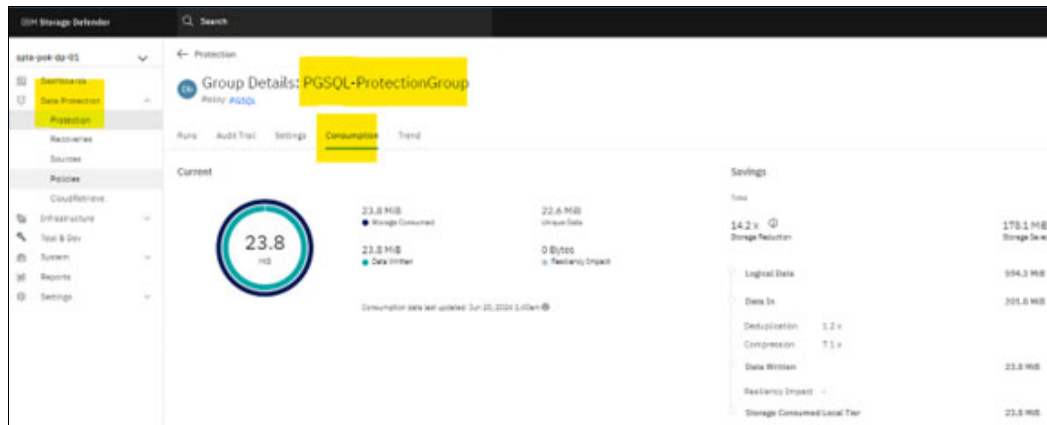


Figure 6-16 Data Protection PostgreSQL Protection Group Consumption

Figure 6-17 shows a screenshot taken from the recovery wizard, indicating the ability to recover to a specific point in time, as we configured the backup to take logs.

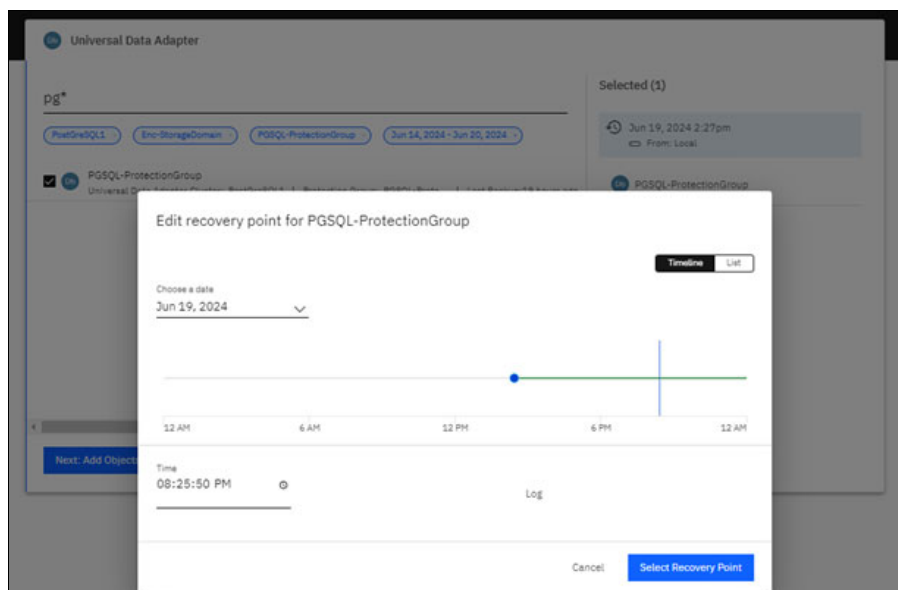


Figure 6-17 Data Protect PostgreSQL Recovery Point In Time selection

The Recovery wizard and workflow is explained in 6.3.4, “Recovery Workflows” on page 93.

## 6.5 Troubleshooting

The following section contains information about the various logs related to the data protection components and the database protection process. Beside the information which are gathered and presented in Defender Data Management Service interface, there is a way for you to investigate with very detailed logs located in the database host itself.

Hereafter are the different logs and their location you can consult when deeper investigations are required.

### **Linux Agent logs**

Agent logs are usually located under the `/var/log/` directory.

The installation path may differ depending on the configuration of the user which deployed the agent.

Local agent creates file named `linux_agent_exec.*` which contains detailed messages regarding the local backup and recovery activities.

### **Linux Universal Adapter Agent logs**

Universal Adapter Agent logs are usually found under the `/var/log/cohesity/uda` directory.

The installation path may differ depending on the configuration of the user which deployed the agent.

In this folder a log file is created for each scheduled backup activity type being handled by the specific adapter (PostgreSQL in this case). The types include full, incremental and log backups.

