


Spectrum Protect Plus

Protecting Database Applications

Julien Sauvanet
Kenneth Salerno
Markus Fehling



Storage



IBM Redbooks

**Spectrum Protect Plus Protecting Database
Applications**

July 2021

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

First Edition (July 2021)

This edition applies to Version 10.1.8 of IBM Spectrum Protect Plus.

This document was created or updated on October 8, 2021.

© Copyright International Business Machines Corporation 2021. All rights reserved.

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

Contents

Notices	v
Trademarks	vi
Preface	vii
Authors	vii
Now you can become a published author, too!	viii
Comments welcome	viii
Stay connected to IBM Redbooks	viii
Chapter 1. Protecting database applications	1
1.1 Database application backup configuration basics	2
1.1.1 Creating an Identity	2
1.1.2 Adding an application server	3
1.1.3 Configuring the database application options	4
1.1.4 Assigning an SLA policy	6
1.2 IBM Spectrum Protect Plus database restore and data reuse	7
1.2.1 Test restore	8
1.2.2 Instant access	9
1.2.3 Production restore	10
1.3 Database protection and vSnap server operations	12
1.3.1 Backup operations	13
1.3.2 Restore operations	14
1.4 Database backup with pre-script and post-script	15
Chapter 2. Protecting Oracle database applications	21
2.1 IBM Spectrum Protect Plus requirements for Oracle	22
2.1.1 Oracle features related to backup	22
2.1.2 Server registration	23
2.1.3 Oracle log backup	26
2.1.4 Backup details	28
2.1.5 Restore details	30
2.1.6 Troubleshooting hints	36
Chapter 3. Backing up and restoring MongoDB databases	39
3.1 IBM Spectrum Protect Plus requirements for MongoDB	40
3.1.1 Fundamental IBM Spectrum Protect Plus requirements for MongoDB	40
3.1.2 MongoDB databases without authentication	41
3.1.3 MongoDB databases with authentication enabled	41
3.1.4 Register a MongoDB server	43
3.2 MongoDB backup and restore with Spectrum Protect Plus	44
3.2.1 MongoDB backup	44
3.2.2 MongoDB restore	46
Chapter 4. Backing up and restoring Db2 databases	51
4.1 IBM Spectrum Protect Plus Db2 features	52
4.2 Prerequisites for Db2 databases	52
4.3 Protecting Db2 databases	55
4.3.1 Registering the Db2 database server	55
4.3.2 Backup Db2 data	56

4.3.3 Restoring Db2 databases	61
Chapter 5. Backing up and restoring SQL Server	69
5.1 IBM Spectrum Protect Plus SQL Server features	70
5.2 Prerequisites for SQL Server databases	70
5.3 Protecting SQL Server databases	72
5.3.1 Register the SQL Server	72
5.3.2 Defining an SQL Server backup job	74
5.3.3 SQL database backups logs	77
5.3.4 vSnap commands used to manage SQL database backups logs	80
5.3.5 Parallel ad-hoc SQL database backups	82
5.3.6 SQL Server global preferences	83
5.4 Restoring SQL Server databases	84
Chapter 6. Backing up and restoring Microsoft Exchange data	91
6.1 Microsoft Exchange server	92
6.1.1 Server roles	92
6.1.2 Stand-alone or availability group databases	92
6.1.3 Mailbox movement	93
6.1.4 Microsoft built-in data loss prevention	93
6.2 Prerequisites for protection in IBM Spectrum Protect Plus	94
6.2.1 Granular restore remote package installation	94
6.3 IBM Spectrum Protect Plus configuration for Exchange	101
6.3.1 Log backup	102
6.3.2 Database Availability Groups	103
6.4 Backup jobs overview	103
6.4.1 Assigning an SLA policy	103
6.4.2 Backup types	103
6.4.3 Scheduled backup	104
6.4.4 Ad hoc backup	108
6.5 Restore jobs	111
6.5.1 Complete Restore	114
6.5.2 Restoring individual items with granular restore	117
Related publications	127
IBM Redbooks	127
Online resources	127
Help from IBM	127

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 <http://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®
Db2®
DB2®

IBM®
IBM Services®
IBM Spectrum®

Redbooks®
Redbooks (logo) ®
Tivoli®

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.

OpenShift, 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

IBM® Spectrum Protect Plus is a data protection solution that provides near-instant recovery, replication, retention management, and reuse for virtual machines, databases, and application backups in hybrid multicloud environments.

This IBM Redpaper publication focuses on protecting database applications. IBM Spectrum® Protect Plus supports backup, restore, and data reuse for multiple databases, such as Oracle, IBM Db2®, MongoDB, Microsoft Exchange, and Microsoft SQL Server. Although other IBM Spectrum Protect Plus features focus on virtual environments, the database and application support of IBM Spectrum Protect Plus includes databases on virtual *and* physical servers.

Authors

This paper was produced by a team of specialists from around the world.



Julien Sauvanet Bert Dufrasne is a Open Group Certified Expert IT Specialist, working at IBM for more than 15 years. He spent more than 10 years working in IBM Services®, helping with Spectrum Protect design, deployment as well as helping clients to solve their challenges around resiliency. Julien focuses on helping customer with their Spectrum Protect/Plus deployments, working as a Technical Advisor in the IBM Systems organization. He has co-authored several IBM Redbooks® publications.



Kenneth Salerno is an Open Group Certified Distinguished Technical Specialist working for IBM Systems in the USA. He has 23 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 and enterprise data center operations. He has contributed code to various open source projects, and is also Linux and Cisco certified.



Markus Fehling joined IBM more than 30 years ago as Dipl. Ing. in Electrical Engineering as a software developer for Hard Disk production, then became developer for Tivoli® Storage Manager. In 2005, he became SAP relationship manager. Markus is an IBM Certified IT specialist, with expertise in AIX®, databases, VMware and Spectrum Protect Plus. Markus is now part of the Storage sales team at the European Storage Competence Center, in IBM Germany, focusing on Spectrum Protect Plus and Red Hat OpenShift.

Thanks to the following people for their contributions to this project:

Bert Dufrasne
IBM Redbooks, San Jose Center

Siddharth Bhatt
Dominic Mueller
Jim Smith
Joerg Walter
Axel Westphal
IBM

Now you can become a published author, too!

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

Find out more about the residency program, browse the residency index, and apply online at:
ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

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

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- Stay current on recent Redbooks publications with RSS Feeds:

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



Protecting database applications

IBM Spectrum Protect Plus supports backup, restore, and data reuse for multiple databases, such as Oracle, IBM Db2, MongoDB, Microsoft Exchange, and Microsoft SQL Server.

Most of the configuration information included in this paper applies to all database applications.

This chapter includes the following topics:

- ▶ Database application backup configuration basics
- ▶ IBM Spectrum Protect Plus database restore and data reuse
- ▶ Database protection and vSnap server operations
- ▶ Database backup with pre-script and post-script

Note: IBM Spectrum Protect Plus offers data reuse functions in addition to backup and restore. You can use the database backup data to create a permanent copy (or clone) of your production database, or to temporarily establish a database copy directly from the vSnap server volumes.

1.1 Database application backup configuration basics

This section describes how to configure and run a database backup in IBM Spectrum Protect Plus, and how to schedule a job to regularly backup the database transaction logs.

Backup, restore, and data reuse handling functions for the supported (relational) databases are all similar in IBM Spectrum Protect Plus.

These are the steps which we cover in this chapter, enabling application's backup:

- ▶ Optional but considered as good practice: create an identity
- ▶ Register the database application
- ▶ Perform an inventory of the application (the system automatically triggers an inventory following the registration process)
- ▶ Assign the application to an SLA

Note: The SLA must have been previously created and must indicate, among other options, the Data backup frequency (**not** the **log backup** frequency) and associated retention (for both data and logs).

- ▶ Configure the application database Options

Note: At this step, you will have to specify how you want the application log to be handled.

- ▶ Start the backup manually or wait for the next schedule as per the SLA instructions

Note: Do not assign more than one application per machine as an application server to a resource group. For example, if Microsoft SQL Server and Microsoft Exchange Server occupy the same machine and both are registered with IBM Spectrum Protect Plus, only one of the applications can be added as an application server to a given resource group. As a reminder Resource Group is used for the Role Based Access Control (RBAC) definition.

For more information about supported databases and environments, see [IBM Spectrum Protect Plus Installation and User's Guide](#).

1.1.1 Creating an Identity

An operating system user is required to register a database application server, and discover the databases that exist on that server. You can enter the user ID and the server-specific data, such as IP name or IP address. However, we advise you to create so-called Identities (user definition entries) in advance, to maintain a customized ordering scheme.

In the IBM Spectrum Protect Plus GUI, select **Accounts** → **Identity** → **Add Identity** to enter the user definition for your specific databases, as shown in Figure 1-1 .

Tip: If you use identical operating system users and passwords for multiple database servers, IBM Spectrum Protect Plus allows you to manage these databases under one identity.

Identity

Identities

Identity Properties

Name: DB administrator - OS user

Username: dbadmin

Password: *****

Cancel Save

Figure 1-1 Creating an identity for an operating system user

With a list of identity entries, as shown in Figure 1-2 , you can see that there are default system identities (such as *serveradmin*) and identities explicitly created for database backup and restore (DB administrator, operating system user, Mongo DB user, and Oracle DBA).

Name	
	DPR-OCP
	ISV-K8S
	ISV-OCP
	Oracle_10.0.240.209
	serveradmin

Figure 1-2 Spectrum Protect Plus Identity page

1.1.2 Adding an application server

In the IBM Spectrum Protect Plus GUI, select **Manage Protection** → **Databases** → **<your database type>**, as shown in Figure 1-3 , then, click **Manage application servers** and **Add application server** buttons.




Figure 1-3 Manage Protection Menu > Databases

Enter the database server's IP name or address, as shown in Figure 1-4 , then enter the database administration user or select an identity that you defined earlier. Click **Get databases** button to start a database discovery job on the server.

Figure 1-4 Manage Application, step to register an application database to SPP

If you save this application server entry, IBM Spectrum Protect Plus automatically starts an inventory job. This job confirms a network connection, adds the application server to the IBM Spectrum Protect Plus database, and then catalogs the instance. You can follow the inventory job or its status in the **Jobs and Operations** menu.

1.1.3 Configuring the database application options

A complete database backup includes the data files, metadata (such as database control files), and the transaction logs. While an IBM Spectrum Protect Plus database backup includes data and metafiles, transactions logs must be backed up more frequently to enable a future database roll forward to a given current point in time. In addition to the backup schedule, IBM Spectrum Protect Plus allows you to automatically create a cron job that regularly starts a transaction log backup.

After you registered the application, an important step is to specify the options. This is where you define how you want the logs to be managed for the database. This is not controlling the retention but the frequency of the log backup as well as whether SPP will delete them from the local system or not, and if yes, when.

Application Options is also where you specify how many streams will be used to perform the database backup, for the applications which are supporting this feature (Oracle is one of those). When enabling multiple streams for database backup, consider the performance impact it can cause to the application server. Moreover, the parallel stream is working at the datafile level, so consider the number of datafiles to set the number of parallel streams. Tuning this performance setting is likely to be an iterative process to find the right spot specific to your environment.

Figure 1-5 shows the possible Application Options. By checking the **Enable Log backup**, additional options are proposed, to let you decide how log backup has to occur.

The screenshot displays the 'Application Database options selection' interface. At the top, there's a breadcrumb trail: 'Instances / its0-oracle / OraDB12Home1'. A status bar indicates 'Last inventory completed May 14, 2021 10:01:53 AM' and a 'Run Inventory' button. Below this is a table with columns: 'Name', 'SLA Policy', and 'Eligible for Log Backup'. Two applications are listed: 'ORATST' and 'SPP', both with 'Oracle' as the SLA Policy and 'Yes' for log backup eligibility. The 'ORATST' checkbox is checked and circled in red. Below the table, there's a 'Clear Selections (1)' button. To the right, there are buttons for 'Run', 'Select an SLA policy', and 'Select Options'. A red arrow points to the 'Select Options' button. The 'Options' section is expanded, showing several settings: 'Enable Log Back up' (checked and circled in red), 'Log Backup Frequency and Start Time' (set to 'Subhourly', '15' minutes, '05/12/2021', '00:00', 'Europe/Paris'), 'Truncate source logs after successful backup' (set to 'Older than a specified number of hours'), 'Primary log retention in hours' (set to '12'), and 'Maximum Parallel Streams per Database' (set to '3'). A 'Save' button is at the bottom left.

Figure 1-5 Application Database options selection - for log backup management & multi stream backup for supported databases

Starting with Spectrum Protect Plus 10.1.8, the log backup can be done below the hour frequency (SubHourly option).

Starting with Spectrum Protect Plus 10.1.8, log backups can be done at frequencies less than an hour (SubHourly option). In addition, a log can be deleted days or hours, or never, from the production system following a Spectrum Protect Plus log backup. Possible options are shown in Figure 1-6.

Truncate source logs after successful backup	Older than a specified number of hours
Primary log retention in hours	Never
	Older than a specified number of days
	Older than a specified number of hours
	Immediately after log back up

Figure 1-6 How applications log are handled on the production database environment

1.1.4 Assigning an SLA policy

After a database instance is defined in IBM Spectrum Protect Plus, assign an SLA policy to that instance. In general we recommend that you create dedicated SLA policies for any single database, or for groups of logically related databases.

To assign an SLA, select the application resource and use the **Select an SLA policy** button. You can do it as host level for all databases or at instance level.

Following the SLA assignment, your application or database is ready for backup. You can decide to let the backup start as per the SLA schedule or manually trigger backup immediately. In the latter, from the Application menu, select the application resource and click the **Run** button, as shown in Figure 1-7 .

The screenshot shows the 'Oracle Backup' section of the IBM Spectrum Protect Plus interface. The left sidebar contains a navigation menu with options like Dashboard, Jobs and Operations, Manage Protection, Policy Overview, File Restore, Virtualized Systems, File Systems, Containers, Cloud Management, Databases, Db2, Exchange, MongoDB, and Oracle. The main panel is titled 'Oracle' and includes a 'Manage application servers' button and a 'Create job' button. Below this is a search bar and a table of instances. The table has columns for Name, Location, SLA Policy, and Eligible for Log Backup. Two instances are listed: 'ORATST' and 'SPP'. The 'ORATST' instance is selected, and a red arrow points to the 'Run' button at the bottom right of the panel.

Name	Location	SLA Policy	Eligible for Log Backup
ORATST	itso-oracle / OraDB12Home1	Oracle	Yes
SPP	itso-oracle / OraDB12Home1	Oracle	Yes

Figure 1-7 Start a manual backup for an application, after registration or at any later time

You can track the progress of a running backup from the Jobs & Operations menu.

Your application is now registered and is part of the Spectrum Protect Plus backup plan, according to the assigned SLA.

1.2 IBM Spectrum Protect Plus database restore and data reuse

IBM Spectrum Protect Plus features a restore wizard that simplifies restore operations for virtual machines and databases. The wizard guides you through the configuration of restore types and parameters, and optionally schedules a job that performs the actual restore.

IBM Spectrum Protect Plus treats data reuse and data recovery as a restore activity. In either case, you must create a restore job. The Databases and the Jobs and Operations menus in IBM Spectrum Protect Plus have a button that is used to start creating a restore job. The parameters that you select during job creation define which activity is performed.

The following list describes the parameters that control the final restore or data reuse activity:

- ▶ Type of Restore:
 - On-Demand Snapshot: one-time restore operation (you choose from On-Demand Point in Time: one-time restore by selecting a point-in-time backup of that database (you specify a point in time or a transaction number)
 - Recurring: repeating point-in-time restore job that runs on schedule
- ▶ Restore Method:
 - A production restore overwrites the original database or creates a database copy with a different database name. In the database copy case, you must specify a new database name and the destination paths.
 - A test restore mounts the vSnap server directories with a database backup to a database server, recovers and opens the database. You can rename the database.
 - An instant access restore also mounts the vSnap server directories with a database backup to a database server, but does not recover or open the database.
- ▶ Destination:
 - Restore to the original instance
 - Restore to an alternative instance

The combination of these selections define which action to perform:

- ▶ Restore a database and optionally overwrite an existing database
- ▶ Establish a copy of a previously backed up database (DevOps)
- ▶ Get access to the database files (data and metadata) of a previous backup and more

The following sections describe examples for these use cases. The sample database is Oracle database 12c.

Note: Consider the following points:

- ▶ For test or instant access restores, IBM Spectrum Protect Plus creates an internal snapshot on the vSnap server to prevent any change to the database backup data. The snapshot directory is then mounted to the selected database server.
- ▶ If you decide to open the database after a restore, you can choose the point-in-time for a database roll forward: either a specific point-in-time or end of backup.

Before you start to create a restore job, you must first select the database and an associated backup to restore, as shown in Figure 1-8 .

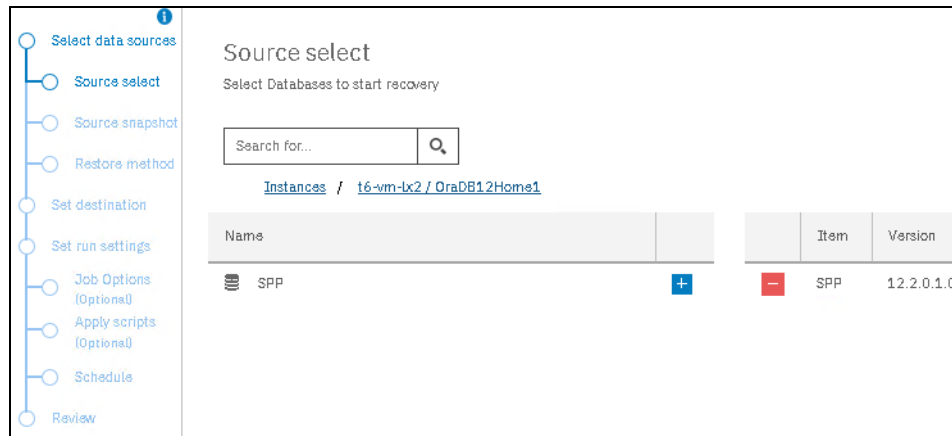


Figure 1-8 Select the source for a database restore

1.2.1 Test restore

This section describes a possible test restore use case: Every Monday morning, a development tester or a particular application user requires a fresh copy of a production database for testing a DevOps scenario.

In the IBM Spectrum Protect Plus restore wizard, you can set up such a user requirement by choosing the following parameter settings:

- ▶ Restore type: On-demand point in time (or On-demand snapshot, depending on the available backups)
- ▶ Restore method: Test
- ▶ Destination: Original or alternative instance

First, select the database instance and an associated database backup, as shown in Figure 1-8 . In addition, select a site and a location for the instance to restore. These settings depend on your specific environment, which can include a cloud or copy location, or a secondary site that you use for replication.

In our example, we chose the primary site and we chose **Recurring** to create a repeating restore job that runs on a schedule, as shown in Figure 1-9 .




Figure 1-9 Selecting a site to create the new database instance

For our use case, we decided to create the test database in an alternative destination (which means not on the original production server) and give the database a new name. Figure 1-10 and Figure 1-11 show the corresponding parameter selections.

Restore method

Select the restore method to be used for the source selections.

☐ Instant Access
 ☐ Production
 ☒ Test

Name	New Database Name
SPP	TQQI

Figure 1-10 Selecting the Test restore method

Set destination

Select which instance you would like to restore for each chosen source.

☐ Restore to original instance
 ☒ Restore to alternate instance

Some instances/groups may be disabled for selection due to version incompatibility.

Instances

	Name	Version
<input checked="" type="radio"/>	t6-vm-lx / OraDB12Home1	12.2.0.1.0
<input type="radio"/>	t6-vm-lx2 / OraDB12Home1	12.2.0.1.0

Figure 1-11 Selecting a restore to an alternative instance

Finally, we define the schedule: every Monday at 8:00 o'clock. Figure Example 1-12 shows an example. If you do not want to wait for the first test run, you can find the scheduled job in the **Jobs and Operations** menu of IBM Spectrum Protect Plus. Select the **Schedule** tab, find the job in the list, and start it manually.

Schedule

Run the restore job now or schedule the time for the restore to run.

DB_test_setup_Monday_at_0800

Frequency: 7 Days

Start Time: 07/01/2019 00

Figure 1-12 Weekly schedule for a database copy

IBM Spectrum Protect Plus does not reflect the new database name in the name of the mounted directory or in the data file names, but it starts the database with the new database identifier (System ID, SID).

The IBM Spectrum Protect Plus test restore job that you started stays active until you manually terminate it. In the **Job and Operations** menu, the job status is shown as “Resource active”. To terminate the job, select it and choose **End instant disk restore**.

1.2.2 Instant access

In Instant access restore mode, IBM Spectrum Protect Plus mounts the volume from the vSnap server repository. The instant access restore wizard is similar to the test restore one, described in “Test restore” on page 8.

In comparison to the test restore, an instant access restore job does not start a database; therefore, you do not need to select a database instance as a restore target.

From the mounted file system you can use the data for custom recovery; for example:

- ▶ Reload individual files such as control files, configuration files, and data files.
- ▶ Rebuild a customized database copy.

As described in 1.2.1, “Test restore”, the instant access job remains active with the “Resource active” state until you terminate it manually. To terminate the job, select it and choose the **End instant disk restore** action.

1.2.3 Production restore

This section describes a common use case for a backup and restore solution: a traditional database restore that overwrites the original database and optionally rolls forward the database to a specific point-in-time.

In the IBM Spectrum Protect Plus restore wizard, the following parameters initiate a traditional database restore that overwrites the existing database:

- ▶ Restore type: On-Demand Point in Time
- ▶ Restore method: Production
- ▶ Destination: Original Instance

First, select the database instance and an associated database backup, as shown in Figure 1-8 . In addition, select a site and a location for the instance to restore. These settings depend on your specific environment, which can include a cloud or copy location, or a secondary site that you use for replication.

In our case, we chose the primary site, as shown in Figure 1-13 .

Source snapshot

Select which instance you would like to restore for each chosen source.

Recurring Site Primary

Item	Version
SPP	12.2.0.1.0

Figure 1-13 Select a site to restore the database

The next two selections indicate what we are trying to achieve; that is, production restore to the original instance, as shown in Figure 1-14 and Figure 1-15 .

Restore method

Select the restore method to be used for the source selections.

☐ Instant Access ☒ Production ☐ Test

Name	New Database Name
SPP	

Figure 1-14 Production restore to the original database

Set destination
Select which instance you would like to restore for each chosen source.

☒ Restore to original instance ☐ Restore to alternate instance

Figure 1-15 Database restore to the original instance

For a restore of a production database, the IBM Spectrum Protect Plus restore wizard assumes a database roll forward to a specific point in time that you can configure in the next menu.

You must also decide whether to overwrite an existing database, as shown in Figure 1-16 . IBM Spectrum Protect Plus provides an auxiliary protection against an unintended data overwrite: If the database still exists and you do not select the overwrite choice box, the restore job fails.

Job Options
Configure the options for this restore job.

Recovery Options

☐ Recover until end of backup ☒ Recover until specific point-in-time

☒ By Time ☐ By SCN

Jun 28, 2019 2:46:00 PM

Application Options

☒ Overwrite existing databases.

Maximum Parallel Streams per Database

Init Params

Figure 1-16 Selecting the database roll forward and overwrite options

Carefully review the job summary that IBM Spectrum Protect Plus displays (Figure 1-17) and, if the information describes what you are trying to achieve, run the restore job.

[Back to Jobs and Operations](#)

Snapshot restore

Review

1

Select data sources

✓

Source type

✓

Source select

✓

Source snapshot

✓

Restore method

✓

Set destination

✓

Set run settings

✓

Job Options (Optional)

✓

Apply scripts (Optional)

✓

Schedule

○

Review

Review

Review your selections before submitting.

Select data sources

Source type:	Oracle
Source select:	SPP
Source snapshot:	SPP - Use Latest
Restore Type:	On-Demand: Point In Time
Restore Source Type:	Site
Restore Source:	Primary
Restore method:	Production

Set destination

Destination:	Restore To Original Instance
--------------	------------------------------

Set run settings

Run cleanup immediately on job failure: Yes

Allow session overwrite: Yes

Continue with restores of other selected databases even if one fails: Yes

Overwrite existing databases: Yes

Recovery type: Pitrecovery

Maximum Parallel Streams per Database: 1

Init Params: Source

Figure 1-17 Restore job summary

The IBM Spectrum Protect Plus restore job verifies whether a database exists or is even up and running. In this case, the Allow database overwrite setting is relevant, as shown in the restore job log shown in Example 1-1.

Example 1-1 Restore job log

```
[t6-vm-lx2] SPP: Another DB with the same name is already running. Proceeding because the overwrite option is enabled.
```

1.3 Database protection and vSnap server operations

This section explains how the database backups are organized inside the vSnap server, how backed up data is stored in the vSnap server, and the different ways of restoring from these backups.

The intent here is to provide technical details about how the process works. For more information about how to configure backup and run recovery operations for specific databases with IBM Spectrum Protect Plus, see the following resources:

- Chapter 2, “Protecting Oracle database applications” on page 21

- ▶ Chapter 3, “Backing up and restoring MongoDB databases” on page 39
- ▶ Chapter 4, “Backing up and restoring Db2 databases” on page 51
- ▶ Chapter 5, “Backing up and restoring SQL Server” on page 69
- ▶ Chapter 6, “Backing up and restoring Microsoft Exchange data” on page 91

At the time of this writing, the following databases are supported:

- ▶ Oracle
- ▶ Microsoft SQL
- ▶ Microsoft Exchange
- ▶ IBM DB2®
- ▶ MongoDB

1.3.1 Backup operations

A complete database backup includes the data files, metadata (such as database control files), and the transaction logs. While an IBM Spectrum Protect Plus database backup includes data and metafiles, transactions logs must be backed up more frequently to enable a future database roll forward to a given current point in time. In addition to the backup schedule, IBM Spectrum Protect Plus allows you to automatically create a cron job that regularly starts a transaction log backup.

For the database that includes log backup enabled, two volumes are created:

- ▶ one volume for the data
- ▶ one volume for the log management - which remains mounted at any time to the target application server

MongoDB: Unlike the other supported databases, MongoDB does not have log management configurable in IBM Spectrum Protect Plus.

The SLA policy that you assigned to your database defines the schedule time for the first backup. If you did not define a schedule, or you do not want to wait for the first automatic backup schedule, click the **Run** button or scroll down to the SLA policy that you provided for this database and select the **Actions** button to start the database backup.

At this point, you also decide to perform a backup of a single database by clicking the **Run** button) or a backup of all applications included in the SLA policy by clicking the **Actions** button (see Figure 1-5).

For each SLA and each database included in that SLA, the following resources are created on the vSnap server to store backups:

- ▶ One primary data vSnap server volume
- ▶ One primary log vSnap server volume

Optionally, the following resources are created:

- ▶ Replication vSnap server volumes (if vSnap server replication is enabled for that SLA)
- ▶ Copy data volume (S3 or Repository Server if copies are enabled for that SLA)
- ▶ Archive data volume (S3 cold, tape, if archive is enabled for that SLA)

These data and log volumes are made available to the database server, using different access method depending on the platform and application, as shown in Figure 1-18 .

Note: Log volumes feature the main characteristics:

- ▶ Log volumes stay mounted.
- ▶ Log volumes are not copied or archived to any S3-mode repositories

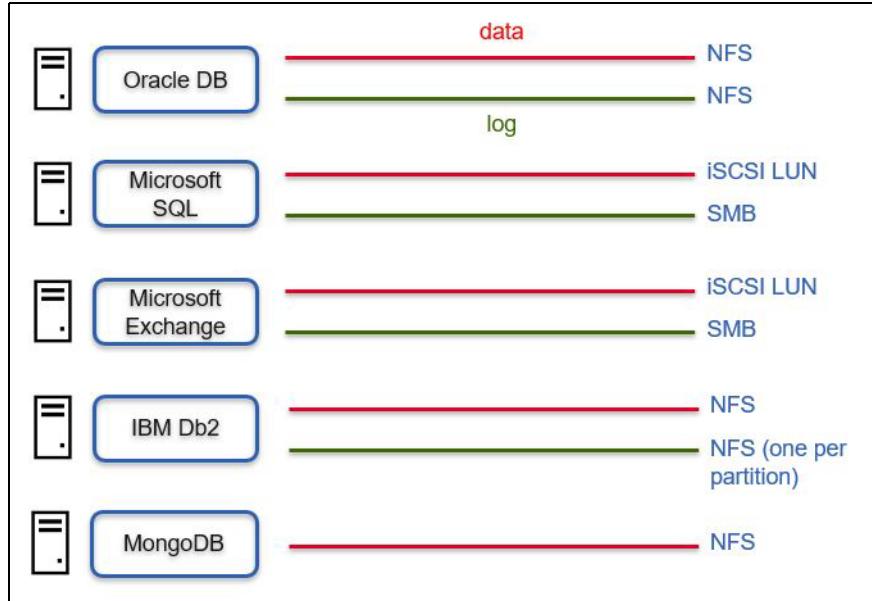


Figure 1-18 How database server access their respective vSnap server volumes

Table 1-1 lists what actions take place for data backup and describes the mechanism that is used by IBM Spectrum Protect Plus to handle the application logs backup.

Table 1-1 Application Data and log backup

Application	Data backup	Log backup trigger
Oracle	RMAN inconsistent incremental level0 for the first backup, level1 for the subsequent, copy all file to share	cron job
Microsoft SQL	VSS snapshot, copy all files to iSCSI LUN	Windows task scheduler
Microsoft Exchange	VSS snapshot, copy all files to iSCSI LUN	Windows task scheduler
IBM Db2	LVM snapshot, copy all files to share	db2 archive log scheduler
MongoDB	LVM snapshot, copy all files to share	Copy journal with data file

1.3.2 Restore operations

The following restore operations are available:

- ▶ Production restore: Replace application data
- ▶ Test restore: Clone new instance of production
- ▶ Instant Access: Access backup data

A point-in-time backup is represented by a volume snapshot in the vSnap server.

A point-in-time Production restore creates a temporary vSnap server clone volume of the last vSnap server data volume before the selected point in time and mounts that clone to the target server. A copy then occurs on the target server, from the clone volume to the production volume. After the copy process completes, the clone volume is dismounted and deleted from the vSnap server.

The next step is to create a temporary clone of the log volume that contains database logs that are created after the selected point in time and mount that clone to the target. This clone contains the log backup with database transactions that occurred after the data was restored in the first step, and allows a roll forward recovery until the specified point in time.

The Test restore works the same way as point in time production restore, but the production data is not copied back. The restore data is provided as a share from the vSnap server.

The Instant Access restore creates a temporary vSnap server clone volume of the selected (point in time) backed up data and mounts that clone to the target application server for access. The same clone and mount operation occurs on the log volume of that same point in time.

These clones allow read/write access, so the application can work with the data. However, when the instant access process completes, the data modifications are not persistent, and any modifications that were made during the instant access are lost. The original backup does not change.

1.4 Database backup with pre-script and post-script

Backup pre-script and post-script are features of IBM Spectrum Protect Plus that can be used and configured as part of an SLA policy options. SLA policy options are unique to an SLA policy. Therefore, if you decide to use SLA options to set pre or post-script actions, any workload that is assigned to that SLA policy uses the same options, including the execution of the script on the configured script server.

Several environments exist in which the database or other application requires their backup to be integrated into a sequence of actions, in a specific order. These tasks are usually managed by a scheduler, such as Tivoli Workload Scheduler.

The intent of this section is to explain which steps and features can be used within IBM Spectrum Protect Plus to synchronize the backup with an external scheduler. Although the backup is still triggered by an IBM Spectrum Protect Plus policy, we can use the pre-script option to have IBM Spectrum Protect Plus run a piece of code on the server host. In our example, we show how to wait for an external scheduler flag before triggering the backup.

Example 1-2 shows the few lines of code that we use as a pre-script. This code is going to loop until a specific file (acting as a flag) is placed in a specific location by an external mechanism. This flag file is the signal for IBM Spectrum Protect Plus to trigger the backup.

Example 1-2 Pre-script to make an IBM Spectrum Protect Plus backup waiting for external signal

```
#!/bin/sh
while [ ! -f /tmp/external_scheduler.flag ]
do
    date >> /tmp/SPP_wait4external_scheduler.log
```

```
echo " Flag /tmp/external_scheduler.flag not there, wait to start backup " >>
/tmp/SPP_wait4external_scheduler.log
sleep 60
done
echo " Flag /tmp/external_scheduler.flag is here, backup time ! " >>
/tmp/SPP_wait4external_scheduler.log
```

The script is named SPPWait4externalscheduler.sh.

Tips: Supported scripts include shell scripts for Linux-based machines and batch and PowerShell scripts for Windows-based machines. Scripts must be created by using the associated file format for the operating system.

Running the **dos2unix** command before uploading the script in IBM Spectrum Protect Plus might help you to ensure suitable format of a shell script if you encounter a formatting problem (that is, ^M at end of line).

The use of the scripts with IBM Spectrum Protect Plus is an easy three-step process:

1. Define the script.
2. Define the script servers.

Update the SLA policy options from the SLA Policy Status page. The first step is to define the script. From the IBM Spectrum Protect Plus GUI, select **System Configuration** → **Script** to define the script.

Defining the script means uploading the script to IBM Spectrum Protect Plus, as shown in Figure 1-19 . Click **Browse** to select and upload the script that you plan to run on the application or database server as part of the backup job.

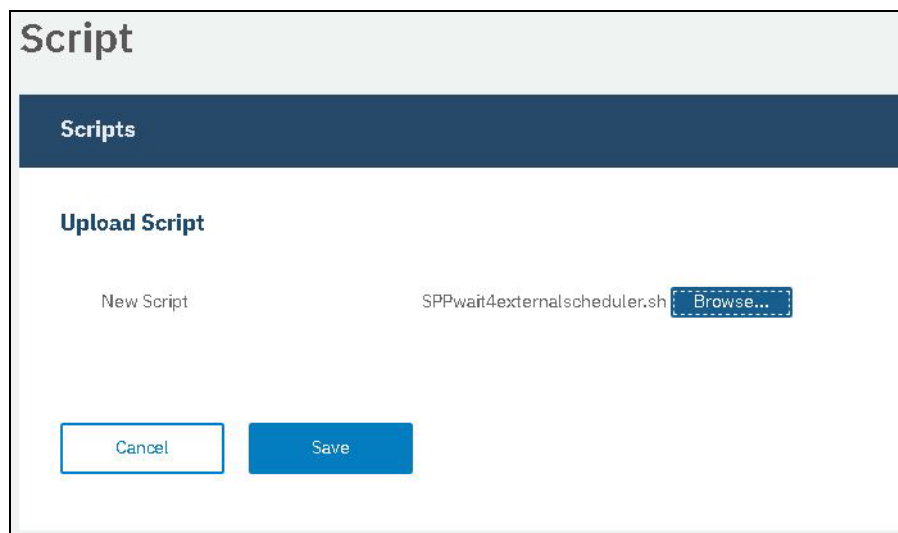
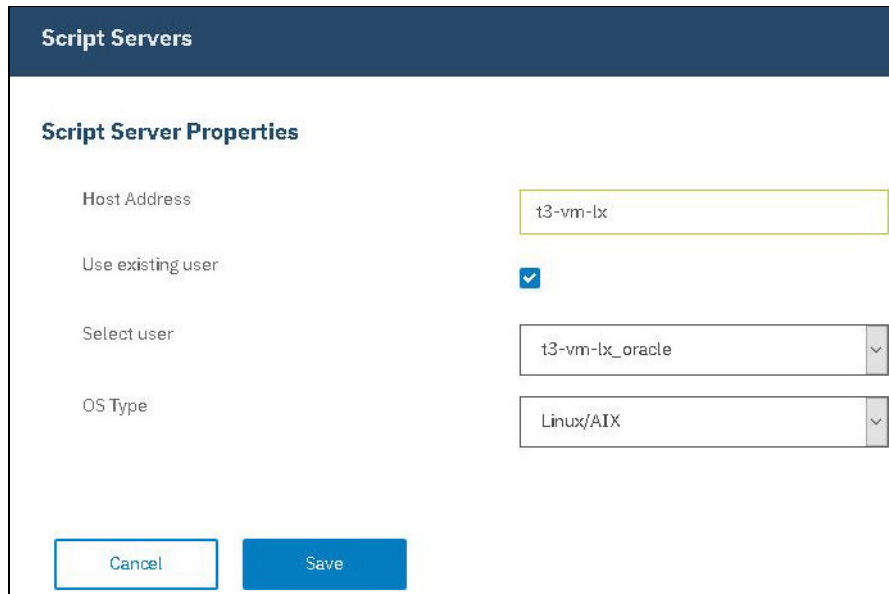


Figure 1-19 Define script by uploading it

The second step is to define the Script Server, as shown in Figure 1-20 . Select **System Configuration** → **Script**. Specify the Host address, login credentials, and operating system type for which you plan to use script.



Script Servers

Script Server Properties

Host Address: t3-vm-lx

Use existing user: ☒

Select user: t3-vm-lx_oracle

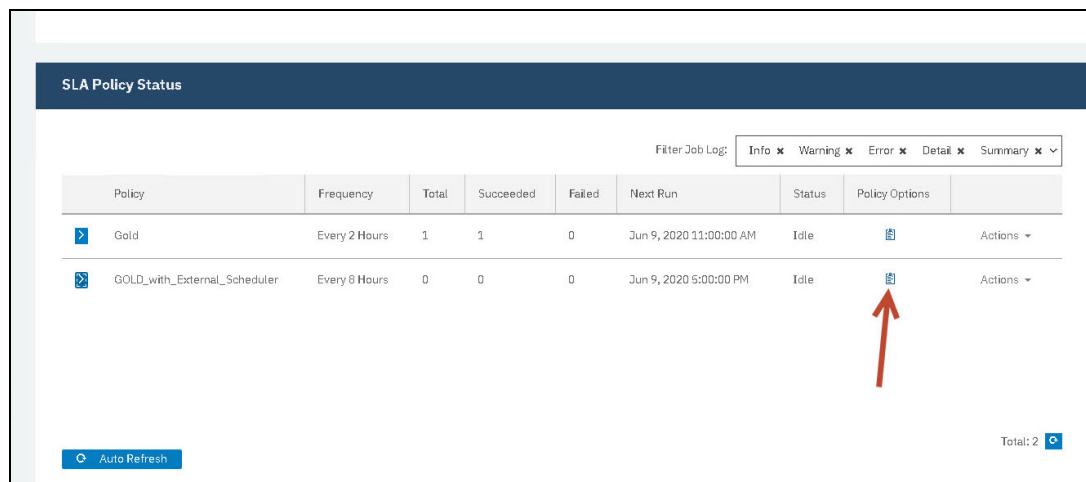
OS Type: Linux/AIX

Buttons: Cancel, Save

Figure 1-20 Define SCript Server by specifying its address and credentials



The third step is to update an SLA Policy or create an SLA Policy and enable the use of pre-script in the policy options. Such an option can be accessed by editing the SLA Policy Options by selecting **Manage Protection** → **<Backup Workload of your choice>** → **SLA Policy Status**.

