Redbooks
ibm.com/redbooks

IBM Spectrum Protect Plus
Practical Guidance for Deployment, Configuration, and Usage

Gerd Becker
Chris Bode
Alberto Delgado Ramos
Bert Dufrasne
Andre Gaschler
Mikael Lindstrom
Peter Minig
Julien Sauvanet
Martin Stuber
Markus Stumpf
Jozef Urica
Joerg Walter
Daniel Wendler
Axel Westphal

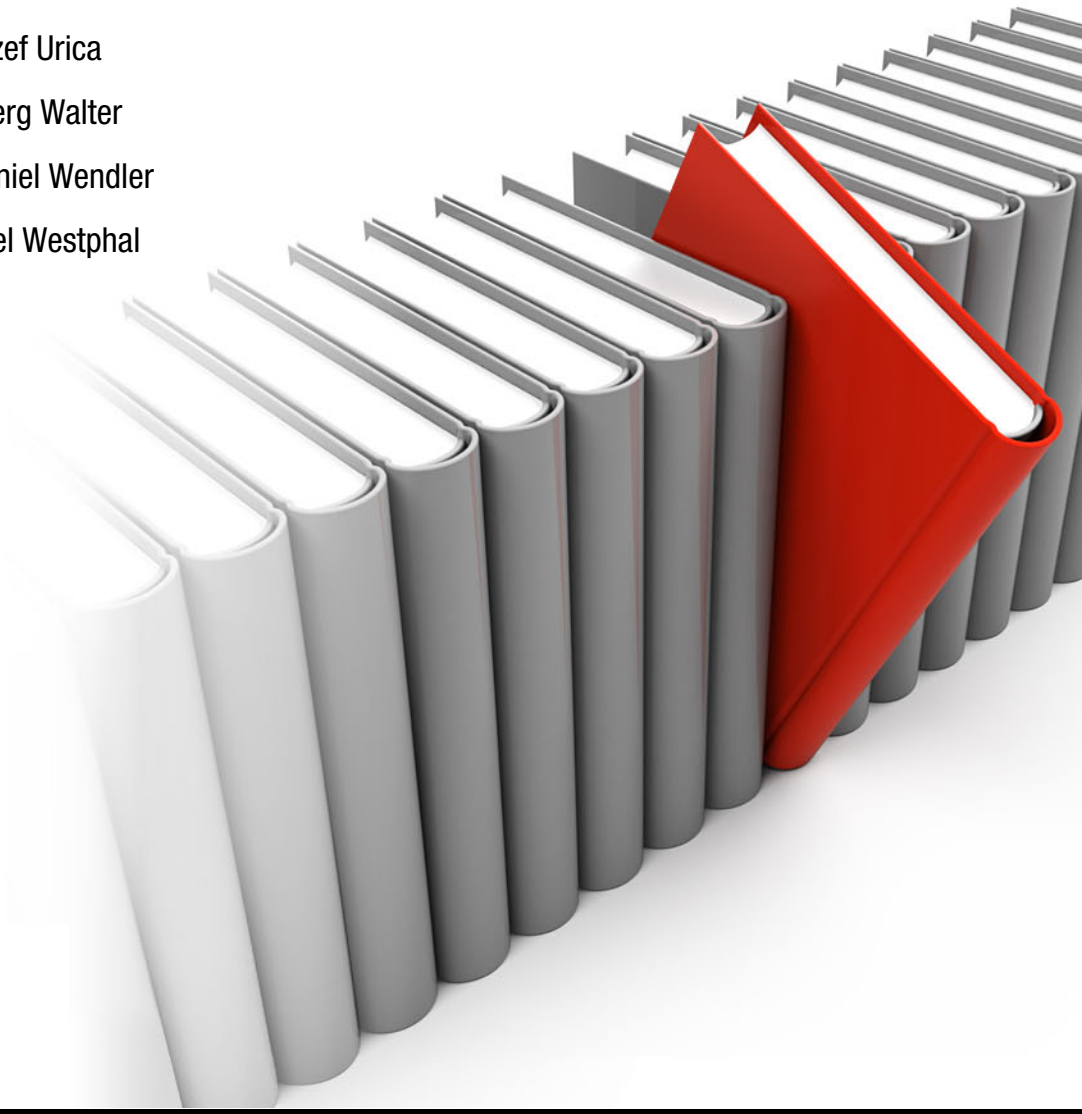

IBM

Redpaper

IBM Redbooks

Spectrum Protect Plus Usage Scenarios Best Practices

December 2020

REDP-5532-01

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

Second Edition (December 2020)

This edition applies to Version 10.1.6 of IBM SPectrum Protect Plus.

This document was created or updated on December 14, 2020.

Contents

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®	IBM Cloud®	Redbooks®
Db2®	IBM Resiliency Services®	Redbooks (logo)  ®
DB2®	IBM Security™	Resilient®
Global Technology Services®	IBM Spectrum®	Storwize®
IBM®	Passport Advantage®	Tivoli®

The following terms are trademarks of other companies:

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

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

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

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

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

VMware, VMware vSphere, 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 applications backups in hybrid multicloud environments.

IBM Knowledge Center for IBM Spectrum® Protect Plus provides extensive documentation for installation, deployment, and usage. In addition, IBM Spectrum Protect Plus Blueprints provide guidance about how to build and size an IBM Spectrum Protect Plus solution.

The goal of this IBM Redpaper® publication is to summarize and complement the available information by providing useful hints and tips that are based on the authors' practical experience in installing and supporting IBM Spectrum Protect Plus in customer environments. Over time, our aim is to compile a set of best practices that cover all aspects of the product, from planning and installation to tuning, maintenance, and troubleshooting.

The paper features the following structure:

- Chapter 1, "IBM Spectrum Protect Plus product architecture and components" on page 1 provides an overview of key components and characteristics of IBM Spectrum Protect Plus version 10.1.6. It is followed by a deeper dive into the architecture, concepts, and operational features, including encryption.
- Chapter 2, "Solution architecture, planning, and design" on page 25 discusses architectural planning considerations, such as the system context, roles and responsibilities, and functional and non-functional requirements. It also reviews architectural decisions that drive the design of the solution for deployment in your environment.
- Chapter 3, "Installation and deployment" on page 47 focuses on installation and deployment, and additional configuration actions that must occur after the initial installation. For more information about the installation, see IBM Knowledge Center or the *Installation and User's Guide*.
- Chapter 4, "Networking" on page 83, provides advice for setting up the network environment and separation, if necessary, between backup data traffic and system management traffic.
- Chapter 5, "Daily operations and maintenance" on page 107, describes daily operations, maintenance, troubleshooting, and monitoring. It discusses specifics of the Role Based Access Control concept and its implementation in IBM Spectrum Protect Plus. It also describes the vSnap Backup Storage server command-line interface, which is used to administer the vSnap server and configure advanced options.
- Chapter 6, "Backing up and restoring virtualized systems" on page 171 - Chapter 11, "Backing up and restoring SQL Server" on page 309 discuss backup, restore, and reuse operations for virtual machines, file systems, and supported databases (Oracle, MongoDB, IBM Db2®, and Microsoft SQL Server).
- Chapter 12, "Backing up and restoring Microsoft Exchange data" on page 329 and Chapter 13, "Backing up and restoring Microsoft 365 data" on page 365, focus on backup and restore operations for Microsoft Exchange and Microsoft 365.
- Chapter 14, "Backing up and restoring containers" on page 381, reviews the concept of containers in a Kubernetes environment and explains how IBM Spectrum Protect Plus can be used to provide protection for containerized workloads.

- Chapter 15, “Replication and additional copies” on page 405 explains options for creating extra copies of data for Disaster Recovery purposes or for long-term archiving on tape or in the cloud.
- Chapter 16, “REST API” on page 461 provides practical information and examples about how to take advantage of the REST API services.

Authors

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

Gerd Becker is a Project Manager for EMPALIS Consulting GmbH, an IBM Business Partner in Germany. He has more than 40 years of IT experience, including over 20 years experience with storage management products, such as DFSMS and IBM Tivoli® Storage Manager. He holds several certifications, including technical and sales, and is an IBM Tivoli Certified Instructor. He has been Chairman of the Guide Share Europe (GSE) Storage-User group for more than 20 years. Gerd has authored many IBM Redbooks® publications. He is also an IBM Champion 2019 and 2020.

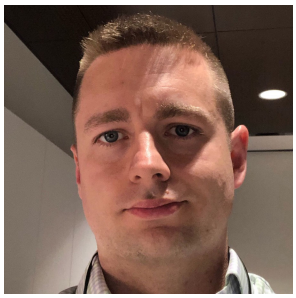

Chris Bode is an Offering Development Architect in IBM Resiliency Services® working from the Research Triangle Park in North Carolina. He has worked for the past 10 years in backup and Disaster Recovery-related XaaS offerings. Before that, he worked within the service delivery arm of IBM’s Backup as a Service business. He has broad expertise in enterprise-scale on-premises and public cloud data protection solution design.

Alberto Delgado Ramos is an IBM Spectrum Protect SME in Madrid, Spain. He joined IBM in 2007, working for IBM Global Services as a Backup Engineer. Since 2016, he is the SPGI Market Process Execution Specialist (PES) for BaaS Process Compliance in Europe. His areas of expertise include supporting customers with designing, planning, and implementation solutions that are based on IBM Spectrum Protect and IBM Spectrum Protect Plus.

Bert Dufrasne is an IBM Certified Consulting IT Specialist and Project Leader for IBM System Storage products at the International Technical Support Organization (ITSO), San Jose Center. He has worked at IBM in various IT areas. He has authored many IBM Redbooks publications, and he has also developed and taught technical workshops. Before Bert joined the ITSO, he worked for IBM Global Services as an Application Architect. He holds a Master’s degree in Electrical Engineering.

Andre Gaschler is an IBM Certified IT Specialist at the IBM EMEA Storage Competence Center (ESCC) in Kelsterbach, Germany. He has been working at IBM for 25 years with more than 15 years of experience with IBM Spectrum Protect. He is a member of IBM's Systems Storage Lab Services team, supporting customers with designing, planning, and implementation solutions that are based on IBM Spectrum Protect and IBM Spectrum Protect Plus. He contributed to the backup chapter of the book *Speichernetze*, and is an author several IBM Redbooks publications.

Mikael Lindstrom is a Senior Technical Staff Member in Resiliency and an IBM Inventor within the IBM Global Services. He has 19 years of IT experience. Mikael advises executive management and offers senior leadership on emerging technology and design elements. Mikael also provides customer-facing expertise for projects worldwide, helping to solve their challenges around resiliency, and to build future-proof Software Defined Resiliency Cloud infrastructure. He has co-authored over 10 IBM Redbooks publications and white papers.

Peter Minig has a degree in electrical engineering and computer sciences in Germany. He joined IBM in 1998 and started by supporting client and server infrastructure, gaining experience in systems management. He is a member of IBM's Systems Storage Lab Services team, supporting customers with designing, planning, and implementing DP&R related products, such as IBM Spectrum Protect, IBM Spectrum Protect Plus, and Tape Storage.

Julien Sauvanet is an IBM Certified Expert IT Specialist, working in the French Infrastructure Services organization. He is also involved in architecture and deployments world-wide for IBM Global Technology Services®, with a focus on IBM Spectrum Protect and IBM Spectrum Protect Plus.
Julien is a subject matter expert in system and storage virtualization. He has co-authored three other IBM Redbooks publications and also published several IBM Tivoli Storage Manager white papers.

Martin Stuber is a Backup and Restore SME and IBM Certified IT Specialist and IT Architect for IBM Global Services Delivery Center, Czech Republic. He has over 17 years of experience in the IT field. He joined IBM in 2008. His area of expertise includes Backup and Archive Solution design, deployment, and data protection implementation. He has written extensively about backup and recovery, installation, operation, and other related architecture topics.

Markus Stumpf is Head of Services and Operations at Empalis Consulting GmbH, an IBM Business Partner in Germany. He handles petabytes of customer data in the service division and manages hundreds of IBM Spectrum Protect Servers. Markus also contributed to several certification tests for IBM Spectrum Protect and IBM Spectrum Protect Plus. He has worked with Tivoli Storage Manager, IBM Spectrum Protect, and IBM Spectrum Protect Plus for more than 16 years and is especially engaged in backing up large database and mail environments.

Jozef Urica is an IBM Certified IT Specialist at the IBM Global Services Delivery Center Czech Republic. He joined IBM in 2010, working for IBM Global Services as a Systems Engineer. His areas of expertise include supporting customers with designing, planning, and implementing IBM Spectrum Protect, IBM Spectrum Protect Plus, and Tape Storage. He holds a Master's degree in Electrical Engineering and Communication in Czech Republic.

Joerg Walter is an IBM Certified IT Specialist at the IBM EMEA Storage Competence Center (ESCC) in Kelsterbach, Germany. Since he joined IBM in 2000, he worked in various technical positions, starting with network engineering and design, but changed focus to storage, data protection, and retention in 2005. Since 2011, Joerg is a member of IBM's Systems Storage Lab Services team, supporting customers with the planning and implementation of DP&R related hardware and software products, such as IBM Spectrum Protect, IBM Spectrum Protect Plus, and IBM Spectrum Protect Snapshot, and physical and virtual tape.

Daniel Wendler has a degree in computer sciences. After completing his studies at the University of Applied Sciences in Wiesbaden, Germany, he joined IBM in 2005. He worked in various technical positions within IBM. He is a member of IBM's EMEA Systems Storage Lab Services pre-sales team, supporting customers with the planning and implementation of DP&R and IBM Security™ Key Lifecycle Manager encryption solutions.

Axel Westphal is an IBM Certified IT Specialist at the IBM EMEA Storage Competence Center (ESCC) in Kelsterbach, Germany. He joined IBM in 1996, working for IBM Global Services as a Systems Engineer. His areas of expertise include setting up and demonstrating IBM System Storage products and solutions in various environments. He has written several storage white papers and co-authored several IBM publications.