The `*PULSE*` log file, will contain detailed messages regarding action and commands that the PostgreSQL agent is executing to perform the given operation. In the example below, a full database backup action is being taken:

#### *Example 6-14 Universal Adapter Agent PostgreSQL PULSE log for full backup*

---

```
[root@jsa-rhel-01 uda]#less
full-backup.5219051150900661-1690887961521-177730.PULSE.log
AgentInput [databases=[mytest], ParallelObjects=6, Concurrency=8, objects=[],
restoreObjectsMap={}, dataView=5219051150900661-65964-177730,
logView=5219051150900661-65964-3683-log, connectorType=POSTGRES,
allowIncrementalBackup=false, opType=FULL_BACKUP, userName=julien,
targetRestoreDir=null, customConnProps=null, truststorePassword=null,
truststorePath=null, startTime=Mon May 13 03:33:14 EDT 2024, host=1.2.3.4,
port=5432, VIP's=[1.2.3.5], s3Endpoint=null, backupHangTimeout=1200,
convertIncrToFullBackupIfError=true, retentionPeriod=0, kerberosConfigFile=null,
pitrTime=0 : Wed Dec 31 19:00:00 EST 1969, createDatabase=false, overwrite=false,
instantRestore=false, startServer=true, dataViewMount=null, useSecureGrpc=false,
certificateConfigPath=null, enableDedupWrite=false, enableDedupRead=false,
maxIOBytes=4194304, ioThreadCount=64, rpcTimeoutMsecs=0,
maxGrpcMessageBytes=41943040, postgresCLIOptions=, applyPermissions=false,
defaultPermission=null, jobDataServicePort=0, archivalDataServicePort=0,
offlineBackup=false, deactivateDatabase=false, activateDatabase=false,
```



```

isRedirectedRestore=false, redirectedRestoreSQLFile=null, startDataService=false,
logArchiveDirectory=null, pruneLogs=false, pruneHours=0, logStagingDirectory=null]
Shutdown hook added
Data source verified
Progress monitoring started
Data movement tasks created
Database weights for progress reporting : {DEFAULT=100.0}
Data movement tasks added for monitoring
Data movement tasks started
BackupTaskInput [allowIncrementalBackup=false, opType=FULL_BACKUP,
connectorType=POSTGRES, databaseName=DEFAULT,
dataView=5219051150900661-65964-177730, retentionPeriod=0, backupHangTimeOut=1200,
convertIncrToFullBackupIfError=true, logView=5219051150900661-65964-3683-log,
concurrencyPerDb=8]
PROCESSING RESULT
QUERY : select pg_is_in_recovery()
QUERY : checkpoint
QUERY : show log_directory
QUERY : show data_directory
Command : /usr/sbin/runuser -l postgres -c /usr/pgsql-16/bin/pg_ctl -VExitcode = 0
QUERY : SELECT system_identifier FROM pg_control_system()
QUERY : select substring(pg_walfile_name(pg_current_wal_lsn()), 1, 8) as timeline
QUERY : show archive_mode
QUERY : show archive_command
QUERY : select spcname,pg_tablespace_location(oid) as location from pg_tablespace
Table spaces : {}
target base path : /5219051150900661-65964-177730/DEFAULT/1715585594339
Sources :
{/var/lib/pgsql/16/data=/5219051150900661-65964-177730/DEFAULT/1715585594339}
Sources :
{/var/lib/pgsql/16/data=/5219051150900661-65964-177730/DEFAULT/1715585594339}
QUERY : show archive_mode
Postgres archive mode is already ON
Command : id postgres -gnExitcode = 0
QUERY : show archive_command
QUERY : show archive_command
Postgres archive command is already set and up-to-date, no changes required.
QUERY : select pg_backup_start('Mon May 13 03:33:14 EDT 2024', true)
List generation done : DataMoverStatus [totalSize=35979052, movedSize=0,
failedSize=0, skippedFiles=0, skippedDirs=0, skippedSize=0, totalFiles=1279,
movedFiles=0, failedFiles=0, totalDirs=26, movedDirs=0, failedDirs=0,
toDeleteSize=0, toDeleteFiles=0, toDeleteDirs=0, deletedSize=0, deletedFiles=0,
deletedDirs=0, failedToDelete=0]
...

```

---

### ***PostgreSQL data protection logs***

Backup logs are located within the PostgreSQL database installation path, under the data/log folder. The database in this example is dumping verbose logs for any of its backup activities to this location.

#### *Example 6-15 Listing PostgreSQL database protection logs*

---

```

[root@jsa-rhel-01 log]#pwd
/var/lib/pgsql/16/data/log
[root@jsa-rhel-01 log]#ls -ltr

```

```
total 420
-rw-----. 1 postgres postgres 58828 May  7 10:09 postgresql-Tue.log
-rw-----. 1 postgres postgres 58828 May  8 10:09 postgresql-Wed.log
-rw-----. 1 postgres postgres 58828 May  9 10:09 postgresql-Thu.log
-rw-----. 1 postgres postgres 58828 May 10 10:09 postgresql-Fri.log
-rw-----. 1 postgres postgres 58828 May 11 10:09 postgresql-Sat.log
-rw-----. 1 postgres postgres 58828 May 12 10:09 postgresql-Sun.log
-rw-----. 1 postgres postgres 58884 May 13 03:33 postgresql-Mon.log
[root@jsa-rhel-01 log]#
```

---





REDP-5730-00

ISBN DocISBN

Printed in U.S.A.

Get connected