In our example, we want to run a script to check whether an Oracle backup ran. We select **Manage Protection** → **Databases** → **Oracle** → **SLA Policy Status**. Then, we select **Policy Options**, as shown in Figure 1-21 .



SLA Policy Status

Filter Job Log: Info x Warning x Error x Detail x Summary x

Policy	Frequency	Total	Succeeded	Failed	Next Run	Status	Policy Options	
Gold	Every 2 Hours	1	1	0	Jun 9, 2020 11:00:00 AM	Idle		Actions
GOLD_with_External_Scheduler	Every 8 Hours	0	0	0	Jun 9, 2020 5:00:00 PM	Idle		Actions

Buttons: Auto Refresh, Total: 2

Figure 1-21 Open the Policy Options to enable the use of a script for that SLA Policy run

When you click **Policy Options**, a pop-up window opens (see Figure 1-22) and you must specify whether you want to enable pre-script or post-script, and which script to run.

Configure Options

☒ Pre-Script

☒ Use Script Server Script

Script Server: t3-vm-lx

☐ Post-Script

☐ Continue job/task on script error

Exclude Resources: e.g. dbname {match} dbname* {wildcard}

Force full backup of resources: e.g. dbname {match} ";" separated

* Forcing a full backup of a resource, runs a once-off, new full base backup of that resource.

Save

Figure 1-22 Configure Policy Options to use a specific script for pre or post backup tasks

Note: Consider the following points:

- ▶ Configuring options for an SLA policy means that any application that is associated with that SLA policy runs with the specific options. In the case of pre-script and post-script, they are triggered on the configured script server, although it might not be the server that is running the backup.
- ▶ In the Policy Options configuration menu, when enabling the use of a script, you can also instruct IBM Spectrum Protect Plus what to do if the script is failing by selecting the **Continue job/task on script error** option.

Whenever the SLA policy is used, it completes the pre-script step before triggering the backup commands. Moreover, if you disable the Continue job/task on script error option, the backup does not run if the pre-script failed.

Figure 1-23 shows the output log that lists the execution of pre-script.

SLA Policy Status			
Filter Job Log: Info x Warning x Error x Detail x Summary x v			
Status	Time ▲	ID	Description
	9:54:19 AM		
Info	Jun 9, 2020 9:54:19 AM	CTGGA1655	This job contains at least two types of tasks: job resolution and protection.
Detail	Jun 9, 2020 9:54:20 AM	CTGGA2434	----- Starting execution task 1 for command preScript on Guest. -----
Detail	Jun 9, 2020 9:54:20 AM	CTGGA2427	Running SPPwait4externalscheduler.sh on Script server t3-vm-lx
Detail	Jun 9, 2020 9:54:20 AM	CTGGA2245	SPP log dir: /data/log/guestdeployer/2020-06-09/2d0b16db-0170-41c2-9c03-7a99f0e3752c

Figure 1-23 Job log showing that pre script action is executed on the target server t3-vm-lx

Figure 1-24 shows the output log when the pre-script completed so the backup job continues and triggers the backup of (in our example) the Oracle database.

SLA Policy Status			
Filter Job Log: Info x Warning x Error x Summary x Detail x v			
Status	Time ▲	ID	Description
Detail	Jun 9, 2020 9:54:53 AM	CTGGA2642	[t3-vm-lx] Script execution completed successfully.
Detail	Jun 9, 2020 9:54:54 AM	CTGGA2435	----- Execution task 1 for command preScript on Guest completed in 34 sec(s) with RC=0 and status COMPLETED. -----
Info	Jun 9, 2020 9:54:57 AM	CTGGA2054	Performing discovery on application server t3-vm-lx
Detail	Jun 9, 2020 9:54:58 AM	CTGGA2245	SPP log dir: /data/log/guestdeployer/2020-06-09/1591689258176/23f43a5c-0a2b-4758-87e3-139c74a433fd

Figure 1-24 Job log showing that backup action is happening after the pre script completed successfully

Tip: For more information about pre-script or SLA policy options in the job log, enable the **Detail** filter of the Job Log.

In this chapter, we demonstrated the use of pre-script with an external scheduler example. However, be aware that schedule might also be triggered externally by using the IBM Spectrum Protect Plus REST API.



Protecting Oracle database applications

This chapter describes Spectrum Protect Plus backup and restore operations for Oracle databases.

This chapter includes the following topics:

- ▶ Basics of Oracle components used for backup
- ▶ Spectrum Protect Plus Oracle backup and restore Details
- ▶ Troubleshooting and logs

Tip: If you are running a standby Oracle database, you can register the standby database for backup instead of the production database to reduce the backup impact on business operations.

2.1 IBM Spectrum Protect Plus requirements for Oracle

This section describes specific IBM Spectrum Protect Plus requirements for Oracle databases. For up to date information related to prerequisites, check the Oracle requirements section of the *IBM Spectrum Protect Plus Installation and User's Guide*, which is available at: <https://www.ibm.com/docs/en/spp/10.1.8?topic=requirements-oracle>.

The steps to define an Oracle database for the backup are :

1. Optional but considered as good practice: create an identity
2. Register the Oracle database
3. Perform an inventory of the application
 - Optional, as it is automatically triggered following the registration process.
4. Assign the application to an SLA
 - The SLA must exist, and contains, amongst other option, the Data backup frequency (**not** the **log backup** frequency) and associated retention (for both data and logs).
5. Configure the application database Options
 - At this step, you will have to specify how you want the Oracle archive redo logs to be handled
6. Optional: Start manually the backup

Before jumping into the details of these steps, and the Spectrum Protect Plus capabilities, let's review some Oracle contextual information.

2.1.1 Oracle features related to backup

This section reviews a few Oracle related concepts that require consideration as they are impacting how Spectrum Protect Plus performs Oracle database backups.

RMAN

IBM Spectrum Protect Plus relies on Oracle RMAN (Recovery MANager) to perform the backups and recovery. Any backup related metadata generated by Spectrum Protect Plus are written to Oracle instance control files. Writing to a remote RMAN catalog is not supported at the time of this writing.

Spectrum Protect Plus builds its own backup catalog. It is stored inside the Spectrum Protect Plus virtual appliance and should be protected via Spectrum Protect Plus catalog backup. It is not possible to use an application specific backup for the catalog.

All commands used by Spectrum Protect can be found in the Job logs, which includes the `command.log` file. To find the RMAN command used, you can download the Job log (From **Jobs and Operations** → **Job History**, select the job you want and click **Download.zip** button).

Once you have extracted that zip file, browse the directory structure and look for the file named **command.log**. In this file, you will have all the commands used for the specific operation. Note that one Job log may have multiple logs (as many as sub-operations) and therefore multiple directories in this structure.

For example in Figure 2-1, you can view an extracted log from the Spectrum Protect Plus server, for an Oracle application backup job ID **1620121641156**, which ran against the Oracle server named 10.0.240.209 (IP address)

Name	Date modified	Type	Size
command.log	5/4/2021 11:58 AM	Text Document	1 KB
diag.tar.gz	5/4/2021 11:58 AM	GZ File	0 KB
input.json	5/4/2021 11:58 AM	JSON File	1 KB
output.json	5/4/2021 11:58 AM	JSON File	1 KB
status.log	5/4/2021 11:58 AM	Text Document	1 KB
version	5/4/2021 11:58 AM	File	1 KB

Figure 2-1 Extracted Job logs for RMAN command review

Oracle Block Change Tracking

IBM Spectrum Protect Plus uses Oracle Block Change Tracking to perform incremental backups.

Note: If Block Change Tracking is not enabled, IBM Spectrum Protect Plus enables it automatically during the first backup

Oracle Compression

IBM Spectrum Protect Plus uses its own compression and deduplication mechanisms. It does *not* use the Oracle Advanced Compression feature (which requires an extra license).

Oracle multi-threading

Oracle 12c introduced the concept of multi-threading. In IBM Spectrum Protect Plus, a multi-threaded database configuration requires Oracle credentials for backup processing.

The discovery process identifies if multi-threading processing is enabled, and prompts the user for the credentials. Enter the credentials for multi-threaded databases at the time of registration. IBM Spectrum Protect Plus passes on the credentials to the Oracle agent during backup, and the agent uses the credentials to log in to the database.

Note: When restoring an Oracle database that was configured for multi-threading at the time of backup, the restored database is non-multithreaded. The restored database must be manually reconfigured to use multi-threading.

2.1.2 Server registration

An Oracle server must be registered in IBM Spectrum Protect Plus by using an operating system user that exists on the Oracle server. This user must belong to the Oracle Inventory group (*oinstall*), have database administration permissions OSDBA (dba) and OSASM (asmadmin), and sudo permissions to perform system operations.

For an Oracle environment it is important that the user used for the registration process has proper Oracle permissions as well as sudo permissions. Before registering the database on the Spectrum Protect Plus server, on the system you want to protect, create a user and configure the proper sudo permissions as shown in Example 2-1.

Example 2-1 Create sppagent user for Oracle server registration on Spectrum Protect Plus

```
itso-oracle:~ # useradd -m -G oinstall,dba,asmadmin sppagent
itso-oracle:~ # passwd sppagent
New password:
Retype new password:
passwd: password updated successfully

itso-oracle:~ # cat /etc/sudoers
...
Defaults:sppagent !requiretty
Defaults:sppagent env_keep+="ORACLE_HOME"
Defaults:sppagent env_keep+="ORACLE_SID"
sppagent ALL=(ALL) NOPASSWD:ALL
```

To reflect this configuration in Spectrum Protect Plus, you can create a Spectrum Protect Plus identity pointing to the user you just created on the machine you need to protect. Follow the generic steps to register a database application as described in 1.1, “Database application backup configuration basics” on page 2.

Specifically, to register an Oracle database, use the **Manage Protect** → **Databases** → **Oracle** page, click the **Manage applications servers** button and enter IP and credentials.

Note that when registering Oracle Real Application Cluster (RAC) nodes, each node must be registered using its physical name or address, **not** the virtual or SCAN address.

When Oracle multithread is enabled, the **Set Credential** button appears at a database instance level in the second registration step, as shown in Figure 2-2. You must specify the credentials of an Oracle database user who has SYSDBA privileges. You can specify the same credentials as those specified to register the application if the privileges associated meet the SYSDBA requirement.

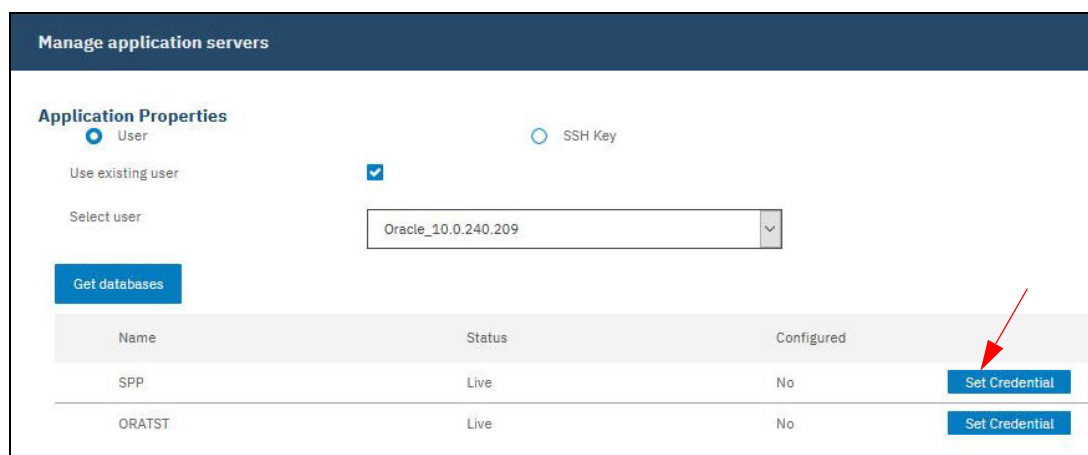


Figure 2-2 Oracle registration in an Oracle multi-thread configuration.

Application **Inventory** process, which is triggered right after the application registration automatically (as well as run on a regular basis - every 24 hours by default), is collecting extensive information about the Oracle database; that information is key to understand how the database can be protected.

Application Server Inventory: this process runs by default every 24 hours and refreshes Application inventory information, including Oracle information. It is advised to **not schedule** this inventory within the backup window. You can configure it in the **Jobs & Operations → Schedule → Application Server Inventory** menu.

Here after are the steps performed by the Application Inventory process:

- ▶ Spectrum Protect Plus (SPP) deploys agents into each Oracle server
- ▶ Agent discovers all Oracle Homes, databases, ASM disk groups (if any) and reports them back to SPP.
 - Collect installation details by reading the `/etc/orainst.loc` which contains the path to `oraInventory`
 - Collect list of Oracle databases
 - Collect database details

A set of SQL queries are executed to retrieve information specific to each database such as DB Name, DB ID, current incarnation, character set, list of data files, list of redo logs, list of control files, list of temp files, FRA location, list of archived log files, SPFILE location, log mode, open mode, size of the database, block change tracking status and current SCN.

Example 2-2 List of SQL queries executed to retrieve information for an Oracle Instance

```
DB Name: select value from v$parameter where name='db_name';
DB ID: select dbid from v$database;
Check if multi-tenant DB: select cdb from v$database;
Current incarnation: select resetlogs_id from v$database_incarnation where
status='CURRENT';
Character set: select value from nls_database_parameters where
parameter='NLS_CHARACTERSET';
List of datafiles: select f.file#, t.name, f.name from v$datafile f, v$tablespace
t where f.ts#=t.ts#;
List of redo logs: select l.thread#, l.group#, l.bytes, l.blocksize, f.member from
v$log l, v$logfile f where l.group#=f.group#;
List of control files: select name from v$controlfile;
List of temp files: select name from v$tempfile where status='ONLINE';
List of archived log locations: select destination, status, error from
v$archive_dest where target='PRIMARY' and status!='INACTIVE';
FRA Location: select name from v$recovery_file_dest;
SPFILE Location: select value from v$parameter where name='spfile';
Log mode: select log_mode from v$database;
Open mode: select open_mode from v$database;
Size of DB: select sum(bytes) from dba_segments;
Block Change Tracking: select status from v$block_change_tracking;
Current SCN: select current_scn from v$database;
```

Based on this inventory job result, Spectrum Protect Plus displays all homes and databases in the GUI and the user can select them for backup.

Following this inventory, some databases may be marked as "not eligible" and will therefore not be selectable in the GUI.

To be eligible for backup with Spectrum Protect Plus, a database must:

- ▶ be MOUNTED and OPEN. If the database is not running, Spectrum Protect Plus cannot determine its details and therefore cannot protect it
- ▶ be in ARCHIVELOG mode.

2.1.3 Oracle log backup

The application log backups have been presented along with screen captures as part of the Application backup options in 1.1.3, “Configuring the database application options” on page 4.

Oracle log backup must be discussed with the database administrator to understand if the log backup and truncation is managed by Spectrum Protect Plus or by another process implemented by the Oracle database administrator.

Note: Do not combine the management of Oracle database log backup between Spectrum Protect Plus and another tool.

Spectrum Protect Plus allows, for Oracle, the following possible scenarios:

- ▶ Backup the Oracle logs and let them on the source system
- ▶ Backup the Oracle logs and delete them from the source following a successful backup and after the specified source retention
- ▶ Do not backup the Oracle logs

Oracle definition of inconsistent backups: Any database backup that is not consistent is an inconsistent backup. A backup that is made when the database is open or after an instance failure or **SHUTDOWN ABORT** command is used is inconsistent. When a database is restored from an inconsistent backup, the Oracle database must perform media recovery before the database can be opened, applying any pending changes from the redo logs.

Spectrum Protect Plus performs an "online incremental update" backup of the Oracle database. While Spectrum Protect Plus is copying the data files, the database accepts updates. Consequently, it is likely that at the end of the backup, the datafiles are not consistent. However, Spectrum Protect Plus cope with this situation by using archive log mode, meaning after restoring a specific backup, the collected log will be applied to recover the database to this backup point in time and make it consistent again.

Therefore, IBM Spectrum Protect Plus requires that the Oracle database to be in ARCHIVELOG mode. If the archive log is disabled, the database will be displayed as **not eligible** for backup in Spectrum Protect Plus.

Log backup management is controlled via crontab; Spectrum Protect Plus creates entries (one per database), as shown in Example 2-3.

Example 2-3 crontab entries example for sub-hourly Oracle log backup

```
sppagent@itso-oracle:~> crontab -l
# DO NOT EDIT THIS FILE - edit the master and reinstall.
# (/tmp/crontab_1621230300 installed on Mon May 17 07:45:00 2021)
# (Cronie version 4.2)
0,15,30,45 * * * * *
/opt/IBM/SPP/logbackup/ORATST/df125510cb9a87dbcbba902ea40f0061/logbackup.sh #
Added by SPP
```

0,55 * * * *

/opt/IBM/SPP/logbackup/SPP/df125510cb9a87dbcbba902ea40f0061/logbackup.sh # Added
by SPP

Spectrum Protect Plus Oracle log backup can be summarized as a six-step process:

1. Mount the vSnap volume to the target application server (only if this is the first log backup or the mount was removed manually)
2. Start RMAN and connect to the target database
3. Verify that the target database is mounted or open
4. Connect to the database
5. Perform the log backup:
 - a. Perform the backup of the archive logs
BACKUP AS COPY ARCHIVELOG FROM SCN ##LATEST_KNOWN_SCN##
 - b. Perform the archive log truncate from the Oracle archive log path - as per the settings you selected
DELETE NOPROMPT FORCE ARCHIVELOG UNTIL TIME ###YOUR SETTINGS ###
6. Disconnect from the database

Step 5.b depends on how you configured the log management from the Spectrum protect Plus Oracle backup configuration options. If you decided to purge the log after 4 hours (see Figure 2-3). Figure 2-4 shows the kind of command that Spectrum Protect Plus will trigger.

Options

☒ Enable Log Back up

Log Backup Frequency and Start Time

Subhourly 15 minute(s)

05/12/2021 00:00 Europe/Paris

Truncate source logs after successful backup

Older than a specified number of hours

Primary log retention in hours

4

Maximum Parallel Streams per Database

3

Figure 2-3 Oracle database backup configuration options for log backup management

Example 2-4 Oracle archive log truncate management - post archivelog backup

```
[2021-06-01 15:29:25] INFO pid:27060 Thread-12 worker_logbackup: purgePrimaryLogs:
AFTER_HOURS
[2021-06-01 15:29:25] INFO pid:27060 Thread-12 worker_logbackup: ORATST: purging
primary logs older than [4] hours
[2021-06-01 15:29:25] JOBL0G pid:27060 Thread-12 writeInputFile: <CTGGF0175>
ORATST: Running command: delete noprompt force archivelog until time
'SYSDATE-4/24' like '/home/ORATST/arch%';
[2021-06-01 15:29:26] JOBL0G pid:27060 Thread-12 worker_logbackup: <CTGGF0219>
ORATST: Log backup worker finished.
```

Note: If you change the archive log management settings from the Oracle database backup configuration options (Figure 2-3), the change you made will be applied on the target system with the next SLA run.

Beyond the 6 step process described above, Spectrum Protect Plus also does metadata management such as:

- ▶ Record the latest SCN in the SPP catalog so that this information will be available for the next log backup.
- ▶ Create Point in time snapshot of two vSnap volumes simultaneously. First vsnap volume contains data files control files, spfile copies. The second vsnap volume contains the oracle log files.

When an SLA backup runs (data files, control file, SPFile backup), Spectrum Protect Plus deals also with archive log cleaning, but this time, for the archive log that has been copied to the vSnap NFS share as part of the crontab archive log backup. Example 2-5 shows the command that is run to cleanup the Oracle archive log no longer required in the log backup destination, that are stored on the vSnap.

Example 2-5 SPP archive log cleanup from archive log backup destination

```
[2021-06-01 15:33:55] INFO pid:28896 MainThread backup_main: Adding DB ORATST to
worker queue for logpurge
[2021-06-01 15:33:55] INFO pid:28896 MainThread backup_main: Waiting for all
workers to finish
[2021-06-01 15:33:55] JOBLLOG pid:28896 Thread-13 worker_logpurge: <CTGGF0220>
ORATST: Log purge worker started.
[2021-06-01 15:33:55] JOBLLOG pid:28896 Thread-13 logParams: ORATST>
(ORACLE_SID=ORATST) (ORACLE_HOME=/opt/oracle/product/12.2.0/dbhome_1)
(OWNER=oracle)
[2021-06-01 15:33:55] JOBLLOG pid:28896 Thread-13 worker_logpurge: <CTGGF0062>
ORATST: Deleting archived logs older than 1 day(s) under log backup destination.
[2021-06-01 15:33:55] JOBLLOG pid:28896 Thread-13 writeInputFile: <CTGGF0175>
ORATST: Running command: delete noprompt force archivelog until time 'SYSDATE-1'
like '/mnt/spp/vsnap/vpool1/fs181/10_0_240_223/ORATST%';
[2021-06-01 15:33:59] JOBLLOG pid:28896 Thread-13 worker_logpurge: <CTGGF0064>
ORATST: Log purge worker finished.
[2021-06-01 15:33:59] JOBLLOG pid:28896 MainThread backup_main: <CTGGF0003>
Completed logpurge operation in 4s. 1 database(s) succeeded and 0 failed.
```

2.1.4 Backup details

This section provides some insight into Oracle database backup process in Spectrum Protect Plus.

Oracle backup method: “Incrementally updated backups”

The Oracle backup strategy used in Spectrum Protect Plus is called "Incrementally Updated Backups". In this strategy, Spectrum Protect Plus will create a level 0 image copy of each data file, and then periodically roll forward this copy by making and then applying a level 1 incremental backup. This approach avoids the overhead of making repeated full image copies of data files, but still provides all of the advantages of full data file copies.

Incrementally updated backups were introduced in Oracle 10g. Using this feature, all changes between the SCN (System Change Number) of the original image copy and the SCN of the incremental backup are applied to the image copy, winding it forward to make the equivalent of a new data file image copy, without the overhead of actually creating a new image copy.

As mentioned as part of the prerequisites, database archive logs are enabled, which implies that local cleaning of these logs is required. IBM Spectrum Protect Plus can perform this task. To control the local archive log cleaning by IBM Spectrum Protect Plus, select the **Enable Log Backup** option and specify the **Primary log retention** parameter in the Backup Options window, as shown in 1.1.3, “Configuring the database application options” on page 4.

Note: When the log backup is enabled, it will be executed as part of the incremental backup

Oracle backup processing

As already mentioned, Spectrum Protect Plus relies on RMAN to protect the Oracle database. Below are the different steps and RMAN commands being used to perform the Oracle database and logs backup as part of the first full or subsequent incremental backup.

1. Mount the vSnap volume to the target application server
2. Start RMAN and connect to the target database
3. Verify that the target database is mounted or open
4. Database backup
 - a. If this is the initial backup of the target database (no data on vSnap at this time)
Run > BACKUP INCREMENTAL LEVEL 0 FOR RECOVER OF COPY WITH TAG 'SPP_BACKUP_####' DATABASE
 - b. If this is an incremental backup of the target database
Run > BACKUP INCREMENTAL LEVEL 1 FOR RECOVER OF COPY WITH TAG 'SPP_BACKUP_####' DATABASE
 - c. then the other objects backup:
Run > BACKUP AS COPY CURRENT CONTROLFILE REUSE TAG 'SPP_BACKUP_####'
Run > BACKUP AS COPY ARCHIVELOG FROM SCN ##### UNTIL SCN #####
create pfile='/mnt/spp/vsnap/vpoo11/fs182/10_0_240_223/SPP/pfile.txt' from spfile;
5. Disconnect from database
6. Create a snapshot on the vSnap volume
7. Catalog the snapshot
8. Unmount the vSnap volume from the target application server
9. Enable parallelism in a single database backup process

For Oracle database backup, you have the ability to enable parallel processing at the data file level. Parallel processing is achieved by opening multiple RMAN channels. This parallel process option is available when you configure the Oracle application options as shown in Figure 2-4 .

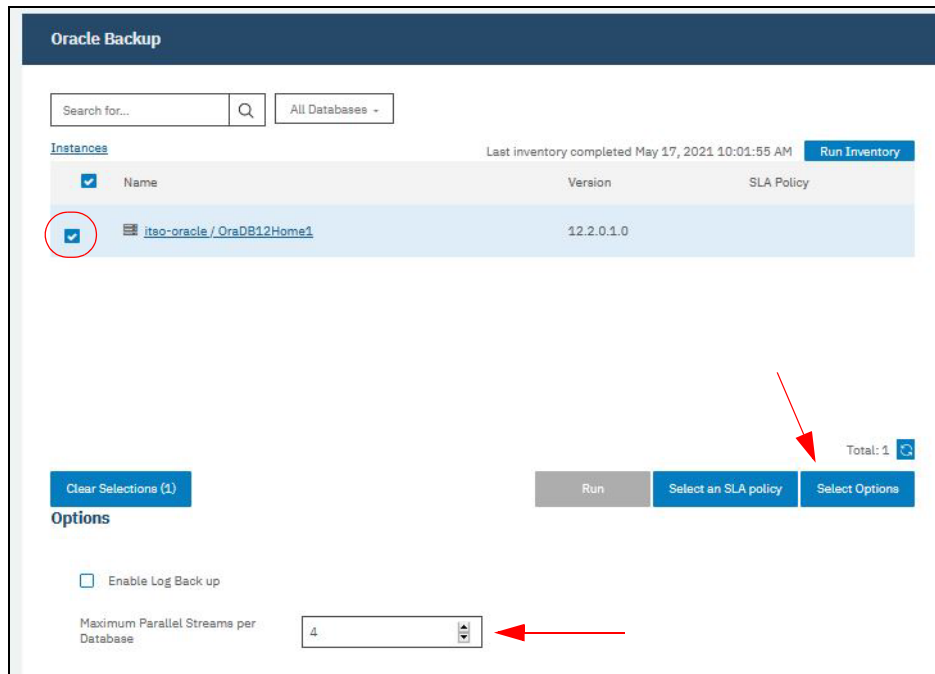


Figure 2-4 Specifying the Parallel Streams for Oracle datafiles backup

Note: Oracle Parallel Streams parameter shown in Figure 2-4 is satisfied for both backup and **simple Production** restore. It is not used by Instant Access, Test restore, and Production restore when changing the Oracle Resource Name and/or Destination Path. In the latter, it means that restore is likely to be slower as it is using only one stream.

2.1.5 Restore details

This section provides some insight into the Oracle database restore process in Spectrum Protect Plus.

Restore modes

There are different restore modes available for Oracle database that can address different use cases. They are explained in 1.2, "IBM Spectrum Protect Plus database restore and data reuse" on page 7.

These restore modes are:

- ▶ Test
- ▶ Production
- ▶ Instant Access

Oracle Test restore processing

When doing a Test restore for Oracle, Spectrum Protect Plus creates a vSnap clone from the version selected by the user and creates an NFS share. The agent mounts the share on the Oracle server where the restore operation is to be performed.

For Oracle RAC, the restore operation is performed on all nodes in the cluster. The agent **spins up a new database using the data files directly from the vSnap volume**. Spectrum Protect Plus performs a point-in-time restore by first mounting the database from the

preceding backup and then applying archived logs to roll forward to the specified intermediate point in time.

In this mode, a temporary database running directly off the vSnap backup repository can be launched without performing any data copy back to production storage.

There are a few options to note about the test restore:

- ▶ Test restore can be done to the original Oracle server or an alternate server.
- ▶ The test database can be brought up with the same name as the original or an alternate name.

Oracle Production restore processing

A Production restore is a permanent database created by copying data from the vSnap backup repository back to production storage and then launching the database.

The production restore database can have its data restored to the same path (directory or ASM disk group) as the original, or to an alternate path (directory or ASM disk group).

Spectrum Protect Plus performs a point-in-time restore by first restoring the database from the preceding backup and then applying archived logs to roll forward to the specified intermediate point in time.

Here are the RMAN steps and instructions being used for Production restore:

1. Mount the vSnap volume that holds the datafiles of the target database
2. Mount the vSnap volume that holds the log files of the target database
3. Copy the datafiles from the vSnap mount to the target directory
4. Recover the database based on archive log:
 - a. Run:

```
>STARTUP NOMOUNT PFILE=' oracle target database location '
```
 - b. Connect to database
 - c. Run the following commands:

```
>CREATE CONTROLFILE REUSE SET DATABASE ' oracle database name ' RESETLOGS  
NOARCHIVELOG  
>SET LOGSOURCE ' vSnap log backup mount '  
>SET AUTORECOVERY ON  
>RECOVER DATABASE USING BACKUP CONTROLFILE UNTIL CANCEL  
>ALTER DATABASE ARCHIVELOG  
>ALTER DATABASE OPEN RESETLOGS
```
5. Disconnect from database

If the database name and path are not modified, the recovery uses RMAN commands to copy the data from vSnap to production system (see Example 2-6) and uses the parallel streams option that was set when configuring the Oracle backup.

Example 2-6 Extract of RMAN logs for Production restore using 3 parallel streams

```
...
Recovery Manager: Release 12.2.0.1.0 - Production on Tue May 4 14:10:18 2021

RMAN> set echo off;
2> connect target *
```

```

3> run {
4> configure controlfile autobackup off;
5> set command id to 'SPP_BACKUP_1056_1620130218';
6> allocate channel spp1 type disk;
7> allocate channel spp2 type disk;
8> allocate channel spp3 type disk;
9> restore database from tag 'SPP_BACKUP_1056';
10> release channel spp1;
11> release channel spp2;
12> release channel spp3;
13> }
14> exit;
echo set off

```

connected to target database: SPP (DBID=2016102274, not open)

using target database control file instead of recovery catalog
old RMAN configuration parameters:
CONFIGURE CONTROLFILE AUTOBACKUP OFF;
new RMAN configuration parameters:
CONFIGURE CONTROLFILE AUTOBACKUP OFF;
new RMAN configuration parameters are successfully stored

executing command: SET COMMAND ID

allocated channel: spp1
channel spp1: SID=130 device type=DISK

allocated channel: spp2
channel spp2: SID=9 device type=DISK

allocated channel: spp3
channel spp3: SID=53 device type=DISK

Starting restore at 04-MAY-21

channel spp1: restoring datafile 00001
input datafile copy RECID=6 STAMP=1071670217 file
name=/mnt/spp/vsnap/vpool1/fs123/10_0_240_223/SPP/data_D-SPP_I-2016102274_TS-SYSTEM_FNO-1_02vu0fqo
destination for restore of datafile 00001: /home/SPP/data/system01.dbf
channel spp2: restoring datafile 00003
input datafile copy RECID=7 STAMP=1071670217 file
name=/mnt/spp/vsnap/vpool1/fs123/10_0_240_223/SPP/data_D-SPP_I-2016102274_TS-SYSAUX_FNO-3_03vu0fqo
destination for restore of datafile 00003: /home/SPP/data/sysaux01.dbf
channel spp3: restoring datafile 00004
input datafile copy RECID=8 STAMP=1071670217 file
name=/mnt/spp/vsnap/vpool1/fs123/10_0_240_223/SPP/data_D-SPP_I-2016102274_TS-UNDOTBS1_FNO-4_04vu0fqo
destination for restore of datafile 00004: /home/SPP/data/undotbs01.dbf
channel spp3: copied datafile copy of datafile 00004
output file name=/home/SPP/data/undotbs01.dbf RECID=0 STAMP=0
channel spp3: restoring datafile 00007


```

input datafile copy RECID=5 STAMP=1071670217 file
name=/mnt/spp/vsnap/vpoo11/fs123/10_0_240_223/SPP/data_D-SPP_I-2016102274_TS-USERS
_FN0-7_06vu0fq
destination for restore of datafile 00007: /home/SPP/data/users01.dbf
channel spp3: copied datafile copy of datafile 00007
output file name=/home/SPP/data/users01.dbf RECID=0 STAMP=0
channel spp2: copied datafile copy of datafile 00003
output file name=/home/SPP/data/sysaux01.dbf RECID=0 STAMP=0
channel spp1: copied datafile copy of datafile 00001
output file name=/home/SPP/data/system01.dbf RECID=0 STAMP=0
Finished restore at 04-MAY-21

released channel: spp1

released channel: spp2

released channel: spp3

Recovery Manager complete.
....

```

However, If either the Instance Name or Path is modified (see Figure 2-5), then the process is using the operating system copy command “cp”. In the latter case, the recovery will use only one parallel copy stream (see Figure 2-7).

Name	New Resource Name
<input checked="" type="checkbox"/> SPP	ORA2

Source Path	Destination Path
/home/SPP/data	/home/ORA2/data
/home/SPP/log2	/home/ORA2/log2
/home/SPP/log1	/home/ORA2/log1
/home/SPP/arch	/home/ORA2/arch

Figure 2-5 Spectrum Protect Plus Oracle production redirect restore wizard

Example 2-7 shows the kind of command you will observe when doing a restore of the SPP instance to an alternate location, here in the /home/ORA2 path, as specified in the Restore wizard shown in Figure 2-5.

Example 2-7 oracle redirect restore using cp command (not RMAN)

```

root      28029 27877  0 14:43 ?          00:00:00 sudo -n cp -p
/mnt/spp/vsnap/vpoo11/fs126/10_0_240_223/SPP/data_D-SPP_I-2016102274_TS-SYSTEM_FN0
-1_02vu0fqo /home/ORA2/data/system01.dbf

```

Note the following options about the test restore:

- ▶ Test restore can be done to the original Oracle server or an alternate server.
- ▶ The production database can be restored using a specific backup, or any intermediate point-in-time (time stamp or SCN).

Note related to Test and Production restore:

It is possible to perform a PIT restore of the Oracle database based on a time stamp or on a SCN. To recover a database PIT both sources are required, the datafiles and the archive log files. SPP performs a roll forward recovery of the target database.

Therefore, the most recent datafiles are needed that were created in the past relative to the PIT (SCN). In addition, the archive log data is needed that contains the PIT (SCN). Because of this dependency, the most recent vSnap snapshot that was taken from the datafiles, in the past relative to the PIT (SCN), is used for recovery and the next vSnap snapshot that was taken from the archive log backup in the future relative to the PIT (SCN).

Oracle Instant Access processing

Spectrum Protect Plus can perform an Instant Access which is in fact an Instant storage mount from the vSnap backup repository to allow the user to access the files stored in any recovery point. The instant access process restores both snapshots taken at the same time, one for the datafiles and one for the archive log.

The instant access restore method is meant to have the ability to verify the backup, but not the recovery of the database.

The user can then perform any custom recovery action using those files, such as manually copy some files or use RMAN catalog to be able to work with these file and do further specific recovery, which you cannot do with other Spectrum Protect Plus restore modes.

A good use case for Instant Access could be to retrieve a specific file, such as control file, as illustrated in Example 2-8, or to execute “CATALOG START WITH <mount point>” to make RMAN scan the contents of the Instant Access mount, then perform any RMAN supported granular restore operation.

Figure 2-6 shows the wizard Review menu when doing an Oracle instant access process. This operation can be done through **Create Job** → **Restore from the Manage Protection** → **Databases** → **Oracle** menu, for example.

[Back to Oracle](#)
Restore - Oracle
Review

☐ Default Setup

Progress bar steps:
☒ Select data sources
☒ Select source
☒ Source snapshot
☒ Restore method
☒ Set destination
☒ Set run settings
☒ Job options
☐ Review

Review

Review your selections, and then click Submit.

Select data sources

Selected source:	ORATST
Source snapshot:	ORATST - Jun 1, 2021 7:47:44 AM
Restore Type:	On-Demand: Snapshot
Restore Source Type:	Site
Restore Source:	Primary
Restore method:	Instant Access

Set destination

Destination:	Restore to original instance
--------------	------------------------------

Set run settings

Run cleanup immediately on job failure:	Yes
Allow session overwrite:	Yes
Continue with restores of other selected databases even if one fails:	Yes

Figure 2-6 Oracle Instant Access wizard - review screen

From the Oracle system side, an NFS share is mounted and can be browsed with regular operating system commands. You can then read or copy any of the desired files back to your environment.

Example 2-8 Instant access allows you to access file directly from the Spectrum Protect Plus backup

```
itso-oracle:~ # df |grep vsnap
10.0.240.223:/vsnap/vpool1/fs181 95060480 195456 94865024 1%
/mnt/spp/vsnap/vpool1/fs181/10_0_240_223
10.0.240.223:/vsnap/vpool1/fs243 96222976 1357952 94865024 2%
/mnt/spp/vsnap/vpool1/fs243/10_0_240_223

itso-oracle:~ # ls -l /mnt/spp/vsnap/vpool1/fs243/10_0_240_223
total 21
drwxr-xr-x 3 oracle oinstall 10 Jun 1 08:06 ORATST
drwxr-xr-x 3 oracle oinstall 10 Jun 1 08:06 SPP

itso-oracle:~ # cp /mnt/spp/vsnap/vpool1/fs243/10_0_240_223/ORATST/controlfile.txt
/tmp/restore_controlefile.txt

itso-oracle:~ # head /tmp/restore_controlefile.txt
-- The following are current System-scope REDO Log Archival related
-- parameters and can be included in the database initialization file.
--
-- LOG_ARCHIVE_DEST=''
-- LOG_ARCHIVE_DUPLEX_DEST=''
--
-- LOG_ARCHIVE_FORMAT=%t_%s_%r.dbf
--
-- DB_UNIQUE_NAME="ORATST"
--
....
```

Restore use cases

Among the three possible restore methods, you must obviously choose the one that best addresses your needs.

Here are the few restore use cases along with the suggested Spectrum Protect Plus restore mode.

Restore for data validation:

For Oracle, as there are no corruption checks being done at backup time, there is an interesting scenario which consists of scheduling a regular “Test mode” recovery. Using this mode mounts the database to any Oracle environment, not necessarily the same as the backup source, allowing the database administrator to run a validation script against the backup copy of the database.

Database duplication

Native RMAN provides a function called RMAN duplicate. Production restore mode of Spectrum Protect Plus can be used to create a PIT clone of an existing Oracle database.

Database build standby

That same production restore mode from Spectrum Protect Plus can also provide an equivalent to the RMAN Target Database for Standby to create a standby database, although the relationship between the production and the standby must be built manually after the restore has finished.

2.1.6 Troubleshooting hints

This section provides some troubleshooting guidance.