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

Jason Basler
Stefan Bender
Andy Cheong
Jason Cooley
Michael Kulessa
Dominic Mueller-Wicke
Gerd Munz
Linda Sandmann
Jim Smith
Tuan Vu
Adam Young
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 Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for 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

1

IBM Spectrum Protect Plus product architecture and components

This chapter starts with a general overview of the IBM Spectrum Protect Plus solution and its key concepts.

The chapter continues with more details about the product architecture and its main components, including the IBM Spectrum Protect Plus server, the vSnap Backup Storage server, and the VMware vSphere Storage APIs for Data Protection (VADP) proxy server. It briefly explains how service level agreement (SLA) policies and sites are used to manage backup data.

This chapter includes the following topics:

1.1 IBM Spectrum Protect Plus overview

IBM Spectrum Protect Plus provides data availability by using snapshot technology for rapid backup, recovery, and data management. It features a secure, self-service, user-facing portal for data protection management and monitoring.

Designed to be used for backup of virtualized systems, databases, file systems, cloud-managed applications, and containers, it also provides data cloning functions to support and automate DevOps workflows. Unlike other data availability solutions, IBM Spectrum Protect Plus uses automated Service Level Agreements (SLAs) to provide backup status and support retention compliance.

Figure 1-1 shows an overview of the IBM Spectrum Protect Plus architecture.

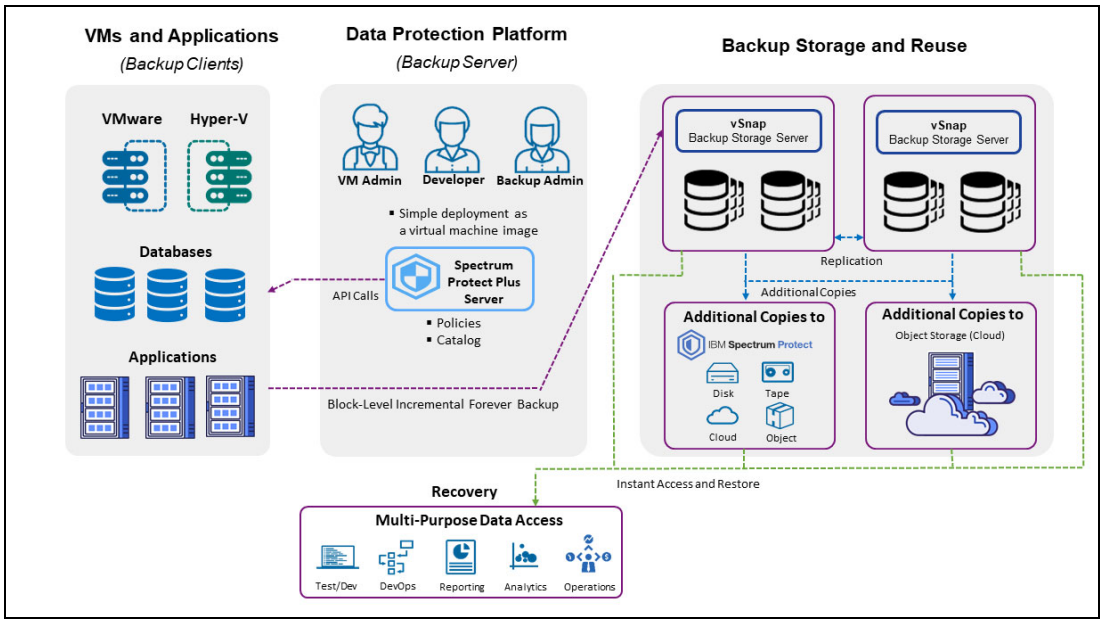

Figure 1-1 IBM Spectrum Protect Architecture Overview

The IBM Spectrum Protect Plus server is the component that manages and orchestrates the entire system. It also is the “brain” that also provides the web interface portal that is used for configuring and operating the solution, and performing centralized scheduling of activities. The vSnap Backup Storage server is the component that is responsible for storing and processing the backup data that is received from production systems.

A typical deployment includes a single backup server (IBM Spectrum Protect Plus server) with multiple backup storage servers (vSnap servers), depending on customer requirements and needs. Because IBM Spectrum Protect Plus is a zero touch data protection solution, a traditional backup client software component is not required.

Figure 1-1 also shows the key component in the middle of Figure 1-1, where the IBM Spectrum Protect Plus server is deployed as a virtual machine. The IBM Spectrum Protect Plus server communicates by way of standardized API calls with the backup clients, which can run on virtual machines, physical servers, or cloud instances. This communication is used in the inventory processes, and the backup processing during the backup period. The IBM Spectrum Protect Plus server is responsible for scheduling, starting, and cataloging the backups.

Backups that are taken as block-level incremental forever backups from virtual machines, databases, and applications are stored according to the backup policies (SLA) as a read/write snapshot in the vSnap Backup Storage server. This vSnap server can run on a virtual or physical server, and can be replicated to a second vSnap server as an operational protection.

If required, data also can be tiered to IBM Spectrum Protect, onto multiple different storage options, such as tape copy or Object Storage, or onto a dedicated Object Storage such as IBM Cloud® Object Storage, for more protection.

The vSnap server manages the backup traffic and workload. The backup metadata is sent to the IBM Spectrum Protect Plus server. To achieve local backup data placement by using SAN or LAN networks and to scale out, multiple vSnap servers are typically used within an IBM Spectrum Protect Plus solution. For efficient performance, proper sizing of the vSnap server is crucial. For more information, see the IBM Spectrum Protect Plus Blueprints that are available at this web page.

Finally, you can then use any backup image to restore back to the virtual environment for recovery, or for any of the multi-purpose use cases, such as instant data recovery boot by way of vSnap repository or data reuse of existing data copies for tests and DevOps. Instant access and restore capabilities are provided by the different backup storage pools, including vSnap pool, IBM Spectrum Protect disk container, and tape and Object Storage, such as IBM Cloud Object Storage.

The built-in role-based access control (RBAC) model enables you to allocate and control access to specific areas and resources within an IBM Spectrum Protect Plus data protection solution, and to dedicated teams. You can use the RBAC role model to assign different roles for VM administrators or database administrators. After these roles are allocated, these users can create backup jobs, restore data, or apply any of the multiple data reuse cases.

1.1.1 Key concepts

This section discusses key concepts in IBM Spectrum Protect Plus.

The primary users for the IBM Spectrum Protect Plus solution are customers running virtualized workloads in a traditional data center, cloud data center, and remote branch offices. IBM Spectrum Protect Plus is also capable of providing multitenancy for service providers. See also Chapter 2, “Solution architecture, planning, and design” on page 25.

- Native data format. Traditional backup products take the data that they obtain from client systems and package it in proprietary formats on disk or tape. Before the data is usable, the backup application must first extract it from the packaged format. IBM Spectrum Protect Plus, by contrast, stores data in the format that is created by the application it is protecting, such as `.vmdk` files for VMware or Hyper-V `.vhdx` files.
- Instant data recovery boot by way of vSnap repository is at the core of the IBM Spectrum Protect Plus solution. It allows for rapid recovery or access to data, without having to wait for traditional data restore operations to move data.
- Zero touch backup and recovery of virtualized systems, databases, file systems, cloud-managed applications, and containers. Traditional backup approaches require the installation of a heavyweight application on each protected machine, which can be time-consuming and intrusive.

  IBM Spectrum Protect Plus uses a modern *zero touch* approach with remote code injection/execution to avoid the need to manually roll out and configure application code across the environment. These benefits also extend beyond implementation because upgrades are no longer required for every client in the environment.

- RESTful API functions that are supported by REST, such as automation, integration, configuration, and data collection. By using the REST API, IBM Spectrum Protect Plus is able to integrate with automation tools like Jenkins, Puppet, and others.
- Encryption at rest. Encryption at rest is a requirement for many service providers and Business Partners bound to the General Data Protection Regulation (GDPR), Europe's new data privacy regulation that went into effect on May 25, 2018. IBM Spectrum Protect Plus offers the option to encrypt data using AES 256-bit encryption in the vSnap server at-rest.
- On-demand data reuse of existing backup copies for Test/Dev, DevOps, analytics, and reporting.
- Scale-out architecture. Some traditional backup products handle larger environments by augmenting the size and performance of a single "island" installation of the product. IBM Spectrum Protect Plus scale-out architecture aligns with future growth needs and handles increases in hardware size with multiple load handlers, known as vSnap servers.
- The data lifecycle is policy-driven by service level agreements (SLA), enabling flexible configurations with local and replica copies, and with copies with on tape or Object Storage. This configuration provides long-term retention protection on low-cost, low-performance storage, each with independently managed retentions. The solution allows for a cloud-friendly scale-out architecture for future growth needs.
- Role-based access control (RBAC) is an easy to use, self-service access capable, web-based user interface where users can run traditional restore operations. Customers can define users and assign them roles, such as that of a hypervisor admin or a database admin so that those users can run actions like restores jobs while having their scope limited to the systems or data types for which they are responsible. Teams can access and reuse existing data copies for Test/Dev, DevOps, analytics, and reporting on demand.
- Data reduction. Compression and Deduplication are a technology to reduce the size of backup data that is stored on disk. Deduplication works by examining incoming backup data streams and looking for pieces of data that is in storage. When the deduplication engine encounters a piece of data that it has seen before, rather than storing it again the engine makes metadata references to where the data is already stored. Data that needs to be stored is first compressed before it is written on disk.
- Incremental forever back up. IBM Spectrum Protect Plus uses a full-once, incremental forever backup technology, which minimizes the amount of redundant data that the system must transfer and process. If you compare this to a backup system that periodically requires a full backup of data that mostly has not changed, more potential exists for data deduplication because more redundant data is present.
- DR and High Availability Replication. In the context of backup replication, this refers to taking backups from the backup client, storing them to a vSnap server, and sending a copy to another vSnap server (normally in a separate physical location). This action protects against the loss of data and safeguards the primary backup server and associated storage infrastructure. When replication is done in backup software rather than at the storage layer, the software can determine the retention for each copy of the data and store it on different types of storage. This action ensures that the replica is properly accounted for and reported on. Restores can be taken from the primary backup target and the replication backup target.

1.2 IBM Spectrum Protect Plus architecture

This section focuses on the IBM Spectrum Protect Plus architecture and its components. Figure 1-2 shows the deployment of IBM Spectrum Protect Plus in two active locations, with the following key components:

- A single IBM Spectrum Protect Plus server
- One or more vSnap servers (multiple in the example)
- One or more VADP proxy servers (deployed as a separate server, or collocated with a vSnap server, multiple in the example). Only a VDAP proxy is used and required with VMware and is responsible for moving data from vSphere data stores. Each site can include one or more VADP proxies.

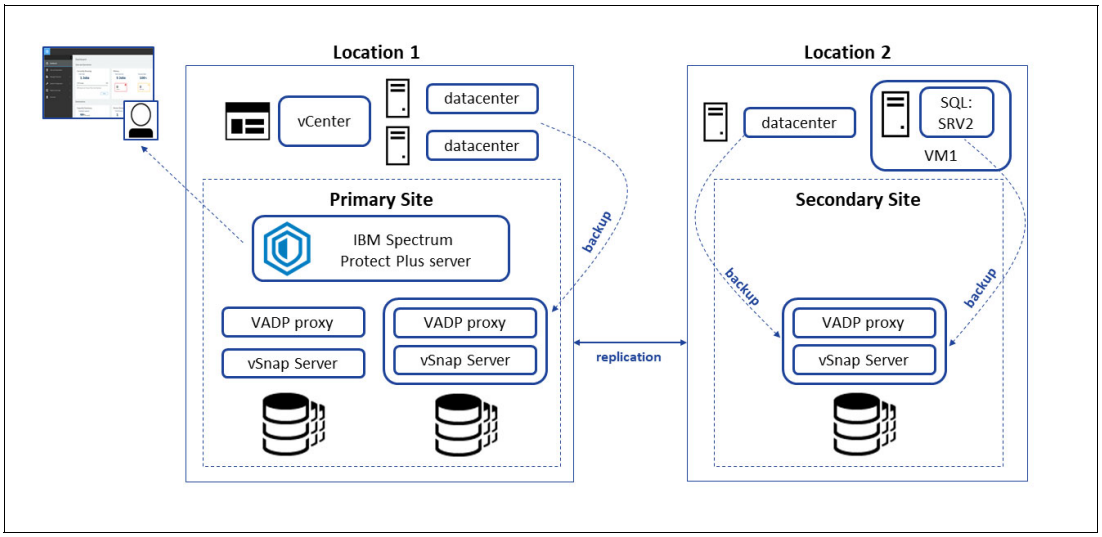

Figure 1-2 IBM Spectrum Protect Plus deployed in two active locations

In the following sections, we explain the roles and features of these key components and how they interact with each other and further elaborate on key concepts Sites and SLA policies.

1.2.1 IBM Spectrum Protect Plus server

This component of the infrastructure manages and orchestrates the entire system. The server consists of several catalogs that track various system aspects, such as recovery points, configuration, access, and customizations. Therefore, all system configurations, backup policies (SLAs), backup schedules, and daily operation aspects are done by using this server.

The server is deployed and pre-configured as a virtual machine, including several pre-configured disks, as shown in Figure 1-3 on page 6. It also provides the web-based graphical user interface (GUI) and a job scheduler.

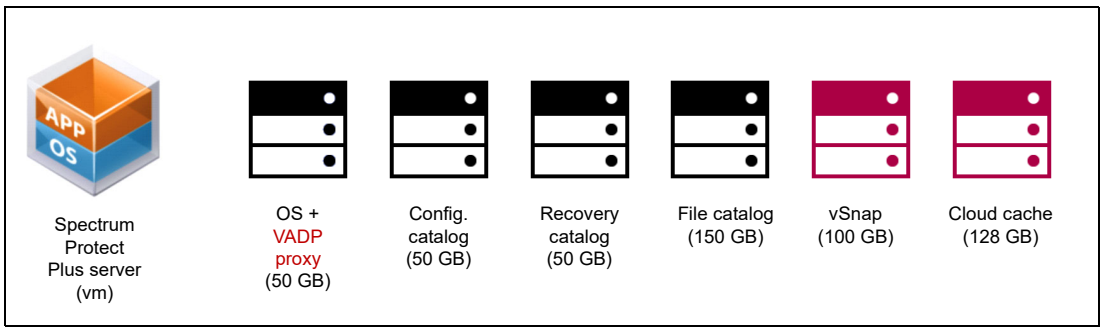

Figure 1-3 IBM Spectrum Protect Plus server: initial storage layout overview

The IBM Spectrum Protect Plus Blueprints provide a basic recommendation to host the vSnap server (including Cloud Cache) and the VADP proxy on dedicated hosts and remove the disks for the vSnap server and Cloud Cache from the IBM Spectrum Protect Plus server. The integrated vSnap server and Cloud Cache are used for demonstration purposes only, but not for production.