Access the Job log

For each job, IBM Spectrum Protect Plus records the commands that it uses to handle the database (including SQL and RMAN commands) in a `command.log` file. The following options are available to access these log files:

- ▶ Select **Download.zip** in the Jobs and Operations menu to download the collection of logs for a specific job. The `.zip` file contains folders that are named `application/<uuid>` where `<uuid>` matches the last portion of the log dir location. Check the `command.log` files in these folders.
- ▶ Check the `/data/log/guestdeployer/<date>` subdirectories on the IBM Spectrum Protect Plus appliance, which also stores the `command.log` files.

For more information about requirements, see *IBM Spectrum Protect Plus Installation and User's Guide*.

NFS troubleshooting during log backups

SPP uses a NFS mount to perform log backup. What happens when the NFS mount fails while the log backup is ongoing?

The NFS mount failure can have multiple root causes. For example:

- ▶ Network issues
- ▶ vSnap goes down for maintenance
- ▶ vSnap error

- ▶ Depending on when the NFS mount failure occurs, the result can vary: When the database processes the log backup, the vSnap share is already invalid (vSnap down). This situation will cause an alert in SPP, but SPP won't lose log backup data.
- ▶ When the database processes the log backup, the vSnap goes down in the middle of the log backup processing, which will cause an alert in SPP and log data can get lost. If the log files or metadata files are damaged or incomplete, SPP loses log backup data. In this case, PIT restores will fail.
- ▶ When the database has finished the log backup the vSnap goes down before it takes a snapshot on the log volume. The data was written to disk and the snapshot will be taken later depending on the next snapshot schedule.

Note: If the vSnap server goes down for any reason, such as maintenance for example, the persistent NFS mount point for log backup becomes invalid and must be unmounted. The next backup (either SLA or log backup) would mount the NFS share again.



Backing up and restoring MongoDB databases

In this chapter, we describe backing and restoring MongoDB databases.

This chapter includes the following topics:

- ▶ 3.1, “IBM Spectrum Protect Plus requirements for MongoDB” on page 40
- ▶ 3.2, “MongoDB backup and restore with Spectrum Protect Plus” on page 44

Note: Although database configuration and handling is widely similar for databases in IBM Spectrum Protect Plus, some differences exist for the supported database systems. We describe that information in Chapter 4, “Backing up and restoring Db2 databases” on page 51, and Chapter 5, “Backing up and restoring SQL Server” on page 69

For more information about generic test restore or DevOps use cases, see Chapter 1., “Protecting database applications” on page 1. This chapter also describes database backup, restore, and DevOps use cases in general, but refers specifically to an Oracle database whenever necessary.

3.1 IBM Spectrum Protect Plus requirements for MongoDB

This chapter describes specific IBM Spectrum Protect Plus requirements for MongoDB databases. For more information about the latest list, check the MongoDB requirements section of the *IBM Spectrum Protect Plus Installation and User's Guide*, which is available at [IBM Documentation](#).

3.1.1 Fundamental IBM Spectrum Protect Plus requirements for MongoDB

This section provides an overview of important IBM Spectrum Protect Plus requirements in MongoDB environments. See *IBM Spectrum Protect Plus Installation and User's Guide* for a complete list of up-to-date support information, which is available at [IBM Documentation](#).

Operating system support

IBM Spectrum Protect Plus supports MongoDB environments on Linux systems, including Red Hat Enterprise Linux (RHEL), CentOS, and SUSE Linux Enterprise Server.

Logical Volume Manager

IBM Spectrum Protect Plus requires that logical volumes of MongoDB data and log paths are managed by Linux Logical Volume Manager (LVM2). LVM2 is used for creating temporary volume snapshots. The database files and the journal must be on a single volume.

Operating system user

The initial discovery of MongoDB databases on an IBM Spectrum Protect Plus server requires an operating system user (called the IBM Spectrum Protect Plus agent user) with the following permissions:

- ▶ Run commands as the root *user* and as the MongoDB software owner user by using *sudo*. IBM Spectrum Protect Plus requires this privilege for tasks, such as discovering storage layouts, mounting and unmounting disks, and managing databases. Example 3-1 shows an appropriate */etc/sudoers* entry for a user named *mosuser*.
- ▶ Read, write, and execute permissions for the database directories. The MongoDB default database directory is */data/db*.

Example 3-1 An entry for mosuser

```
Defaults:mosuser !requiretty
mosuser ALL=(ALL) NOPASSWD:ALL
```

Current restrictions

In IBM Spectrum Protect Plus version 10.1.6, MongoDB is configured as a stand-alone instance or replica set. Currently, IBM Spectrum Protect Plus does not support backup operations of MongoDB shared cluster instances. A backup always includes all databases in the instance.

3.1.2 MongoDB databases without authentication

After installing the MongoDB software, you can immediately start the MongoDB daemon (mongod) or service on your operating system and access a (default) database. On Linux operating systems, a default database is created on the data path /data/db. However, such a database is open to anybody in your network, or even the internet. Therefore, we strictly recommend that you secure your MongoDB databases more carefully.

IBM Spectrum Protect Plus offers a two-stage process to access a MongoDB database. First, you register the database server with an IP name or address, an operating system user, and a corresponding password. IBM Spectrum Protect Plus initiates a database discovery job on this server. If you run your MongoDB without authentication, the database registration in IBM Spectrum Protect Plus is complete at this point.

Also, if you secured your databases on the database level, you specify more user credentials for each secured database that IBM Spectrum Protect Plus discovered.

3.1.3, “MongoDB databases with authentication enabled” on page 41 describes how to enable MongoDB authentication.

3.1.4, “Register a MongoDB server” on page 43 describes the MongoDB registration in IBM Spectrum Protect Plus.

3.1.3 MongoDB databases with authentication enabled

This section describes more configuration steps for a MongoDB database that runs with authentication enabled.

If your MongoDB database is configured without credentials, you should secure it. There are many MongoDB databases open on the internet, providing the opportunity for massive data breaches.

For more information about available authentication options, see the MongoDB manuals, which are available at [this website](#).

MongoDB authentication requires the definition of at least one MongoDB user. If database authentication is enabled, IBM Spectrum Protect Plus must provide a user name and a password to run backup and restore activities.

For each MongoDB user that you plan to use for backup and restore with IBM Spectrum Protect Plus, specify MongoDB access roles by using the **db.grantRolesToUser()** command, as shown in Example 3-2.

Example 3-2 Grant permissions to an existing MongoDB user

```
> use admin
switched to db admin

> db.grantRolesToUser("mdbuser",
[ { role: "hostManager", db: "admin" },
{ role: "clusterMonitor", db: "admin" } ] )

> db.grantRolesToUser("mdbuser",
[ { role: "clusterManager", db: "admin" } ] )
```

The MongoDB hostManager and clusterMonitor roles provide access to MongoDB commands that IBM Spectrum Protect Plus requires to monitor, read the state of, and handle the databases including:

- ▶ getCmdLineOpts
- ▶ serverVersion
- ▶ replSetGetConfig
- ▶ replSetGetStatus
- ▶ shutdown

The clusterManager role is required only for running test restore operations of replica sets.

If you decide to create a new or dedicated user for backup and restore purposes, you can use the **db.createUser()** command, as shown in Example 3-3. According to the MongoDB manuals, the ClusterAdmin role includes the clusterManager, clusterMonitor, and hostManager roles.

Example 3-3 Create a MongoDB user with the permissions required by IBM Spectrum Protect Plus

```
> show dbs
admin    0.000GB
config   0.000GB
local    0.000GB
> use admin
switched to db admin
> db.createUser(
  {
    user: "mdbuser",
    pwd: "mypasswd",
    roles: [ "readWrite", "dbAdmin","clusterAdmin" ]
  }
)
```

Use the **db.getUsers()** command to display users and their permissions.

Note: Enhanced database administration permissions are required to create users and grant roles. The roles that are required for backup and restore with Spectrum Protect Plus are not sufficient.

For MongoDB authentication to take effect, restart the MongoDB daemon (mongod) with the “--auth” option. Example Example 3-4 on page 42 shows how to start the daemon on a Linux command line.

Example 3-4 Starting mongod on Linux

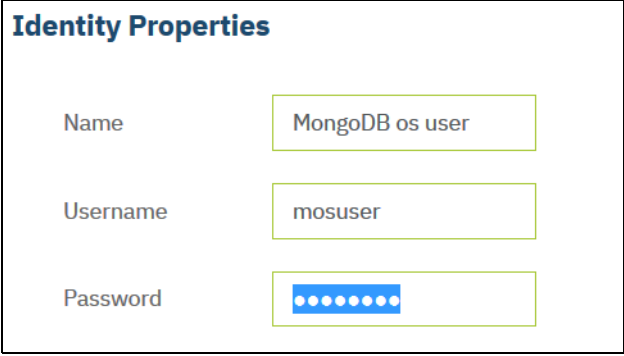
```
mongod --bind_ip_all --auth &
```

3.1.4 Register a MongoDB server

This section describes the tasks required to register a MongoDB server.

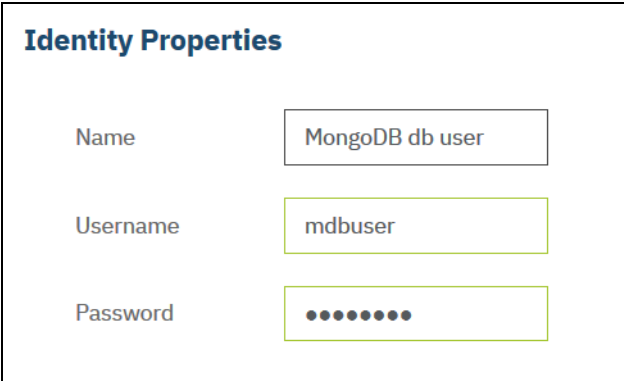
Create identities

Based on your decision to run your MongoDB database with or without authentication, one or two user definitions are required: an operating system user and optionally a MongoDB user. You can specify the users in the **Add application server** menu, but we recommend explicitly creating a so-called **Identity** with a customized name first. Figure 3-1 and Figure 3-2 show Identities for an operating system user and a MongoDB user. The two user names can be identical.



The figure shows a form titled "Identity Properties" with three fields: Name, Username, and Password. The Name field contains "MongoDB os user", the Username field contains "mosuser", and the Password field is masked with blue dots.

Figure 3-1 Identity definition for an operating system user



The figure shows a form titled "Identity Properties" with three fields: Name, Username, and Password. The Name field contains "MongoDB db user", the Username field contains "mdbuser", and the Password field is masked with grey dots.

Figure 3-2 Identity definition for a MongoDB user

Add an application server

In the IBM Spectrum Protect Plus GUI, select **Manage Protection** → **Databases** → **MongoDB**. Then, click **Manage application servers**, and finally, click **Add application servers** to register the database server. Enter the database server IP name or address and select an existing identity. Alternatively, enter a user name and a password.

If you want to start a database discovery job on the server, click **Get Instances**. If IBM Spectrum Protect Plus discovers databases, it shows the connection data for these databases: IP name or address, and IP port.

Important: If you run your MongoDB database without authentication, the registration procedure is complete. However, you should secure your database. If your database is secured, you must specify more user credentials to access the database. The IBM Spectrum Protect Plus GUI provides a **Set Credential** option for the discovered databases (see Figure 3-3).

For more information about handling of MongoDB databases with authentication, see section 3.1.3, “MongoDB databases with authentication enabled”.

Manage application servers

Application Properties

Host Address: t6-vm-lx

☒ User ☐ SSH Key

Use existing user: ☒

Select user: MongoDB os user

Get Instances

Name	Status	Configured
t6-vm-lx Connection: 't6-vm-lx:27017'	Live	No

Set Credential

Figure 3-3 Add a MongoDB server

For more information about required configuration steps and parameters, see *IBM Spectrum Protect Plus Installation and User's Guide*, which is available at [IBM Knowledge Center](#).

3.2 MongoDB backup and restore with Spectrum Protect Plus

In this chapter, we describe MongoDB database backup and restore. The sample restore in this chapter is a restore to the original destination.

For more information about the configuration of other use cases, see 1.2, “IBM Spectrum Protect Plus database restore and data reuse” on page 7.

3.2.1 MongoDB backup

This section describes the tasks that are required to register and back up a MongoDB server.

Assigning an SLA policy

After a MongoDB instance is defined in IBM Spectrum Protect Plus, assign an SLA policy to the instance. In general, we recommend creating dedicated SLA policies for single databases or groups of logically related databases.

After you set up an SLA policy for your MongoDB backup job, you can choose to configure extra options for that job. More SLA options include running scripts, and forcing a full base backup.

For more information, see 1.4, “Database backup with pre-script and post-script” on page 15.

Starting the database backup

The SLA policy that you assigned to your database (see Figure 3-4) defines the schedule time for the first backup. If you did not define a schedule or do not want to wait for the first automatic backup schedule, click **Run** or scroll down to the SLA policy that you provided for this database and select **Actions** to start the database backup.

Now, you also decide whether to perform a backup of a single database (click **Run**), use the Create Job wizard, or perform a backup of all applications that are included in the SLA policy (click **Actions**).


MongoDB Backup			
Search for...		Q	All Databases ▾
Instances		Last Inventory at Jul 9, 2019 1:16:20 PM Run Inventory	
<input checked="" type="checkbox"/>	Name	Version	SLA Policy
<input checked="" type="checkbox"/>	 t6-vm-lx Connection: 't6-vm-lx:27017'	4.0.9	Mongo_SLA

Figure 3-4 MongoDB instance discovered by IBM Spectrum Protect Plus with an SLA policy assigned

Wait until a backup is automatically scheduled or scroll down to the SLA policy section in the window and select **Actions** → **Start** to manually start a backup. This process is IBM Spectrum Protect Plus standard handling, and not specific to MongoDB environments.

To run an on-demand backup job for multiple MongoDB databases that are associated with an SLA policy, click **Create job**, select **Ad hoc backup**, and follow the instructions.

Note: Do not run inventory jobs at the same time that MongoDB backup jobs are scheduled.



SLA Policy Status								
					Filter Job Log: Info x Warning x Error x Summary x ▾			
Policy	Frequency	Total	Succeeded	Failed	Next Run	Status	Policy Options	
 Mongo_SLA		0	0	0		Idle		<div>Actions ▾ Start Pause Sche</div>

Figure 3-5 Manually start a database backup

IBM Spectrum Protect Plus mounts a vSnap server directory to the database server to copy the backup data (see Example 3-5). During the initial backup operation, IBM Spectrum Protect Plus creates a vSnap server volume and NFS share.

Example 3-5 A vSnap server directory mounted on the database server

```
t6-vm-lx:~ # df -h
```

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/sda2	40G	14G	26G	35%	/
devtmpfs	1.9G	8.0K	1.9G	1%	/dev
...					
/dev/mapper/mongovg-mongo1v	15G	410M	14G	3%	/data
10.0.250.48:/vsnap/vpool1/fs11	49G	128K	49G	1%	/mnt/spp/vsnap/vpool1/fs11

During incremental backups, the created volume is reused. The IBM Spectrum Protect Plus MongoDB agent mounts the share on the MongoDB server where the backup is performed.

Switch to the Jobs and Operations menu to display the job protocol and optionally download the job logs and command files.

3.2.2 MongoDB restore

IBM Spectrum Protect Plus offers several restore methods for databases: Production restore, test restore, and instant access. You can select between restore to the original or an alternative destination with or without overwriting an existing database. These features are available for all supported databases.

In this chapter, we demonstrate a MongoDB database restore to the original destination.

For more information about the configuration of other use cases, see 1.2, “IBM Spectrum Protect Plus database restore and data reuse” on page 7.

Restoring a MongoDB database to the original destination

In the IBM Spectrum Protect Plus restore wizard, the following parameters start a traditional database restore that overwrites the existing database:

- ▶ Type Restore: On-Demand Snapshot
- ▶ Restore method: Production
- ▶ Destination: Original Instance

First, select the database instance and an associated database backup, as shown in Figure 3-6. Select the available Source Snapshot that needs to be restored, as shown in Figure 3-7.

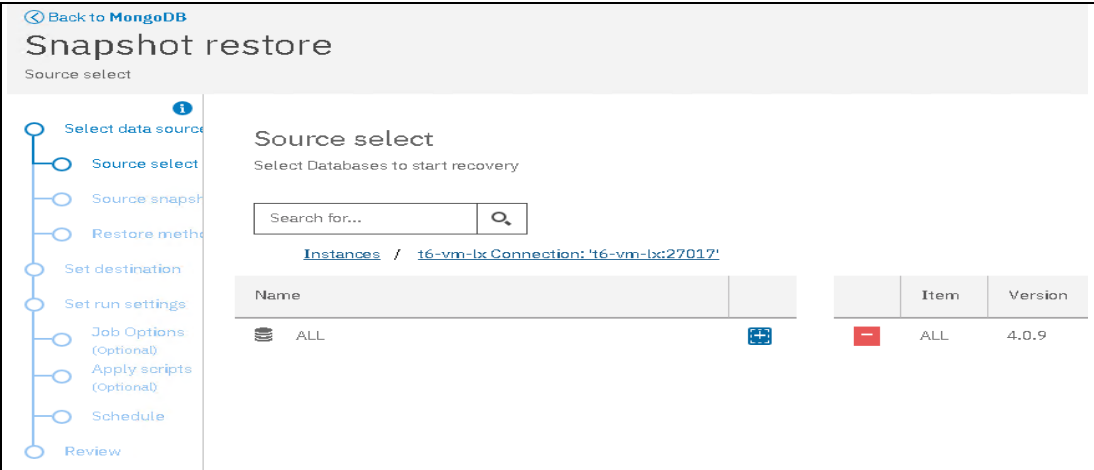


Figure 3-6 Select the source for a database restore

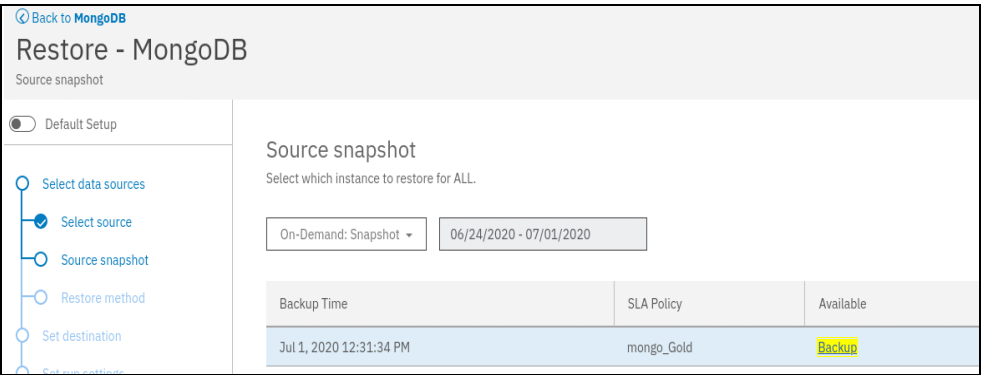


Figure 3-7 Select a site from which to restore the database

The next two selections express what we are trying to achieve: A production restore to the original instance, as shown in Figure 3-8 and Figure 3-9 on page 48.

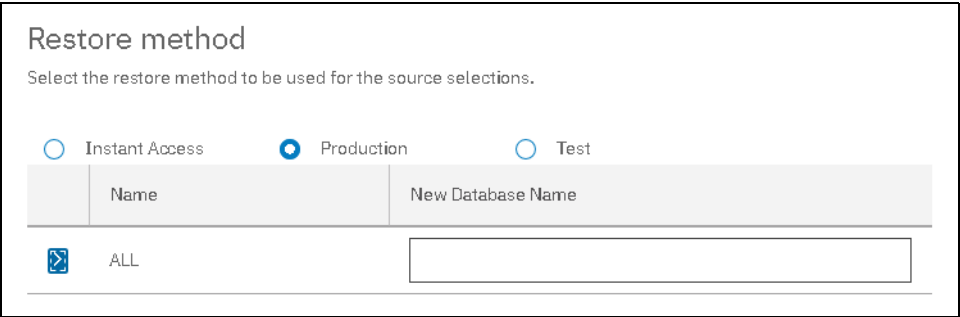


Figure 3-8 Select the restore method

Set destination

Select which instance you would like to restore for each chosen source.

☒ Restore to original instance
 ☐ Restore to alternate instance

Figure 3-9 Select the restore destination

For an on-demand snapshot restore of a production database, the IBM Spectrum Protect Plus restore wizard assumes a subsequent database roll forward to the end of logs included in the backup (see Figure 3-10).

Job Options

Configure the options for this restore job.

Recovery Options

☒ Recover until end of backup

Application Options

☒ Overwrite existing databases.

Maximum Parallel Streams per Database:

Advanced Options

☒ Run cleanup immediately on job failure.

☒ Allow session overwrite

☒ Continue with restores of other selected databases even if one fails.

Figure 3-10 Select database overwrite and other restore options

You must also decide about overwriting a database. IBM Spectrum Protect Snapshot provides an auxiliary protection against an unintended data overwrite; that is, if the database still exists and you do not select the overwrite option, the restore job fails.

In IBM Spectrum Protect Plus, an on-demand snapshot restore is not scheduled. Spectrum Protect Plus runs it only once, as shown in Example 3-11.

Schedule

Run the restore job now or schedule the time for the restore to run.

On-Demand restore selected - job will run on completion of the snapshot restore wizard.

Figure 3-11 Information about a job schedule

Carefully review the job summary that IBM Spectrum Protect Plus displays. If the information describes what you trying to achieve, run the restore job.

Finally, switch to the Job and Operations menu to check the job results (see Figure 3-12).

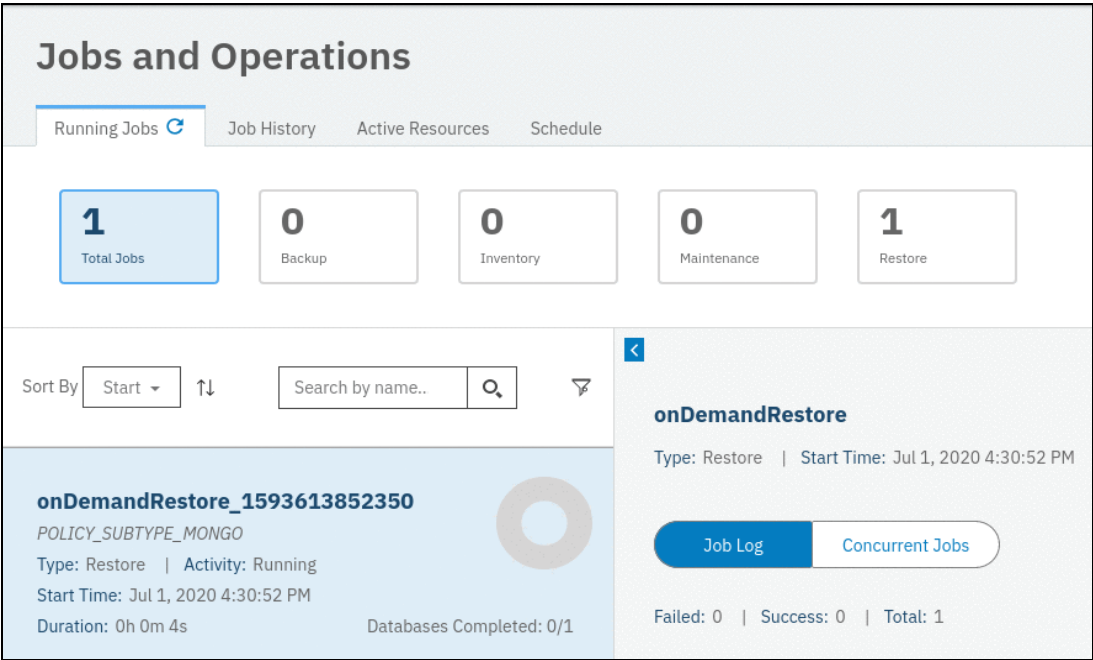


Figure 3-12 Job status view



Backing up and restoring Db2 databases

This chapter describes the management of Db2 databases with IBM Spectrum Protect Plus. Backup, restore, and recovery of single-partitioned and multi-partitioned Db2 databases are supported.

This chapter includes the following topics:

- ▶ 4.1, “IBM Spectrum Protect Plus Db2 features” on page 52
- ▶ 4.2, “Prerequisites for Db2 databases” on page 52
- ▶ 4.3, “Protecting Db2 databases” on page 55

4.1 IBM Spectrum Protect Plus Db2 features

IBM Spectrum Protect Plus supports the following features with Db2 databases:

- ▶ Automatic discovery of Db2 installations on registered machines in IBM Spectrum Protect Plus.
- ▶ Backup, restore, and recovery of single- and multi-partitioned Db2 databases.
- ▶ IBM Spectrum Protect Plus is performing software snapshot-based *online* backups using LVM2 or journaled file system (JFS2) and the Db2 Advanced Copy Services (ACS) interface.
- ▶ IBM Spectrum Protect Plus uses a custom incremental-copying algorithm for data movement from snapshot to vSnap server repository. This algorithm is effective for incremental forever backups.
- ▶ Multiple restore methods are available. Production Restore (database is restored by copying data), Test Restore (database is restored in-place without data movement) and Instant Access (IBM Spectrum Protect Plus only mounts the backup volume).
- ▶ IBM Spectrum Protect Plus supports continuous Db2 archive log backup. This feature can be used optionally for a backup.
- ▶ For Db2, various recovery (transaction roll forward) modes are available in IBM Spectrum Protect Plus, which includes point-in-time recovery by using the archive logs.
- ▶ Note that Spectrum Protect Plus Db2 support is not HADR aware; it will perform the backup on the Db2 instance that has been registered but does not support a failover scenario with regard to performing an automatic switching.

4.2 Prerequisites for Db2 databases

The supported operating systems for IBM Spectrum Protect Plus with Db2 are:

- ▶ On PowerPC: IBM AIX 7.1, 7.2, and later fixpack and modification levels (64-bit kernel)
- ▶ On Linux x86_x64: Red Hat Enterprise Linux 6.8, 7, 11.0 SP4, and 12.0 SP1; SUSE Linux Enterprise Server 11.0 SP4 and 12.0 SP1 and later maintenance and modification levels
- ▶ On Linux on Power Systems (little endian): Red Hat Enterprise Linux 7.1, SUSE Linux Enterprise Server 12.0 SP1 and later maintenance and modification levels.

IBM Db2 Version 10.5, 11.1, 11.5 and later maintenance levels: Enterprise Server Edition are supported at the time of this writing.

To manage Db2 databases with IBM Spectrum Protect Plus the following prerequisites must be met:

- ▶ Define a dedicated IBM Spectrum Protect Plus agent user, for example *sppagent*, on every Db2 server with the required privileges for sudo, as shown in Example 4-1.

Example 4-1 A sudoers file with sppagent user

```
Defaults:sppagent !requiretty
sppagent ALL=(ALL)
NOPASSWD:ALL
```

- ▶ Db2 archive logging is activated and Db2 is in recoverable mode, which requires that at least LOGRETAIN is enabled. The archive logs are continuously backed up to vSnap by Db2

autonomously (the Spectrum Protect Plus Db2 agent just configured logarchmeth1 and logarchmeth2 accordingly with a vSnap NFS mounted volume as the target). Note that the existing setting on logarchmeth1 can coexist with the logarchmeth2 setting of Spectrum Protect Plus.

- Logical volumes holding IBM Db2 table spaces (data and temporary table spaces), the local database directory, and IBM Db2 log files are managed by Logical Volume Management system (LVM2) on Linux and by the Journaled File System (JFS2) on AIX. LVM2 on Linux and JFS2 on AIX are used for creating temporary volume snapshots. Ensure that there is at least 10% free capacity for logical volume snapshots.
- Each Db2 host has to be registered in IBM Spectrum Protect Plus. In a Db2 DPF environment with multiple hosts, every Db2 host has to be registered in IBM Spectrum Protect Plus.

In this publication, the Db2 database example consists of a multi-partitioned Db2 Database Partitioning Feature (DPF) database version 10.5 that is running on two Red Hat Enterprise Linux Server hosts, as shown in Figure 4-1.

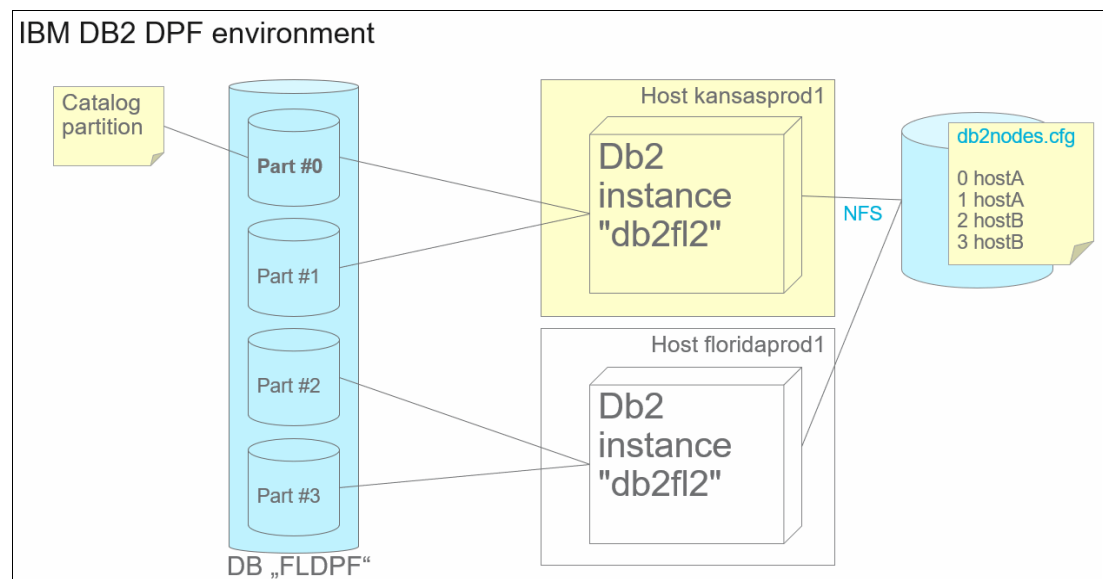


Figure 4-1 Db2 DPF environment

In our example, the Db2 partitions 0, 1, 2, and 3 are spread over the two servers kansasprod1 and floridaprod1, as shown in the db2nodes.cfg file in Example 4-2.

Example 4-2 The db2nodes.cfg file

```
bash-4.1$ cat sqllib/db2nodes.cfg
0 kansasprod1 0
1 kansasprod1 1
2 floridaprod1 0
3 floridaprod1 1
```

To be able to manage the Db2 DPF database with IBM Spectrum Protect Plus the *parallel backup mode*, as shown in Figure 4-2 has to be enabled. To run parallel backup processing of partitions in your Db2 environment, ensure that one of the following prerequisites is met:

- The Db2 registry variable **DB2_PARALLEL_ACS** is set to **YES**, for example: **db2set DB2_PARALLEL_ACS=YES**

- In earlier versions of Db2, the backup mode is determined by the Db2 registry variable **DB2_WORKLOAD**. To enable parallel backup mode, run the Db2 command **db2set DB2_WORKLOAD=SAP**. Check with the Db2 command **db2set -a11 DB2_WORKLOAD**.

Note: Db2 serial backup mode is not supported with IBM Spectrum Protect Plus because of the fact that logs included in the backup can be inconsistent across partitions.

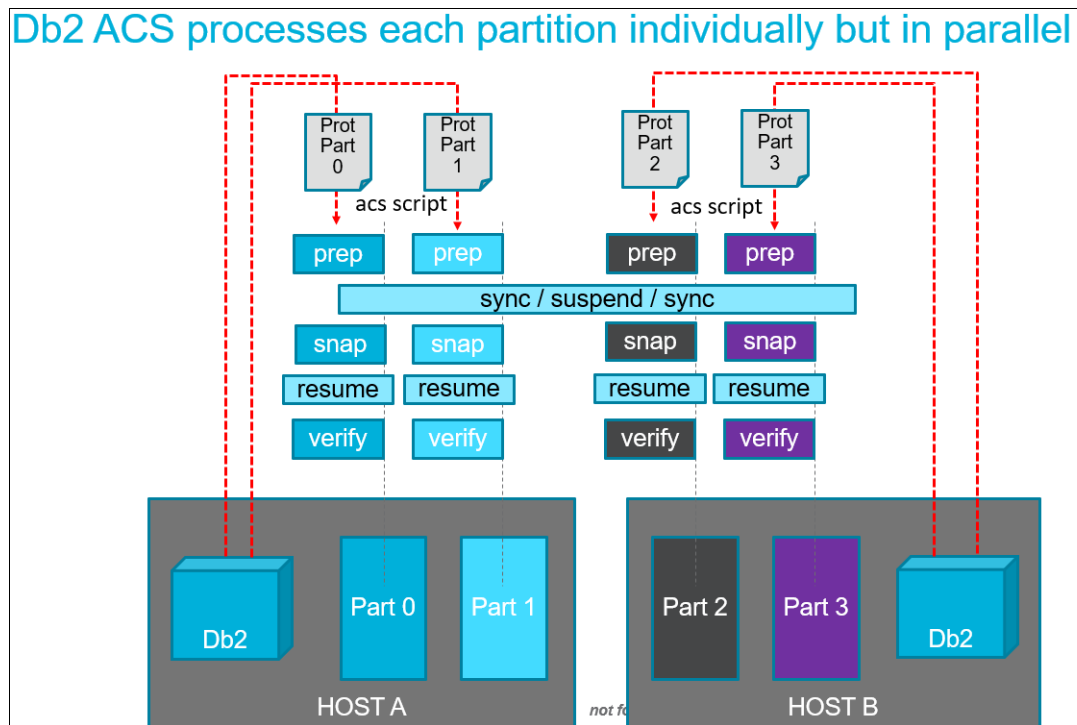


Figure 4-2 Parallel backup mode with Db2 Advanced Copy Services (ACS)

IBM Spectrum Protect Plus triggers the Db2 agent once per host, and if there is more than one partition on the host, Db2 will trigger ACS for each partition individually. A dedicated protocol file is available per partition that is later stored on the vSnap server volume. The Db2 agent can handle the multiple invocations of its ACS scripted part through Db2.

In parallel backup mode, which is the default mode for an SAP Db2 database, all partitions are suspended before Db2 issues snapshot requests. The requests are then performed in parallel on all partitions, as shown in Figure 4-2. IBM Spectrum Protect Plus runs the Db2 backup command on the Db2 catalog partition. The main Db2 ACS processes are:

1. Prepare phase: The write operations of the database are suspended; that is, **WRITE SUSPEND** is set automatically on the database. Db2 prepares the file systems, checks space requirements in the storage system and does other things to keep the database consistent.
2. Snapshot phase: Db2 instructs the Db2 agent to perform a software snapshot on each partition in parallel. The snapshot request is done by taking software snapshots of the corresponding volumes.
3. Verify phase: Db2 checks if the snapshot was taken successfully. If the snapshot is correct, the data is moved to the vSnap server by the Db2 agent.

For more information about updates to the Db2 database prerequisites, see *Spectrum Protect Plus- All Requirements*, which is available at [this web page](#).

4.3 Protecting Db2 databases

To protect Db2 with IBM Spectrum Protect Plus, the database servers have to be registered so that IBM Spectrum Protect Plus can discover the Db2 databases. To start the backup, the Db2 database always has to be assigned to an SLA policy.

4.3.1 Registering the Db2 database server

Before IBM Spectrum Protect Plus can manage the Db2 database, the Db2 servers have to be registered in IBM Spectrum Protect Plus. To register a Db2 database server, complete the following steps:

1. In the navigation pane, click **Manage Protection** → **Databases** → **Db2**.
2. Click **Manage Application Servers** → **Add Application Server**. Enter the required login credentials for the Db2 server, as shown in Figure 4-3.

Note: Pre-define the sppagent username as an Identity in **Accounts** → **Identity** → **Add Identity** before you enter the login credentials of the db2 server. Otherwise, IBM Spectrum Protect Plus will append the ip-address or FQDN to the sppagent username to make it a dedicated user. Especially if you have to change the sppagent password, it makes it easier for the IBM Spectrum Protect Plus admin when the sppagent user can be reused for multiple Db2 servers.

The screenshot shows the 'Manage application servers' interface. Under the 'Application Properties' section, the following fields are visible:

- Host Address:** kansasprod1.spp.ibm.com
- Authentication:** Two radio buttons are present. The 'User' option is selected (indicated by a blue dot), and the 'SSH Key' option is unselected.
- Use existing user:** An unchecked checkbox.
- UserId:** sppagent
- Password:** A field with masked characters (dots).

Figure 4-3 Add Db2 application server pane

Test connection to a Db2 server

The IBM Spectrum Protect Plus test function verifies communication with the Db2 host and tests Domain Name System (DNS) settings between IBM Spectrum Protect Plus and the

host. It also tests that certain services are enabled, and that the specified user has sudo privileges. To start the test, select the host and click **Actions** → **Test**. A pop-up window displays, as shown in Figure 4-4 on page 56.

Test result of kansasprod1.boeblingen.de.ibm.com

1. Physical - Basic Test for physical host network configuration			
Name	Description	Status	Message
Host FQDN Resolvable Test	Host FQDN must be resolvable to an IPv4 address	✓	
Socket Connection Test	Must allow socket connection on port 22 for Linux	✓	
2. Remote - Remote executor test for session creation and remote agent deployment			
Name	Description	Status	Message
Remote Session Test	Latest remote agent must be installed on host, SSH and SFTP service must be installed on Linux host, and port must be open to create session to SSH service.	✓	
Remote Agent Execute Test	Remote agent must be configured correctly using user credentials with sufficient privileges.	✓	
3. LINUX - Basic Linux prerequisites for file and volume operations			
Name	Description	Status	Message
Sudo Privileges	User must have password-less sudo privileges	✓	

Figure 4-4 Test result pop-up

4.3.2 Backup Db2 data

Before starting a backup of Db2, the Db2 database must be assigned to one or more SLA policies.

Defining a Db2 backup job

Assign the selected Db2 database to a SLA policy to create a backup job. Db2 backups run in a “Base-Once-Incremental-Forever” scheme. During the initial base (full) backup, IBM Spectrum Protect Plus creates a vSnap server volume and mounts it to the Db2server by using NFS.

After assigning the Db2 database to an SLA policy, as shown in Figure 4-5, you can optionally click the **Select Options** button, to enable Log Backup, as shown in Figure 4-6, “Select options to enable log backup of Db2” on page 58. With log backup enabled, IBM Spectrum Protect Plus will automatically create a log backup volume and mount it to the application server.

Clear Selections (1)

Run

Select an SLA policy

SLA Policy

	SLA Policy	Frequency	Retention
<input type="checkbox"/>	Gold	Every 4 Hours	1 Weeks
<input type="checkbox"/>	Silver	Every 1 Days at 11:06:29 PM	1 Months
<input checked="" type="checkbox"/>	Bronze	Every 1 Days at 11:06:29 PM	1 Weeks
<input type="checkbox"/>	Demo	Every 1 Days at 11:06:36 PM	1 Months

Figure 4-5 Assign a SLA policy to the database

Enable Log backup

Archived logs for databases contain committed transaction data. This transaction data can be used to run a roll forward data recovery when you are running a restore operation. The use of archive log backups enhances the recovery point objective for your data.

For IBM Spectrum Protect Plus, the Db2 archive logging must be enabled and Db2 must be in recoverable mode. If log backup is enabled in IBM Spectrum Protect Plus, one of the Db2 parameters, LOGARCHMETH1 or LOGARCHMETH2, is updated with the path of the vSnap pool for the log files, as shown in Example 4-3. Therefore, it is important that one of the LOGARCHMETH parameters includes the value OFF and can be used for a vSnap log volume assignment.

Example 4-3 Log backup enabled in IBM Spectrum Protect Plus

```
[db2inst1@spp-db2-01 ~]$ db2 get database configuration for SPPDB | grep LOGAR* -i
First log archive method          (LOGARCHMETH1) = DISK:/mnt/spp/vsnap/vpool1/fs20/
Archive compression for logarchmeth1 (LOGARCHCOMPR1) = OFF
Options for logarchmeth1          (LOGARCHOPT1) =
Second log archive method         (LOGARCHMETH2) = DISK:/mnt/spp/vsnap/vpool1/fs148/192_168_5_234/
Archive compression for logarchmeth2 (LOGARCHCOMPR2) = OFF
Options for logarchmeth2          (LOGARCHOPT2) =
```

Note: To successfully enable Db2 log backup in Spectrum Protect Plus, the Db2 agent expects (and verifies) that all partitions have unique settings for logarchmeth1 and logarchmeth2.

In the Db2 Backup window, select the Db2 database and click **Select Options** → **Enable Log Backup** → **Save**, as shown in Figure 4-6, to allow roll forward recovery when you set up a backup job or SLA policy. When selected for the first time, you must run a backup job for the SLA policy to activate log archiving to Spectrum Protect Plus on the database.

Figure 4-6 Select options to enable log backup of Db2

vSnap commands used to manage Db2 Logs

IBM Spectrum Protect Plus agent creates a separate volume on the vSnap server repository, which is mounted by NFS shared persistently on the Db2 application server. The backup process updates the LOGARCHMETH1 or LOGARCHMETH2 parameters to point to that volume for log archiving purposes. The volume is kept mounted on the Db2 server unless the Enable Log Backup option is cleared and a new backup job is run.

Log backup transaction files are copied to this share according to the schedule created for log backup.

If the DB2 backup job is running, we can see an NFS share on the file system that is associated with the SLA. As shown in Example 4-4, running the vSnap CLI command **vsnap share show** lists the active share, in which the Volume ID 3671 and the share name `/vsnap/vpool1/fs148` can be identified.

Example 4-4 Active share

```
[serveradmin@vsnap fs114]$ vsnap share show
```

ID	TYPE	PARENT VOL	PARTNER ID	NAME
2733	smb	81	N/A	vpool1_fs81
3557	nfs	113	N/A	/vsnap/vpool1/fs113
3558	nfs	114	N/A	/vsnap/vpool1/fs114
3671	nfs	148	N/A	/vsnap/vpool1/fs148

```
[serveradmin@vsnap fs114]$ vsnap share show --id 3671
```

```
ID: 3671
NAME: /vsnap/vpool1/fs148
SHARE TYPE: nfs
VOLUME ID: 148
PARTNER ID: N/A
CREATED: 2020-06-30 11:58:46 UTC
UPDATED: 2020-06-30 11:58:46 UTC
SHARE OPTIONS:
```

ALLOWED HOSTS:
192.168.122.1
192.168.5.94
READ ONLY: No

The share is used to transfer the backup data from the database to the vSnap server.

After the backup of the log completes, log backup transaction files are copied to this share according to the schedule that was created for log backup, as shown in Example 4-5.

Example 4-5 Log backup transaction files copied into NFS shared vSnap in DB2 guest spp-db2-01

```
[root@spp-db2-01 ~]# df -h
Filesystem                                Size  Used Avail Use% Mounted on
/dev/mapper/rhel-root                     38G   9.8G   28G   27% /
devtmpfs                                 1.9G     0   1.9G    0% /dev
tmpfs                                     1.9G   12K   1.9G    1% /dev/shm
tmpfs                                     1.9G   25M   1.9G    2% /run
tmpfs                                     1.9G     0   1.9G    0% /sys/fs/cgroup
/dev/sda1                                1014M  143M   872M   15% /boot
/dev/mapper/rhel-home                     19G   844M   18G    5% /home
/dev/mapper/db2-data                      15G   240M   14G    2% /db2_data
/dev/mapper/db2-log                       9.8G   116M   9.1G    2% /db2_log
tmpfs                                     380M   12K   380M    1% /run/user/42
192.168.5.234:/vsnap/vpool1/fs148 898G 128K 898G 1% /mnt/spp/vsnap/vpool1/fs148/192_168_5_234
tmpfs                                     380M     0   380M    0% /run/user/0

[root@spp-db2-01 ~]# cd /mnt/spp/vsnap/vpool1/fs148/192_168_5_234
[root@spp-db2-01 192_168_5_234]# ls -lrt
total 1
drwxr-x---. 3 db2inst1 db2iadml 3 Jun 30 07:59 db2inst1
[root@spp-db2-01 192_168_5_234]# cd db2inst1
[root@spp-db2-01 db2inst1]# ls -lrt
total 1
drwxr-x---. 3 db2inst1 db2iadml 3 Jun 30 07:59 SPPDB
[root@spp-db2-01 db2inst1]# cd SPPDB
[root@spp-db2-01 SPPDB]# ls -lrt
total 1
drwxr-x---. 3 db2inst1 db2iadml 3 Jun 30 07:59 NODE0000
[root@spp-db2-01 SPPDB]# cd NODE0000/
[root@spp-db2-01 NODE0000]# ls -lrt
total 1
drwxr-x---. 3 db2inst1 db2iadml 3 Jun 30 07:59 LOGSTREAM0000
[root@spp-db2-01 NODE0000]# cd LOGSTREAM0000/
[root@spp-db2-01 LOGSTREAM0000]# ls -lrt
total 1
drwxr-x---. 2 db2inst1 db2iadml 4 Jun 30 08:01 C0000000
[root@spp-db2-01 LOGSTREAM0000]# cd C0000000/
[root@spp-db2-01 C0000000]# ls -lrt
total 3
-rw-r-----. 1 db2inst1 db2iadml 12288 Jun 30 08:00 S0002034.LOG
-rw-r-----. 1 db2inst1 db2iadml 12288 Jun 30 08:01 S0002035.LOG
```

Performing a single Db2 Backup

Start the Db2 SLA Policy Backup by clicking **Run** in the Db2 backup window, as shown in Figure 4-7. The Db2 backup of the selected database then starts.

Note: The Run button is enabled only for a single database backup. Also, the database must have an SLA policy applied.

To run an on-demand backup job for multiple Db2 databases that are associated with an SLA policy, click **Create job**. Then, select **Ad hoc backup** and follow the instructions.

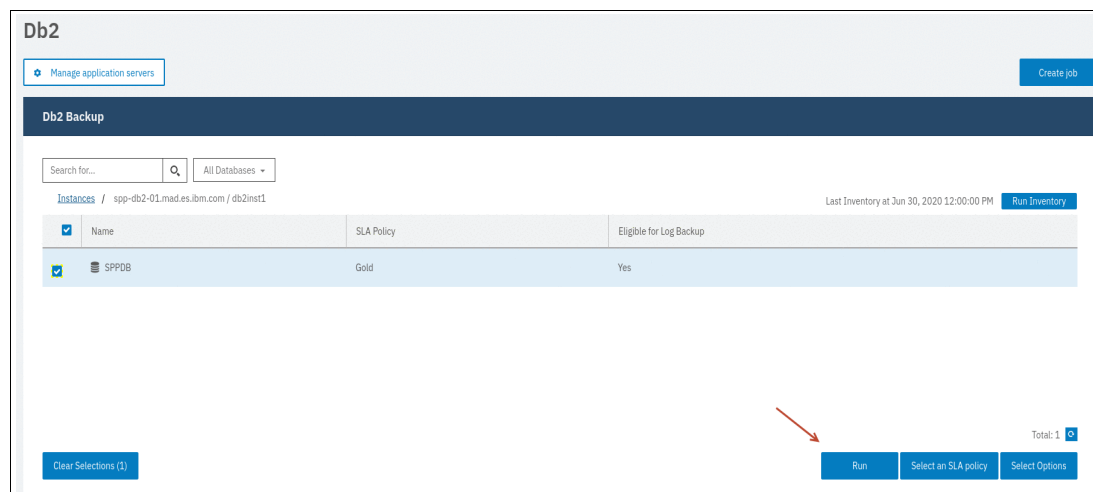


Figure 4-7 Start the Db2 backup

Log in to one of the Db2 database server by using SSH and check where the backup is created. Run the **df -h** command, as shown in Example 4-6, and review the vSnap server volumes.

Example 4-6 vSnap server volumes for data and log backup

```
[root@spp-db2-01 C0000000]# df -h
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/rhel-root      38G       9.8G   28G   27% /
devtmpfs                  1.9G         0 1.9G    0% /dev
tmpfs                     1.9G      12K 1.9G    1% /dev/shm
tmpfs                     1.9G      26M 1.9G    2% /run
tmpfs                     1.9G         0 1.9G    0% /sys/fs/cgroup
/dev/sda1                 1014M     143M   872M   15% /boot
/dev/mapper/rhel-home      19G      989M    18G    6% /home
/dev/mapper/db2-data       15G      240M    14G    2% /db2_data
/dev/mapper/db2-log        9.8G     116M    9.1G    2% /db2_log
tmpfs                     380M      12K   380M    1% /run/user/42
192.168.5.234:/vsnap/vpool1/fs148 898G    128K   898G    1% /mnt/spp/vsnap/vpool1/fs148/192_168_5_234
tmpfs                     380M         0   380M    0% /run/user/0
tmpfs                     380M         0   380M    0% /run/user/1003
192.168.5.234:/vsnap/vpool1/fs113 898G     25M   898G    1% /mnt/spp/vsnap/vpool1/fs113/192_168_5_234
```

One vSnap server log volume is used for multiple Db2 partitions. A single log archive volume on vSnap server is sufficient because the log paths are orthogonal because of the Db2 NODEXXX element in each of the log paths. Log volumes stay mounted on the Db2 application server. When the backup completes, you see the status **Completed**, as shown in Figure 4-8 on page 61.

SLA Policy Status										
Filter Job Log: Info Warning Error Detail Summary										
Policy	Frequency	Total	Succeeded	Failed	Next Run	Status	Policy Options			
Gold	Every 12 Hours	1	1	0	Jul 1, 2020 12:30:00 AM	Idle	Actions			
Log	Start Time	End Time	Duration	Status	Sub Policy Type	Total	Succeeded	Failed		
	Jun 30, 2020 5:41:48 PM	Jun 30, 2020 5:47:35 PM	0hr 5mins 47secs	Completed	Backup	1	1	0		

Figure 4-8 Db2 backup job completed

After you select an SLA policy for your Db2 backup job, you can choose to configure extra options for that job. Other SLA options include running scripts and forcing a full base backup.

For more information, see 1.4, “Database backup with pre-script and post-script” on page 15.

IBM Spectrum Protect Plus automatically deletes older transactional logs after a successful database backup. This action ensures that the capacity of the log archive volume is not compromised by retention of older log files. These truncated log files are stored in the vSnap server repository until the corresponding backup expires and is deleted. The retention period of database backups is defined in the assigned SLA policy.

4.3.3 Restoring Db2 databases

IBM Spectrum Protect features a restore wizard (see Figure 4-9) that simplifies the restore for virtual machines (VMware and Hyper-V) and application data (Db2, Exchange, MongoDB, Oracle, and SQL) to ensure that you can meet all of your recovery and reuse scenarios. Start the restore wizard by clicking **Jobs and Operations** → **Create Job** → **Restore** → **Db2**.

Db2 database restore with IBM Spectrum Protect Plus supports several restore methods that are explained in the following sections.

The following parameters control the restore or data reuse activity:

- ▶ **Type of Restore:**
 - On-Demand Snapshot
 - On-Demand Point in Time
 - Recurring
- ▶ **Restore Method:**
 - A production restore overwrites the original database or creates a database copy on an alternate host. Production is the only restore method that is available for restore operations to the original location.
 - A test restore mounts the vSnap server directories with a database backup to an alternative database server, recovers and opens the database. You can choose to rename the database.
 - An instant access restore also mounts the vSnap server directories with a database backup to a database server, but does not recover or open the database
- ▶ **Destination:**
 - Restore to the original instance on original host.
 - Restore to the original instance on alternate host, optionally with a new database name.

- Restore to an alternate instance on an alternate host with alternate database name. It is required to specify a new name for the database and the original instance must exist on the target host.

Important: For all restore operations, Db2 must be at the same version level on the source and target hosts. In addition to that requirement, you must ensure that an instance with the same name as the instance that is being restored exists on each host. This requirement applies when the target instance has the same name, and when the names are different. In order for the restore operation to succeed, both instances must be provisioned, one with original name and the other with the new name.

The combination of these selections define which action to perform, including the following examples:

- ▶ Restore a database restore and optionally overwrite an existing database
- ▶ Establish a copy of a previously backed up database (DevOps)
- ▶ Get access to the database files (data and metadata) of a previous backup

In this example scenario, a production restore is performed on a multi partitioned Db2 database version 10.5. In Figure 4-9, the first page of the restore wizard is displayed and the user has to choose **Db2**.

Note: When you are restoring a multi-partitioned database to an alternate location, ensure that the target instance is configured with the same partition numbers as the original instance. All of those partitions must be on a single host.

For more information about database examples that show a test restore or instant access, see “IBM Spectrum Protect Plus database restore and data reuse” on page 7.

As shown in Figure 4-9, the user must select the Db2 database that requires a restore.

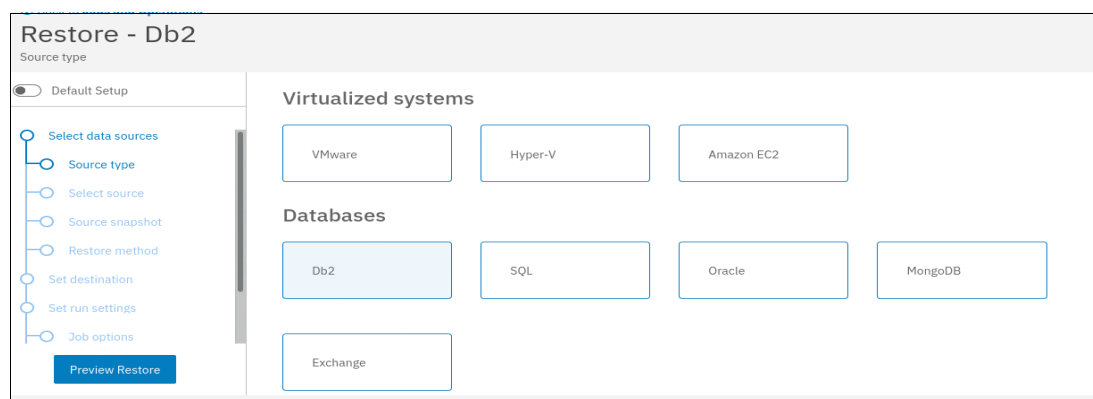


Figure 4-9 Spectrum Protect Plus restore wizard

Important: The Spectrum Protect Plus Db2 agent will not overwrite existing active logs unless you have chosen recovery mode to end of backup. For the other modes (no recovery, recovery to end of logs, and recovery to PIT) it will not restore logs from the snapshot backup.

Speaking in terms of Db2 restore commands, only the **recover to end of backup** mode adds the option `logtarget include force`. For the other modes, it is omitted and it defaults to `LOGTARGET EXCLUDE`.

By selecting the blue plus sign, a backup is associated with the database, as shown in Figure 4-10.

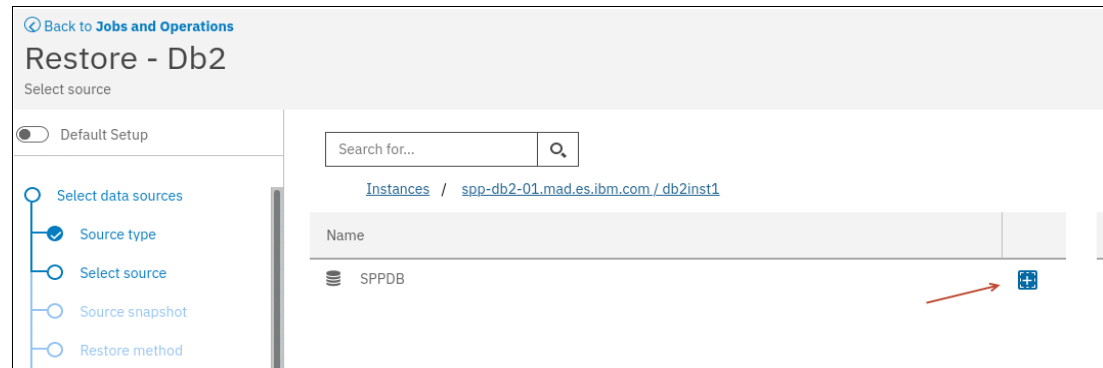


Figure 4-10 Db2 source database

In the IBM Spectrum Protect Plus restore wizard, the following parameters must be selected to start a traditional database restore that overwrites the database:

- ▶ Restore type: On-Demand Point in Time
- ▶ Restore location type and location can vary. Here, we use Site and Primary.
- ▶ Restore method: Production
- ▶ Destination: Restore to original Instance
- ▶ Restore Method: Production
- ▶ Job options: Overwrite existing database

The next step is to select the type of restore, as shown in Figure 4-11. Here, On-demand Point in Time was selected.

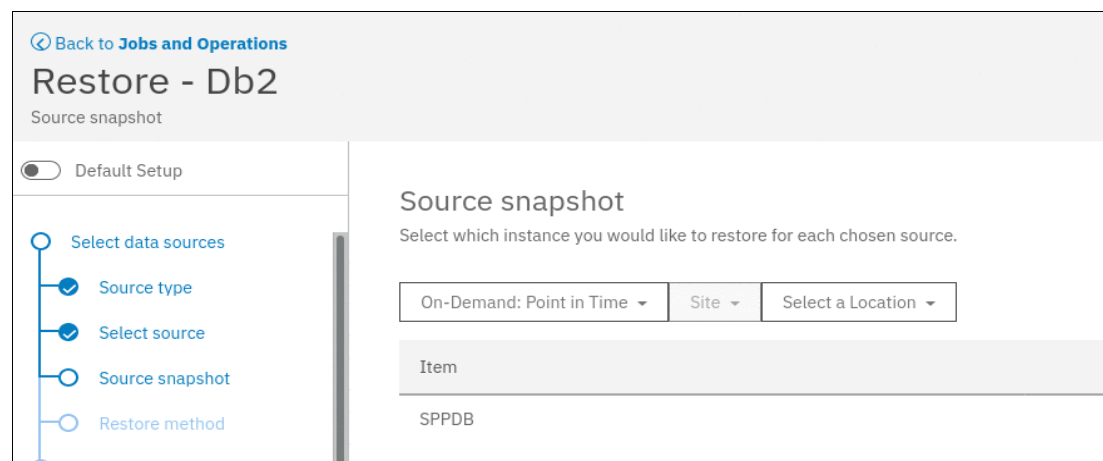


Figure 4-11 Select type of restore: On-Demand: Point in Time

Then, choose a restore location, as shown in Figure 4-12. These settings depend on your specific environment, which can include an object storage or vSnap server location, or a secondary site that you use for replication. In our example, we chose Site and Primary, as shown in Figure 4-12.

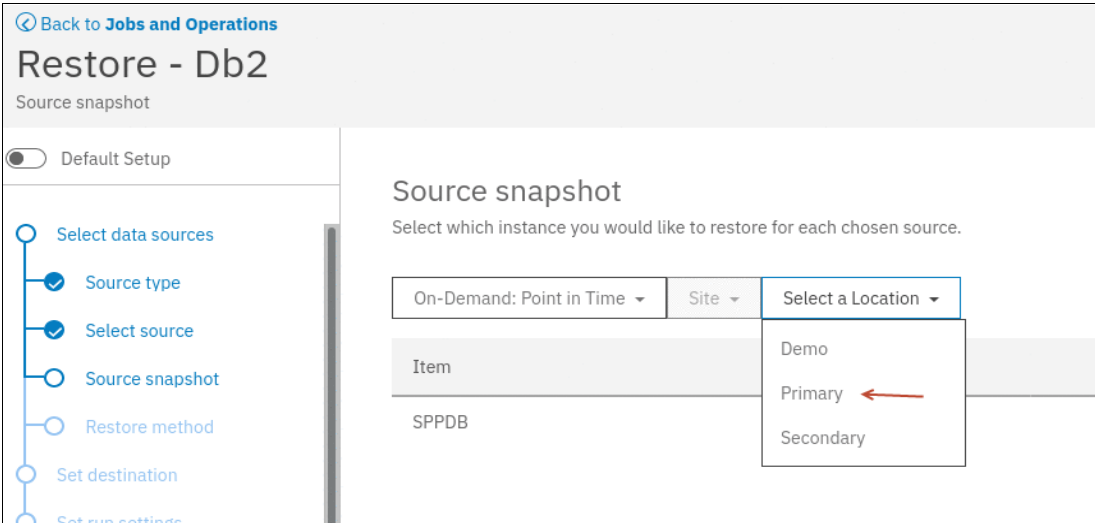


Figure 4-12 Select restore location type

There are three restore methods available, as shown in Figure 4-13. In our scenario we are choosing a *Production* restore.

Production restore

A production restore either overwrites the original database or creates a database copy on an alternate host and optionally in an alternate database instance.

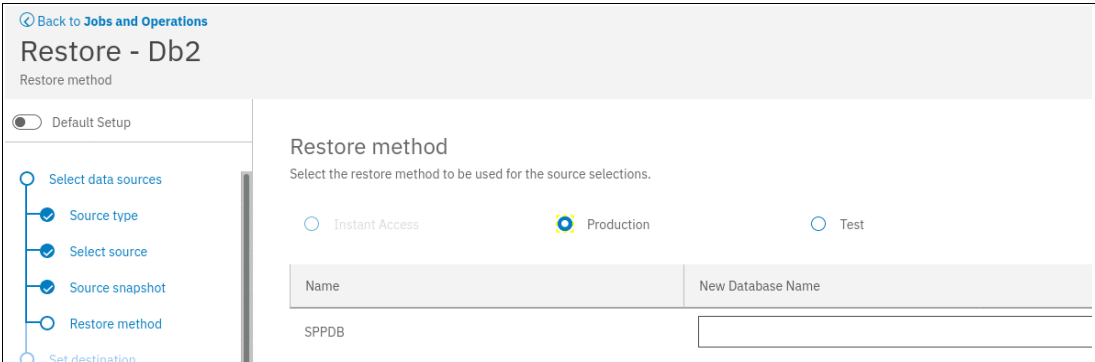


Figure 4-13 Select restore method

As shown in Figure 4-14, click **Restore to original instance** to restore to the Db2 production server.

The screenshot shows the 'Restore - Db2' wizard with the 'Set destination' step selected in the left-hand navigation pane. The main content area is titled 'Set destination' and includes the instruction 'Select which instance you would like to restore for each chosen source.' There are two radio button options: 'Restore to original instance' (which is selected) and 'Restore to alternate instance'. Below these options, a note states: 'Restoring to original location will override any custom database names that have been defined.'

Figure 4-14 Select destination for the restore

For Devops scenarios, it is possible to create a restore job that runs periodically at a specific time. In our scenario, we create an on-demand restore job that runs only once. As a further restore job option, we select **Overwrite existing databases**, as shown in Figure 4-15.

The screenshot shows the 'Restore - Db2' wizard with the 'Job options' step selected in the left-hand navigation pane. The main content area is titled 'Job options' and includes the instruction 'Configure the options for this restore job.' Under the 'Recovery Options' section, there are three radio button options: 'No Recovery', 'Recover until end of backup', and 'Recover until end of available logs' (which is selected). Below these, there are three time selection boxes (each showing '00') and a text box containing 'Europe/Madrid'. Under the 'Application Options' section, the checkbox 'Overwrite existing databases.' is checked, and a red arrow points to this checkbox.

Figure 4-15 Specify restore job options: Overwrite existing databases

Another option, as shown in Figure 4-16, is to provide pre- and post-scripts that perform specific actions before and after the Db2 restore. Those scripts must be uploaded to IBM Spectrum Protect Plus before creating the restore job.

Back to **Jobs and Operations**

Restore - Db2

Apply scripts

☐ Advanced Setup

- ✓ Source type
- ✓ Select source
- ✓ Source snapshot
- ✓ Restore method
- ✓ Set destination

Apply scripts

Select Pre or Post scripts to run along with this restore job.

- ☐ Pre-Script
- ☐ Post-Script
- ☐ Continue job/task on script error

Figure 4-16 Specify scripts for the restore job

Finally, the Review page is displayed and after checking all values, the on-demand restore job can be submitted. See Figure 4-16. To start the on-demand restore job, click **Submit**.

Back to **Jobs and Operations**

Restore - Db2

Review

☐ Advanced Setup

- ✓ Source type
- ✓ Select source
- ✓ Source snapshot
- ✓ Restore method
- ✓ Set destination
- ✓ Set run settings
- ✓ Job options
- ✓ Apply scripts (Optional)
- Review

Review

Review your selections, and then click Submit.

Select data sources

Source type:	Db2
Selected source:	SPPDB
Source snapshot:	SPPDB - Use Latest
Restore Type:	On-Demand: Point in Time
Restore Source Type:	Site
Restore Source:	Primary
Restore method:	Production

Set destination

Destination:	Restore to original instance
--------------	------------------------------

Preview Restore

Figure 4-17 Review of restore job parameters.

The restore job can be monitored by selecting **Jobs and Operations** → **Running Jobs**, as shown in Figure 4-18. When the restore job finishes, it is removed from the Running Jobs list.

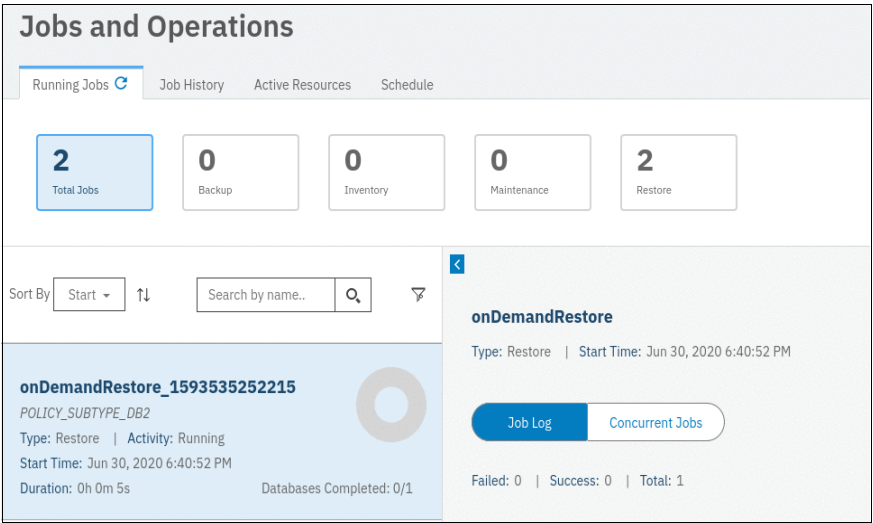


Figure 4-18 Monitor the Db2 restore job



Backing up and restoring SQL Server

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

This chapter includes the following topics:

- ▶ 5.1, “IBM Spectrum Protect Plus SQL Server features” on page 70
- ▶ 5.2, “Prerequisites for SQL Server databases” on page 70
- ▶ 5.3, “Protecting SQL Server databases” on page 72
- ▶ 5.4, “Restoring SQL Server databases” on page 84

5.1 IBM Spectrum Protect Plus SQL Server features

In this section, we describe the features of IBM Spectrum Protect Plus with Microsoft SQL Server. As of July 2020, the following features are supported:

- ▶ Backup, restore, and recovery of stand-alone/failover cluster and AlwaysOn Availability Groups (AAGs)
- ▶ Incremental forever database and log backups, including log truncation
- ▶ Automatic discovery of SQL installations on registered servers
- ▶ Parallel ad-hoc SQL database backups
- ▶ Production restore (database is restored by copying data):
 - To original location
 - To alternative location (that is, alternative source path)
- ▶ Test restore (database is restored in-place without data movement)
- ▶ Instant access restore (database is restored, but not opened)
- ▶ Restore to alternate instance and / or database name
- ▶ No recovery (does not require log backup being enabled)
- ▶ Recover to specific point-in-time (requires log backups enabled)
- ▶ Recover until end of backup (does not require log backup enabled)
- ▶ Recover standby mode (requires log backups enabled)
- ▶ Microsoft SQL Server restore with file renaming

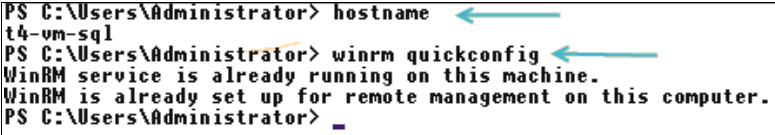
5.2 Prerequisites for SQL Server databases

Before protecting the SQL Server environment with IBM Spectrum Protect Plus, check that all the prerequisites for IBM Spectrum Protect Plus are fulfilled. The main prerequisites must be met:

- ▶ A supported Microsoft SQL Server versions (Standalone and Enterprise editions):
 - SQL Server 2008 R2 SP3
 - SQL Server 2012
 - SQL Server 2012 SP2
 - SQL Server 2014, SQL Server 2016, SQL Server 2017, SQL Server 2019
- ▶ A supported version of the Windows operating system: Windows Server 2012 R2, Windows Server 2016, Windows Server 2019

The following conditions and settings are also important prerequisites:

- ▶ The Windows Remote Management (WinRM) must be enabled by running the command **winrm quickconfig** in a Windows command line session on the guest Microsoft SQL Server system, as shown in Figure 5-1.



```
PS C:\Users\Administrator> hostname
t4-vm-sql
PS C:\Users\Administrator> winrm quickconfig
WinRM service is already running on this machine.
WinRM is already set up for remote management on this computer.
PS C:\Users\Administrator>
```

Figure 5-1 Windows Remote Shell configured into Microsoft SQL Server

- A Microsoft iSCSI Initiator service must be enabled and running on the Microsoft SQL server system, as shown in Figure 5-2.

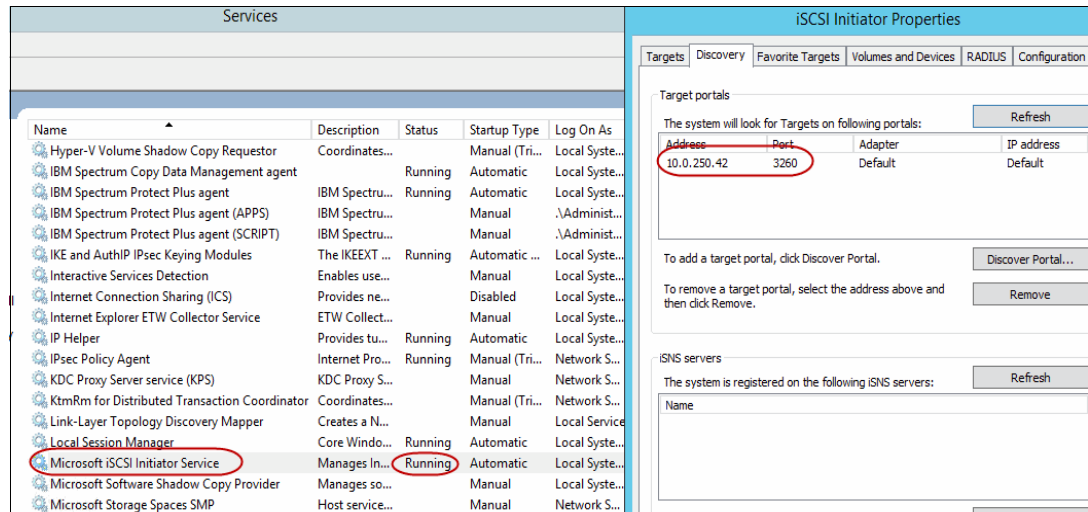


Figure 5-2 Microsoft iSCSI Initiator running on Microsoft SQL Server

- An IBM Spectrum Protect Plus agent user must have “Log on as a service” rights on the SQL application server.
- The login credentials must have public and sysadmin permissions enabled, plus permission to access cluster resources in a SQL Server AAGs environment.
- To perform log backups, the SQL Server agent service user must be a local Windows administrator and must have the sysadmin permission enabled to manage SQL Server agent jobs.
- The host name of the IBM Spectrum Protect Plus appliance should be resolvable from the SQL application servers.
- The Microsoft SQL Server Guest Network Adapter Backup must have the option “Client for Microsoft Networks” enabled to prevent CIFS share issues, when Databases SQL Backup Logs are defined and configured, as shown in Figure 5-3.

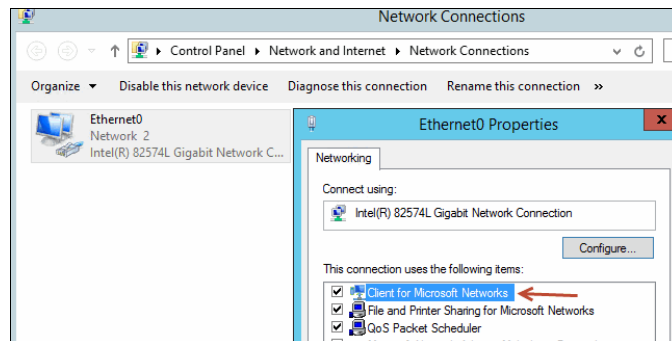


Figure 5-3 Client for Microsoft Networks option enabled into Microsoft SQL Server

For more information about the SQL Server database prerequisites, see *IBM Spectrum Protect Plus- All Requirements*, which is available at [this web page](#).

5.3 Protecting SQL Server databases

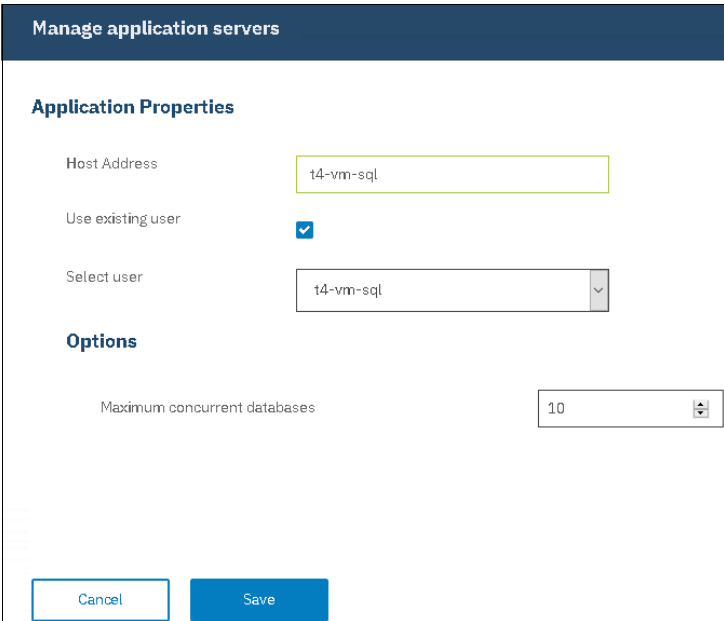
To protect SQL Server with IBM Spectrum Protect Plus the database server has to be registered in IBM Spectrum Protect Plus so that IBM Spectrum Protect Plus can discover the SQL Server databases. To start the backup, the SQL Server database always has to be assigned to an SLA policy.

5.3.1 Register the SQL Server

Before IBM Spectrum Protect Plus can manage a SQL Server database, the SQL application server has to be registered in IBM Spectrum Protect Plus. To register a SQL application server, complete the following steps:

In the IBM Spectrum Protect Plus GUI navigation pane, click **Manage Protection** → **Databases** → **SQL** → **Manage Application Servers** → **Add Application Server**.

Enter the required login credentials for the SQL application server, as shown in Figure 5-4. In this example, the IBM Spectrum Protect Plus admin includes predefined the SQL Server Admin in **Accounts** → **Identity** → **Add Identity**.



The screenshot shows the 'Manage application servers' dialog box. It has a title bar 'Manage application servers'. The main content area is divided into two sections: 'Application Properties' and 'Options'. In the 'Application Properties' section, there are three fields: 'Host Address' with the value 't4-vm-sql', 'Use existing user' which is checked, and 'Select user' which is a dropdown menu showing 't4-vm-sql'. In the 'Options' section, there is a field for 'Maximum concurrent databases' with the value '10'. At the bottom of the dialog are two buttons: 'Cancel' and 'Save'.

Figure 5-4 Register SQL application server

Perform a configuration test of the newly assigned SQL Server in IBM Spectrum Protect Plus, as shown in Figure 5-5 on page 73.

If SQL application servers are attached to a domain, a user name in the format domain\Name must be used. If a user is a local administrator, the format .\local administrator> must be used.

For failover clusters and AAGs, each node must be registered by name or IP address. If fully qualified domain names are used, they must be resolvable and routeable from IBM Spectrum Protect Plus.

Test result of t4-vm-sql

1. Physical - Basic Test for physical host network configuration

Name	Description	Status	Message
Socket Connection Test	Host must allow socket connection on port 5985 for Windows Server	✓	

2. Remote - Remote executor test for session creation and remote agent deployment

Name	Description	Status	Message
Remote Session Test	Latest remote agent must be installed on host, port must be open to create a session to WinRM service, and remote agent must be running on host with administrative privileges.	✓	
Remote Agent Execute Test	Remote agent must be configured correctly using user credentials with sufficient rights including log on as a service privilege.	✓	

3. WINDOWS - Basic Windows pre-requisites for file and volume operations

Name	Description	Status	Message
Local Administrator Privilege	User must have local administrator privilege	✓	
HTTPS connection to SPP appliance	Test HTTPS connection from the Windows server to SPP appliance	✓	

4. SQL - SQL pre-requisites for log backup

Name	Description	Status	Message
SQL pre-requisites for log backup	SQL log backup requires group policy configuration: [LmCompatibilityLevel="Not Defined"]	✓	
SQL pre-requisites for log backup	SQL log backup requires group policy configuration: [DisableDomainCreds = "0"]	✓	

5. SQL - SQL pre-requisites for application operations

Name	Description	Status	Message
Test Operating System of SQL Server	Windows 2012R2 or later is required to protect a SQL Server	✓	
SQL Instance(s) Test	SQL Instance(s) have been discovered	✓	
SQL Inventory Permissions of instance MSSQLSERVER	User must be assigned with sysadmin role	✓	

Figure 5-5 SQL Server configuration test results

5.3.2 Defining an SQL Server backup job

Before starting a backup of an SQL Server database, the SQL Server database has to be assigned to one or more SLA policies. There are four predefined policies (*Demo*, *Gold*, *Silver*, and *Bronze*) available for selection. You can use these policies or specify new policies that meet specific requirements.