A sample configuration of this recommended setup is shown in Figure 1-2 on page 5 where vSnap server and VADP proxy are separated from the IBM Spectrum Protect Plus server. For more information about the required installation and configuration steps for a IBM Spectrum Protect Plus server, vSnap server, and VADP proxy, see Chapter 3, “Installation and deployment” on page 47.

Next, we describe the user interface, catalogs, and job scheduler of the IBM Spectrum Protect Plus server.

User interfaces

The server provides two web-based graphical user interfaces (GUIs):

- The Day-to-day GUI with access to the Dashboard, Backup Configuration, and so on
- The Administrative Console GUI with access to license and certificates management, Network and System Settings, and to apply software updates.

In addition to the GUI, a command-line interface (CLI) is provided to access advanced networking and system settings. A REST API interface can also be used for system automation. The CLI is further explained in Chapter 3, “Installation and deployment” on page 47 and 5.4, “vSnap server CLI” on page 156. The REST API interface is discussed in the Chapter 16, “REST API” on page 461.

After logging in to the regular web-based GUI by using https (for example, `https://spp-server/`) the Dashboard view is loaded, as shown in Figure 1-4 on page 7. The Dashboard view provides a direct overview and links to the following information:

- Jobs and Operations: Currently running jobs, and a history of jobs in the categories of Failed, Warning, and Success, are displayed together with a job success rate for a specific time frame.
- Destinations: Backup capacity overview is provided, together with a device status (inactive and full) and overall Data Reduction statistics.
- Coverage: Providing an overview of how many discovered resources are protected using different policies.

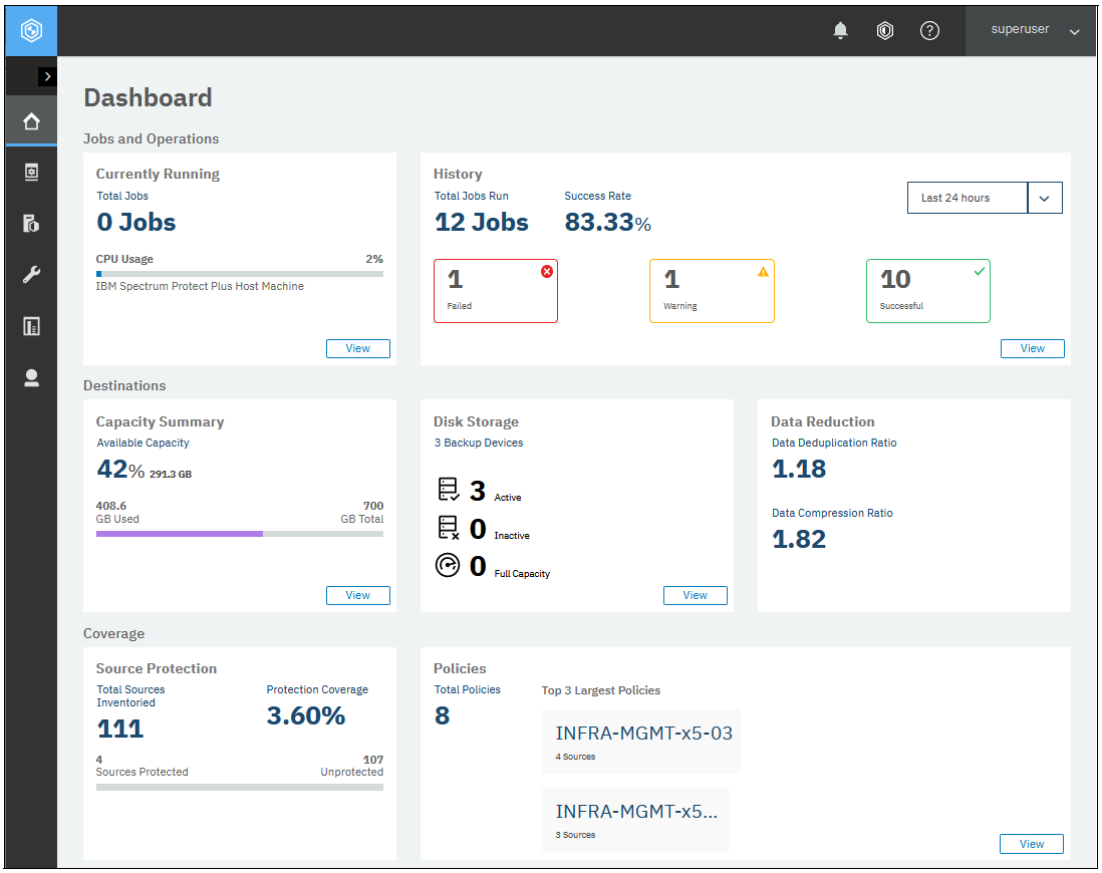

Figure 1-4 IBM Spectrum Protect Plus Dashboard

Since version 10.1.6, the IBM Spectrum Protect Plus Dashboard can be integrated into the IBM Spectrum Protect Operations Center to provide single pane for the status of an enterprise backup environment. Further on, a link to the IBM Spectrum Protect Operations Center can be integrated into the IBM Spectrum Protect Plus interface in the upper right corner. For more information about the configuration of both integrations, see 5.2.1, “IBM Spectrum Protect Plus in Spectrum Protect Operations Center” on page 117.

The left menu bar provides one-click access to the following areas:

- Jobs and Operations with details about running and previous jobs, and their corresponding logs. On the Schedule tab, an overview of all scheduled jobs is provided, and the schedule of system jobs can be changed. Furthermore, on the Active Resources tab, the Instant Access Restore and options for applications, hypervisor, and file systems are listed. Find more information, see Chapter 5, “Daily operations and maintenance” on page 107 and IBM Knowledge Center.
- Manage Protection provides access to the Backup Policy configuration, Single File Restore, and backup configuration for virtualized systems, file systems, containers, cloud-managed applications, and databases. In addition, the catalog backup and the protection of IBM Spectrum Protect Plus are configured here.

  For more information about protecting virtualized systems, see Chapter 6, “Backing up and restoring virtualized systems” on page 171 and in IBM Knowledge Center.

  For more information about protecting Windows file systems, see Chapter 7, “Backing up and restoring Windows file system data” on page 227 and in IBM Knowledge Center.

For more information about protecting Kubernetes Container, see Chapter 14, “Backing up and restoring containers” on page 381 and in IBM Knowledge Center.

For more information about Cloud Managed Applications, such as Microsoft 365, see Chapter 13, “Backing up and restoring Microsoft 365 data” on page 365 and in IBM Knowledge Center.

For more information about protecting applications and databases, see Chapter 8, “Backing up and restoring databases” on page 255, Chapter 12, “Backing up and restoring Microsoft Exchange data” on page 329, and IBM Knowledge Center.

- System Configuration includes configuration items for the Backup Storage Options Disk (vSnap Backup Storage server), Object Storage (for example, Amazon S3), and Repository Server (IBM Spectrum Protect). Further on the VADP proxy, Sites, LDAP, and SMTP can be configured. At the Global Preferences area, system preferences for your IBM Spectrum Protect Plus environment can be set. For more information, see IBM Knowledge Center.
- Reports and Logs enables you to configure and customize the built-in reports, and provides access to the Audit Log. For more information, see IBM Knowledge Center.
- Accounts is used to configure user access by using a role-based access control model. It also enables you to manage IBM Spectrum Protect Plus that is used to identify access to different resources. For more information, see 5.1, “Role-based access control overview” on page 108 and IBM Knowledge Center.

Note: For more information about system configuration, reports and logs, and accounts, see Chapter 5, “Daily operations and maintenance” on page 107.

The web-based Administrative Console GUI can be entered by using a web browser on port `8090` (for example, `https://spp-server:8090`) that uses an IBM Spectrum Protect Plus user account or the System ID `serveradmin`. In the console, you can perform the following actions:

- Get more information about the installed product versions
- Manage and install the licenses
- Manage and install certificates; for example, Active Directory LDAP certificates
- Apply and install IBM Spectrum Protect Plus software updates
- Perform System Actions, such as Start/Stop the server, restart the virtual machine, and configure the time zone
- Adjust the network configuration

Data catalogs

The IBM Spectrum Protect Plus server maintains several data catalogs, all running on the server appliance. Recovery points of the backed-up entities (for example, virtualized systems, and databases) are tracked in addition to the information for file-level recovery. System configuration is also tracked in these catalogs.

The default sizes provisioned for the catalogs should be sufficient for most workloads, but can be expanded, if necessary. It is recommended by the IBM Spectrum Protect Plus Blueprints to use the default catalog sizes. In general, place the IBM Spectrum Protect Plus appliance on high-performing, flash storage to optimize catalog performance.

Figure 1-5 shows the different catalogs and their corresponding disks and sizes.

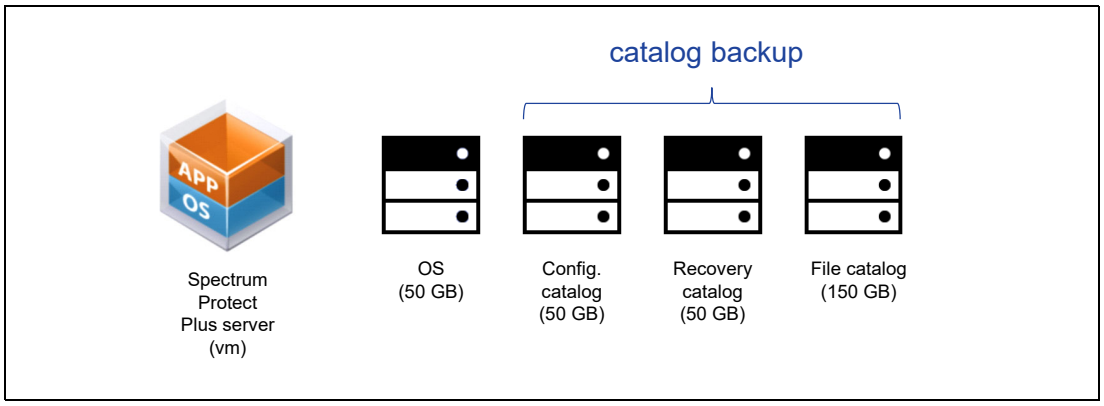

Figure 1-5 The catalogs maintained by the IBM Spectrum Protect Plus server

To protect your IBM Spectrum Protect Plus installation, the catalogs and all required information must be backed up daily. This backup is configured under **Manage Protection** → **IBM Spectrum Protect Plus** → **Backup**.

Backup Schedule and Retention are managed through an SLA and backups are stored on a vSnap Backup Storage server from where it can also be restored. Ensure that at least one copy of the catalog backup is targeted outside of your fault domain by backing up directly to another site, replicating to another site, or creating a copy to an IBM Spectrum Protect server or Cloud Storage.

Find more information about the configuration, see 3.4.5, “Adding an SLA for IBM Spectrum Protect Plus catalog backup” on page 76, and IBM Knowledge Center.

Scheduler (Jobs and Operations)

Figure 1-6 shows the overview page of the Scheduler, which can be accessed by selecting **Jobs and Operations** on the **Schedule** tab.

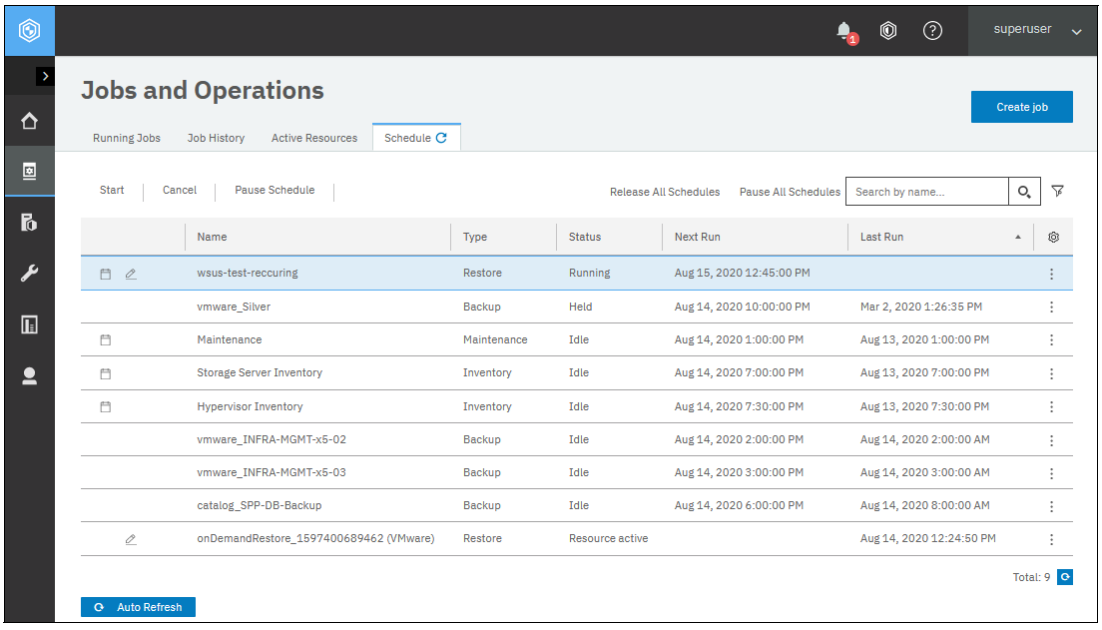

Figure 1-6 IBM Spectrum Protect Plus Schedule tab

In this view, all backup, restore, and system-defined (Maintenance and Inventory) jobs from IBM Spectrum Protect Plus are listed with the information about the last and next run time. The visible columns in the view can be configured by clicking **View Configuration Wheel** in the end of the header row.