Assign the selected SQL Server database to an SLA policy to create a backup job. SQL Server backups run in a “Base-Once-Incremental-Forever” scheme. During the initial base (full) backup, IBM Spectrum Protect Plus creates a vSnap server volume and mounts it to the SQL application server over iSCSI.

Note: An iSCSI route must be enabled between the SQL Server and vSnap server. For more information, see [this web page](#).

Optionally, the SQL Server admin can click the *Select Options* button to enable Log Backup, as shown in Figure 5-6. With log backup enabled, IBM Spectrum Protect Plus manages the log backup by using the SQL Server agent service.

To complete log backups, the SQL Server Agent service user must be a local Windows administrator and must have the sysadmin permission enabled to manage SQL Server agent jobs. Also, the SQL VSS Writer service running on the local SQL Server system must be started from a local system user.

The agent uses the administrator account to enable and access log backup jobs. The IBM Spectrum Protect Plus SQL Server agent service user must also be the same as the SQL Server service and SQL Server agent service account for every SQL Server instance to be protected.

Enable Log Backup

Log Backup Frequency and Start Time

30 Minutes

0 00

☒ Backup database files one at a time using parallel streams. ☐ Backup database files in parallel using parallel streams.

Maximum Parallel Streams per Database

2

Save

Figure 5-6 Enable SQL Server log backup

Set the maximum number of data streams per database to the backup storage. This setting applies to each database in the job definition. Databases can be backed up in parallel if the value of the option is set to 1. Multiple parallel streams might improve backup speed, but high bandwidth consumption might affect overall system performance.

The SQL Server backup job status can be monitored in the **Jobs and Operations** → **Schedule** panel, as shown in Figure 5-7.

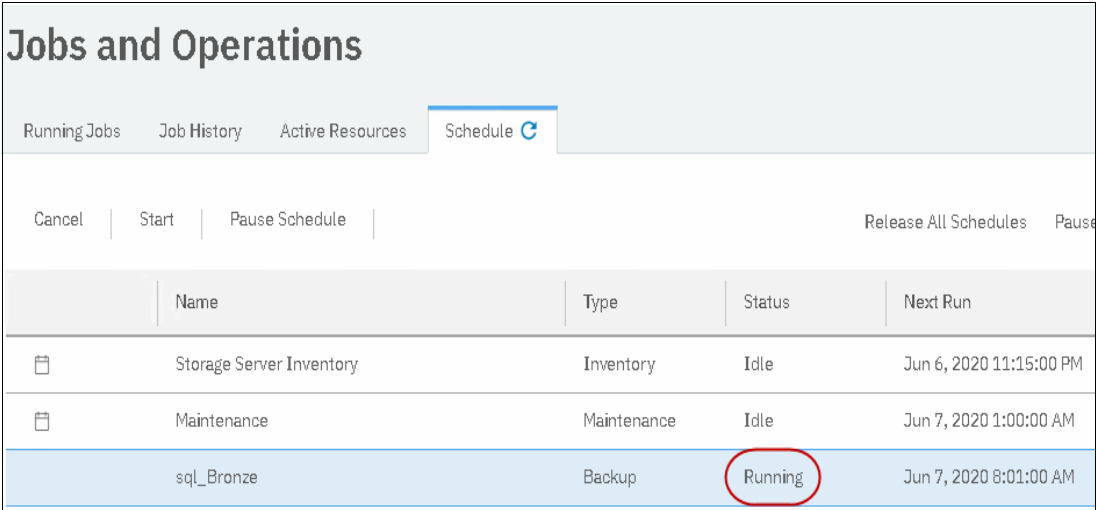


Figure 5-7 SQL Server backup job with Status: Running

It also can be monitored by selecting **Jobs and Operations** → **Running Jobs** → **Progress**, as shown in Figure 5-8.

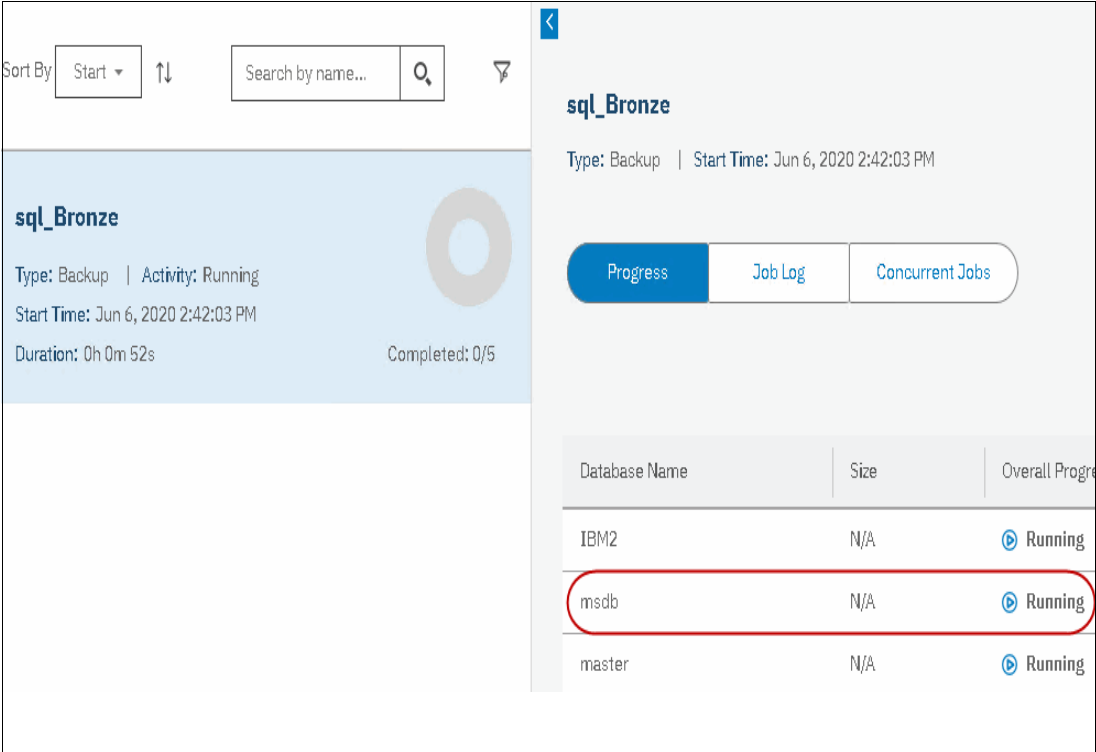


Figure 5-8 SQL Server Databases Backup Status: Running

Note: The Microsoft SQL Server agent sets the VSS backup type to COPY_ONLY for all database backups.

SQL Server backup workflow

Some readers might be interested in the detailed workflow of a SQL Server base backup. Here are the internal steps of the backup workflow:

1. Discover the SQL Server client to get the current SQL Server instance, database information, cluster information, availability group information (for AlwaysOn), disk, and volume information.
2. Request the SQL Server iSCSI initiator information. Create an iSCSI LUN on vSnap server, map the LUN to the SQL Server client iSCSI initiator.
3. Prepare vSnap server LUN for Backup:
 - a. Rescan SQL Server.
 - b. Identify the iSCSI LUN provisioned as backup target.
 - c. Clear the readonly flag.
 - d. Bring disk online.
 - e. Initialize the disk.
 - f. Create GPT partition table.
 - g. Create a primary partition.
 - h. Bring the partition online.
 - i. Quick format the volume.
 - j. Label the volume with "SPPB_*".
 - k. Collect the volume GUID, serial number information for cataloging.
 - l. Mount the backup target volume to a volume mount point on
C:\ProgramData\SPP\mnt\subfolder.
4. Check and enable USN Journaling for block level incremental capability.
5. Backup: VSS Snapshots
 - a. Start a VSS backup request & get VSS writer metadata.
 - b. Collect the source volume information of the selected SQL Server databases.
 - c. Add the instance and database to the application backup list. Add covering the volumes to the snapshot set.
 - d. Commit snapshot set.
 - e. Copy the database files from the VSS shadow copy to vSnap server iSCSI backup target.
 - f. Notify the writer of the backup status, save the backup document.
 - g. Report the backup status and backup metadata.
6. For incremental backups, use USN Journal to identify changed blocks since the last successful snapshot. Copy changed blocks from the VSS shadow copy to vSnap server iSCSI backup target.
7. Merge those changes into the last snapshot in the vSnap server.
8. Unmount "C:\ProgramData\SPP\mnt\subfolder" and Unmap the iSCSI LUN.
9. Rescan on SQL server and ensure cleanup was successful.
10. Take vSnap server snapshot of backup volume and log share volume (if applicable).
11. Catalog the backup metadata to the IBM Spectrum Protect Plus Server.

5.3.3 SQL database backups logs

IBM Spectrum Protect Plus version 10.1.6 allows archiving of log files for databases that contain committed transactions log data. Transactions data can be used to run a roll forward recovery process as part of a restore operation. The use of archive log backups enhances the recovery point objective for data.

Depending on what type of SQL backup log is required, it can be configured by using one of the following methods:

- With Truncate SQL Logs option activated on Virtualized Systems wizard

Note: If you multiple backup solutions are performing log truncation, you can establish discontinuity in the log chain. It must be ensured that the log truncation occurs only once during a backup.

With this option activated, logs might be truncated during the VM Backup as a result of log clearing. In this case, you can restore a VM only; a roll forward of the transaction log data cannot be performed.

The option to truncate SQL logs can be defined under **Manage Protection** → **Virtualized Systems** → **VMware or Hyper-V** → **Select VM Server** → **Select Options**, as shown in Figure 5-9.

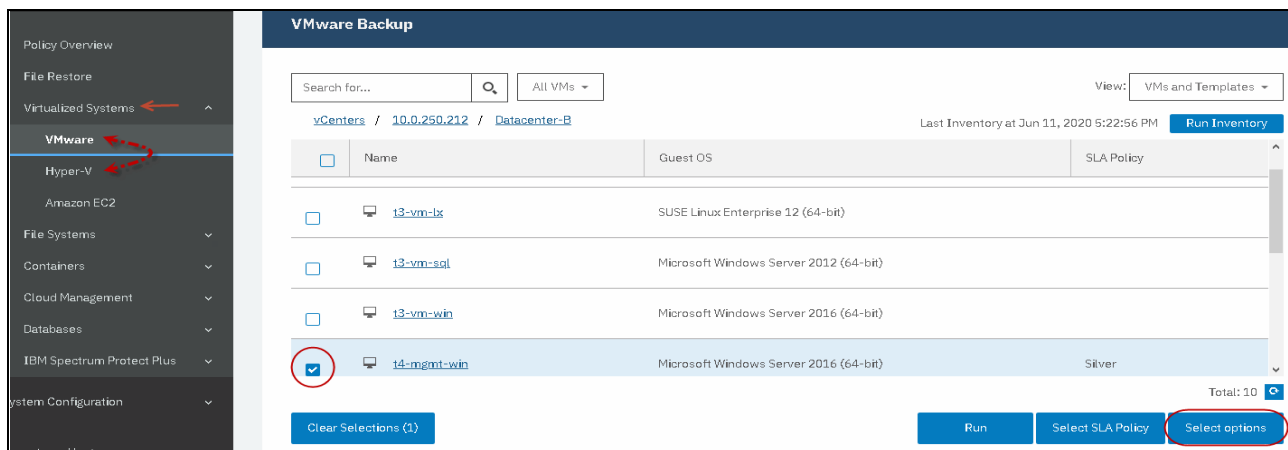


Figure 5-9 Selected options for VM Server

Under **Agent Options**, select the **Truncate SQL Logs** option, as shown in Figure 5-10. Click **Save**.

Agent Options

☒ Truncate SQL logs

☒ Catalog file metadata

Exclude Files

C:\Program Files\C:\Program Files (x86);C:\Windows;C:\winnt;C:\Drivers\tmp\usr\bin\bin\sbin

☒ User

☐ SSH Key

Use existing user ☒

Select user

escc:t4admin_a25b9e25c02t

Save Test

Figure 5-10 Truncate SQL Logs Option Enabled

Note: For more information about how to enable Log Truncation, see *Protecting Virtualized Systems - Backing up VMware / Hyper-V data Guides*, which is available at IBM Knowledge Center:

- ▶ [Backing up VMware data](#)
- ▶ [Backing up Hyper-V data](#)

▶ Enable Log Backup option

You can configure log backups by using database SQL backup. The use of archive log backups enhances the recovery point objective for your data. Enabling this option allows roll forward recovery when you restore Microsoft SQL Server data.

The option to enable SQL log backup can be defined under **Manage Protection** → **Databases** → **SQL** → **Select SQL Instance** → **Select Options**, as shown in Figure 5-11.

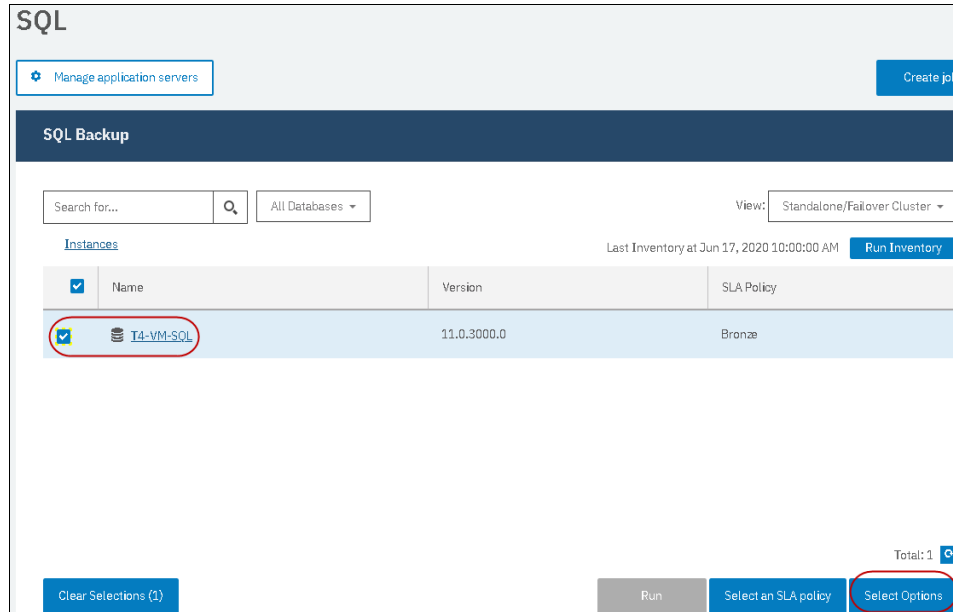


Figure 5-11 Select Options for SQL Instance

Under Options, select **Enable Log Backup** and define a Log Backup Frequency, as shown in Figure 5-12.

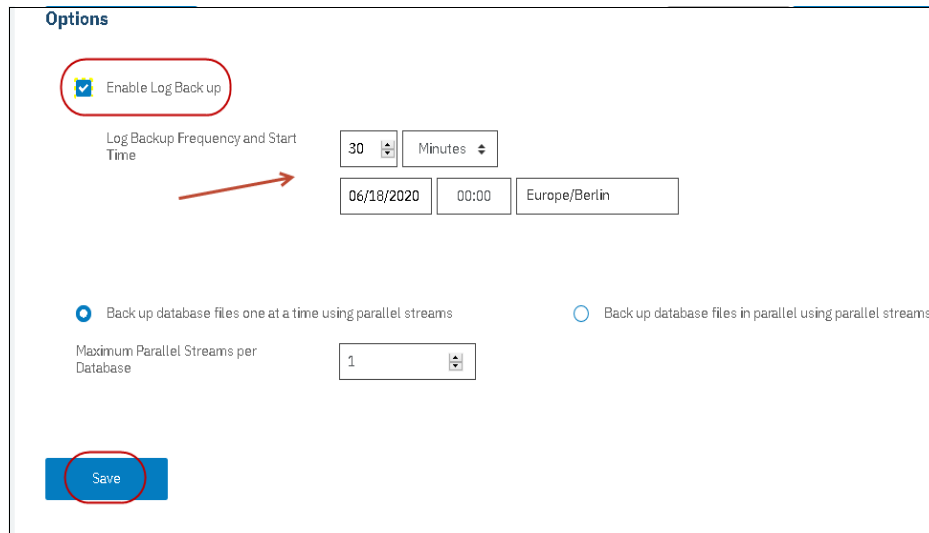


Figure 5-12 Enable Log Backup option enabled

The enabled SQL log backup schedule option can also be reviewed in the Microsoft SQL Server system, as shown in Figure 5-13 on page 80 under **Task Scheduler** → **Task Scheduler Library** → **IBM** → **SPP Windows Agent**.

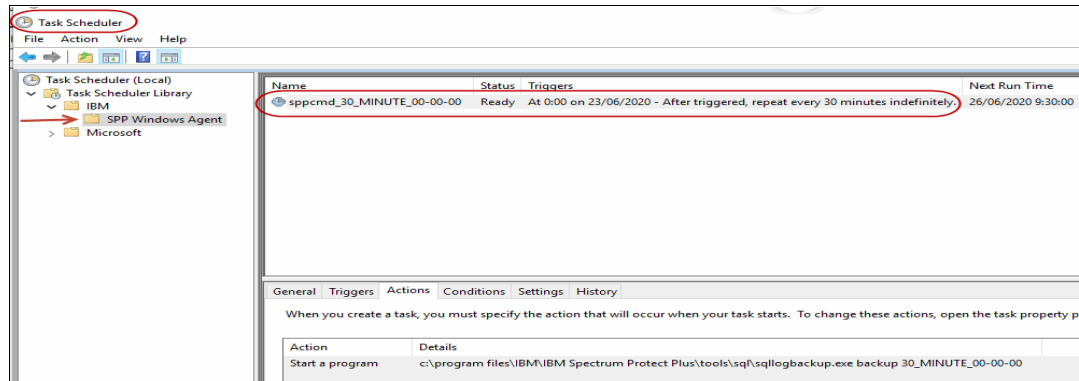


Figure 5-13 SQL Log Backup Scheduler in Microsoft SQL Server

Note: To run the Windows log backup task, the IBM Spectrum Protect Plus agent user must have the Log On As Batch Job assignment privilege.

Note: For more information about how to enable Log backup, see *Backing Up SQL Server Data Guide*, which is available at [IBM Knowledge Center](#).

5.3.4 vSnap commands used to manage SQL database backups logs

The IBM Spectrum Protect Plus agent maps the LUN to the SQL server and mounts the NTFS volume to perform the backup. If log backups are enabled, IBM Spectrum Protect Plus creates a separate vSnap server volume and creates a CIFS share on that volume. Log backup transaction files are copied to this share according to the schedule that was created for the log backup.

If the SQL backup SLA job is running, you can see a share smb on the file system, which is associated with the SLA. The vSnap CLI command **vsnap share show**, as shown in Example 5-1, lists the active share where the Volume ID 1 and the file system name /vpool1_fs2 can be identified.

Example 5-1 Active Share

```
[vsnapadmin@t4-spp-vsnap vsnap]$ vsnap share show
ID | TYPE | PARENT VOL | PARTNER ID | NAME
-----
1 | smb | 2 | N/A | vpool1_fs2

[vsnapadmin@t4-spp-vsnap vsnap]$ vsnap share show --id 1

ID: 1
NAME: vpool1_fs2
SHARE TYPE: smb
VOLUME ID: 2
PARTNER ID: N/A
CREATED: 2020-06-03 12:06:15 UTC
UPDATED: 2020-06-18 06:00:57 UTC
SHARE OPTIONS:
  ALLOWED HOSTS:
    10.0.250.46
```


The shared volume is used to transfer the backup data from the database to the vSnap server.

After the log backup completes, log backup transaction files are copied to this share, as shown in Example 5-2.

Example 5-2 Log Backup Transaction Files Copied into vSnap

```
[vsnapadmin@t4-spp-vsnap vsnap]$ df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                  16G         0    16G   0% /dev
tmpfs                     16G         0    16G   0% /dev/shm
tmpfs                     16G      1.1G    15G   7% /run
tmpfs                     16G         0    16G   0% /sys/fs/cgroup
/dev/mapper/lg_os-lv_root   35G      3.5G    31G  11% /
/dev/mapper/lg_data-lv_home 997M      33M   965M   4% /home
/dev/mapper/vsnapdata-vsnapdata1v 126G      33M  126G   1% /opt/vsnap-data
/dev/mapper/lg_os-lv_tmp    9.8G      33M   9.8G   1% /tmp
/dev/sda1                 969M     166M   737M  19% /boot
/dev/mapper/lg_os-lv_var    12G     116M    12G   1% /var
/dev/mapper/lg_os-lv_var_log 3.0G     145M   2.8G   5% /var/log
/dev/mapper/lg_os-lv_var_log_audit 497M      61M   436M  13% /var/log/audit
/dev/mapper/lg_os-lv_var_tmp 997M      33M   965M   4% /var/tmp
tmpfs                     3.1G         0    3.1G   0% /run/user/1001
vpool1                    79G      128K    79G   1% /vsnap/vpool1
vpool1/fs1                79G      128K    79G   1% /vsnap/vpool1/fs1
vpool1/fs2                79G      128K    79G   1% /vsnap/vpool1/fs2
vpool1/fn3                79G       15M    79G   1% /vsnap/vpool1/fn3
vpool1/fs6                97G       18G    79G  19% /vsnap/vpool1/fs6
[vsnapadmin@t4-spp-vsnap vsnap]$ cd /vsnap/vpool1/fs2

[vsnapadmin@t4-spp-vsnap fs2]$ ls -lrt
total 3
drwxrwxrwx. 2 vsnap vsnap 10 Jun 18 09:59 T4-VM-SQL_ESCC
drwxrwxrwx. 2 vsnap vsnap 10 Jun 18 09:59 T4-VM-SQL_IBM2
[vsnapadmin@t4-spp-vsnap fs2]$ cd T4-VM-SQL_ESCC

[vsnapadmin@t4-spp-vsnap T4-VM-SQL_ESCC]$ ls -lrt
total 32
-rw-r--r--. 1 vsnap vsnap 86528 Jun 17 09:59 T4-VM-SQL_ESCC_log_z_1592388001.trn
-rw-r--r--. 1 vsnap vsnap 86528 Jun 17 13:59 T4-VM-SQL_ESCC_log_z_1592402401.trn
-rw-r--r--. 1 vsnap vsnap 86528 Jun 17 17:59 T4-VM-SQL_ESCC_log_z_1592416801.trn
-rw-r--r--. 1 vsnap vsnap 86528 Jun 17 21:59 T4-VM-SQL_ESCC_log_z_1592431201.trn
-rw-r--r--. 1 vsnap vsnap 86528 Jun 18 01:59 T4-VM-SQL_ESCC_log_z_1592445601.trn
-rw-r--r--. 1 vsnap vsnap 86528 Jun 18 05:59 T4-VM-SQL_ESCC_log_z_1592460001.trn
-rw-r--r--. 1 vsnap vsnap 86528 Jun 18 09:59 T4-VM-SQL_ESCC_log_z_1592474401.trn
-rw-r--r--. 1 vsnap vsnap   73 Jun 18 09:59 lsn.json
[vsnapadmin@t4-spp-vsnap T4-VM-SQL_ESCC]$
```

5.3.5 Parallel ad-hoc SQL database backups

IBM Spectrum Protect Plus version 10.1.6 introduces a new ad-hoc backup wizard that simplifies the process of performing individual backups of one database without starting the complete SLA job. The wizard guides you through the backup selection. It shows the new settings options that allow you to select one or more databases with the same SLA and do their backup concurrently without collision, and synchronizes the operations at various steps.

One ad-hoc job can be started from **Manage Protection → Databases → SQL → Create Job → Ad hoc Backup → Select SLA Policy → Select Source**, as shown in Figure 5-14.

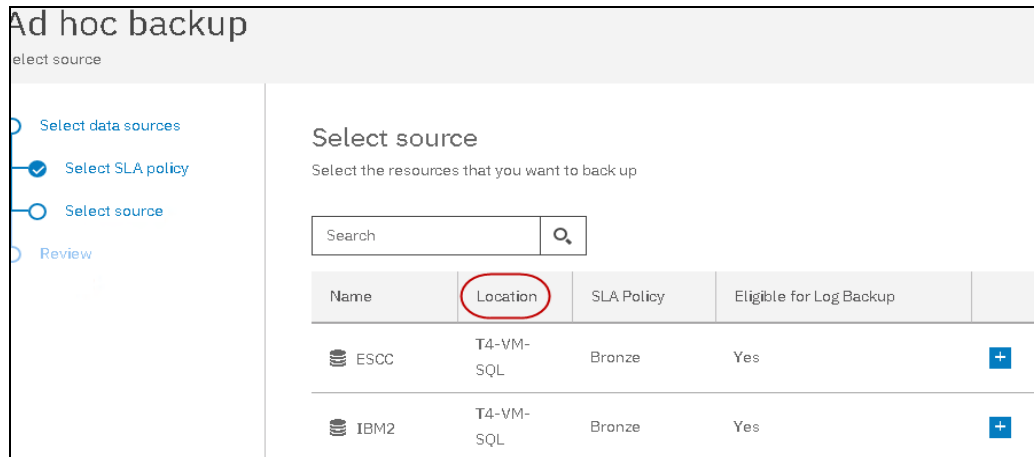


Figure 5-14 Ad Hoc SQL Backup showing Name, Location and SLA Policy

Multiple sessions from the same SLA policy can be started from **Manage Protection → Databases → SQL → Create Job → Ad hoc Backup → Select SLA Policy → Select Source**. The sessions can be monitored from **Jobs and Operation → Running Jobs**, as shown in Figure 5-15.

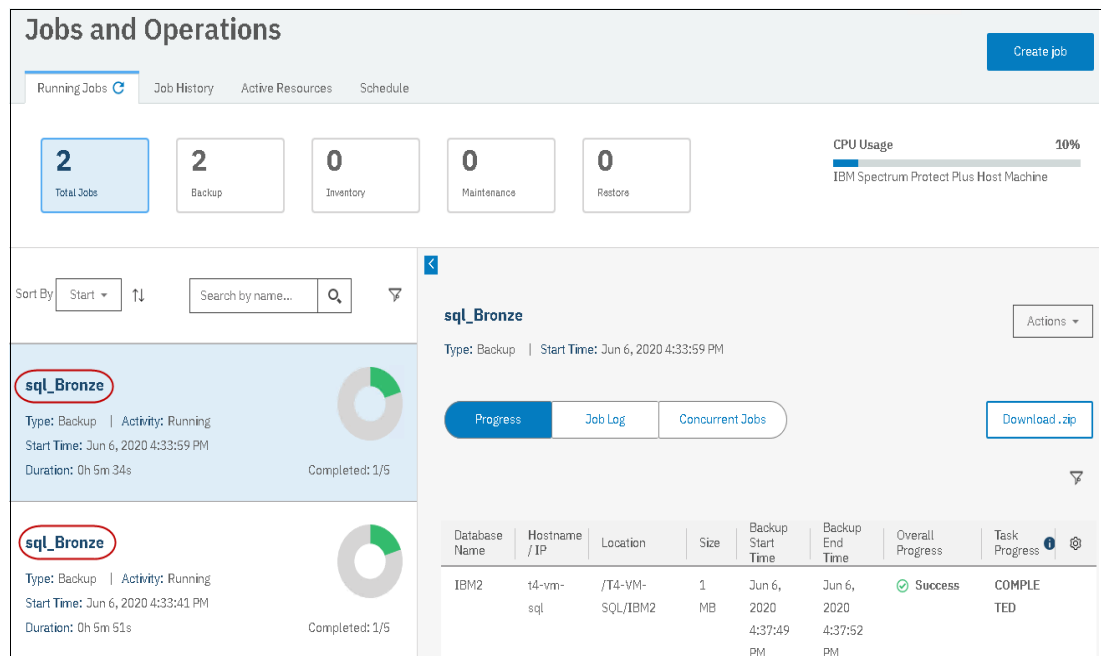


Figure 5-15 Multiple running sessions from the same SLA policy

5.3.6 SQL Server global preferences

The Global Preferences panel contains default values for parameters that apply to all IBM Spectrum Protect Plus operations. To change parameters, select **System Configuration** → **Global Preferences** → **Application**, as shown in Figure 5-16

The following options are available for SQL Server:

Note: Only users with administrator credentials can manage global preferences.

- ▶ Enable SQL Server databases restored in test mode eligible for backup
When this option is selected, SQL Server databases that were restored in test mode are available for selection in the SQL Backup pane or ad hoc backup wizard.
- ▶ Allow SQL database backup when transaction log backup chain is broken
Run a database SLA backup job when IBM Spectrum Protect Plus detects a break in the log backup chain for a database.
- ▶ Rename SQL data and log files when database is restored in production mode with new name
This options allows to rename SQL database and log files files during a production or test restore job. This field applies only when a new database name is provided during an SQL database restore job.

Application	
Enable SQL Server databases restored in test mode eligible for backup	<input checked="" type="checkbox"/>
Maximum volume size for backup target LUNs on Windows (TB)	256
Maximum backup retries(kBs)	3
Maximum concurrent servers running backups	0
Allow SQL database backup when transaction log backup chain is broken	<input checked="" type="checkbox"/>
Rename SQL data and log files when database is restored in production mode with new name	<input checked="" type="checkbox"/>

Figure 5-16 Global Preferences SQL Application Options

5.4 Restoring SQL Server databases

IBM Spectrum Protect Plus features a restore wizard that simplifies restores for virtual machines and databases. The wizard guides you through the configuration of restore types and parameters and optionally schedules a job that performs the restore.

IBM Spectrum Protect Plus treats data reuse and data recovery as a restore activity. In both cases, you must create a restore job. A restore job can be started by making one of the following selections in IBM Spectrum Protect Plus:

- ▶ **Manage Protection → Databases → SQL → Create Restore Job**
- ▶ **Jobs and Operations → Create Restore Job → Restore**

The parameters that you select during backup job creation define which is performed.

The following main parameters control the final restore or data reuse activity:

- ▶ **Type of Restore:**
 - On-Demand Snapshot
 - On-Demand Point in Time
 - Recurring
- ▶ **Restore Method:**
 - A production restore either overwrites the original database or creates a database copy with a different database name. In the latter case you must specify a new database name and the destination paths.
 - A test restore mounts the vSnap server directories with a database backup to a database server, recovers and opens the database. You can choose to rename the database.
 - An instant access restore also mounts the vSnap server directories with a database backup to a database server, but does not recover or open the database. An instant access restore of an Always On database is restored to the local destination instance.

Note: The SQL Server system databases (master, msdb, model) can be restored only with Instant Access mode in IBM Spectrum Protect Plus.

- ▶ **Destination:**
 - Restore to the original instance
 - Restore to an alternate instance

The combination of these selections define which action to perform, including the following examples:

- ▶ Perform a database restore and optionally overwrite an existing database
- ▶ Establish a copy of a previously backed up database (DevOps)
- ▶ Get access to the database files (data and metadata) of a previous backup

In the first example scenario, a Production restore of a SQL Server stand-alone database is performed by using SQL Server version 2012. As shown in Figure 5-17 on page 85, the databases ESCC and IBM2 are selected for the restore. By selecting the blue plus sign, a backup is associated with the database.

For more information about database examples that show a test restore or instant access, see “IBM Spectrum Protect Plus database restore and data reuse” on page 7.

Back to SQL

Restore - SQL

Select source

☐ Default Setup

Select data sources

Select source

Source snapshot

Restore method

Set destination

Set run settings

Job options

Select source

Select the databases to recover

Search for...

View: Standalone/Failover Cluster

[Instances](#) / [T4-VM-SQL](#)

Name		Item
ESCC		IBM2
IBM2		ESCC

Figure 5-17 Select the SQL Server backup source

Other restore parameters that must be specified are shown in Figure 5-18:

- ▶ **Restore type: On-Demand: Snapshot**
Runs a one-time restore job from a database snapshot. The restore job starts immediately upon the completion of the wizard.
- ▶ **Restore location type: Site**
The site where snapshots were backed up. The site is predefined in IBM Spectrum Protect Plus.
- ▶ **Location = Primary**
The primary site location from which to restore snapshots.

Restore - SQL

Source snapshot

☒ Advanced Setup

Select data sources

Select source

Source snapshot

Restore method

Set destination

Set run settings

Job options

Source snapshot

Select which instance you would like to restore for each chosen source.

On-Demand: Snapshot Site Primary 06/10/2020 - 06/17/2020

Item	Restore Point
ESCC	Jun 17, 2020 8:01:00 AM
IBM2	Jun 17, 2020 8:01:00 AM

Figure 5-18 Restore parameters

In production mode, the agent first restores the files from the vSnap server volume back to primary storage and then creates the new database by using the restored files. Select **Production**, as shown in Figure 5-19 and then, click **Next**.

Restore - SQL

Restore method

☐ Default Setup

Select data sources

Select source

Source snapshot

Restore method

Set destination

Set run settings

Job options

Review

Restore method

Select the restore method to be used for the source selections.

☐ Instant Access
☒ Production

	Name
<input checked="" type="checkbox"/>	IBM2
<input checked="" type="checkbox"/>	ESCC

Figure 5-19 Select the restore method

When selecting production mode, you can also specify a new folder for the restored database by expanding the database section and entering a new folder name.

In our setup, we perform an on-demand restore to the original instance, as shown in Figure 5-20.

[Back to SQL](#)

Restore - SQL

Set destination

☐ Default Setup

Select data sources

Select source

Source snapshot

Restore method

Set destination

Select which instance you would like to restore for each chosen source.

☒ Restore to original instance
☐ Restore to alternate instance

Figure 5-20 Select the restore destination

Enable the restore job to overwrite the selected database. By default, this option is not enabled, as shown in Figure 5-21 on page 87.

Note: Before you run restore operations in an SQL Server Always On environment by using the production mode with the Overwrite existing databases option, ensure that the database is not present on the replicas of the target availability group. As a prerequisite, manually clean up the original databases (to be overwritten) from all replicas of the target availability group.

Restore - SQL
Job options

Advanced Setup

☒ Select data sources
☒ Select source
☒ Source snapshot
☒ Restore method
☒ Set destination
☐ Set run settings
☐ Job options
☐ Apply scripts (Optional)
☐ Review

Preview Restore

Job options
Configure the options for this restore job.

Recovery Options

☐ No Recovery
☒ Recover until end of backup
☒ By Time
☐ By Transaction ID

☐ Europe/Berlin

Application Options

☒ Over write existing databases.
 Maximum Parallel Streams per Database:

Figure 5-21 Select restore job options

In the Review page, check all entered restore job parameters, as shown in Figure 5-22. Click **Submit** to start the on-demand restore job.

Restore - SQL
Review

Default Setup

☒ Select data sources
☒ Select source
☒ Source snapshot
☒ Restore method
☒ Set destination
☒ Set run settings
☒ Job options
☐ Review

Review
Review your selections, and then click Submit.

Select data sources

Selected source: IBM2, ESCC
 Source snapshot: IBM2 - Jun 7, 2020 8:01:00 AM
 ESCC - Jun 7, 2020 8:01:00 AM
 Restore Type: On-Demand: Snapshot
 Restore Source Type: Site
 Restore Source: Primary
 Restore method: Production

Figure 5-22 SQL Server restore summary

IBM Spectrum Protect Plus mounts the vSnap server backup volume at the SQL application server and copies the backup data to the source.

In our example that is shown in Figure 5-23, the vSnap server backup volume is mounted as Disk1 during the restore job.

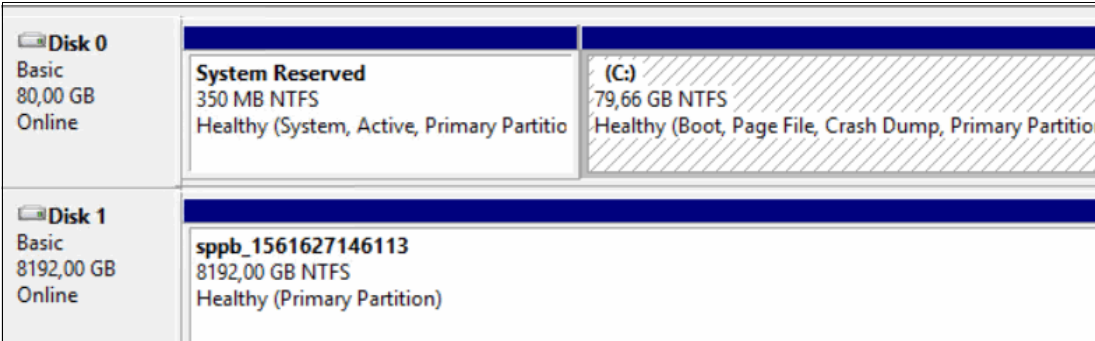


Figure 5-23 Mount of vSnap server volume for the restore on the SQL application server

In the following example scenario, a Production restore of a SQL Server stand-alone database is performed with the new Standby mode. After the restore, IBM Spectrum Protect Plus keeps the database in read-only mode. This option permits:

- Uncommitted transactions are saved in an undo file
- The undo file can be used for bringing the database online

As shown in Figure 5-24, we select the new Standby mode Job Options and perform the restore similar to the previous example.

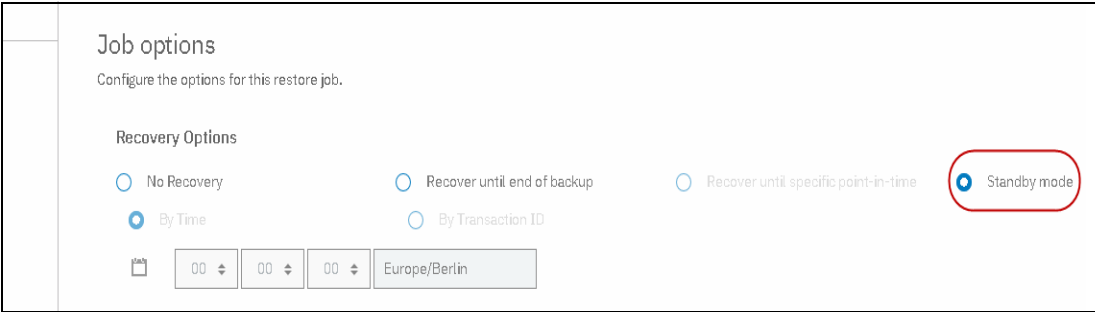


Figure 5-24 Select restore Standby mode job options

IBM Spectrum Protect Plus mounts the vSnap server backup volume at the SQL application server and copies the backup data to the source. Figure 5-25 shows that the new Database IBM2_TEST was restored with Standby/Read-Only mode.