By using the **Line Item Action** ⋮ at the end of a row, a job can be started, paused (deactivated), or canceled when it is running. The same functions for the job can be chosen above the header row for a selected job, as shown in Figure 1-7.

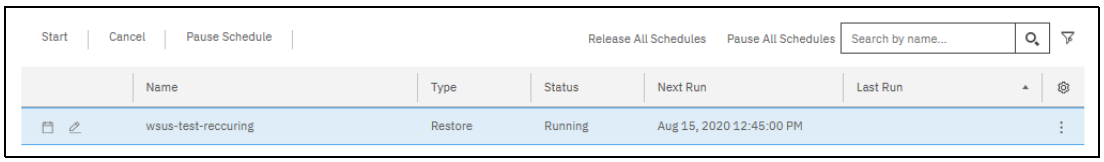

Figure 1-7 Action Options for Schedules

Close to the search bar, all schedules can be paused or released, which is helpful when maintenance activities are performed.

An adjustment of the scheduled start time for system-defined and recurring restore jobs can be performed by clicking the calendar icon in front of the job name. The schedule time for backup jobs can be adjusted by selecting **Manage Protection** → **Policy Overview** and modifying the associated SLA Policy.

An on-demand or recurring restore job can be modified by clicking the pencil icon. This allows (for example) to adjust and correct a failed on-demand restore job and start it again without redefining all settings of the restore job.

Finally, by using the icon, on-demand or recurring restore jobs can be deleted.

In addition to the backup and restore jobs, one maintenance and several inventory jobs are predefined in IBM Spectrum Protect Plus (system-defined jobs).

The maintenance job is predefined with system installation, and runs typically once a day to remove expired and deleted entries from the catalog and the system configuration. This cleanup procedure reclaims space on backup storage devices, cleans up the IBM Spectrum Protect Plus catalog, and removes related snapshots. The maintenance job also removes cataloged data that is associated with deleted jobs.

An inventory job is automatically created when you add a resource to IBM Spectrum Protect Plus. The following inventory jobs are available:

- Hypervisor Inventory is created when the first Virtualized System (VMware, Hyper-V, or Amazon EC2) is added to IBM Spectrum Protect Plus. It scans daily the associated resources on the Hypervisors, such as VMs, storage, and network configuration.
- Application Server Inventory is created when the first database, application, or file system is added to IBM Spectrum Protect Plus. It scans daily the available databases and application instances on the added systems.
- Storage Server Inventory is predefined with the system installation, and scans the configured vSnap Backup Storage server resources daily.

Identities, keys, and certificates

Some features in IBM Spectrum Protect Plus require credentials to access your resources. For example, IBM Spectrum Protect Plus connects to Oracle servers as the local operating system user that is specified during registration to complete system tasks, such as cataloging, data protection, and data restore. Furthermore, IBM Spectrum Protect Plus components, such as a vSnap Backup Storage server and a VADP proxy server, require credentials to access the resources.

It is a best practice to add user names and passwords for your resources through the **Accounts** → **Identity** pane. Then, when a feature is used in IBM Spectrum Protect Plus that requires credentials to access a resource, select **Use existing user**, and then, select an identity from the drop-down menu. A single identity can be reused for accessing multiple resources. For example, a single configured identity can be used to connect multiple Oracle instances.

The following resources require credentials:

- vSnap Backup Storage servers
- VADP proxy servers
- LDAP configuration
- Hypervisor servers
- Database instances
- Application instances
- File system instances
- Virtual machine for file metadata discovery or for specific data reuse cases where VM internal configurations need to be changed

Figure 1-8 on page 11 shows the Identity window that appears when you click (in the right corner) **Accounts** → **Identity** → **Add Identity**. You can enter a common name that best describes your entry in the Name field, and add the Username and Password in the respective fields.

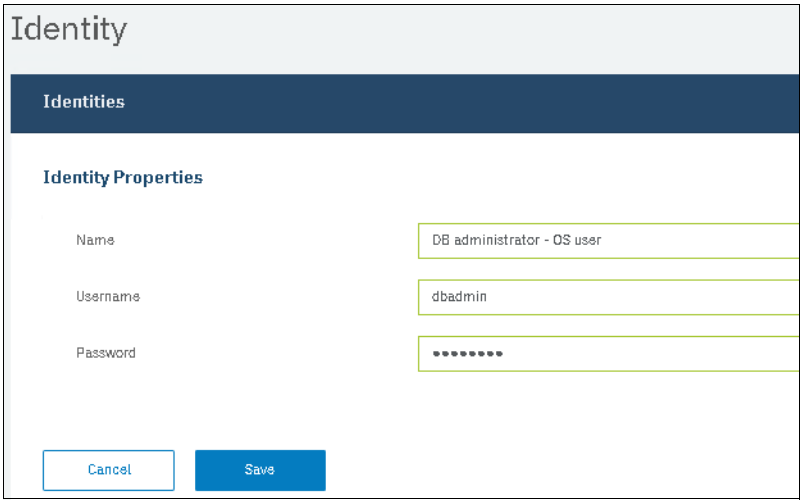

Figure 1-8 Create an identity for an operating system user

You have a list of identity entries, as shown in figure Figure 1-9. Some are default system identities (vsnapadmin, LocalvSnapadmin, and serveradmin), and some are users that were created for database backup and restore (DB administrator, operating system user, Mongo DB user, and Oracle DBA).

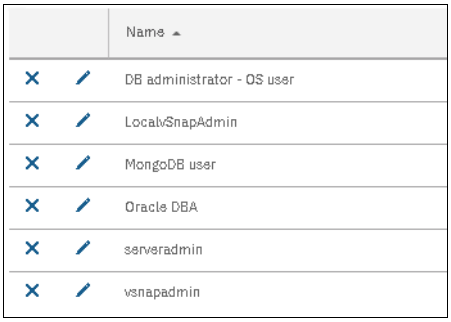

Figure 1-9 Identity entries

Cloud resources require Access Keys and Certificates to be connected with IBM Spectrum Protect Plus. In addition to an identity, a Linux instance can also use an SSH key. Access keys, certificates, and SSH keys can be added and managed by selecting **System Configuration** → **Keys and Certificates**.

1.2.2 Site

A *site* is an IBM Spectrum Protect policy construct that is used to manage data placement in the environment. It can be *physical*, such as a data center, or *logical*, such as a department or organization. IBM Spectrum Protect Plus components are assigned to sites to localize and optimize data paths. A deployment always has at least one site per physical location.

By default, the IBM Spectrum Protect Plus environment has a primary site, a secondary site, and a demo site. You can change the site name and other options for the default Primary and Secondary site, and a demo site. The demo site is available only for the on-board vSnap server. You cannot use this site with any other vSnap server.

Note: The demo site is predefined with the deployment of IBM Spectrum Protect Plus, and is intended to be used for product demonstration purposes only. It should *not* be used for a production backup and can therefore be deleted in such installations.

Sites in IBM Spectrum Protect Plus are configured and maintained by selecting **System Configuration** → **Site** as shown in Figure 1-10 on page 13. Here, you can adjust the names or add sites. A Throttle Rate can be configured to change and limit the throughput for site replication and copy operations so that you can manage your network activity on a defined schedule.

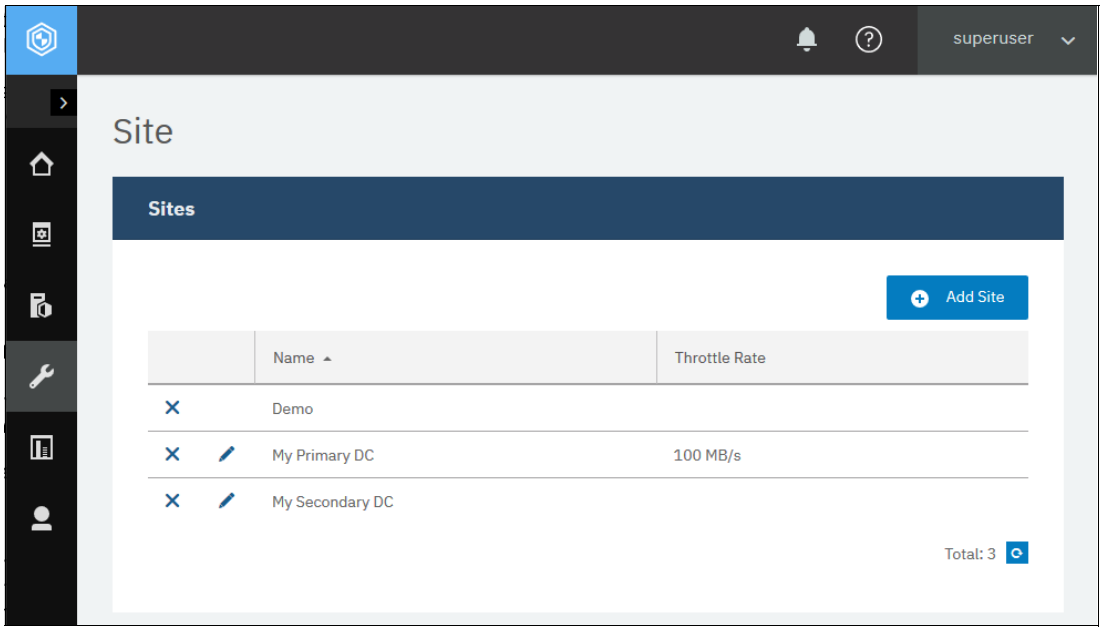

Figure 1-10 IBM Spectrum Protect Plus Site Overview

The general philosophy is to localize data movement to the sites by placing vSnap Backup Storage servers and VADP proxies (VMware only) together in the sites. It is advised to always have at least one site per physical location, and at least one vSnap server per site, including at least one VADP proxy per site if you are protecting VMware vSphere in this site.

The placement of backup data to a site is governed by the SLA policies. Therefore, you specify in which site your backup data is stored or replicated to, instead of specifying a dedicated vSnap Backup Storage server.

Cases where you want to define multiple sites in a physical location and further considerations together with replication are discussed in the IBM Spectrum Protect Plus Blueprints in *Sites and vSnap server distribution* and *Replication considerations*, which are available at this web page.

1.2.3 vSnap Backup Storage server

The vSnap Backup Storage server, or vSnap server, consists of a pool of disk storage and is the primary backup destination for IBM Spectrum Protect Plus. It receives data from production systems for the purposes of data protection or reuse, and is responsible for storing, managing, and maintaining this backup data.

A vSnap Backup Storage server provides compression, deduplication, and encryption functions. It can be deployed as a virtual machine, or installed on a physical server. Physical installations can use SAN networks for data transfers.

Note: A built-in vSnap Backup Storage server is predefined with the deployment of IBM Spectrum Protect Plus, and is intended to be used only for product demonstration purposes. It should not be used for a productive backup, and can therefore be deleted in such installations, as shown in Figure 1-3 on page 6.

IBM Spectrum Protect Plus requires at least one vSnap Backup Storage server that can be scaled up by adding disks to increase capacity. Multiple vSnap Backup Storage servers can be deployed in the following ways:

- Per site to scale out for an overall performance increase
- In different sites to achieve local data placement

The backup data that is stored by a vSnap server can be replicated to a second vSnap server as an operational protection. For more protection, backup data can be copied to an IBM Spectrum Protect server or to Cloud Storage. For all data copy operations to cloud Object Storage or IBM Spectrum Protect (backup, recovery, and data reuse), each vSnap server requires a disk cache area (referred to as the cloud cache, as shown in see Figure 1-11 on page 14) to perform the following functions:

- As a temporary staging area for objects that are pending upload to the cloud Object Storage endpoint during data copy operations.
- To cache downloaded objects and store any temporary data that might be written into the restore volume during restore operations.

Instant access, restore, and data reuse capabilities are provided by these different backup storage providers. For more information, see Chapter 15, “Replication and additional copies” on page 405.

Note: A vSnap Backup Storage server is similar to an IBM Spectrum Protect storage pool in terms of data placement and handling.

Figure 1-11 shows the storage layout of a virtual or physical vSnap Backup Storage server. The vSnap metadata is on the operating system disk (`/etc/vsnap`) and used to track information about the system, replication, and copy configuration, including the encryption key. If the VADP proxy is installed, it is also placed on these disks.

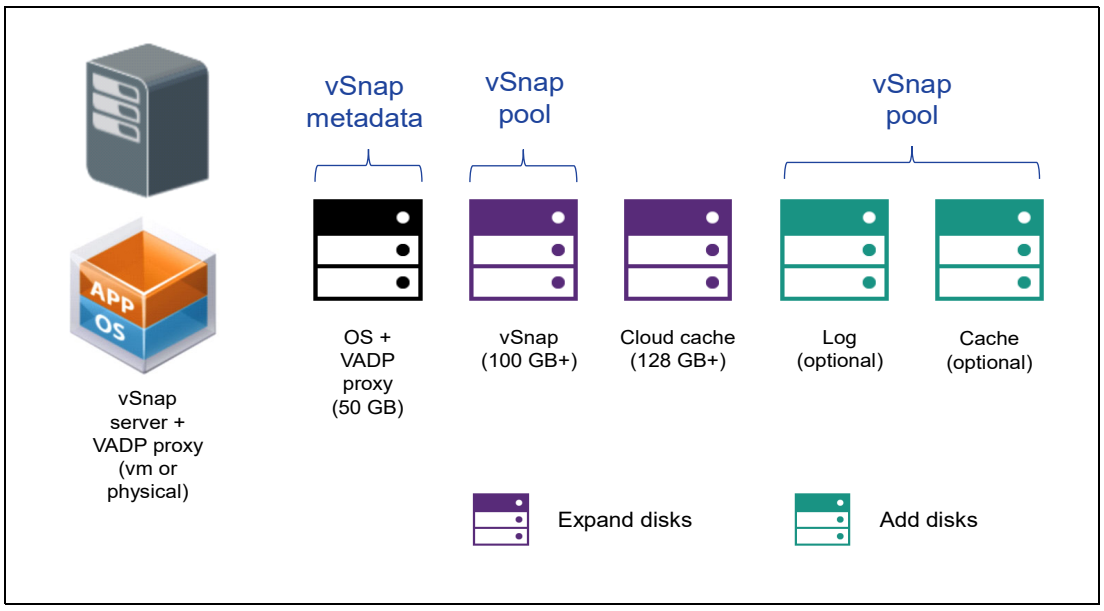

Figure 1-11 Storage layout overview of IBM Spectrum Protect Plus vSnap server

The sizing of the vSnap Backup Storage servers is crucial. The IBM Spectrum Protect Plus Blueprints that provide the required guidance are available at this web page.

vSnap pool

The disks that are used by a vSnap Backup Storage server are organized within a vSnap pool, as shown in Figure 1-11 on page 14. By adding disks, the vSnap pool capacity is extended. The pool provides optional log and cache functions that are managed on separate disks.