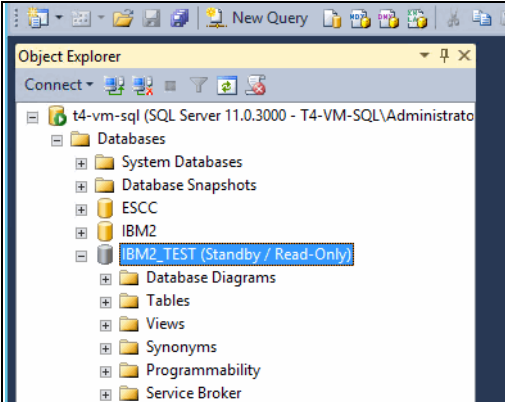


Figure 5-25 Database in standby mode after the restore



Backing up and restoring Microsoft Exchange data

Microsoft Exchange is a widely used mailing solution. IBM Spectrum Protect Plus can be used to protect Microsoft Exchange data, which provides restore functions at a database or single item (mail, contact, or calendar entry) level.

This chapter describes how to set up IBM Spectrum Protect Plus to protect Microsoft Exchange Servers, and explores common scenarios and best practices. It includes the following topics:

- ▶ 6.1, “Microsoft Exchange server” on page 92
- ▶ 6.2, “Prerequisites for protection in IBM Spectrum Protect Plus” on page 94
- ▶ 6.3, “IBM Spectrum Protect Plus configuration for Exchange” on page 101
- ▶ 6.4, “Backup jobs overview” on page 103
- ▶ 6.5, “Restore jobs” on page 111

6.1 Microsoft Exchange server

Microsoft Exchange is an enterprise Groupware and Mail Transport product. Most Exchange servers use database availability groups (DAG) to replicate the mailbox databases between different servers or sites. This approach ensures that every mailbox database has more than one copy to avoid data loss because of a server outage or corruption.

For more information about DAG, see [this web page](#).

6.1.1 Server roles

Depending on the Microsoft Exchange release, different server roles are available that must be protected, as listed in Table 6-1.

Table 6-1 Microsoft Exchange Server roles

Role or Version	Exchange 2013	Exchange 2016	Exchange 2019
Edge/Transport	X	X	X
Client Access	X		
Mailbox	X	X	X

The Edge/Transport role is used to transport mail from external sources into the Exchange infrastructure. A server with installed Edge/Transport is usually placed in a specific secured firewall zone because it is directly connected to the internet. If this role is the only role that is installed on the server, the server needs no Exchange-specific protection (because it has no persistent user data; it acts only as a proxy).

When implemented as a VMware or Hyper-V virtual server, it can be protected by hypervisor backup in IBM Spectrum Protect. When implemented as a physical server, it can be protected with the Windows File System backup component of IBM Spectrum Protect Plus.

The Client Access role is a separate role in Exchange 2013 and was merged into the mailbox role in Exchange 2016 and 2019. If a server is installed with Client Access Server role only, the same type of protection applies that is used for Edge/Transport only servers.

Only Microsoft Exchange servers with installed Mailbox role are protected by IBM Spectrum Protect Plus Backup and Restore for Microsoft Exchange. These servers are usually called *mailbox servers*. In the IBM Spectrum Protect Plus GUI, they are referred to as *Application Servers*.

6.1.2 Stand-alone or availability group databases

Every Mailbox database in Exchange is created and hosted on at least one Exchange server. A Mailbox database is used to store user or service account mailboxes. Every Exchange account is served by only one mailbox database. The maximum size of a Mailbox database is 2 TB.

To provide high availability at a database level, Mailbox databases can be configured in availability groups.

Database Availability Groups (DAG) are a group of Mailbox Servers in the same Exchange domain that share multiple copies of Mailbox databases. Up to 16 copies of a single Exchange database can exist. However, only one copy is active, meaning that users are working on this copy and changes are applied to this copy. The other copies are updated by shipping the committed Exchange log files from the active copy to the other copies. The inactive copies are showing a healthy status if the log replication is working.

A `ReplayLagTime` and `TruncationLagTime` can be defined for every copy to ensure that the copy does not commit or truncate the replicated logs before the `ReplayLagTime` and `TruncationLagTime` are reached. The default value of these two parameters is 0 seconds and the maximum value is 14 days.

A database copy with default settings is a nearly real-time copy (there is always the gap of the active log file, which is not shipped to inactive copy yet) of the active copy.

For example, a database copy with a `ReplayLagTime` of 7 days is a copy that lags the active copy by 7 days. A lagged copy ensures that if the active database copy becomes corrupted, a working copy (7 days back in time) is still available that can be used to fix the corruption or be used as a new base to apply the logs until the corruption occurred.

6.1.3 Mailbox movement

Every Exchange mailbox is hosted in a single Exchange mailbox database and, if applicable, on corresponding copies of this database (Database Availability Groups). Nevertheless, the Exchange Administrator can move Exchange mailboxes from one Mailbox database to another. Common use cases are to move a mailbox to a Mailbox database on faster or more reliable storage when the current database is hitting the recommended maximum size of 2 TB per database, or the user is switching to a different department when location and Mailbox databases are defined by department or location rules.

6.1.4 Microsoft built-in data loss prevention

Microsoft Exchange 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 still 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 20 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 Spectrum Protect Plus adds data protection capabilities that can be used whenever the built-in solutions are not satisfying or in case of a disaster.

6.2 Prerequisites for protection in IBM Spectrum Protect Plus

Ensure that all prerequisites for your Microsoft Exchange application are met before you start protecting Exchange databases with IBM Spectrum Protect Plus.

IBM Spectrum Protect Plus is a zero touch data protection product; therefore, no installation on the Exchange Mailbox Servers is needed. However, some requirements must be met to enable IBM Spectrum Protect Plus to access Exchange Mailbox Servers and perform backup or restore tasks.

For more information about these requirements, see [IBM Documentation](#).

6.2.1 Granular restore remote package installation

To perform granular Mailbox restore requests, an installation of the Spectrum Protect Plus Microsoft Management Console (MMC) GUI on a Microsoft Windows system where Outlook 2016 or later 32-bit edition is required.

This Windows system can be one of the Exchange Mailbox servers, but Microsoft advises against installing Outlook on an Exchange Mailbox Server. Therefore, it is best to use a separate Windows server.

To use the remote management features, you must first install and enable Windows PowerShell 3.0, or later, on all IBM Spectrum Protect Plus protected Exchange servers and the remote server from which you intend to run the IBM Spectrum Protect Plus MMC GUI.

To download, install, and enable the software, follow the instructions in Microsoft Windows Management Framework 3.0 Downloads. The remote server and Application server must be in the same domain.

This installation is called Granular remote package. The installation steps be found in the readme file for the Spectrum Protect Plus MMC GUI, which we included here for convenience.

Installation steps

Deploy the granular restore package to a remote server that has Microsoft Outlook installed.

The following installation steps are performed only once. After the granular restore package is installed, you can continue to use it to perform later granular restore operations:

1. Copy the granular restore package, which is in `C:\Program Files\IBM\IBM Spectrum Protect Plus\tools\exchange\imr\<version>TIV-TSMEXC-Win.exe`, from the Application Server to the remote server from where you manage the granular restore operations. Also, note that `<version>` indicates the version.
2. On the remote server, run the following commands to install the package (these commands assume that you copied to the `C:\temp` directory):
 - a. Create the installation diagnostic folder:

```
mkdir C:\temp\diag
```
 - b. Install MMC GUI and granular components:

```
C:\temp\imr\install_imr.bat *-TIV-TSMEXC-Win.exe 10.1.7 c:\temp
```

Where `*` is the MMC GUI version.

3. Configure the remote connection between the remote server and Application server:
 - a. Verify that the Windows Firewall allows inbound connections on the remote server.
 - b. Set the hostnames for the remote server and respective for the Application server.

The Application server runs the Exchange server and the remote server performs the granular restore operation.

```
$remote_server_host_name = "outlook1.domain.org"
$app_server_host_name = "exchange1.domain.org"
```

4. Enable remote management for the MMC GUI that is deployed with IBM Spectrum Protect Plus entering the following Windows PowerShell command:

```
Enable-PSRemoting -Force
```

Depending on your environment, you might need to add trusted hosts to the Exchange Server and server where the MMC GUI is deployed:

- a. Add the Application Server and remote server to the trusted hosts list by running the following command on each system:

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value
"$remote_server_host_name,$app_server_host_name" -Force
```

- b. Restart the winrm service by running the following command:

```
Restart-Service winrm
```

5. Enable the Windows PowerShell Remoting feature with Credential Security Support Provider (CredSSP) authentication. Complete the following steps:
 - a. On the remote server, run the following command to enable the Windows PowerShell Remoting feature with CredSSP:

```
Enable-WSManCredSsp -Role Client -DelegateComputer $app_server_host_name
-Force
```

- b. On the Application Server that runs the granular restore operation, run the following command to enable the Windows PowerShell Remoting feature with CredSSP:

```
Enable-WSManCredSsp -Role Server -Force
```

6. Verify that the Windows PowerShell Remoting feature is configured by using one of the following methods: (use the **Test-WSMan** cmdlet to test whether the WinRM service is running on the remote computer):

- a. On the remote server, run the following cmdlet to verify that the Windows PowerShell Remoting feature is configured correctly:

```
Test-WSMan $app_server_host_name
```

- b. On the Application Server, run the following cmdlet to verify that the Windows PowerShell Remoting feature is configured correctly:

```
Test-WSMan $remote_server_host_name
```

Optionally, for more remote configuration verification, complete the following steps:

1. Set the credentials object you used. Usually, this credential is a domain administrator:

```
$creds = Get-Credential
```

- a. On the Application Server and remote server, run the following cmdlet to verify basic remote connection:

```
Invoke-Command -ComputerName $remote_server_host_name -ScriptBlock { pwd }
-Credential $creds
```

```
Invoke-Command -ComputerName $app_server_host_name -ScriptBlock { pwd }
-Credential $creds
```

- b. On the Application Server and remote server, run the following cmdlet to verify (CredSSP) authentication is enabled:

```
Invoke-Command -ComputerName $remote_server_host_name -ScriptBlock { pwd }
-Credential $creds -Authentication CredSsp
Invoke-Command -ComputerName $app_server_host_name -ScriptBlock { pwd }
-Credential $creds -Authentication CredSsp
```

For our example, we show the commands that run in our test environment, which consists of the following servers:

- ▶ Windows 10 server ("windows10.xxxxxxx.lab,192.168.111.66"), as shown in Example 6-1
- ▶ Exchange server ("epc-exchange.xxxxxxx.lab,192.168.111.167"), as shown in Example 6-2 on page 97

Example 6-1 PowerShell commands on the Windows 10 server

```
PS C:\Users\Administrator> Set-Item WSMAN:\localhost\Client\TrustedHosts -Value
"epc-exchange.xxxxxxx.lab,192.168.111.167"
```

WinRM Security Configuration.

This command modifies the TrustedHosts list for the WinRM client. The computers in the TrustedHosts list might not be authenticated. The client might send credential information to these computers. Are you sure that you want to modify this list?

[Y] Yes [N] No [S] Suspend [?] Help (default is "Y"): y

```
PS C:\Users\Administrator> Get-Item WSMAN:\localhost\Client\TrustedHosts
WSManConfig: Microsoft.WSMan.Management\WSMan::localhost\Client
```

Type	Name	SourceOfValue	Value
----	----	-----	----
System.String	TrustedHosts		
	epc-exchange.xxxxxxx.lab,192.168.111.167		

```
PS C:\Users\Administrator> Restart-Service winrm
PS C:\Users\Administrator> Test-WSMan 192.168.111.167
```

```
wsmid          : http://schemas.dmtf.org/wbem/wsman/identity/1/wsmanidentity.xsd
ProtocolVersion : http://schemas.dmtf.org/wbem/wsman/1/wsman.xsd
ProductVendor   : Microsoft Corporation
ProductVersion  : OS: 0.0.0 SP: 0.0 Stack: 3.0
```

```
PS C:\Users\Administrator> Test-WSMan epc-exchange.xxxxxxx.lab
```

```
wsmid          : http://schemas.dmtf.org/wbem/wsman/identity/1/wsmanidentity.xsd
ProtocolVersion : http://schemas.dmtf.org/wbem/wsman/1/wsman.xsd
ProductVendor   : Microsoft Corporation
ProductVersion  : OS: 0.0.0 SP: 0.0 Stack: 3.0
```

```
PS C:\Users\Administrator> enable-wsmancredssp -role client -delegatecomputer
epc-exchange.xxxxxxx.lab
```

CredSSP Authentication Configuration for WS-Management

CredSSP authentication allows the user credentials on this computer to be sent to a remote computer. If you use CredSSP authentication for a connection to a malicious or compromised computer, that computer will have access to your user name and password. For more information, see the Enable-WSManCredSSP Help topic. Do you want to enable CredSSP authentication?
[Y] Yes [N] No [S] Suspend [?] Help (default is "Y"): y

```

cfg      : http://schemas.microsoft.com/wbem/wsman/1/config/client/auth
lang     : en-US
Basic    : true
Digest   : true
Kerberos : true
Negotiate : true
Certificate : true
CredSSP  : true

```

```

PS C:\Users\Administrator> invoke-command -computername epc-exchange.xxxxxxx.lab
-scriptblock {pwd}

```

Path	PSComputerName
----	-----
C:\Users\Administrator.xxxxxxx\Documents	epc-exchange.xxxxxxx.lab

```

PS C:\Users\Administrator> $cred = get-credential

```

```

cmdlet Get-Credential at command pipeline position 1
Supply values for the following parameters:
Credential

```

```

PS C:\Users\Administrator> invoke-command -computername epc-exchange.xxxxxxx.lab
-Authentication Credssp -credential $cred -scriptblock {pwd}

```

Path	PSComputerName
----	-----
C:\Users\spp\Documents	epc-exchange.xxxxxxx.lab

Example 6-2 PowerShell commands on the Exchange Server

```

PS C:\Windows\system32> Set-Item WSMan:\localhost\Client\TrustedHosts -Value
"windows10.xxxxxxx.lab,192.168.111.66"
WinRM Security Configuration.
This command modifies the TrustedHosts list for the WinRM client. The computers in the
TrustedHosts list might not be
authenticated. The client might send credential information to these computers. Are you
sure that you want to modify
this list?
[Y] Yes [N] No [S] Suspend [?] Help (default is "Y"): y
PS C:\Windows\system32> Get-Item WSMan:\localhost\Client\TrustedHosts

```

```

WSManConfig: Microsoft.WSMan.Management\WSMan::localhost\Client

Type      Name                                     SourceOfValue  Value
----      -

```

```
System.String TrustedHosts
windows10.xxxxxxx.lab,192.168.111.66
```

```
PS C:\Windows\system32> Test-WSMan 192.168.111.66
```

```
wsmid          : http://schemas.dmtf.org/wbem/wsman/identity/1/wsmanidentity.xsd
ProtocolVersion : http://schemas.dmtf.org/wbem/wsman/1/wsman.xsd
ProductVendor   : Microsoft Corporation
ProductVersion  : OS: 0.0.0 SP: 0.0 Stack: 3.0
```

```
PS C:\Windows\system32> Test-WSMan windows10.xxxxxxx.lab
```

```
wsmid          : http://schemas.dmtf.org/wbem/wsman/identity/1/wsmanidentity.xsd
ProtocolVersion : http://schemas.dmtf.org/wbem/wsman/1/wsman.xsd
ProductVendor   : Microsoft Corporation
ProductVersion  : OS: 0.0.0 SP: 0.0 Stack: 3.0
```

```
PS C:\Windows\system32> enable-wsmancredssp -role server
```

```
CredSSP Authentication Configuration for WS-Management
CredSSP authentication allows the server to accept user credentials from a remote computer.
If you enable CredSSP authentication on the server, the server will have access to the user name and password of the client computer if the client computer sends them. For more information, see the Enable-WSManCredSSP Help topic.
Do you want to enable CredSSP authentication?
[Y] Yes [N] No [S] Suspend [?] Help (default is "Y"): y
```

```
cfg          : http://schemas.microsoft.com/wbem/wsman/1/config/service/auth
lang         : en-US
Basic        : false
Kerberos     : true
Negotiate    : true
Certificate  : false
CredSSP      : true
CbtHardeningLevel : Relaxed
```

```
PS C:\Windows\system32> invoke-command -computername windows10.xxxxxxx.lab -scriptblock {pwd}
```

```
Path                PSComputerName
----                -
C:\Users\spp\Documents windows10.xxxxxxx.lab
```

```
PS C:\Windows\system32> $cred = get-credential
```

```
cmdlet Get-Credential at command pipeline position 1
Supply values for the following parameters:
Credential
```

```
PS C:\Windows\system32> invoke-command -computername windows10.xxxxxxx.lab -Authentication Credssp -credential $cred -scriptblock {pwd}
```