The vSnap server log is used to optimize write performance for application log backups and for data reuse operations. The vSnap server cache is used with a memory cache to optimize backups when deduplication and performance is used for data reuse scenarios.

Redundant Array of Independent Disks (RAID) technology can be used to protect vSnap pool against data loss that is caused by a hardware failure of a disk. This protection is by using software RAID inside the vSnap server or by using hardware or software features of the underlying storage system that provides the capacity to the vSnap pool.

Protect vSnap server configuration (metadata)

The vSnap Backup Storage server configuration and metadata is on the operating system disk (`/etc/vsnap`), as shown in Figure 1-11 on page 14. This vSnap server information about system, replication, and copy configuration (including the encryption key) can be backed up by using the CLI interface. This backup can be used to recover the vSnap server where the vSnap pool data is intact and valid but the configuration or metadata information is lost or not available (for example, operating system disk failure).

At a minimum, a backup should be created once a day after all replication and copy tasks are completed. The backup file is created locally and should be securely copied to a central location, such as the IBM Spectrum Protect Plus server.

For more information about a configuration example, see 3.4.6, “Backing up the vSnap server system configuration” on page 77, and *Appendix D: Protecting vSnap System Configuration* chapter of the IBM Spectrum Protect Plus Blueprints, which are available at this web page.

User interfaces

Two user interfaces for the vSnap Backup Storage servers are provided. Graphical access to the vSnap server configuration is provided as part of the web-based IBM Spectrum Protect Plus standard GUI. Under **System Configuration** → **Backup Storage** → **Disk**, all available vSnap Backup Storage servers with their status and capacity are listed, as shown in Figure 1-12.

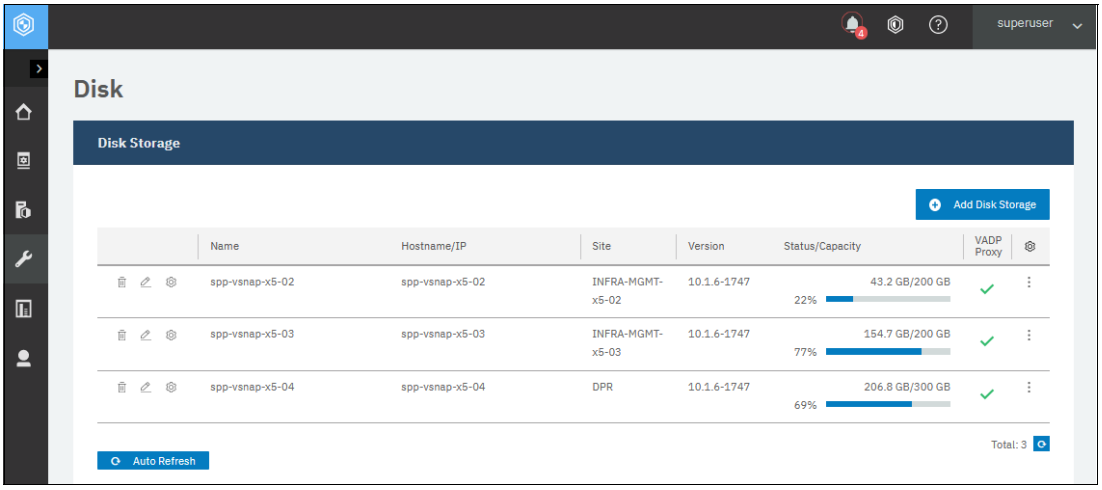

Figure 1-12 vSnap Backup Storage Overview in the IBM Spectrum Plus GUI

The configuration of the vSnap server can be viewed and changed by clicking the **Edit** icon  , which opens the vSnap server configuration page.

The second user interface is a command-line-interface (CLI), which can be reached through SSH, by using the default serveradmin UserID, as shown in Example 1-1.

Example 1-1 vSnap Backup Storage CLI

```
Using username "serveradmin".
serveradmin@9.155.114.248's password:
Last login: Wed Aug 12 10:43:20 2020 from 9.211.79.204
------------------------------------------------------------------
Be sure to adhere to vSnap server hardware and memory requirements
 as described in IBM Spectrum Protect Plus Blueprints
 accessible from the IBM Spectrum Protect Plus Knowledge Center.

------------------------------------------------------------------
[serveradmin@spp-vsnap-x5-02 ~]$ vsnap
Usage: vsnap [OPTIONS] COMMAND [ARGS]...

Options:
  --json     Show output in JSON format.
  --summary  Show output in summary (tabular) format.
  --detail   Show output in detail (multiline) format.
  --help     Show this message and exit.

Commands:
  archive       Manage archive resources.
  cloud         Manage cloud resources.
  disk          Manage disks.
  host          Manage volume host mappings.
  maint         Manage maintenance sessions.
  network       Manage network interfaces.
  partner       Manage partner servers.
  pool          Manage storage pools.
  relationship  Manage replication relationships.
  repair        Manage vSnap repairs.
  session       Manage replication sessions.
  share         Manage volume shares.
  snapshot      Manage volume snapshots.
  system        Manage vSnap system.
  target        Manage storage targets.
  task          Functions to create and monitor tasks.
  throttle      Manage throttling events.
  user          Manage vSnap users.
  volume        Manage storage volumes.
[serveradmin@spp-vsnap-x5-02 ~]$
```

1.2.4 VADP proxy server

The VADP proxy server is the component that is responsible for moving data from the VMware vSphere data stores to the vSnap Backup Storage server to protect VMware virtual machines, as shown in Figure 1-13 on page 17. As the name implies, it is required only for backing up VMware virtual machines. A VADP proxy can be installed on a virtual or physical machine.

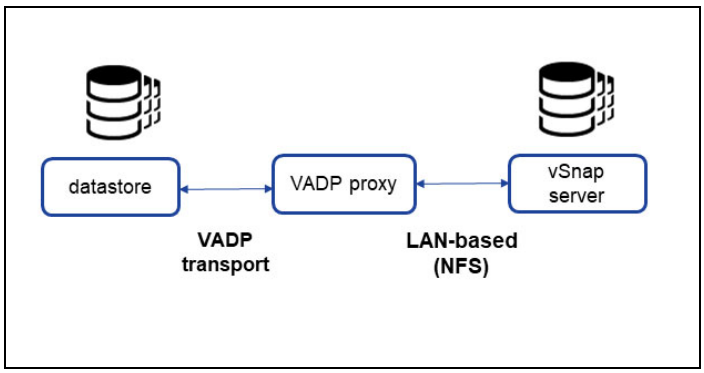

Figure 1-13 VADP proxy connectivity

You can install the vSnap server (backup storage provider) and VADP proxy on the same physical or virtual system. If so, IBM Spectrum Protect Plus optimizes data movement by eliminating an NFS mount over LAN when these two systems are colocated. The vSnap Backup Storage overview that is shown in Figure 1-12 on page 15 shows if a VADP proxy is enabled on a vSnap server.

At least one VADP proxy component is always required for each site, in case VMware virtual machines need to be backed up in this site. Based on sizing needs, more proxies can be required. VADP proxies support the following VMware transport modes: File, SAN, HotAdd, NBDSSL, and NBD. For more information about VMware transport modes, see the *vSphere Documentation Center* at this web page.

The overview and configuration of the VADP proxies can be accessed by selecting **System Configuration** → **VADP Proxy**, as shown in Figure 1-14.

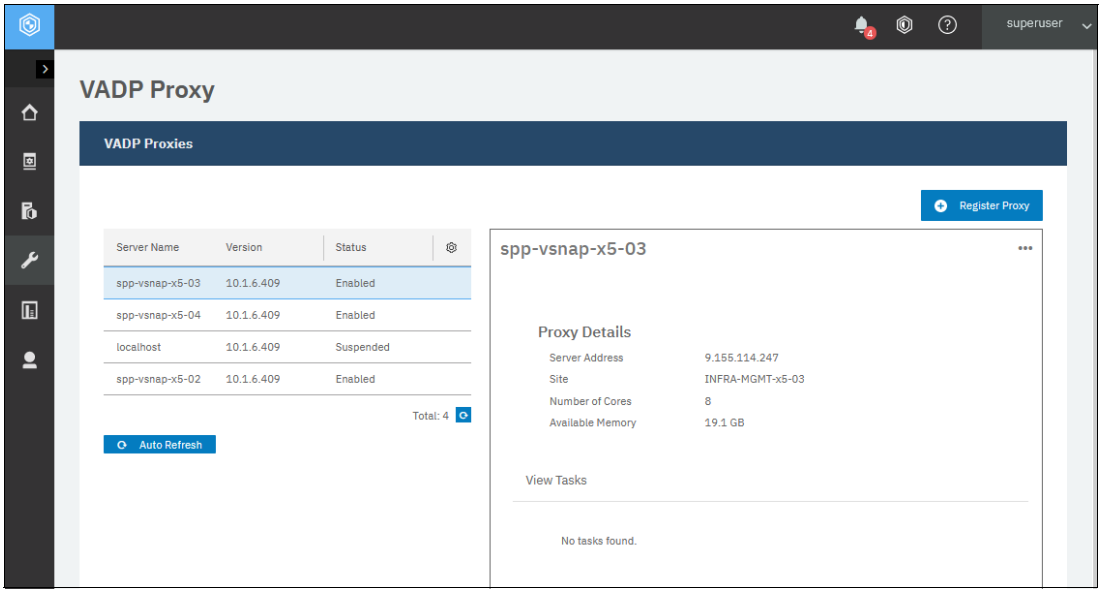

Figure 1-14 VADP proxy overview in the GUI

Note: A built-in VADP proxy is predefined with the deployment of IBM Spectrum Protect Plus, and is intended to be used only for product demonstration purposes. It is advised to suspend the `localhost` VADP proxy, and not use it for production backups.

The detailed sizing and configuration of VADP proxies is discussed under Configuring VADP Proxies in the IBM Spectrum Protect Blueprints (see this web page).

For more information about installing and deploying a VADP proxy, see Chapter 3, “Installation and deployment” on page 47.

1.2.5 Data flows

IBM Spectrum Protect Plus offers incremental forever backups that can be further compressed, deduplicated, and encrypted. Figure 1-15 shows an example of data flow for an incremental VMware backup, which is stored compressed, deduplicated, and encrypted on the vSnap repository. When the backup data is replicated, it is decrypted, rehydrated, and sent compressed through an SSH tunnel to the vSnap server replica where it is again deduplicated and encrypted. If the backup data is copied further on to an Object Storage, it is decrypted and rehydrated and sent compressed through a TLS tunnel to the Object Storage, where it is stored according to Object Storage native capabilities.

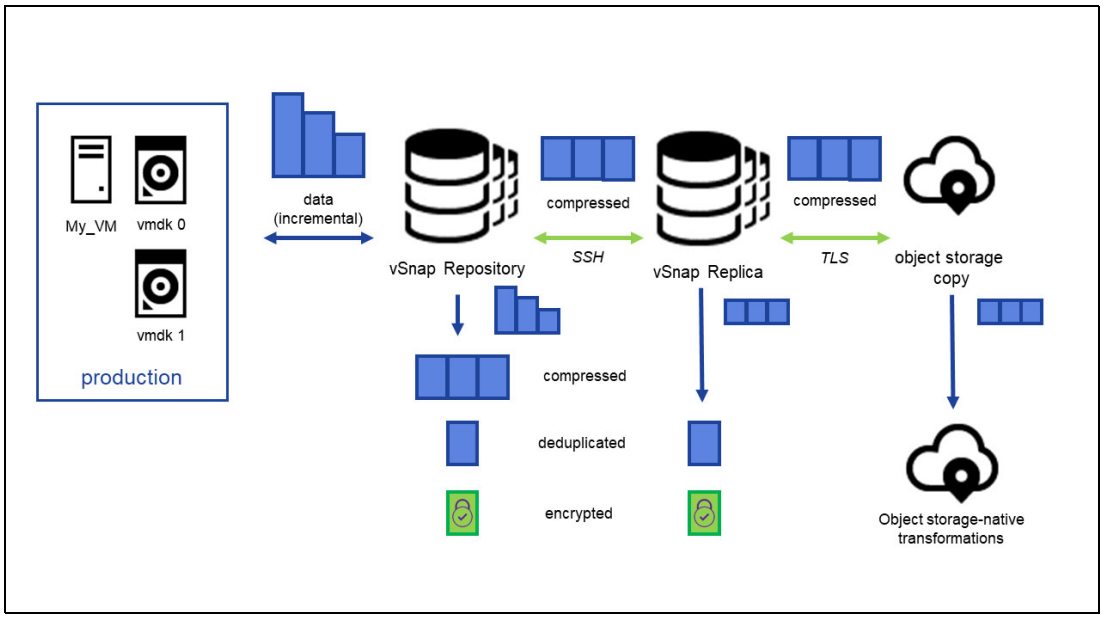

Figure 1-15 Data flows between the IBM Spectrum Protect Plus components

1.3 SLA backup policies

Backup policies, which are also referred to as *service level agreement* (SLA) policies, define parameters that are applied to backup jobs. The overview and configuration of SLA policies can be accessed in the GUI by selecting **Manage Protection** → **Policy Overview**, which displays a Protection Summary in the first half of the window, as shown in Figure 1-16.

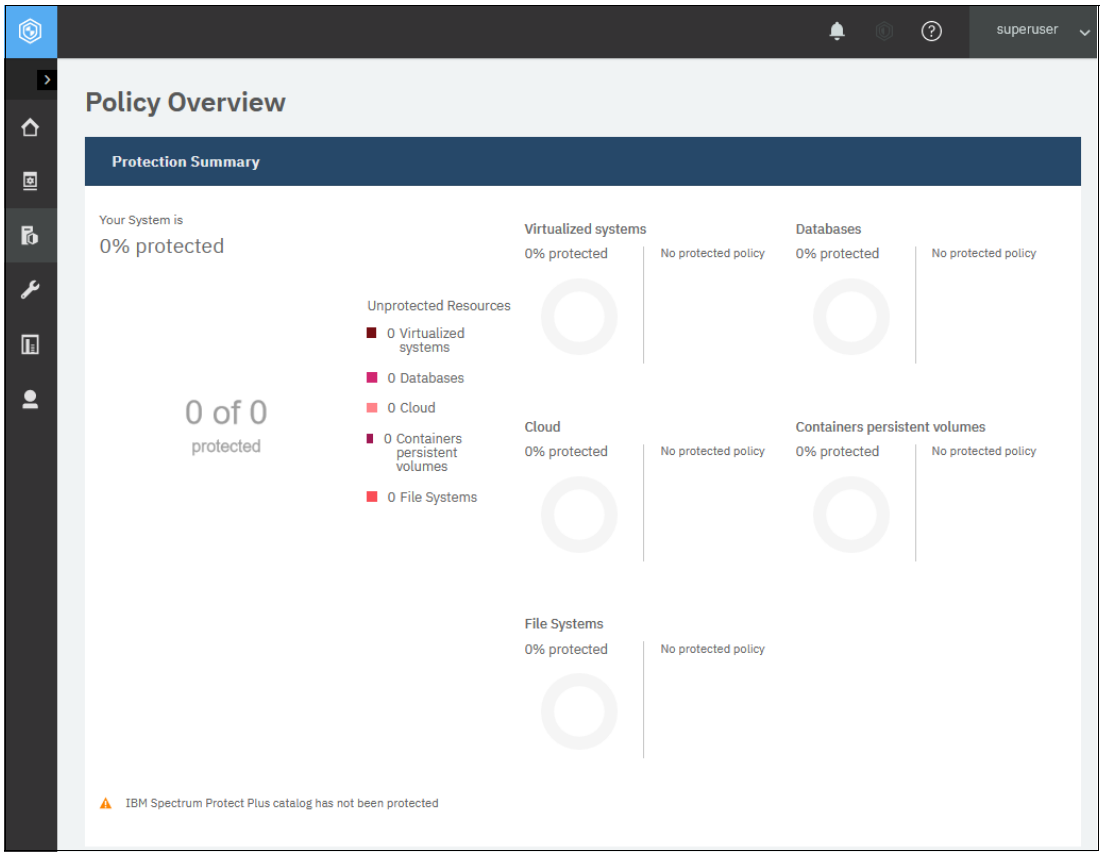

Figure 1-16 SLA Policy Protection Overview

Note: At the bottom of Figure 1-16, information is displayed about missing catalog protection of IBM Spectrum Protect Plus. For more information about the IBM Spectrum Protect Plus catalog, see “Data catalogs” on page 8.

For more information about the configuration of the catalog backup, see 3.4.5, “Adding an SLA for IBM Spectrum Protect Plus catalog backup” on page 76.

In the bottom half of the window, the configured and available SLA policies are listed, as shown in Figure 1-17. The following pre-configured SLA policies are available:

- Gold, Silver, and Bronze policies can be used for VMware, Hyper-V, Exchange, Microsoft 365, SQL, Oracle, IBM DB2®, MongoDB, and Windows file systems.
- EC2 policy can be used for Amazon EC2 backups.
- Container policy is used for Kubernetes Container support.

Note: The demo SLA is predefined with the deployment of IBM Spectrum Protect Plus and is intended to be used only for product demonstration purposes. It should not be used for a productive backup; therefore, it can be deleted in such installations.

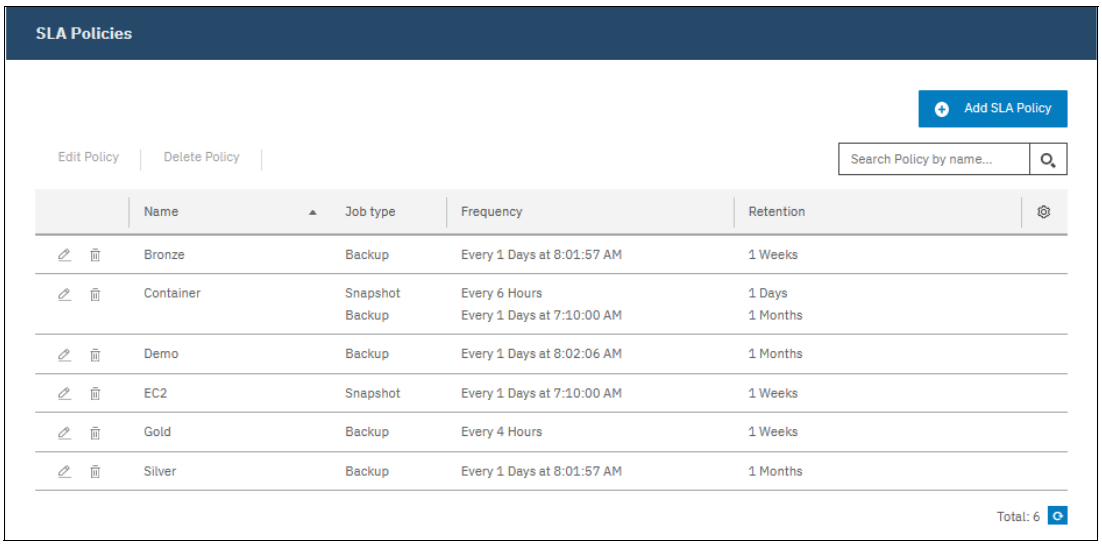

Figure 1-17 SLA Policies Overview with six default policies

You can use these policies as they are, modify the policies, or create custom SLA policies by using the **Add SLA Policy** option in the bottom in the right corner of the GUI.

Important: An SLA cannot be renamed. Consideration this issue when you are specifying the name for an SLA.

For more information about creating and configuring an SLA, see IBM Knowledge Center.

1.3.1 Backup parameters and retention management

SLA customization is done in a single configuration page and defines the following backup parameters, which are applied to the *Backup Clients* that are backed up with this SLA:

- How often and when to back up data (Backup Schedule)
- How long to retain back up data (Backup Retention)
- Where to store back up data (Backup Target)
- How to protect back up data (Replication or extra Copy Target)

Note: Consider that backup retention changes in an SLA take effect for future and existing backups that are associated with that SLA.

The SLA template is divided into three configuration sections, as shown in Figure 1-18:

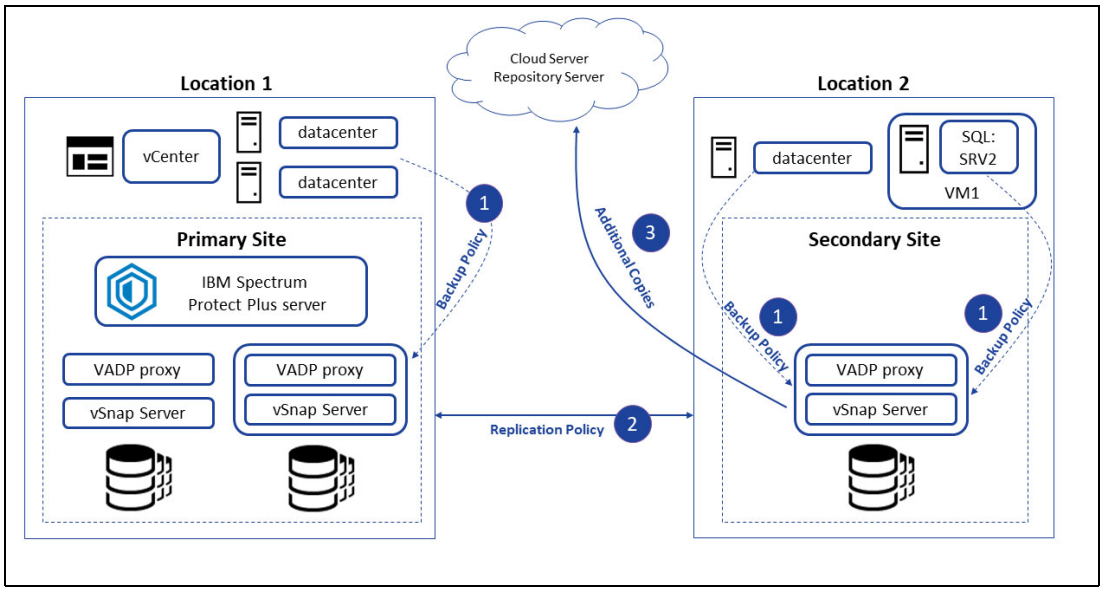

Figure 1-18 The three configuration sections of an SLA policy

1. *Backup policies* define the main backup schedule and retention for the incremental backups that are taken with this SLA. For example, you can configure that a backup is taken once a day at PM and retained for one week (7 days). LAter, it is configured in which Site these backups are stored and if an encrypted vSnap Backup Storage must be used. Figure 1-19 shows these configuration details for the backup policy part.

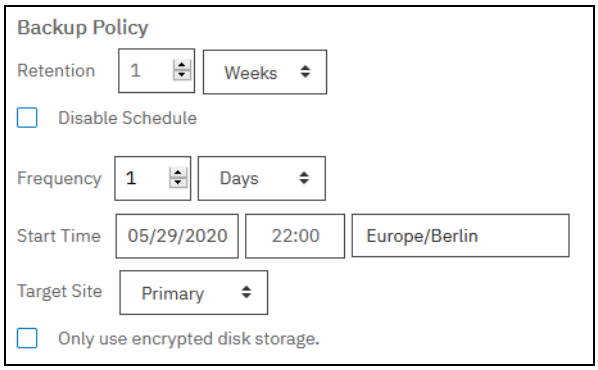

Figure 1-19 SLA template - Backup Policy configuration details

2. *Replication policies* define a dedicated backup replication schedule to protect backup data to another Site. By default the Backup Storage Replication is not enabled and must be enabled to use it. Later, it cab be specified whether an encrypted vSnap Backup Storage must be used in the target site. Finally, you can decide whether the same backup retention is used for the replication copy of the backup data or if another retention should be used (*dissimilar policies*). Figure 1-20 on page 22 shows these configuration details for the replication policy part.

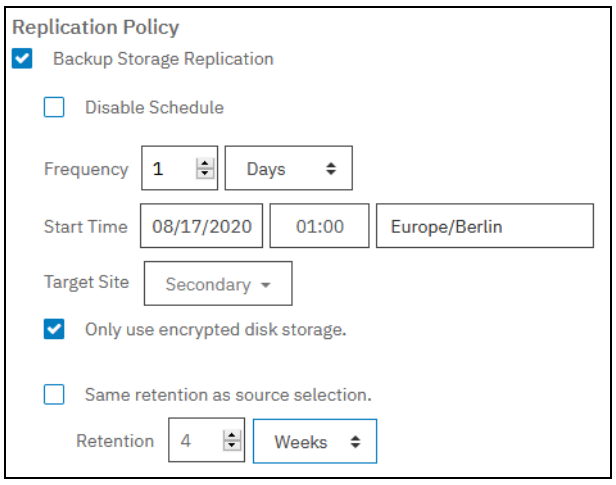

Figure 1-20 SLA template - Replication Policy configuration details

3. *Additional copies* are divided into the following parts:
   - *Standard Object Storage* defines the incremental copy of backup data to a cloud storage provider (Cloud Server) or to IBM Spectrum Protect (Repository Server). At the IBM Spectrum Protect server, a container storage pool must be used in an object domain. A dedicated schedule to perform the copy operation is defined. For the retention, you can decide whether the same retention as for the source is used or if another retention should be applied (*dissimilar policies*).
   - *Archive Object Storage* defines the full copy of backup data to a cloud storage provider (Cloud Server) or to IBM Spectrum Protect (Repository Server). At the IBM Spectrum Protect server, a tape storage pool must be used in an object domain. A dedicated schedule and retention is defined.

     Note: The Frequency and Retention for the archive Object Storage can be set to Weeks, Months, and Years only. Minutes and Hours is not available for this part.

In Figure 1-21 on page 23, you can see the configuration details for the Additional Copies section. Both options, the Standard and the Archive Object Storage, allow to select if data is copied from the primary backup target or from the Replication site.

Note: For more information about replicating backup data to another vSnap Backup Storage server and copying data to a Cloud Storage Provider or to IBM Spectrum Protect, see Chapter 15, “Replication and additional copies” on page 405.