```
Path                PSComputerName
----                -
C:\Users\Administrator\Documents windows10.xxxxxxx.lab
```

Running the MMC GUI on the remote server to perform granular restore

To perform a granular restore by using the MMC GUI, complete the following steps:

1. Start the MMC GUI application:

C:\Program Files\Tivoli\FlashCopyManager\FlashCopyManager.exe

2. To add the Application Server in MMC GUI, click **Actions** → **Manage Computers** to open Manage Computers window.

Click the plus-sign icon (+) in the Computers pane and enter an Application Server name. Click **Set Account** and enter user credentials for the Application Server, as shown in Figure 6-1.

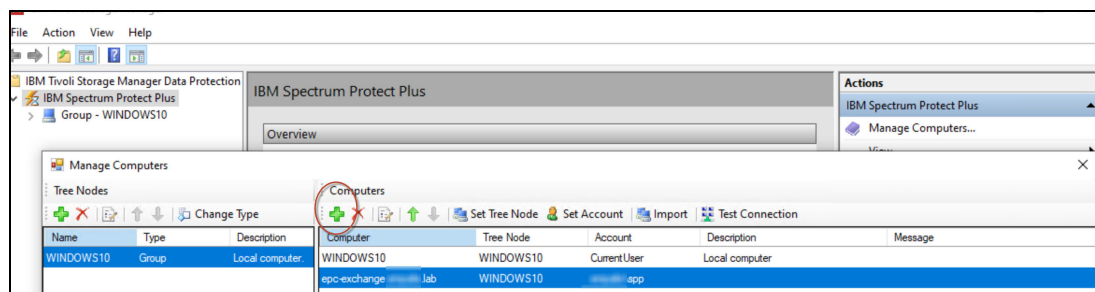


Figure 6-1 Adding the Application Server

Use the same credentials as used by Spectrum Protect Plus GUI. Select **Manage Protection** → **Databases** → **Exchange** → **Backup** → **Manage Application Servers** → **Add Application Server**.

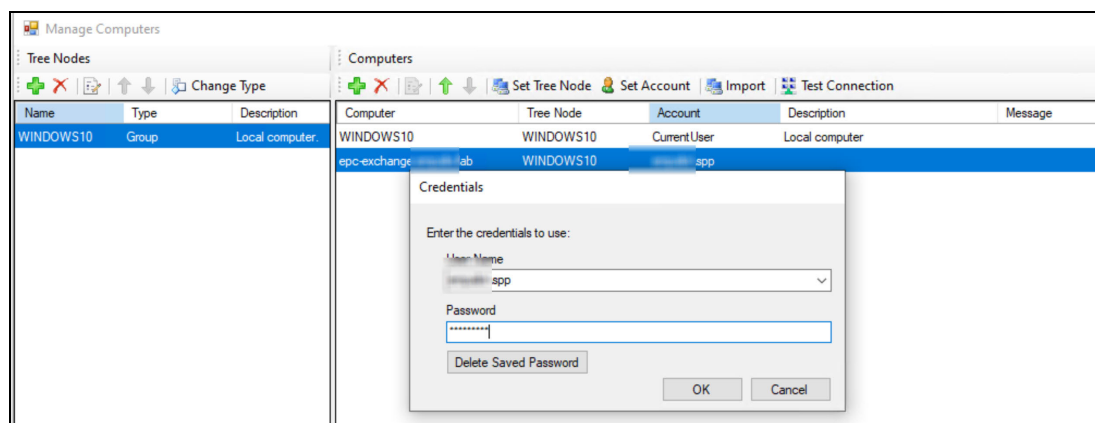


Figure 6-2 Managing credentials

3. Test the connection to the Application Server. Select **Application Server node** and then, click the **Test Connection** tab.




Figure 6-3 Testing the connection

The verify connection status of “Connected successfully” is shown in the Message column. More information can be displayed by clicking the hyperlink, as shown in Figure 6-4.

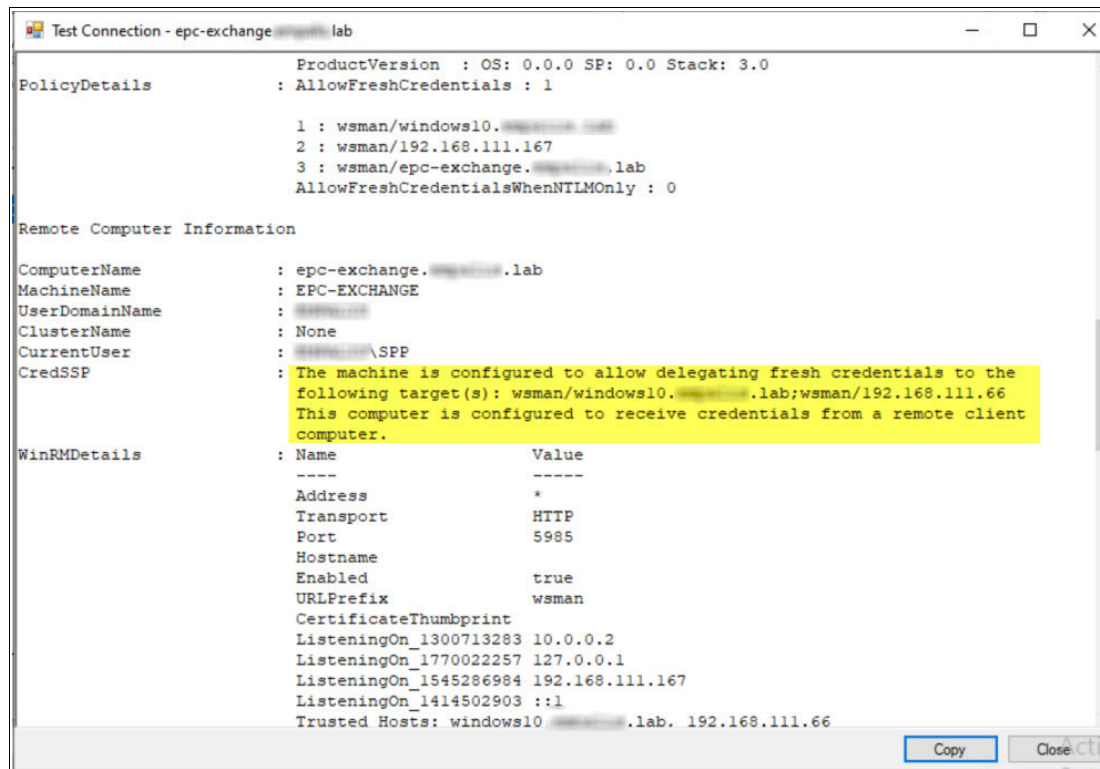


Figure 6-4 Successful connection details

From this output, consider the information about CredSSP, which indicates that the configuration is successful and a connection is possible.

4. Click **OK** to close the Manage Computers window.

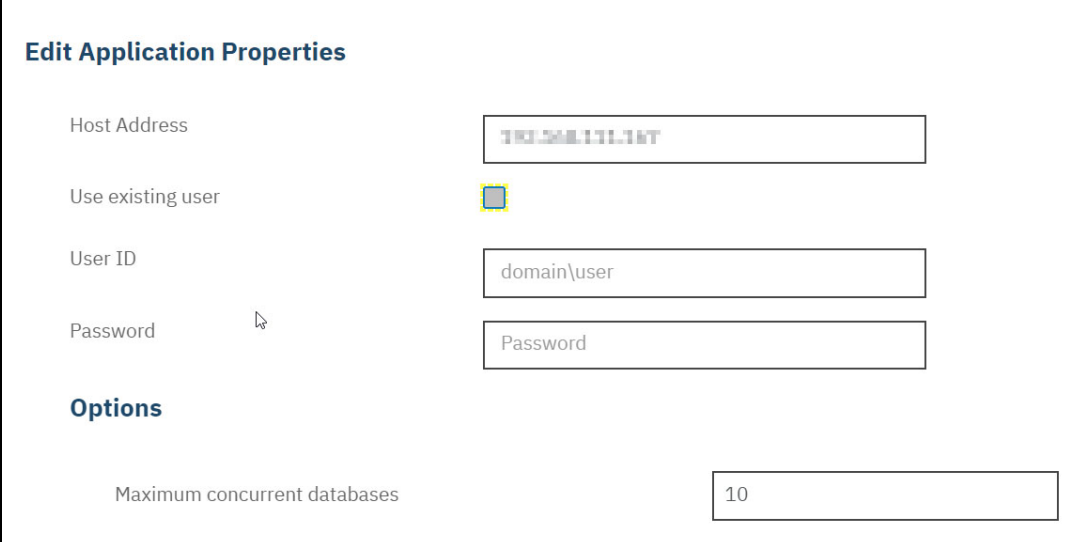
6.3 IBM Spectrum Protect Plus configuration for Exchange

Every Exchange Mailbox server is referred to as Application Server in IBM Spectrum Protect Plus.

To configure an exchange server in IBM Spectrum Protect Plus, start the IBM Spectrum Protect Plus Server GUI and log in to the dashboard.

Select **Manage Protection** → **Databases** → **Exchange**. Click **Manage Application Server** → **Add Application Server**

Figure 6-5 shows how to add or edit an Application Server. The Host address can be the Server name or IP address. The User ID must be entered for the first Exchange Server with the domain or user ID and password. For any other Application Servers, the same user ID can be used by selecting the **Use existing user** option and then, selecting it from the drop-down menu.



Edit Application Properties

Host Address

Use existing user

User ID

Password

Options

Maximum concurrent databases

Figure 6-5 Edit Application Properties window

The Maximum concurrent database number (default: 10) is used to reduce or raise the number of mailbox databases that are backed up concurrently. In production environments, the default of 10 is a good starting point and is raised incrementally only to avoid overloading the Microsoft Exchange Server.

After the Application Server is registered, the Mailbox databases on the server can be browsed and SLAs can be assigned.

Figure 6-6 shows an example of a Microsoft Exchange server with two Mailbox databases that are not part of a Database Availability Group (DAG). These databases can be selected and assigned to an SLA individually.

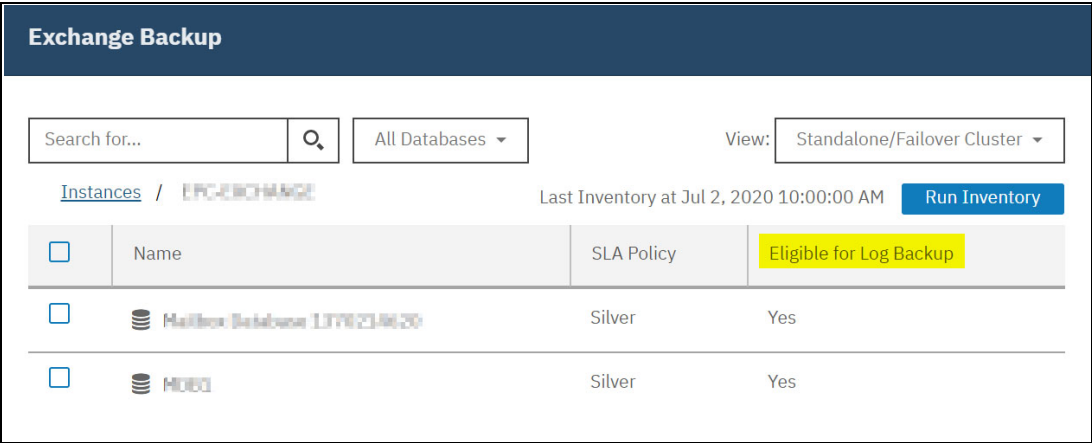


Figure 6-6 Exchange Backup menu

The Run Inventory button can be used to immediately query the Microsoft Exchange Server for a list of databases and their status. The list also indicates whether the Mailbox database is using circular logging or not. Databases with disabled circular logging are flagged with Yes in the Eligible for Log backup column. Databases can also be filtered by using the Search box or the view can be switched from Standalone/Failover Cluster to a list of DAG enabled Mailbox databases.

6.3.1 Log backup

By clicking **Select Options**, the Log Backup menu opens, as shown in Figure 6-7. In this window, a periodic log backup of the Exchange Mailbox Database log can be defined. This option ensures that multiple restore points are used during the day without the need to back up the Exchange Mailbox database. A log backup also enables the Microsoft Exchange server to purge the backed-up log files and free up space in the log directory of this mailbox database.

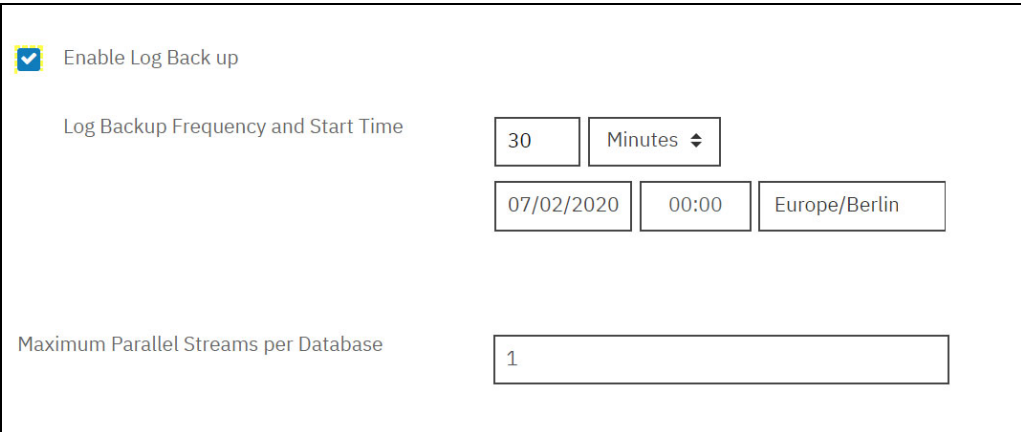


Figure 6-7 Log Back up menu

6.3.2 Database Availability Groups

DAG-enabled Mailbox databases can be protected by backing up only one database copy.

To configure the protection of DAG-enabled Mailbox databases, the view must be changed to “Database Availability Groups”.

By default, the backup is performed on the active copy, which might interfere with the Exchange user workload. To switch the protection to a passive copy in Options, select the **Backup preferred node**. This option can be pointed to the Exchange Mailbox Server with the lowest activation preference. This setting ensures that the backup is performed on the passive copy, which is the last copy to take over the Active copy role in the cluster.

6.4 Backup jobs overview

This section describes how to protect Microsoft Exchange Mailbox databases with IBM Spectrum Protect Plus.

6.4.1 Assigning an SLA policy

Before you can run a backup job, you must define an SLA policy. You can use an existing policy or define specific policies.

Generally, it is preferred to create dedicated SLA policies for single databases or for groups of logically related databases.

6.4.2 Backup types

The following backup jobs are available for Microsoft Exchange Applications:

- ▶ Scheduled
- ▶ Ad hoc

6.4.3 Scheduled backup

IBM Spectrum Protect Plus supports single or multiple Exchange databases per Exchange backup job. Multiple database backup jobs run sequentially.

In the navigation pane, select **Manage Protection** → **Databases** → **Exchange** as shown in Figure 6-8.

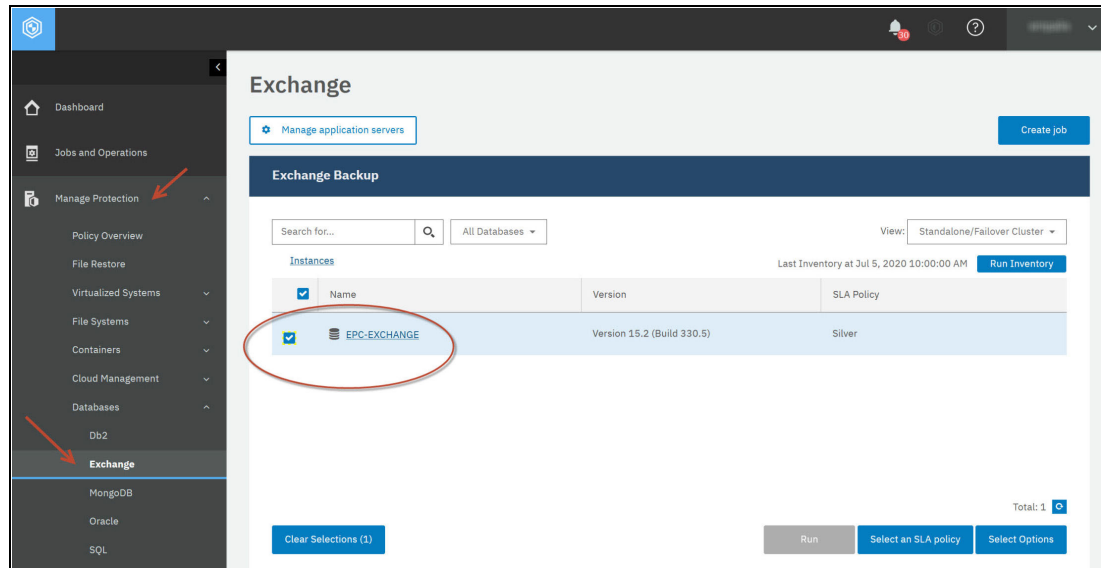


Figure 6-8 Defining a backup job

Select an Exchange instance to back up all the data in that instance. Optionally, you can click an instance name and then, select individual databases that you want to back up.

Three choices are available: Run, Select an SLA policy, and Select options, as shown in Figure 6-9.

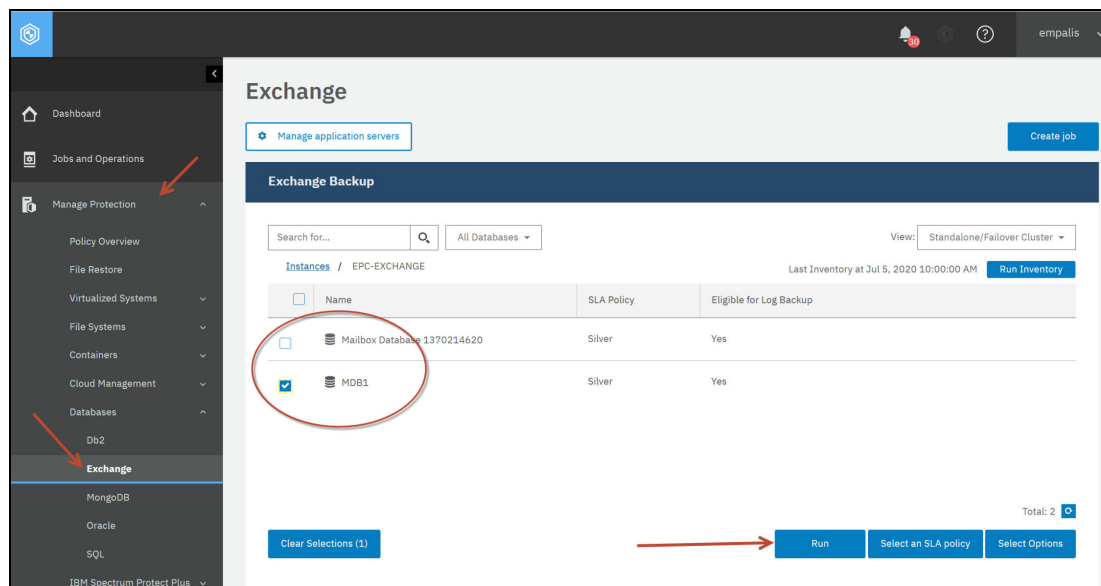


Figure 6-9 Selecting an instance or database

Click **Select an SLA Policy**. Predefined choices are: Gold, Silver, and Bronze. Each choice includes different frequencies and retention rates, as shown in Figure 6-10.

	SLA Policy	Frequency	Retention
<input checked="" type="checkbox"/>	Silver	Every 1 Days at 8:01:57 AM	1 Months
<input type="checkbox"/>	Bronze	Every 1 Days at 8:01:57 AM	1 Weeks
<input type="checkbox"/>	Demo	Every 1 Days at 8:02:06 AM	1 Months
<input checked="" type="checkbox"/>	Exchange_Silver	Every 1 Days at 12:00:00 AM	15 Days

Figure 6-10 Custom SLA policy Exchange_Silver

Gold is the most frequent with the shortest retention rate. You can also create a custom SLA policy or edit a policy, as we did by selecting the **Exchange_Silver** SLA policy, as shown in Figure 6-10. Click **Save** to confirm your choice.

Now, the SLA selection can be verified and options can be defined for the scheduled backup job by clicking **Select Options**, as shown in Figure 6-11.

	Name	SLA Policy	Eligible for Log Backup
<input type="checkbox"/>	Mailbox Database 1370214620	Silver	Yes
<input checked="" type="checkbox"/>	MDB1	Exchange_Silver	Yes

Figure 6-11 Checking SLA selection and selecting options

You can define options for your backup, such as enabling log backups for future recovery and specifying the parallel streams to reduce the time necessary to back up large databases (see Figure 6-12). Click **Save**.

Figure 6-12 Selecting options

Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table, as shown in Figure 6-13.

Policy	Frequency	Total
Silver	Every 1 Days at 8:01:57 AM	1
Exchange_Silver	Every 1 Days at 12:00:00 AM	1

Log	Start Time	End Time
Jul 5, 2020 3:35:10 PM	Jul 5, 2020 3:53:52 PM	

Status	Time	ID	Description
Summary	Jul 5, 2020	CTGGA2398	Starting job for policy

Figure 6-13 Configuring SLA policy

After clicking the icon, a pop-up window appears, as shown in Figure 6-14, in which you can configure more policy options.

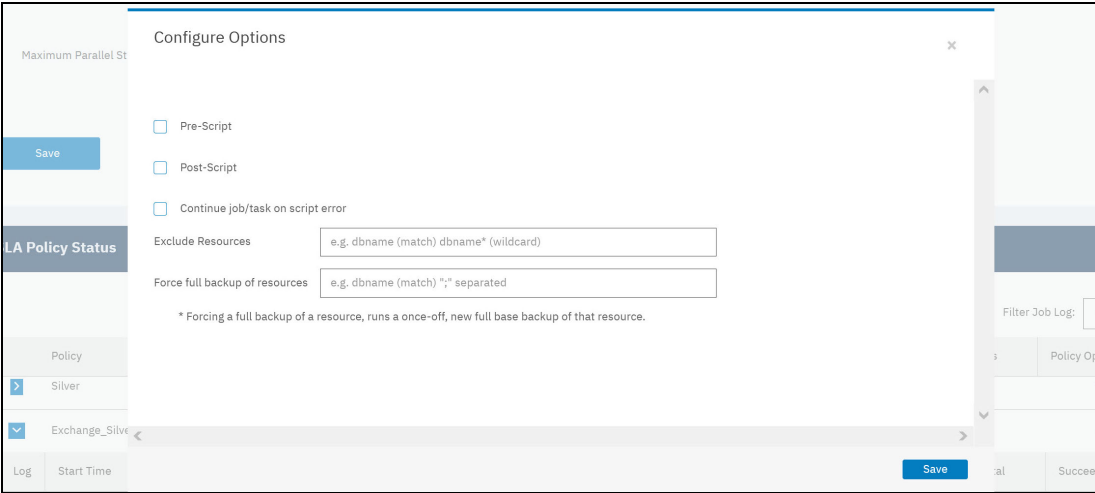


Figure 6-14 Configuring SLA options

To run the policy outside of the scheduled job, select the instance or database and then, click **Actions** → **Start**.

The status changes to Running for your chosen SLA. To pause the schedule, click **Actions** → **Pause Schedule**. To cancel a job after it starts, click **Actions** → **Cancel** (see Figure 6-15).

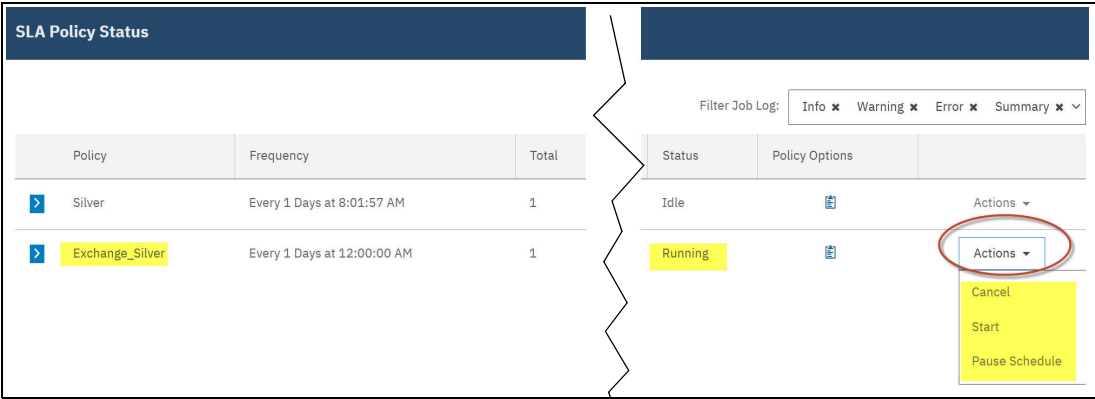


Figure 6-15 Running Jobs and actions

6.4.4 Ad hoc backup

An ad hoc backup is performed from the **Jobs and Operations** window.

Complete the following steps:

1. Click **Create Job**, as shown in Figure 6-16.

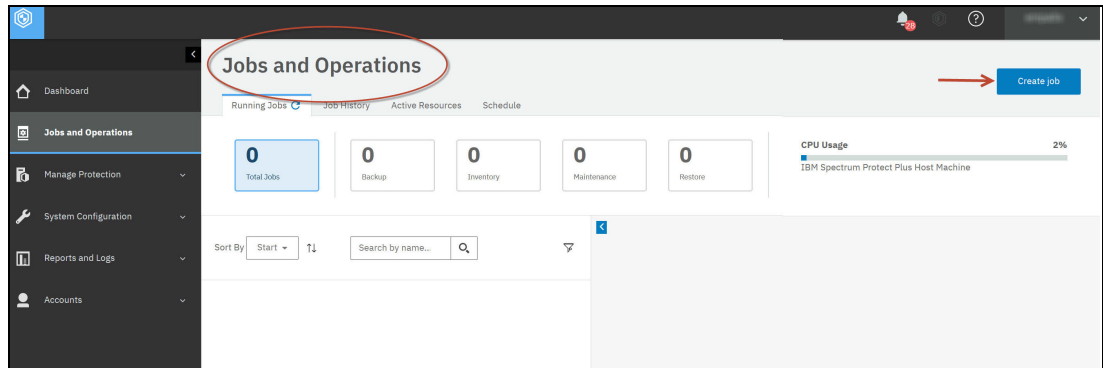


Figure 6-16 Creating a Job

2. You are presented with a choice for Ad hoc backup or Restore. Select **Ad hoc backup** (see Figure 6-17).

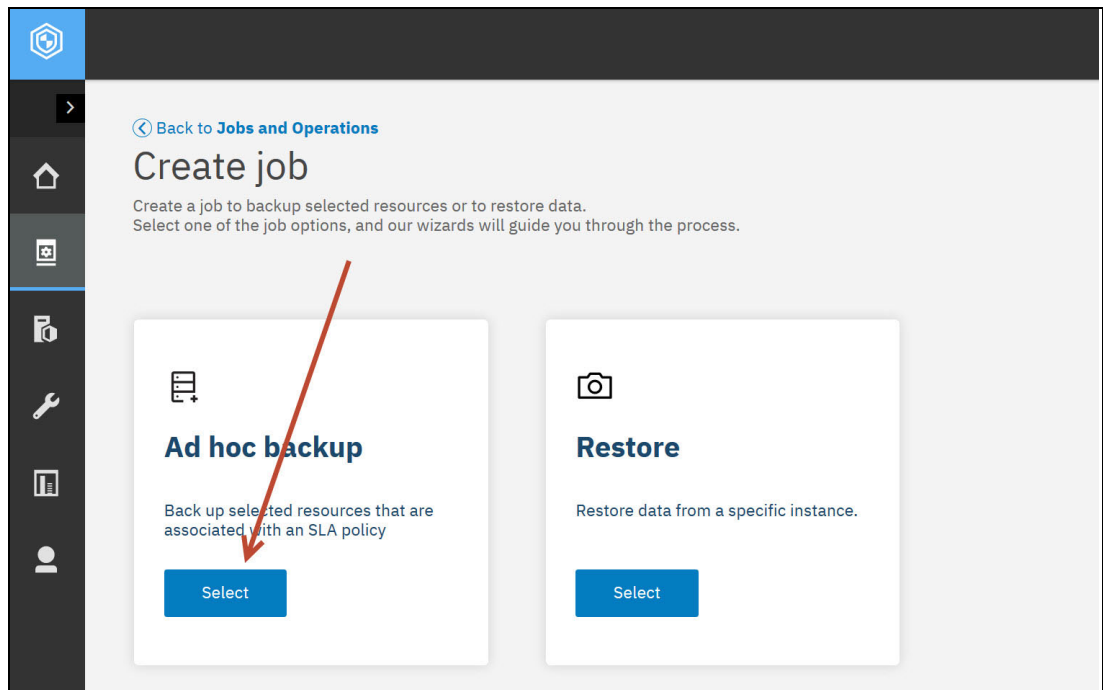


Figure 6-17 Ad hoc backup

3. In the Database selection, select **Exchange**, as shown in Figure 6-18.

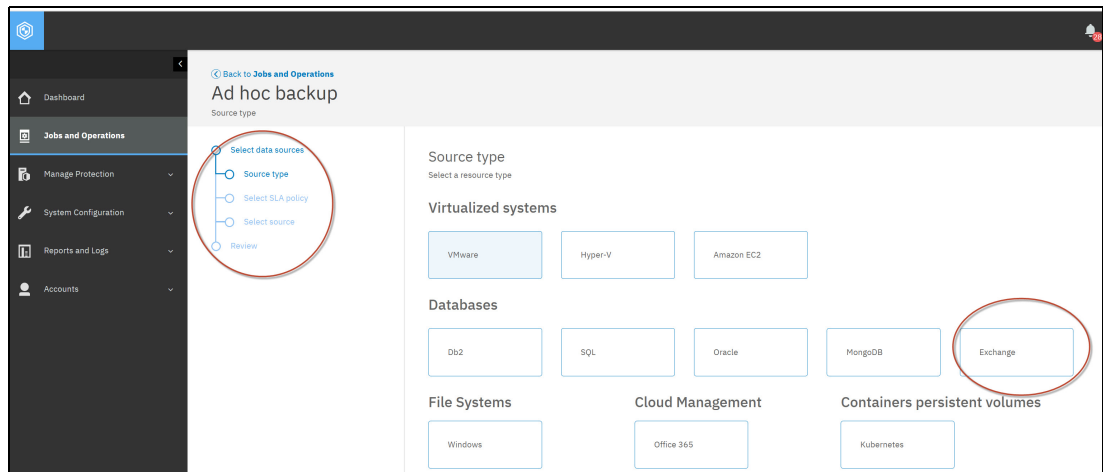


Figure 6-18 Definition diagram and selection

4. Select a predefined SLA policy. After clicking the SLA policy, the defined values for that policy are shown (see Figure 6-19).

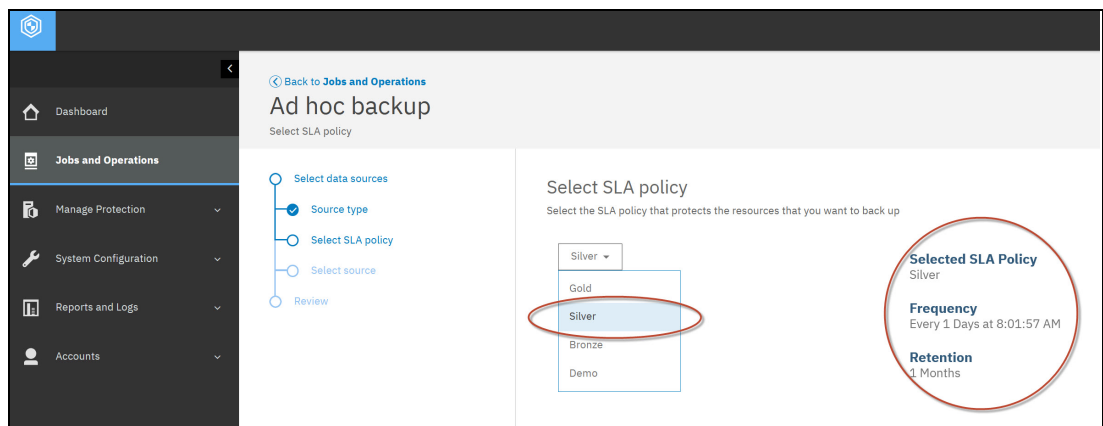


Figure 6-19 Selected SLA policy

5. Select the database to back up. If many databases are available, use the search function to easily find the wanted database. Now, the database can be added to the backup job list by clicking the blue plus sign (+), as shown in Figure 6-20.

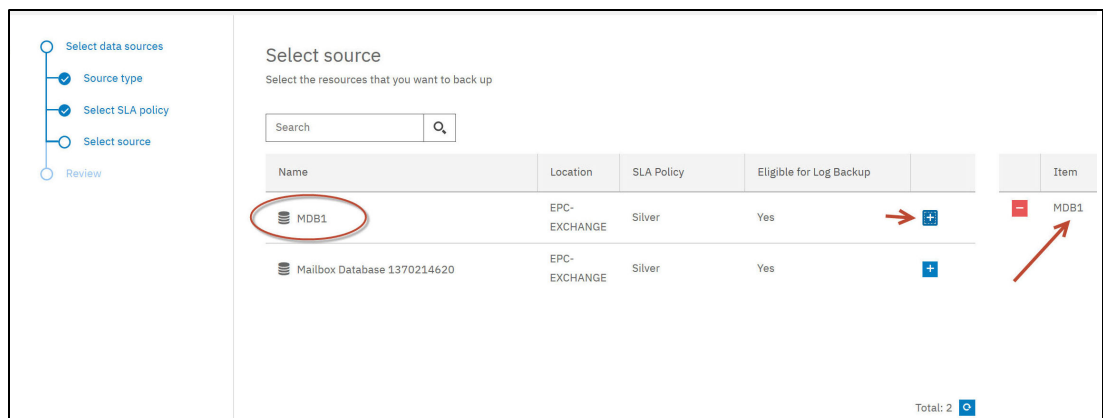


Figure 6-20 Choosing databases for ad hoc backup

6. Review the backup job options. Then, click **Submit** to start the job, as shown in Figure 6-21.

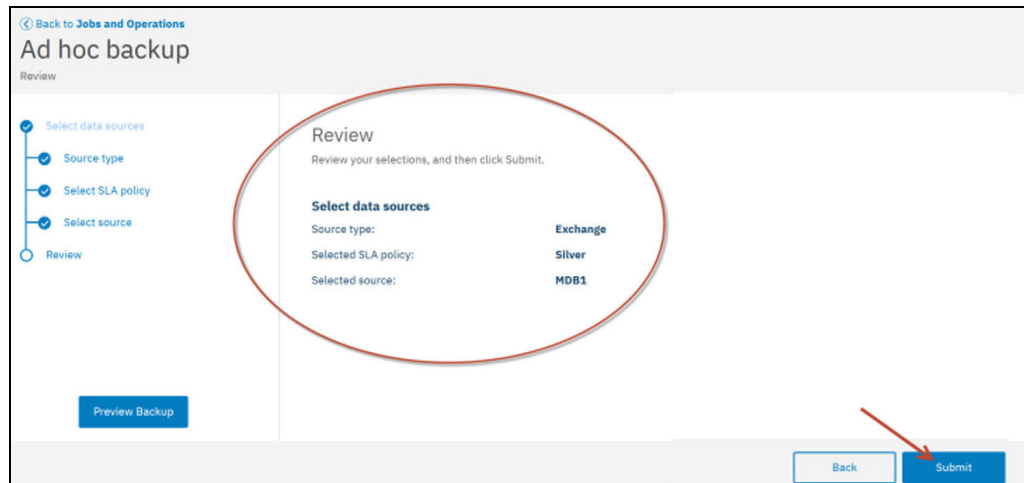


Figure 6-21 Ad hoc backup preview

7. As shown in Figure 6-22, a message is displayed to confirm that the job was submitted. Click **OK** to close the message.




Figure 6-22 Job submission confirmation

The Ad hoc backup job can be monitored under the **Jobs and Operations** pane, as shown in Figure 6-23.

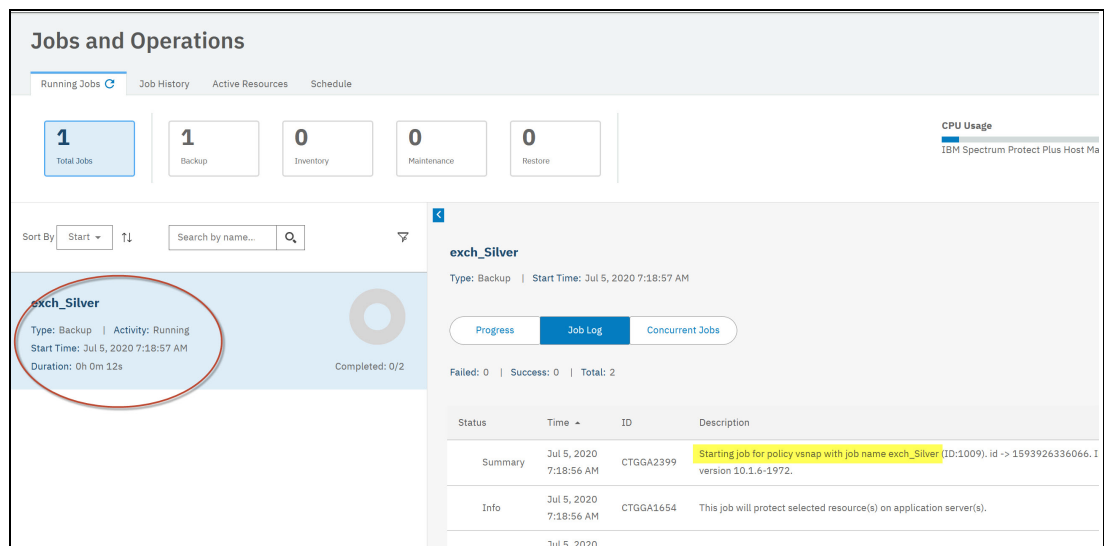


Figure 6-23 Running ad hoc backup

6.5 Restore jobs

The following types of restore jobs available in IBM Spectrum Protect Plus:

- ▶ On-demand: Snapshot: Runs a one-time restore operation. The restore job starts immediately upon the completion of the wizard.
- ▶ On-demand: Point in Time: Runs a one-time restore job from a point-in-time backup of a database. The restore job starts immediately upon the completion of the wizard.
- ▶ Recurring: Creates a repeating point-in-time restore job that runs on a schedule.

Two options are available to restore Microsoft Exchange data. It is possible to recover a complete Exchange Database into any database or Recovery Database (RDB) or to recover individual items, such as mailboxes or individual emails.

In both cases, you find the entry point for the procedure in the Jobs and Operations panel or the Manage Protection panel, as shown in Figure 6-24.

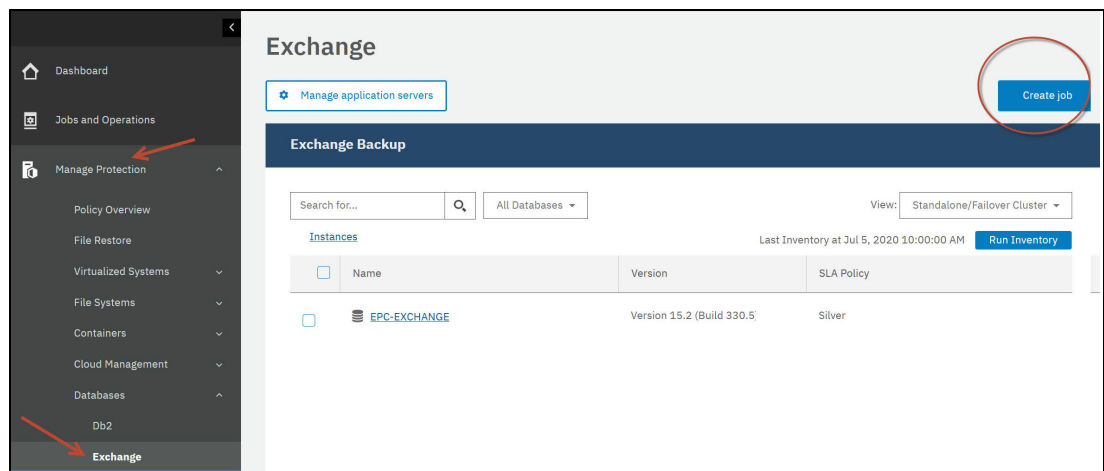


Figure 6-24 Creating Restore job

To create the restore job, select **Create Job** in the **Manage Protection - Exchange** menu. Then, select **Restore**, as shown in Figure 6-25.

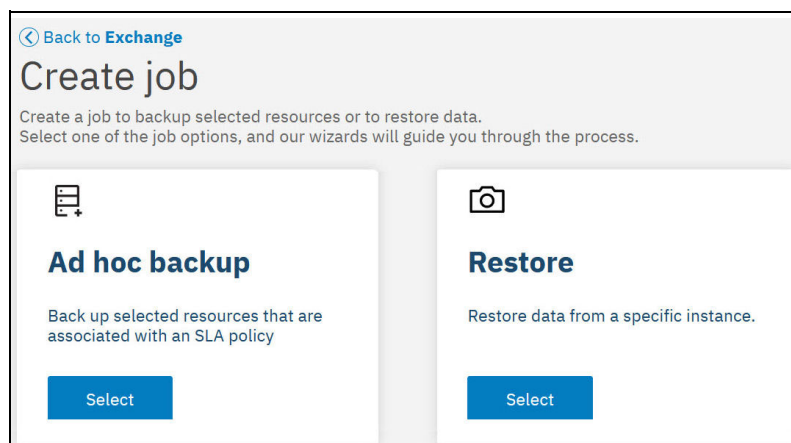


Figure 6-25 Selecting Restore

The restore dialog opens, as shown in Figure 6-26.

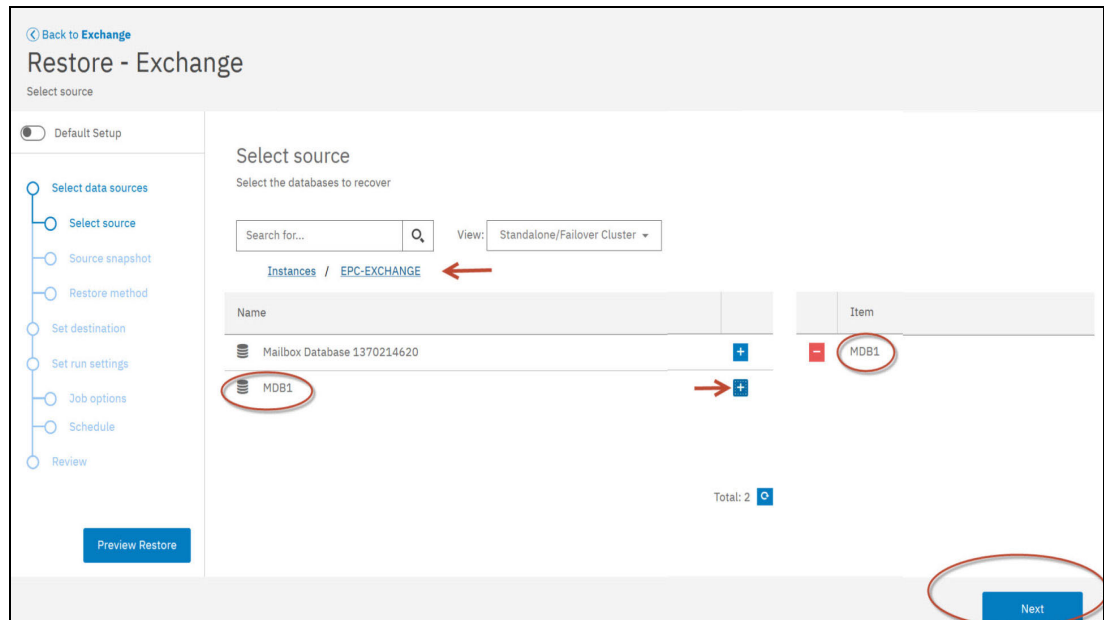


Figure 6-26 Selecting Restore items

Complete the following steps:

1. Select the object to restore. The Microsoft Exchange instance is displayed with the relevant databases for selection. By clicking the blue plus sign (+), the corresponding object is placed on the item list, also known as “job list”. Click **Next** to continue.
2. In the next panel, as shown in Figure 6-27, the type of restore can be selected.

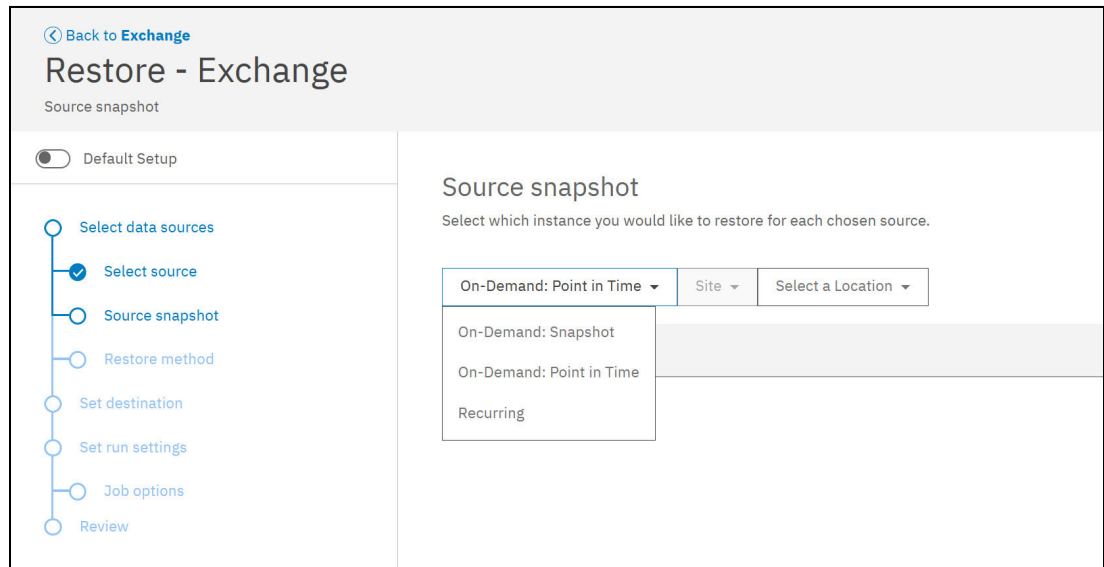


Figure 6-27 Selecting source snapshot

The following types are available:

- On-Demand Snapshot

In this case, select a Date/Time from the list of available backups, as shown in Figure 6-28.

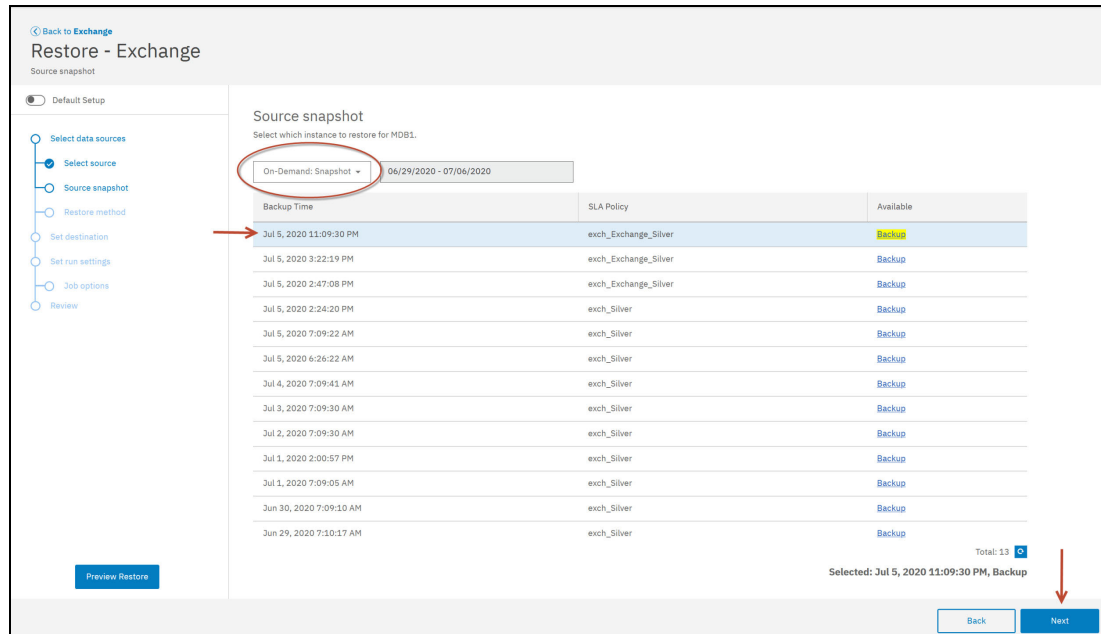


Figure 6-28 Setting date and time

- On-Demand Point-in-Time

In this case, the available restore location sites are: Demo, Primary, and Secondary (see Figure 6-29).

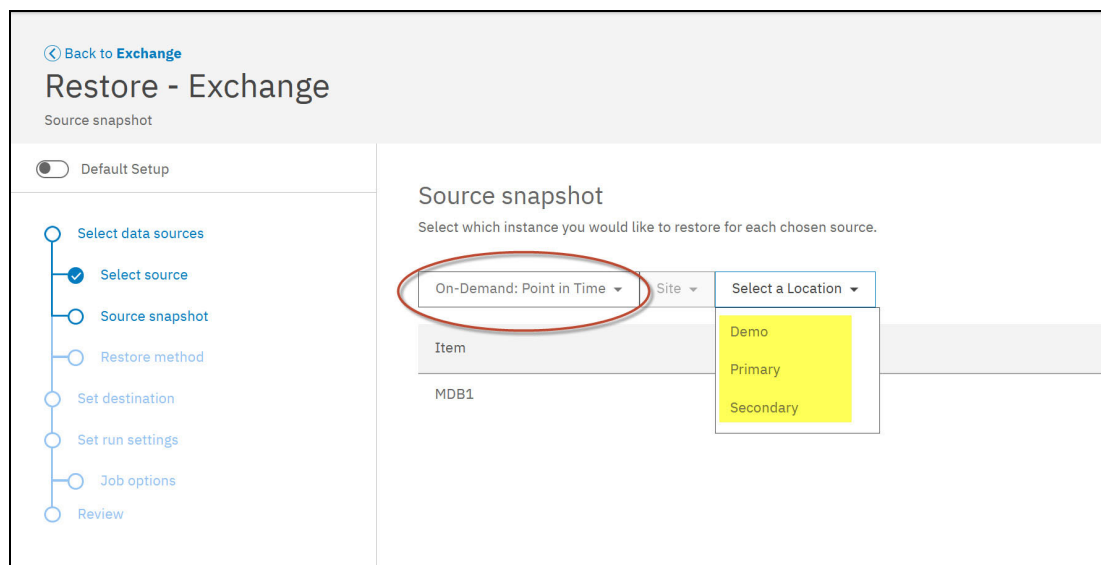


Figure 6-29 Selecting a location for On-Demand Point-in-Time

– Recurring

Choose a restore location type. As shown in Figure 6-30, different location types from where the restore can be taken are also available: Site, Cloud service, Repository server, Cloud service archive, and Repository server archive.

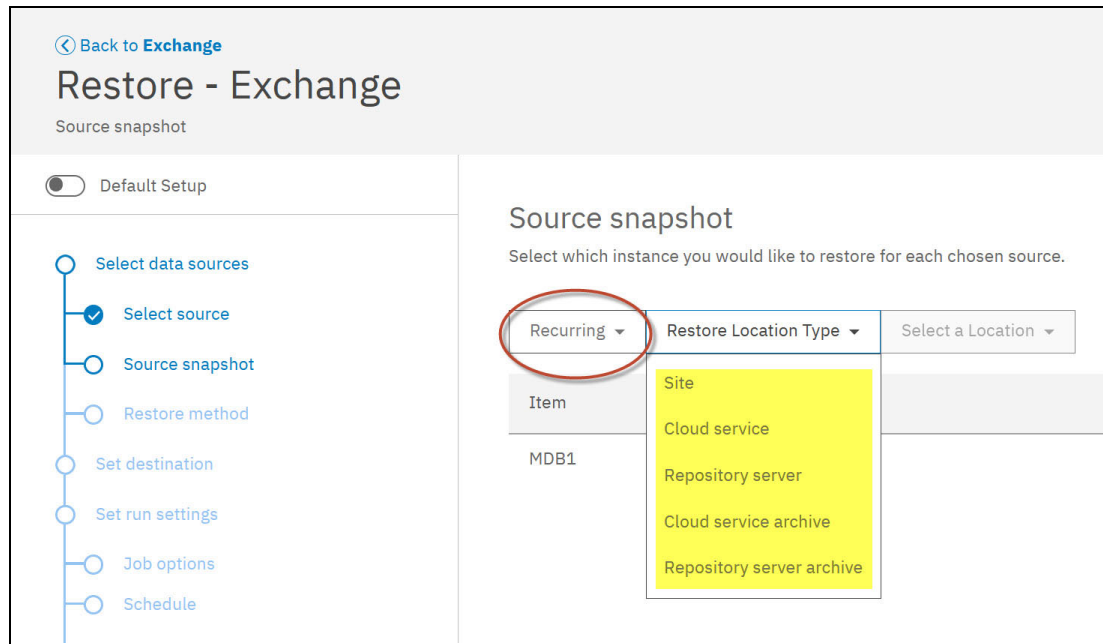


Figure 6-30 Recurring Restore location types

3. For our scenario, we select **On-Demand: Snapshot**. After selecting one of the available backups, click **Next**. You are presented with the next step, which is to select amongst two restore methods: Complete Restore or Item Recovery.

6.5.1 Complete Restore

This section describes how to restore the complete mailbox database and use it in instant access, production and test (see Figure 6-31).

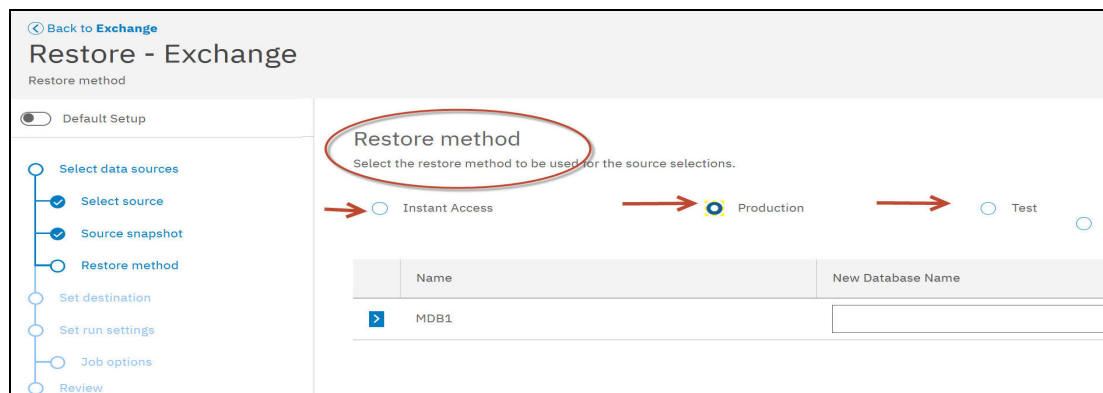


Figure 6-31 Restore method

In Production or Test mode, enter a new database name. In the panel for production, the destination path can be changed, as shown in Figure 6-32.

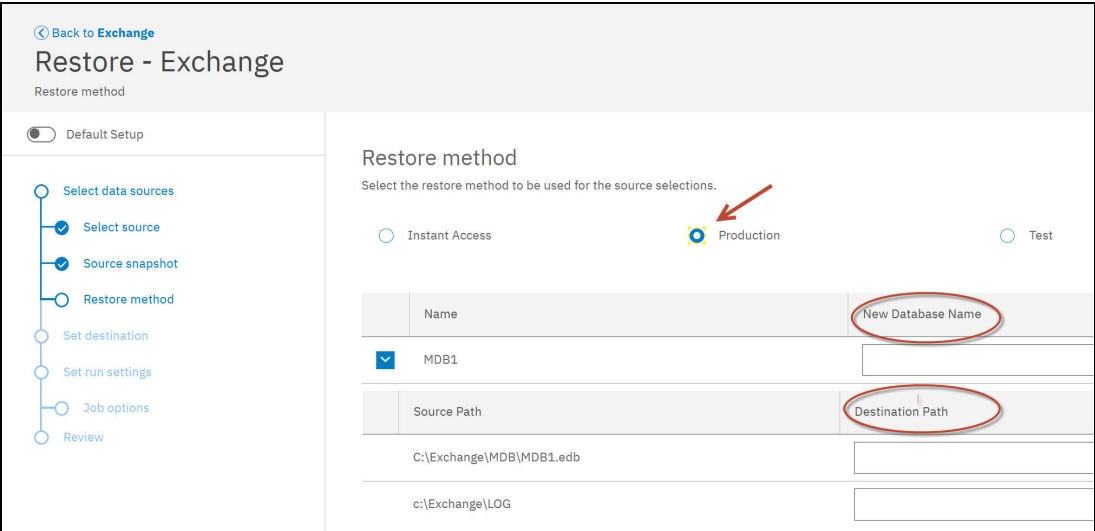


Figure 6-32 Restore method: Production

In our example, we proceed with the restore into production, as can be the case in a situation where the source database is corrupted and must be replaced.

After clicking **Next**, set the destination and choose restore into the original instance, as shown in Figure 6-33.

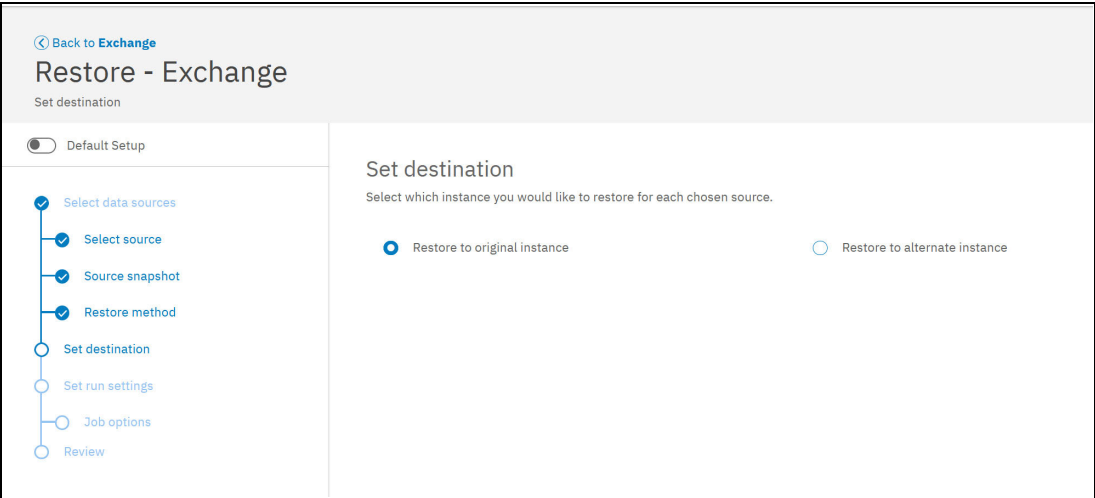


Figure 6-33 Setting destination

By clicking **Next**, other job options are available that are necessary for the recovery. The choice here is: No Recovery or Recover until end of backup.

The options **Recover until end of available logs** and **Recover until specific point-in-time** are not available for this type of restore because no log backups are available.

Only the following other options are available:

- ▶ Maximum Parallel Streams per Database
- ▶ Run cleanup immediately on job failure

The job options are shown in Figure 6-34.

[Back to Exchange](#)

Restore - Exchange

Job options

Default Setup

- Select data sources
- Select source
- Source snapshot
- Restore method
- Set destination
- Set run settings
- Job options
- Review

Job options

Configure the options for this restore job.

Recovery Options

☐ No Recovery ☒ Recover until end of backup ☐ Recover until end of available logs. ☐ Recover until specific point-in-time

☐ By time

00 00 00 Europe/Berlin

Application Options

☐ Overwrite existing databases.

Maximum Parallel Streams per Database 1

Advanced Options

☒ Run cleanup immediately on job failure

☐ Allow session overwrite

☒ Continue with restores of other selected databases even if one fails

[Preview Restore](#)

[Back](#) [Next](#)

Figure 6-34 Job options

Click **Next** to proceed. The last panel displays a summary for review, as shown in Figure 6-35. Click **Submit** to start the restore job.

[Back to Exchange](#)

Restore - Exchange

Review

Default Setup

- Select data sources
- Select source
- Source snapshot
- Restore method
- Set destination
- Set run settings
- Job options
- Review

Review

Review your selections, and then click Submit.

Select data sources

Selected source: MDB1

Source snapshot: MDB1 - Jul 5, 2020 11:09:30 PM

Restore Type: On-Demand: Snapshot

Restore Source Type: Site

Restore Source: Primary

Restore method: Production

Set destination

Destination: Restore to original instance

Set run settings

Run cleanup immediately on job failure: Yes

Allow session overwrite: No

Continue with restores of other selected databases even if one fails: No

Overwrite existing databases: No

Recovery type: recovery

Maximum Parallel Streams per Database: 1

[Preview Restore](#)

[Back](#) [Submit](#)

Figure 6-35 Review and submit

Under **Jobs and Operations**, you can follow the progress of the operation, as shown in Figure 6-36.

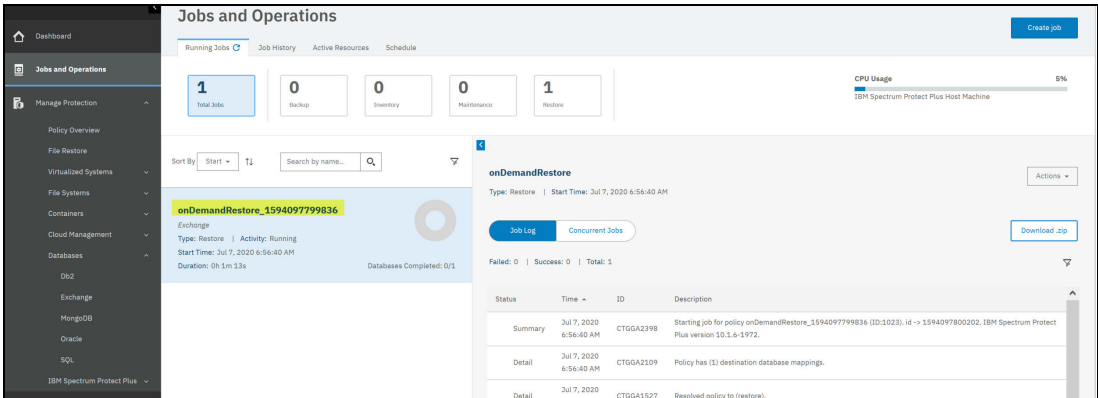


Figure 6-36 Monitoring running job

After completion, the job information is moved into the Job History window, as shown in Figure 6-37.

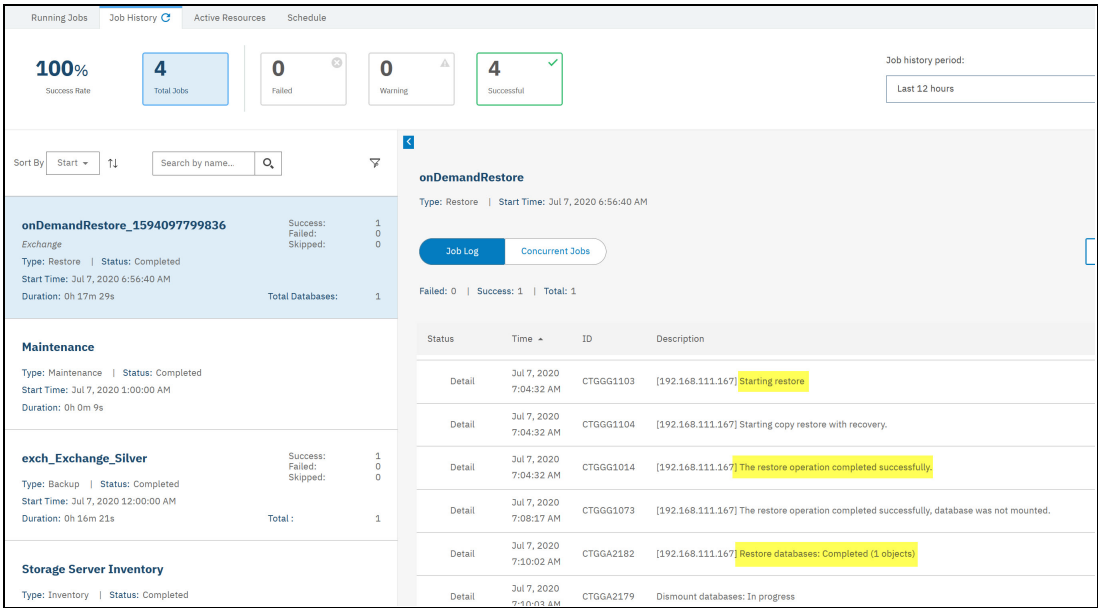


Figure 6-37 Restore Job history

6.5.2 Restoring individual items with granular restore

To recover single mailboxes or single mailbox items, such as individual mails, two methods are available: granular restore by using an Exchange Server or granular restore by using a Remote System.

Refer to the prerequisites, described in 6.2.1, “Granular restore remote package installation” on page 94.

In both cases, the restore procedure is started as described in 6.5, “Restore jobs” on page 111.

Select **Granular Restore** as shown in Figure 6-39.

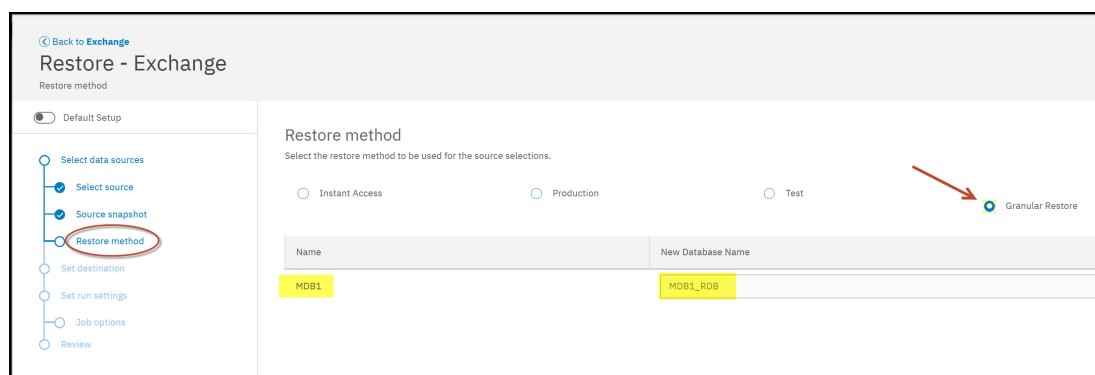


Figure 6-38 Granular restore method

This type of restore uses a Recovery Database (in our case, MDB1.RDB). The rest of the procedure is similar to what is described in 6.5.1, “Complete Restore” in successively setting destination and job options and then submitting the job.

After the restore job is started, it can be monitored in the Jobs and Operations panel.

The recovery database is created and the snapshot is mounted as Recovery Database (RDB). An excerpt from the job log in Example 6-3 shows the steps that are performed.

Example 6-3 Restore job process