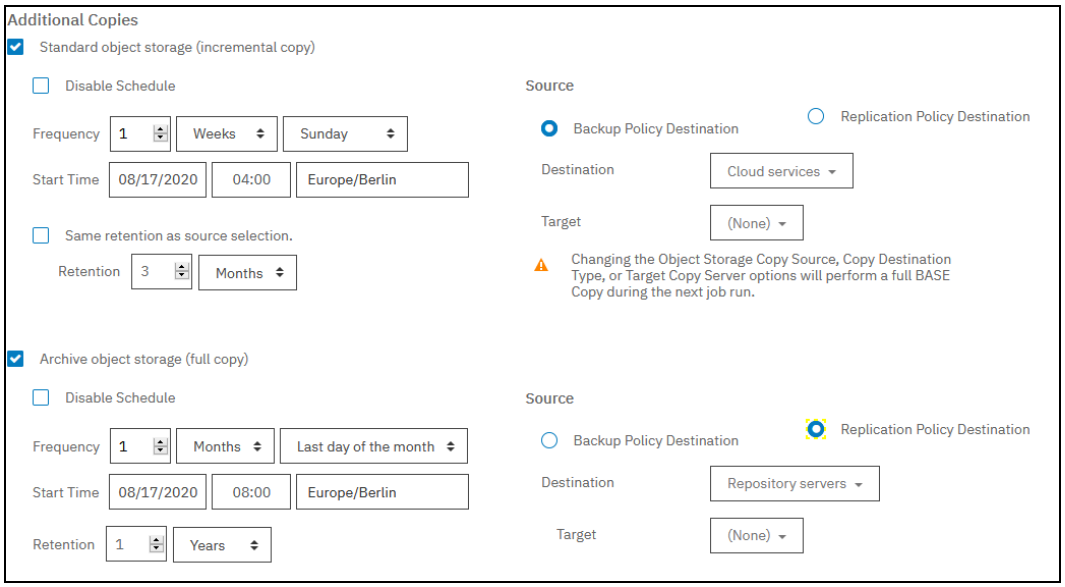

Figure 1-21 SLA template: Additional Copies configuration details

1.3.2 Backup jobs: Associate Backup Clients to an SLA

To back up virtual machines, databases, applications, and so on, associate them with a suitable SLA, as shown in 6.1.3, “Assigning an SLA policy” on page 174.

This association creates a corresponding backup job. For example, if a VMware virtual machine is associated with the Bronze SLA, the `vmware_Bronze` backup job is created. If a Hyper-V virtual machine is associated with the Bronze SLA, the `hyperv_Bronze` backup job is created. Both backup jobs for VMware and Hyper-V run at the same time and use the same configured backup parameters.

If a virtual machine is associated with multiple SLA policies, ensure that the policies are not scheduled to run concurrently. Schedule the SLA policies to run with a significant amount of time between them, or combine them into a single SLA policy.

To delete an SLA policy, ensure that no Backup Clients are associated with it. The associated backup instances remain available for recovery until their expiration, as defined previously by the SLA policy.

1.3.3 SLA to Site to vSnap server relationship

An SLA is configured to place backup data into a defined site, or to replicate backup data from the source site to the target site. Therefore, you cannot specify exactly on which vSnap Backup Storage server the backup data is stored. This relationship between an SLA and a vSnap server is defined when the Backup Policy or the Replication Policy in an SLA is run the first time and it cannot be changed later.

For example, assume that we have a configuration as shown in Figure 1-12 on page 15, in which three vSnap Backup Storage servers are used, each associated to a different site. The build relationship is simple because only one vSnap server per site is available and therefore selected if the SLA points to this site. Usually in larger installations, multiple vSnap servers are available per site.

For more information about guidelines and configuration aspects of distributing backup data on multiple vSnap servers, see 6.1.6, “Distributing VM backups to multiple vSnap servers” on page 180.

For more information about spreading backup data over multiple sites, see 15.2.4, “Dual-site backup using multiple SLAs” on page 408.

2

Solution architecture, planning, and design

This chapter is meant to give you a broad understanding about how to design, plan, and size an IBM Spectrum Protect Plus solution. It helps you to visualize the software interactions with the virtualization layer, operating systems, disk storage, network components, tape storage, and other products. It also discusses the impact of such deployment within IT teams and the overall IT environment, and how data protection operations can be distributed across the various teams.

This chapter includes the following topics:

2.1 Solution design introduction

In Chapter 1, “IBM Spectrum Protect Plus product architecture and components” on page 1, we discussed the different components in an IBM Spectrum Plus environment. In this chapter, we describe building a solution. Building a data protection solution includes different phases, as shown in Figure 2-1.

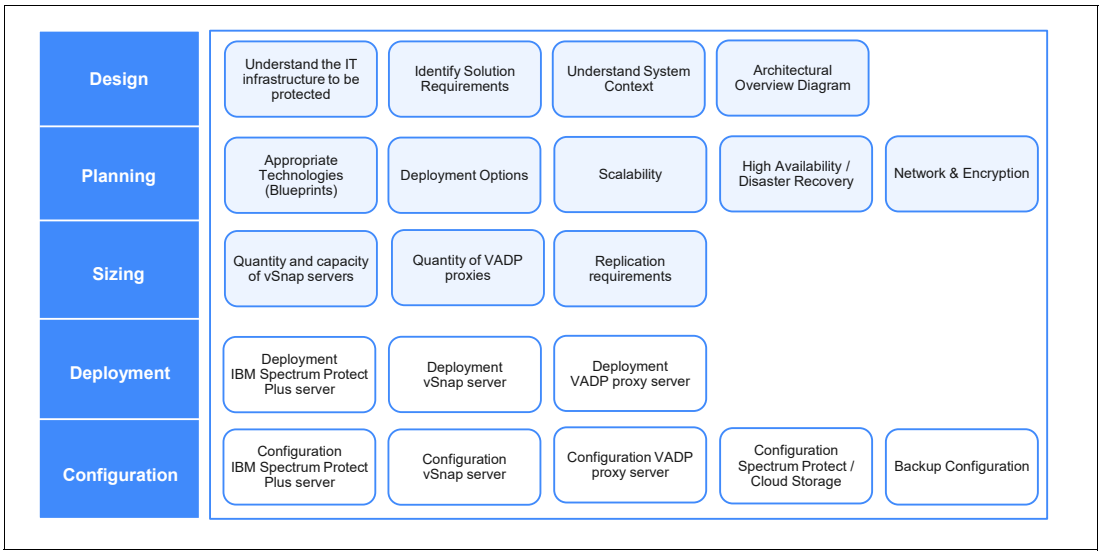

Figure 2-1 The different phases of building a data protection solution

The following phases are discussed in this chapter:

- Design
- Planning
- Sizing

The following phases are described Chapter 3, “Installation and deployment” on page 47 and cover:

- Deployment
- Configuration

2.2 Design

A solution design for a new modern data protection environment consists of the following multi-step approach:

- Understand the IT infrastructure components that must be protected by the solution
- Identify the functional and non-functional solution requirements
- Understand the overall system context
- Create a solution architecture, including an overview diagram, which is further developed in the planning and sizing phase

2.2.1 Understand the IT infrastructure to be protected

To create a solution design for a data protection environment, it is essential to understand the IT infrastructure that must be protected by this new environment.

Understanding the infrastructure is important to know to identify which components of IBM Spectrum Protect Plus are required. For example, to back up VMware virtual machines, we need the VADP proxy component of IBM Spectrum Protect Plus; whereas for Hyper-V virtual machines, this information is not required. However, this information is required when sizing the environment.

Backup clients

The systems that must be protected are typically called *backup clients*. Because we need to know the details of these clients for the solution design, we must ask the following questions:

- Are they running on physical servers?
- Are they running as virtualized systems and which Hypervisor are they using?
- Which operating system are they using?
- Which applications or databases are installed on these systems?
- Are they running on a public cloud and which cloud provider are they using?
- Are they running on a container, which container platform are they using, and what type of storage is being used to persist data?
- Which front-end capacity are these clients using? Front-end capacity is the size of the primary data that is being protected for client applications, virtual machines, and systems.

  Front-end capacity details are required by *Backup Client* category. For example:

  – How much front-end capacity for VMware?
  – How much front-end capacity for Microsoft SQL?

  The information about the front-end capacity is needed for sizing of an IBM Spectrum Protect Plus solution, as discussed in 2.4, “Sizing” on page 45.

For more information about whether all gather backup clients are supported by IBM Spectrum Protect Plus, see IBM Support’s IBM Spectrum Protect Plus - All Requirements Doc.

Infrastructure and environment

After we know the details about the backup clients, we need to gather information about in which locations these clients are running and the details of theses locations:

- Are there multiple data centers or computer rooms per location?
- What is the network connection between these locations, data centers, computer rooms?

The next piece of information we need to gather are parameters about the backup data; for example, how long data must be retained and how often a backup should run. This information is described as non-functional requirements of a data protection environment, which we discuss next.

2.2.2 Solution requirements

To design a data protection solution, you need the functional and non-functional requirements.

The functional requirements describe what the system must do; that is, which functions the solution must provide relative to the business needs.

Typical functional requirements include the following examples:

- VMware and Hyper-V back up and restore

  The system must:

  - Perform full or incremental VM backups of VMware and Hyper-V virtual machines.
  - Restore a virtual machine back into production to replace a lost or damaged production VM.
  - Perform test restores of individual virtual machines into an isolated environment that is separated from production.

- VM and database cloning

  The system must clone individual databases or virtual machines for data reuse for test and development environments with different attributes, such as the name and network location.

- Restore granularity

  The system must provide single file restore for backed up file systems and item level recovery for databases and applications.

- Replication of backup data

  The solution must:

  - Replicate backups offsite at least once a day and restore at the destination once data is transferred.
  - Store a third backup copy in a private or public cloud.

The non-functional requirements for a data protection system often are the capacity, performance, and service levels that must be achieved.

The non-functional requirements address those aspects of the system that, although not directly affecting the functionality of the system as seen by the users, can have a profound effect on cost. They also affect how the data protection system is accepted by the users and the system administrators, while not directly impacting the functionality of the system as seen by the users.

The following non-functional requirements in a data protection solution are most important:

- Recovery Time Objectives (RTO)

  Defines how long it might take until a system is restored and operational.

- Recovery Point Objectives (RPO)

  Defines the maximum time (usually in hours) for which a data loss is accepted. That is, it defines the maximum time frame between two backups (recovery points). For example, an RPO of 24 hours requires that a backup must be taken once per day and the maximum accepted data loss is the last 24 hours. In the case of an RPO of 4 hours, every 4 hours a backup must be taken and the maximum accepted data loss is the last 4 hours. The pre-configured Gold SLA is configured for a backup every 4 hours.

- Backup Retention Policy

  Defines how long a created backup version is kept. That is, it describes the age of the oldest previous state. For example, a retention of 1 week means that you can go back 7 days if a restore is required. Therefore, you have 7 days to notice that, for example, a file was deleted and you must restore it. Having a one week retention and for example, an RPO of 24 hours (daily backup), results in 7 backup versions. In the case of an RPO of 4 hours, 42 backup versions are created.

- Backup window

  Defines the time window in which backups can be performed. For example, a backup window from 10 PM - 6 AM defines a maximum backup duration of 8 hours. During this 8 hours, all backups must be completed.

- Daily Change Rate

  Defines how much percent of the data is changed daily.

- Annual Growth

  Defines how much new data is stored in percent per year.

These most important requirements define the backup data parameters, such as backup frequency, backup retention, and the recovery constraints that lead to a consistent implementation of a data protection solution. These parameters are key inputs for sizing an IBM Spectrum Protect Plus solution, as described in 2.4, “Sizing” on page 45.

Other typical non-functional requirements include the following examples:

- Availability

  The system must be available 99.9% of the time, excluding planned downtime for maintenance.

- Durability

  The system must avoid data loss because of a single component failure, such as a power supply or single hard disk drive.

- Disaster recovery

  The system must recover from an entire site loss for customers with a requirement for offsite recovery.

- Scalability

  The solution must scale in place as the customer grows without having to dismantle and replace components with larger hardware. The service can accomplish this goal by scaling up or scaling out, or a combination of the two.

- Performance:
  - The system must have well-defined performance characteristics that provide consistent speeds that grow with capacity as part of the scalability model.
  - Alerting can occur; for example, by way of RESTful API, email, or syslog.

- Capacity management

  The system must provide a mechanism for assessing capacity usage, and generating alarms as the system approaches full capacity.

- Alerting

  The system must provide a mechanism for sending alerts about events, such as performance issues, and errors to typical external alerting tools by way of standard protocols (for example, syslog or SNMP).

- Provide awareness and security

  The system must provide encryption at rest functionality and encrypt data that is transmitted over the WAN.

- Billing

  Provide usage-based billing based on the amount of data that is stored in the system; for example, by using API calls to the IBM Spectrum Protect Plus server to generate billing volume data.

- Network separation

  Provide network separation as an extra security layer so that back ups are not using the management network or adversely affecting business applications.

- Network bandwidth

  Provide the required bandwidth between the hypervisors, application servers, VADP proxies, and the vSnap servers.

- Firewall

  Specific communication paths through a firewall must be open for the involved components that are placed in secured network environments.

2.2.3 System context

In engineering, a *system context diagram* is a diagram that defines the boundary between the system, or part of a system, and its environment. It also shows the entities that interact with the system.

The diagram that is shown in Figure 2-2 on page 31 shows the system context of a typical IBM Spectrum Protect Plus solution. The diagram records the interactions with the external systems in the IT infrastructure, and with users that might be distributed among several teams.