```

Detail Jul 8, 2020 6:53:34 AM CTGGA2179 Granular restore databases: In progress
Info Jul 8, 2020 6:53:34 AM CTGGA1618 Granular restore for databases (MDB1)...
Detail Jul 8, 2020 6:54:04 AM CTGGA2245 SPP log dir:
/data/log/guestdeployer/2020-07-08/1594183634044/1594184014247/192.168.111.167
Detail Jul 8, 2020 6:54:26 AM CTGGG0000 [192.168.111.167] IBM Spectrum Protect
Plus Exchange Agent
Detail Jul 8, 2020 6:54:26 AM CTGGG1125 [192.168.111.167] The Exchange agent is
running as user ????\administrator, in group ???\Domain-User .
Detail Jul 8, 2020 6:54:26 AM CTGGG1103 [192.168.111.167] Starting restore
Detail Jul 8, 2020 6:54:26 AM CTGGG1104 [192.168.111.167] Starting granular
restore with recovery.
Detail Jul 8, 2020 6:54:26 AM CTGGG1014 [192.168.111.167] The restore operation
completed successfully.

```

Information is summarized in the Active Resources tab, as shown in Figure 6-39.

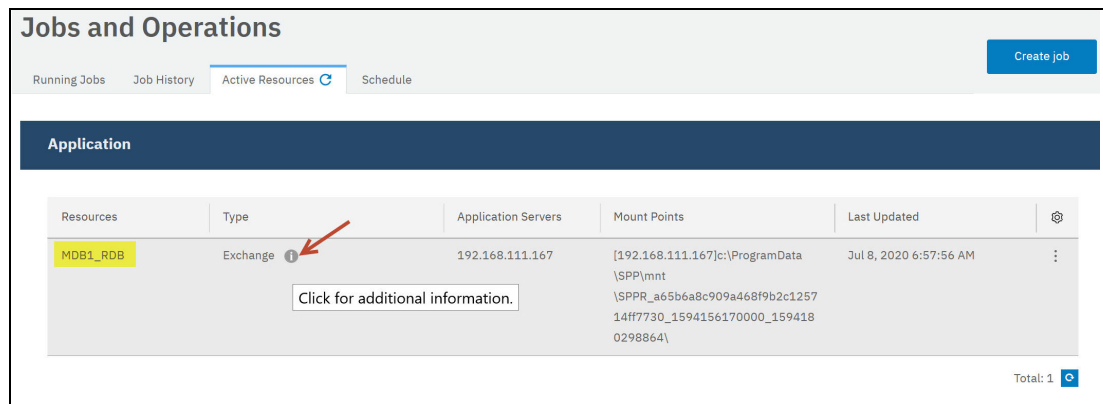


Figure 6-39 Active Resources tab

Clicking the information icon (as indicated by the arrow in Figure 6-39) in the Type column provides more information about how to start the IBM Spectrum Protect Plus MMC GUI. This GUI is automatically installed during the restore procedure on the Exchange server (see Figure 6-40).

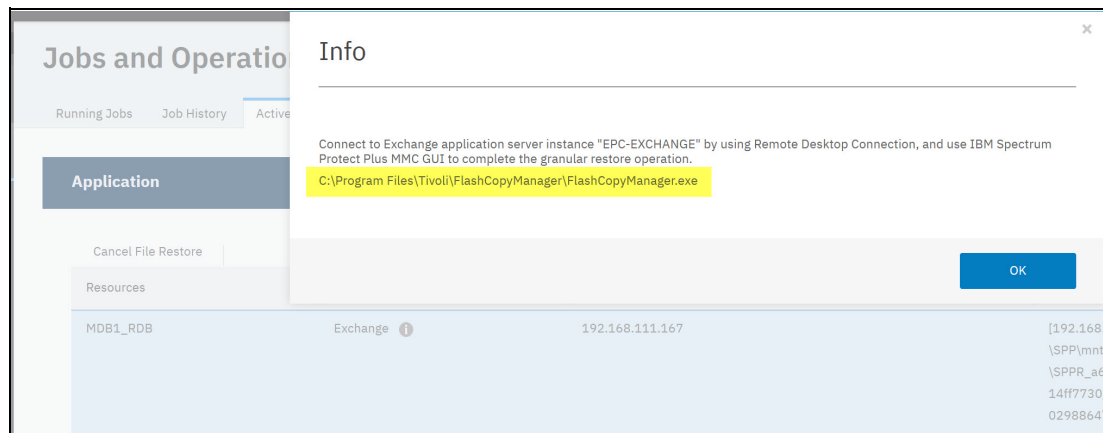


Figure 6-40 Information how to start the IBM Spectrum Protect MMC GUI

You must decide which target to use to proceed with the item recovery: install the MMC GUI in combination with Outlook 2016 on the Exchange Server or run it on a separate server.

Item Recovery by using an Exchange Server

To proceed with the recovery of individual mailbox items, complete the following steps:

1. Log on to the Exchange Server with a user IFD that has the suitable permissions.
2. Open a command window and enter the string that is provided by IBM Spectrum Protect Plus. Usually, the command must be put in quotation marks, as shown in Figure 6-41.

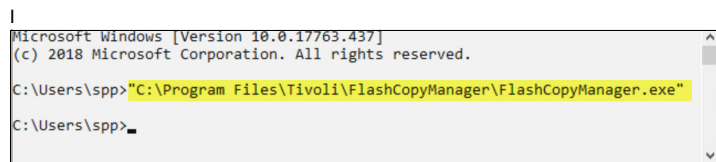


Figure 6-41 Opening IBM Spectrum Protect Plus MMC GUI

3. The IBM Spectrum Protect MMC GUI opens. The next step is to check with the wizard that all prerequisites are fulfilled. Click **IBM Spectrum Protect Plus** → **Dashboard** → **Manage** → **Configuration** and start the Configuration Wizard, as shown in Figure 6-42.

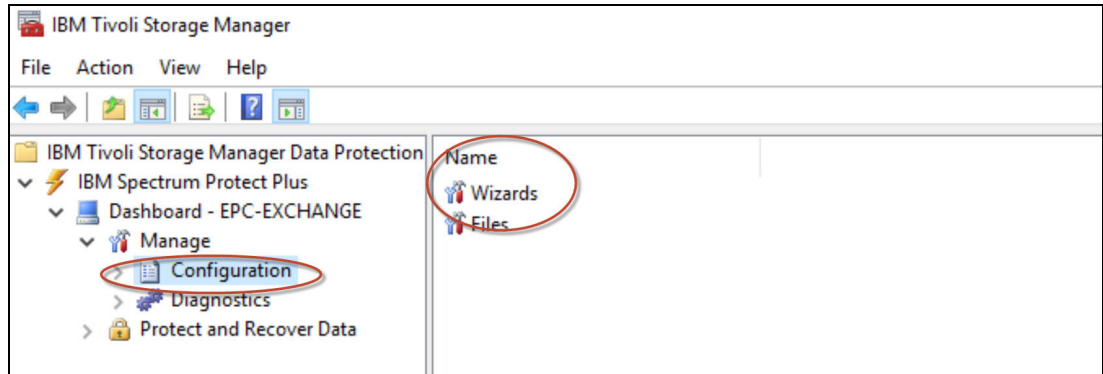


Figure 6-42 Starting the Configuration wizard

4. Click **Wizards** and the configuration option IBM Spectrum Protect Plus configuration is shown. Click **Start** to run the wizard. The result should display failed: 0, as shown in Figure 6-43.

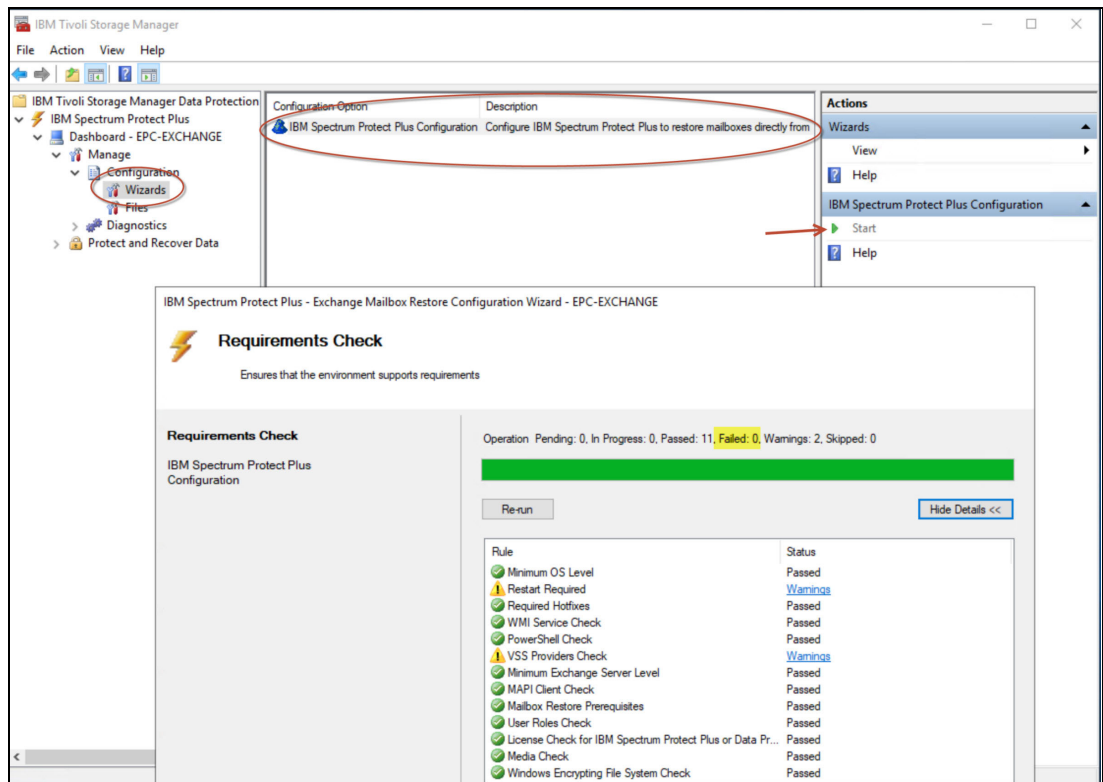


Figure 6-43 IBM Spectrum Protect Configuration wizard

5. Click **Next** and the wizard proceed and completes the process.

The warning about VSS Provider Check can safely be ignored because no IBM VSS Hardware Provider is installed and it is not necessary when restoring from IBM Spectrum Protect Plus.

The restart required warning often is caused by a pending restart (most likely after a Windows patch update on the operating system).

6. After completing the configuration wizard, proceed to recover single mailbox items.
7. Expand the **Protect and Recover Data** tab and select the Exchange server. On the right side of the display, three tabs are available: Protect, Recover, and Automate. Click the **Recover** tab.

A Configuration Error appears, as shown in Figure 6-44. This error is shown because it is not recommended to perform the recovery with the exchange server.

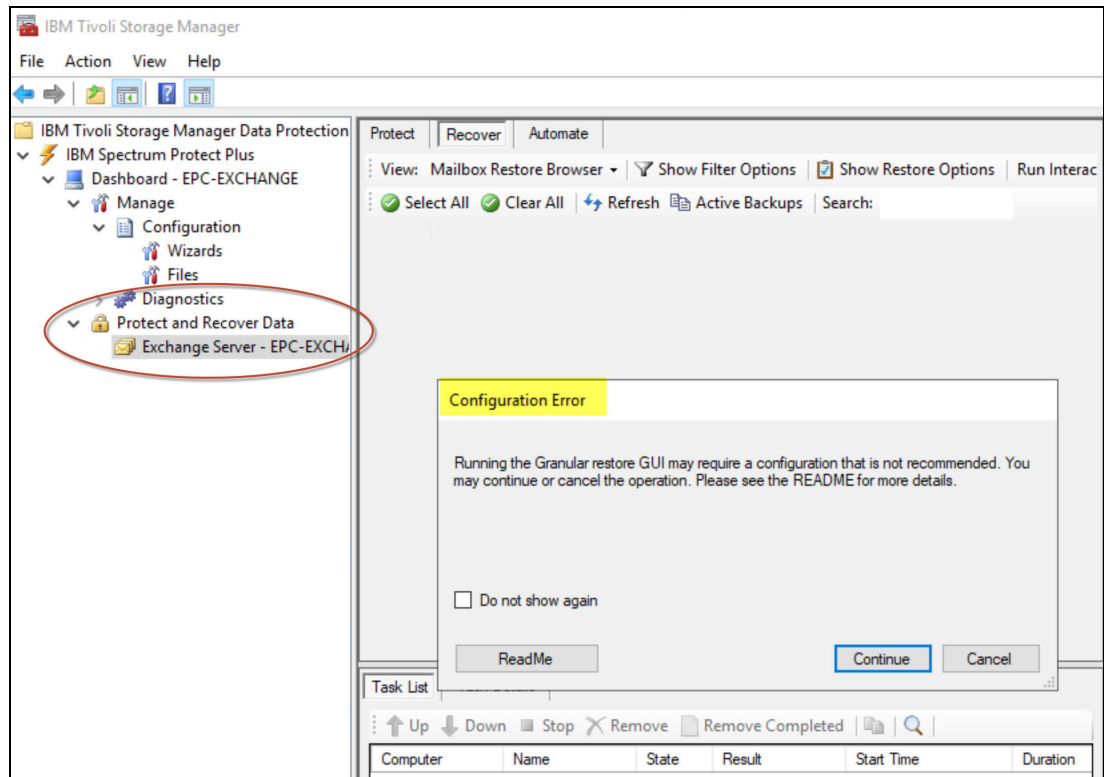


Figure 6-44 Protect and Recover Data

- Click **ReadMe** to see information that is similar to the information that was described in 6.2.1, “Granular restore remote package installation” on page 94 (see Figure 6-45).

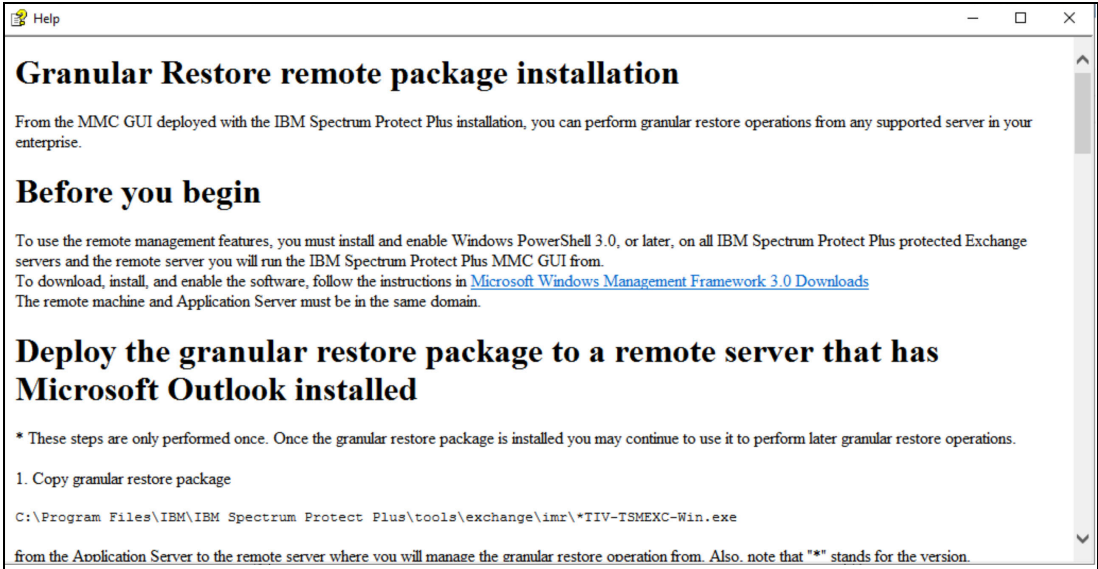


Figure 6-45 Installing ReadMe Granular Restore remote package

- The Recovery Database (RDB) opens, but no mailbox is selected. The mailboxes appear as closed. Proceed with the recovery by selecting the **Mailbox Restore Browser** view, as shown in Figure 6-46.

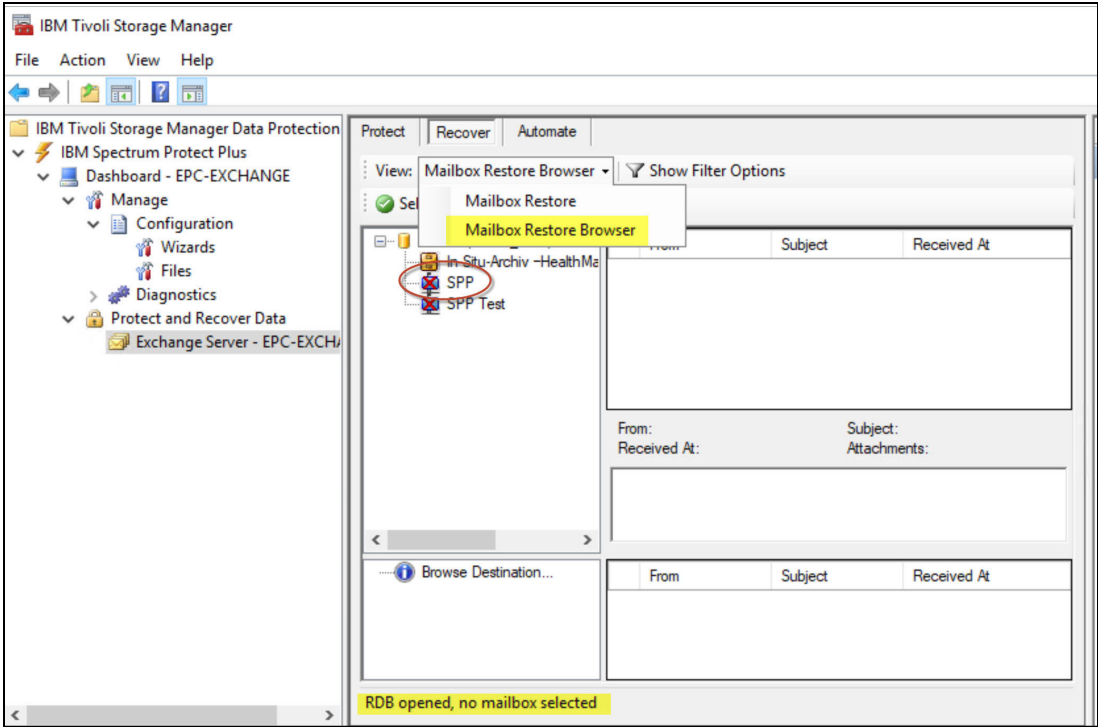


Figure 6-46 Mailbox Restore Browser

By clicking the mailbox icon (in our example, SPP), the mailbox is populated and the items are provided for recovery. This process can take some time.

10. The populated mailbox now shows the mailbox items, such as inbox (see Figure 6-47). Click the inbox and all mail objects are shown. By selecting individual mail items, the content is shown in the middle part of the window.

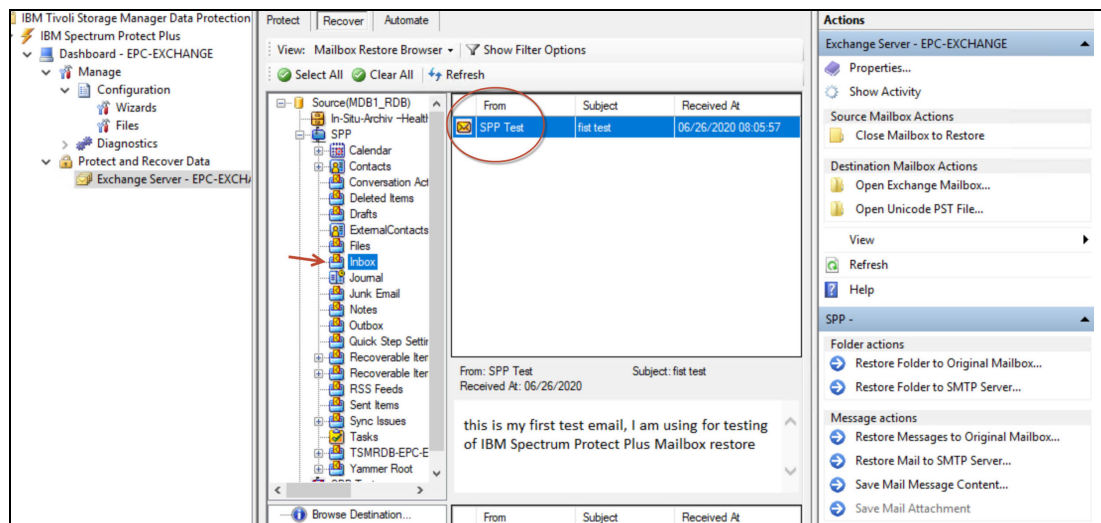


Figure 6-47 Item select

In the Actions column on the right side of the display, the choices for the recovery are listed. The column is divided into Folder Actions and Message Actions sections. We can recover folders or single messages.

11. Click the **Restore Messages to Original Mailbox** entry. The restore from the Recovery Database (RDB) goes done into the active database. The restore progress and the result are displayed in a separate window, as shown in Figure 6-48.

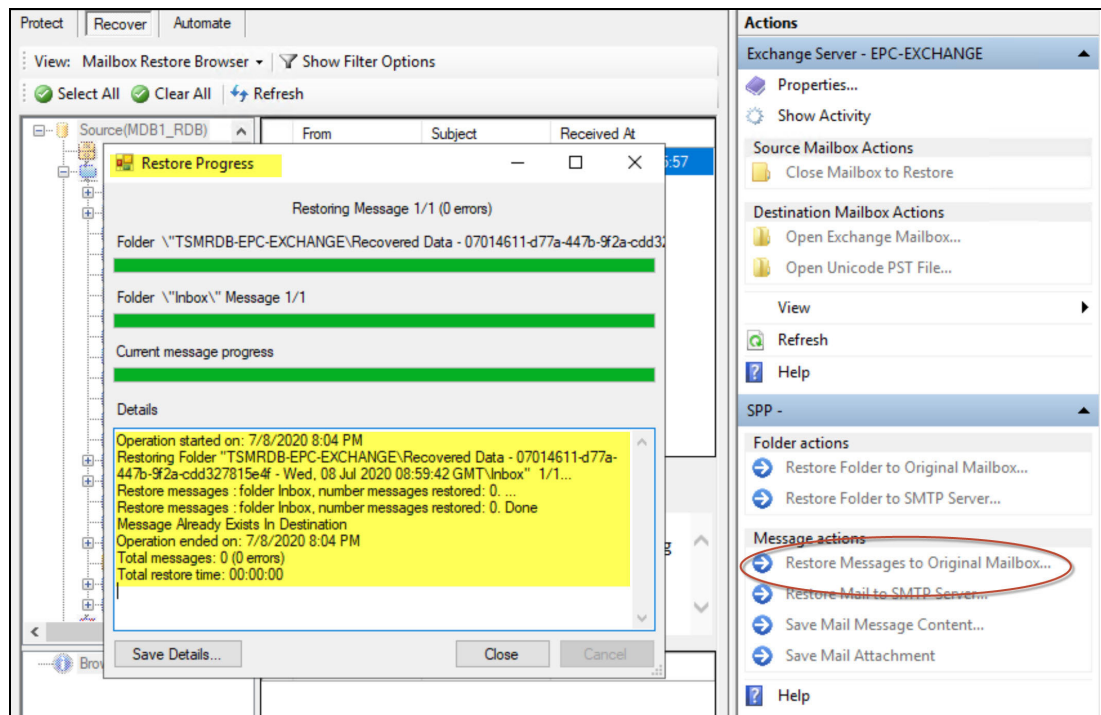


Figure 6-48 Restore progress: Restoring message

After successful recovery, a cleanup procedure must be completed on the IBM Spectrum Protect Plus Server. This cleanup can be done in **Jobs and Operations** → **Active Jobs**, by canceling the running job, as shown in Figure 6-49.

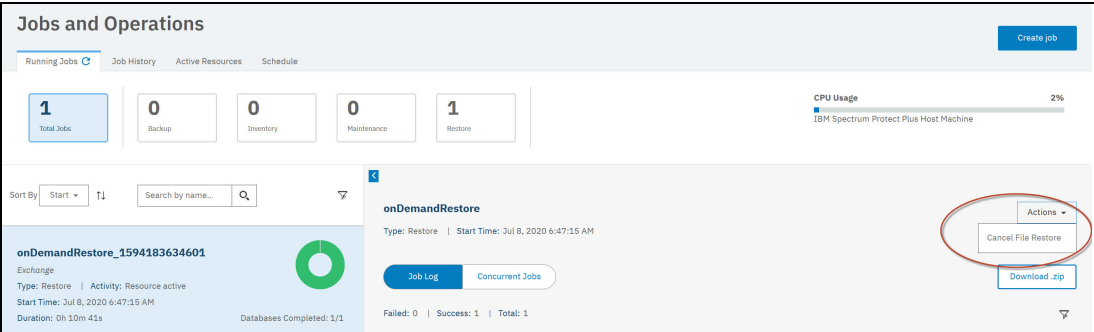


Figure 6-49 Cleanup in Jobs and Operations: Running Jobs

Item Recovery restore job also can be stopped is in **Jobs and Operations** → **Active Resources** by clicking the three vertical dots and selecting **Cancel job**, as shown in Figure 6-50.

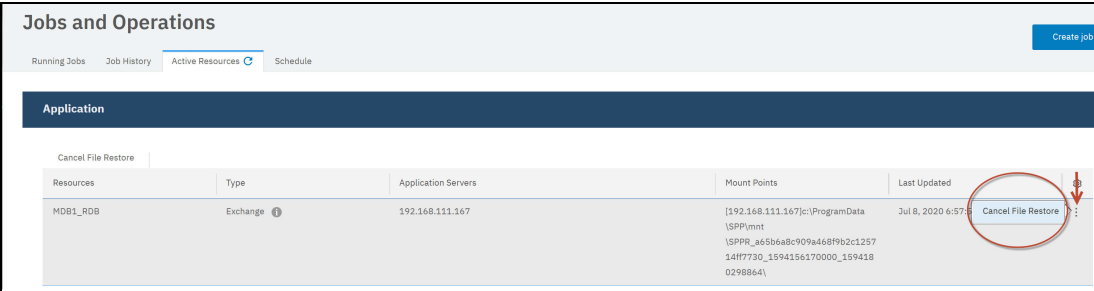


Figure 6-50 Clean up in Jobs and Operations: Active Resources

The Job History Job Logs includes the detailed log of the cleanup procedure and is confirmed with a success message, as shown in Figure 6-51.

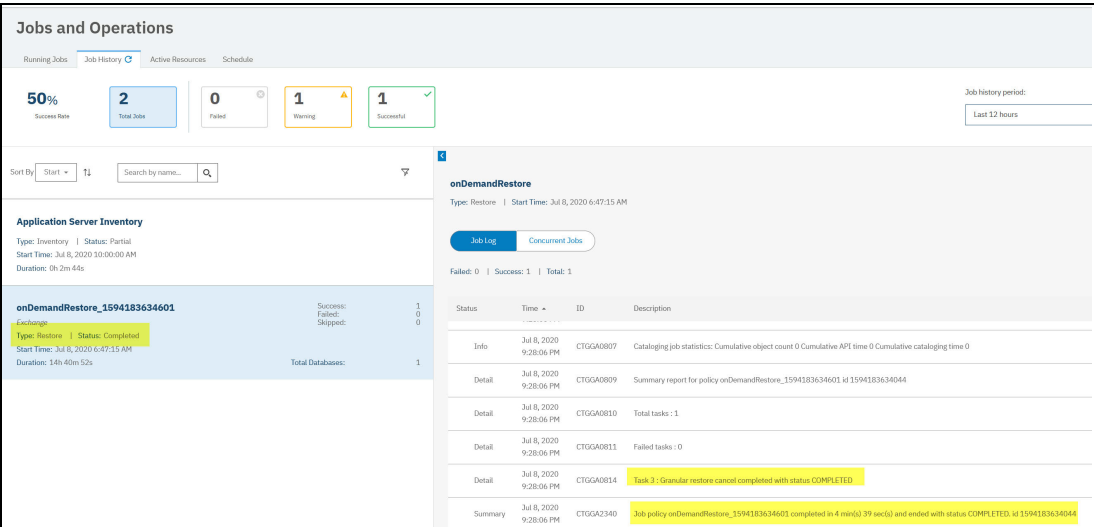


Figure 6-51 Job History cleanup

Item Recovery by using a Remote System

The restore of individual items with the IBM Spectrum Protect MMC GUI is done on a separate Windows system, which is called *Remote System*.

The Exchange server must be added as a managed computer so that it appears in the Group → Dashboard view, as shown in Figure 6-52.

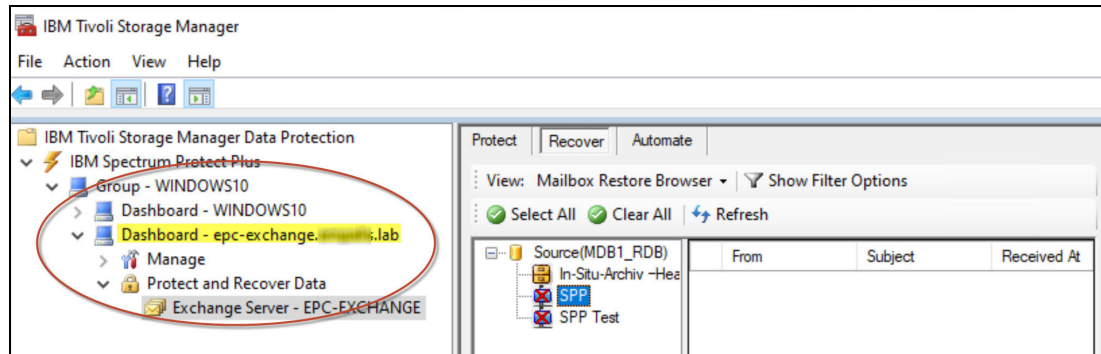


Figure 6-52 Group Dashboard view

After expanding the Protect and Recover Data entry, the Mailbox Restore Browser shows the available mailbox items in Recovery Database (RDB) that are connected to the Exchange Server and provided through the PowerShell communication.

The recovery procedure on a remote system is identical to the recovery procedure on the Exchange Server, as described in “Item Recovery by using an Exchange Server” on page 119.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *IBM Spectrum Protect Plus Practical Guidance for Deployment, Configuration, and Usage*, REDP-5532.

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

The following websites are also relevant as further information sources:

- ▶ Featured Documents for IBM Spectrum Protect Plus:
<https://www.ibm.com/support/pages/featured-documents-ibm-spectrum-protect-plus>
- ▶ IBM Spectrum Protect Plus BluePrints:
<https://ibm.biz/IBMSpectrumProtectPlusBlueprints>
- ▶ IBM Spectrum Protect Plus Documentation:
<https://www.ibm.com/docs/en/spp>
- ▶ IBM Spectrum Protect Plus Support:
https://www.ibm.com/mysupport/s/topic/0T050000000IQWtGA0/spectrum-protect-plus?language=en_US&productId=01t50000004uZGc
- ▶ IBM Spectrum Protect Plus - All Requirements Doc:
<https://www.ibm.com/support/pages/ibm-spectrum-protect-plus-all-requirements-doc>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5640-00

ISBN 0738459852

Printed in U.S.A.

Get connected