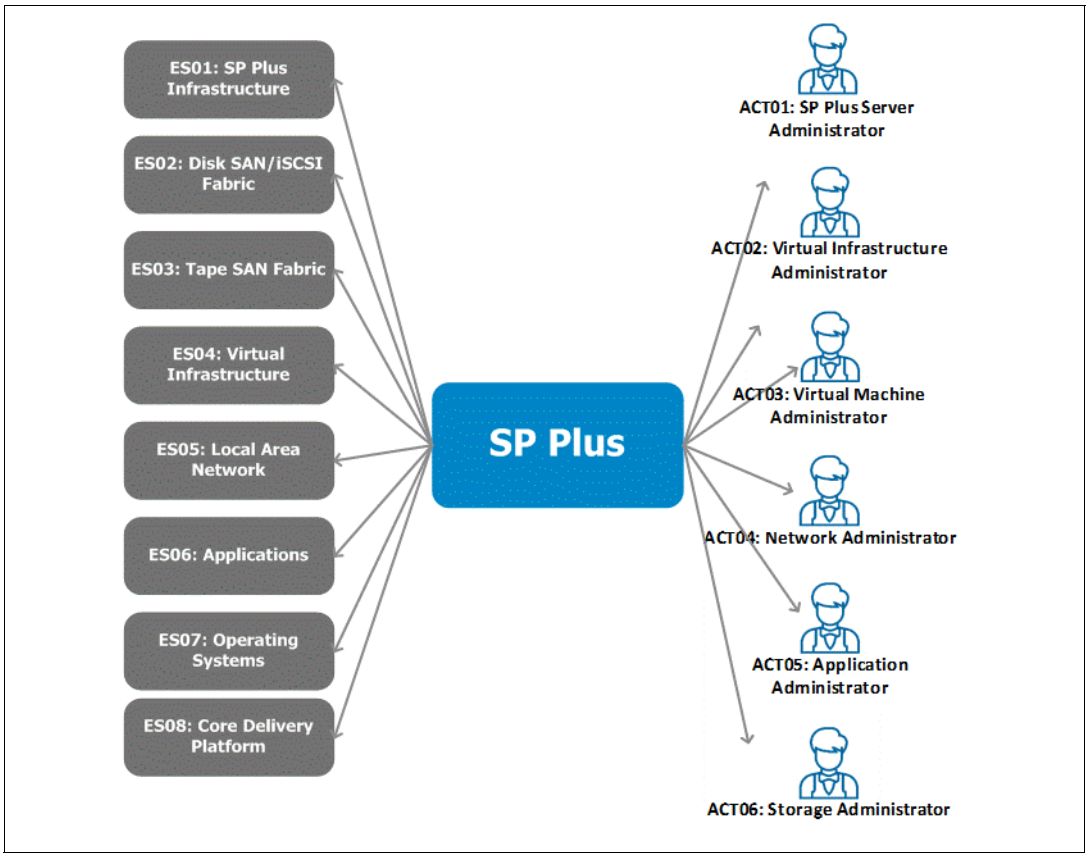

Figure 2-2 IBM Spectrum Protect Plus: System context

Table 2-1 lists how each user interacts with the IBM Spectrum Protect Plus solution in performing their different tasks.

Table 2-1 Users that are involved in an IBM Spectrum Protect Plus solution deployment

Type ID	User/System	Description
ACT01	IBM Spectrum Protect Plus administrator	The IBM Spectrum Protect Plus administrator manages the IBM Spectrum Protect Plus infrastructure (ES01). ACT01 provides essential services in managing, operating, and configuring infrastructure; scheduling backups, and defining data retention policies. Depending on the IT team in an organization, guest access to protect, restore, or reuse data can also be delegated to users (for example, ACT02, ACT03, and ACT05) by using IBM Spectrum Protect Plus role-based access control (RBAC).
ACT02	Virtual Infrastructure administrator	ACT02 is the main actor of IBM Spectrum Protect Plus to use it as the backup and restore solution for all of the guests that are hosted in a virtual infrastructure (ES04). ACT02 also provides access to the IBM Spectrum Protect Plus solution for vCenter and the underlying virtual infrastructure components.
ACT03	Virtual machine administrator	ACT03 is an IBM Spectrum Protect Plus solution stakeholder whose role is to protect and restore the guests for whom they are responsible.

Type ID	User/System	Description
ACT04	Network administrator	The network administrator is responsible for the network infrastructure element (ES05), and must provide adequate infrastructure support for the backup and restore requirements.
ACT05	Application administrator	ACT05 is another stakeholder in the solution who must be aware of how it works and what it can do for their specific database application, such as data reuse.
ACT06	Storage administrator	ACT06 is responsible to ensure that the IBM Spectrum Protect Plus components have appropriate access to storage components where snapshot copies are retrieved and stored.

Roles and responsibilities across functions

Depending on the structure, the separation of duties, roles, and responsibilities between the teams can vary. The distribution of roles and tasks is shown in Figure 2-3.

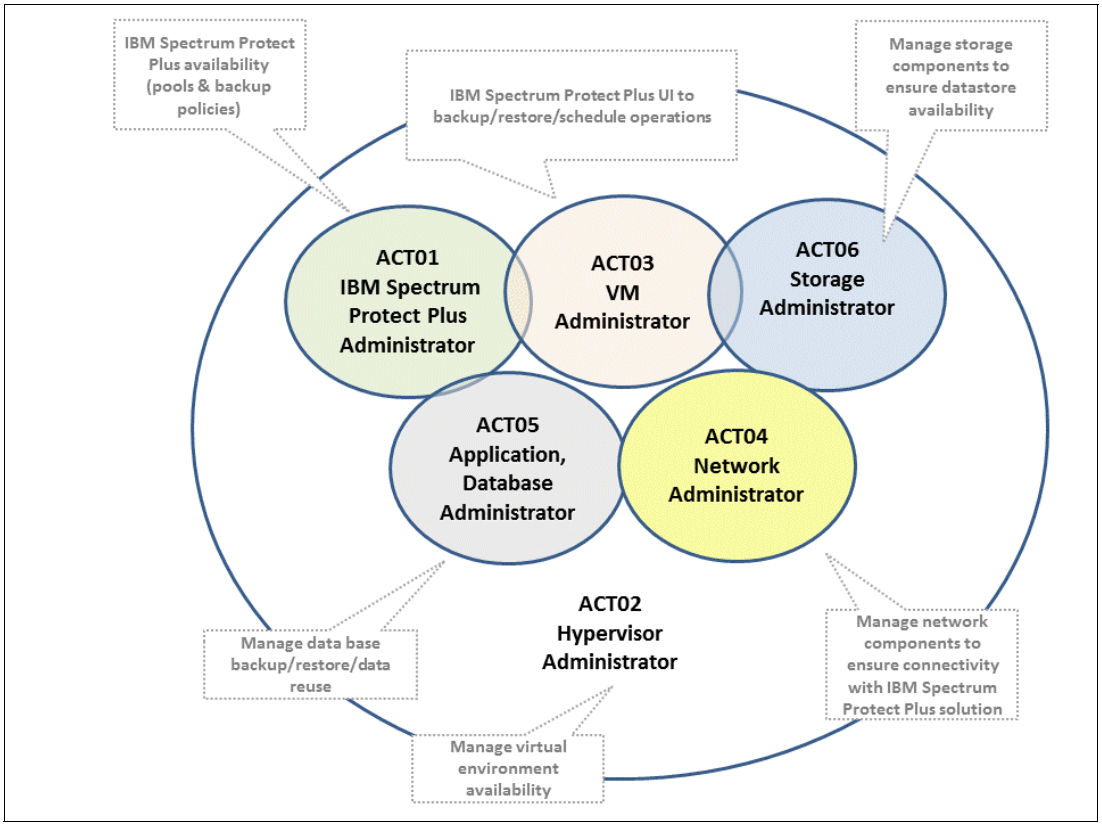

Figure 2-3 Tasks distribution across IT competencies

In Figure 2-3, bubble intersections represent an operation that can be done by IBM Spectrum Protect administrators, virtual machine administrators, or application administrators, depending on your IT organization. Therefore, a bubble intersection indicates that there should be team communication. Each person who is responsible for a task can inform, or request assistance from, other peer users.

Relations within the overall solution infrastructure

Table 2-2 lists how the solution interacts with external systems.

Table 2-2 External systems involved in the IBM Spectrum Protect Plus solution deployment

ID	Name	Description
ES01	IBM Spectrum Protect Plus Infrastructure	The IBM Spectrum Protect Plus infrastructure external service is a requirement for the solution to function. It contains several components, such as the IBM Spectrum Protect Plus server, VADP proxy, vSnap server, and data retention policy. Without ES01, an IBM Spectrum Protect Plus solution cannot exist.
ES02	Disk SAN/Object Storage/iSCSI Fabric	This external service is where snapshot copies are stored. In some cases, it can also include storage where backups are stored before it is migrated to ES03.
ES03	Tape SAN Fabric	The tape SAN external service is where snapshot copies can be stored to provide a physical air gap. It can be accessed indirectly through the IBM Spectrum Protect server when a second backup copy to physical tape occurs or is accessed.
ES04	Virtual Infrastructure	This external service is the main interface for the IBM Spectrum Protect Plus solution that enables data to be retrieved. It includes components, such as the hypervisors, guests, and logical data stores where guest data is hosted.
ES05	Local area network	LAN is the main transport for data being backed up and restored by the IBM Spectrum Protect Plus solution. This service most likely determines the overall performance of the solution.
ES06	Applications, Database	Applications and database constitute the part of the data that is protected by the IBM Spectrum Protect Plus solution.
ES07	Operating Systems	Operating Systems are the part of the data that is protected by the IBM Spectrum Protect Plus solution. One of the main benefits of the solution is the ability to start from a backup, which provides quick and efficient recovery when operating system failures occur.
ES08	Core Delivery Platform	This external service is to provide centralized management, monitoring, alerting reporting, billing, security, compliance, and other IT services with which the IBM Spectrum Protect Plus solution must directly or indirectly interface.

Note: The users that are listed in Table 2-1 on page 31 must be aware of and involved in the data protection product implementation.

2.2.4 Architectural Overview Diagram

At the end of a design phase, an Architectural Overview Diagram (AOD) should be created to visualize the design. This diagram is in the planning and sizing phase and is developed to create a complete solution architecture of a modern data protection environment. In our example for an AOD, we have two locations: multiple vSnap servers and a connection to Amazon, as shown in Figure 2-4 on page 34.

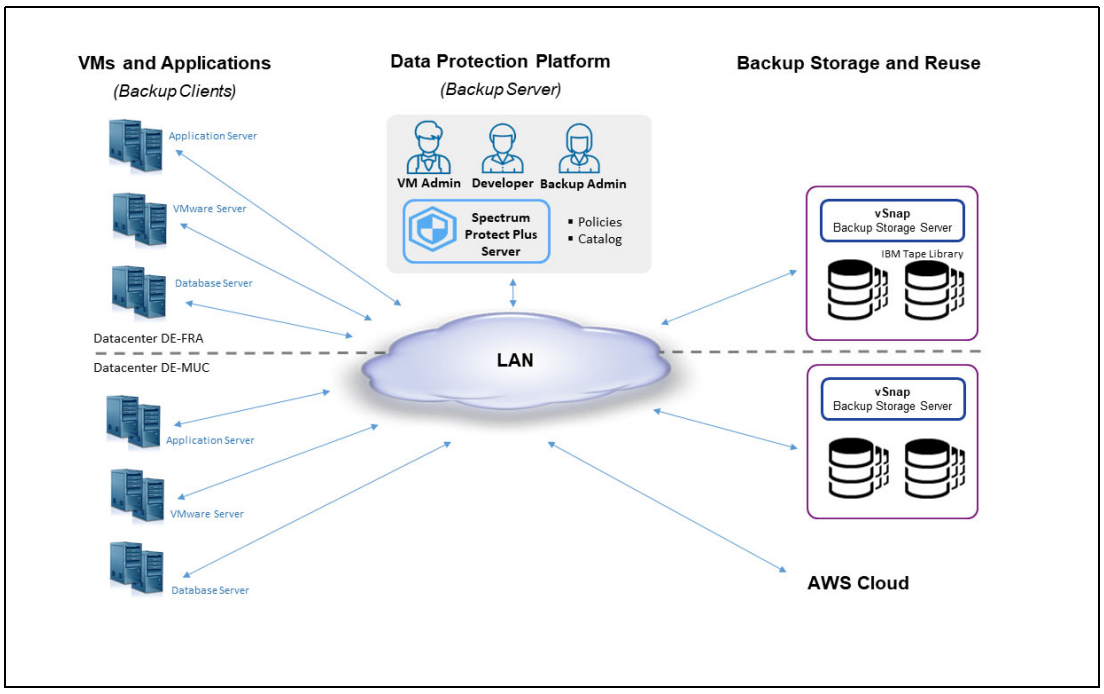

Figure 2-4 Architectural Overview Diagram example

2.3 Planning the solution

The intent of this section is to provide examples of key decisions to be made when designing a data protection solution with IBM Spectrum Protect Plus. This information is not intended to replace the analysis of your environment, and is not a complete set of decisions to make. However, it can help you to sketch a design of what your solution might look like, and what must be considered to build a data protection solution that meets your requirements.

One of the most important steps in designing a data protection solution is to validate that the architectural decisions meet requirements, including networking, security, and various technology considerations. In the first part of this section, we explain different technical options for a deployment of IBM Spectrum Protect Plus. Then, we provide templates to document architectural decisions for a solution design.

The IBM Spectrum Protect Plus Blueprints discuss the planning and sizing of an IBM Spectrum Protect Plus data protection environment. In Chapter 2, “Choosing the Appropriate Technologies” of IBM Spectrum Protect Plus Blueprints, topics are described, such as physical or virtual vSnap server deployment, disk technology, compression, and deduplication with their different options. Therefore, the Blueprints are a *must read* when planning an IBM Spectrum Protect Plus installation. The latest versions are available at this IBM Support web page.

In this chapter, we discuss more aspects of an IBM Spectrum Protect Plus deployment that are outside the information that is included in the Blueprints.

2.3.1 Deployment options for IBM Spectrum Protect Plus

All on-premises environment

In this configuration, the IBM Spectrum Protect Plus server and the vSnap server are deployed on an on-premises environment:

- IBM Spectrum Protect Plus server:

  This server is deployed as a virtual machine image on a VMware or Hyper-V hypervisor. A shared or dedicated hypervisor infrastructure can be used. A dedicated environment for a production data protection environment is preferred to separate the primary systems from the backup environment. At a minimum, the used storage is separated between primary and backup systems.

- vSnap server:

  This server is deployed as a virtual machine image on a VMware or Hyper-V hypervisor, or as a physical server. When deploying as virtual machine, it is recommended to a use a dedicated hypervisor infrastructure.

Hybrid environment

In this configuration, the IBM Spectrum Protect Plus server and a vSnap server are installed and maintained on-premises, while another Snap server is deployed on IBM Cloud or on AWS on an existing or new VPC. For more information about the cloud solutions, see the following web pages:

- IBM Spectrum Protect Plus on IBM Cloud
- IBM Spectrum Protect Plus on AWS Marketplace

A hybrid environment might benefit IBM Spectrum Protect Plus users who want to continue protecting workloads that are running on-premises and cloud. In addition to backup and recovery operations, you can use a hybrid environment to replicate and reuse data between your on-premises location and the cloud for more data protection. For example, you might want to use on-premises data that is replicated on the cloud for DevOps, quality assurance, and testing purposes.

Figure 2-5 shows a hybrid environment example with AWS.

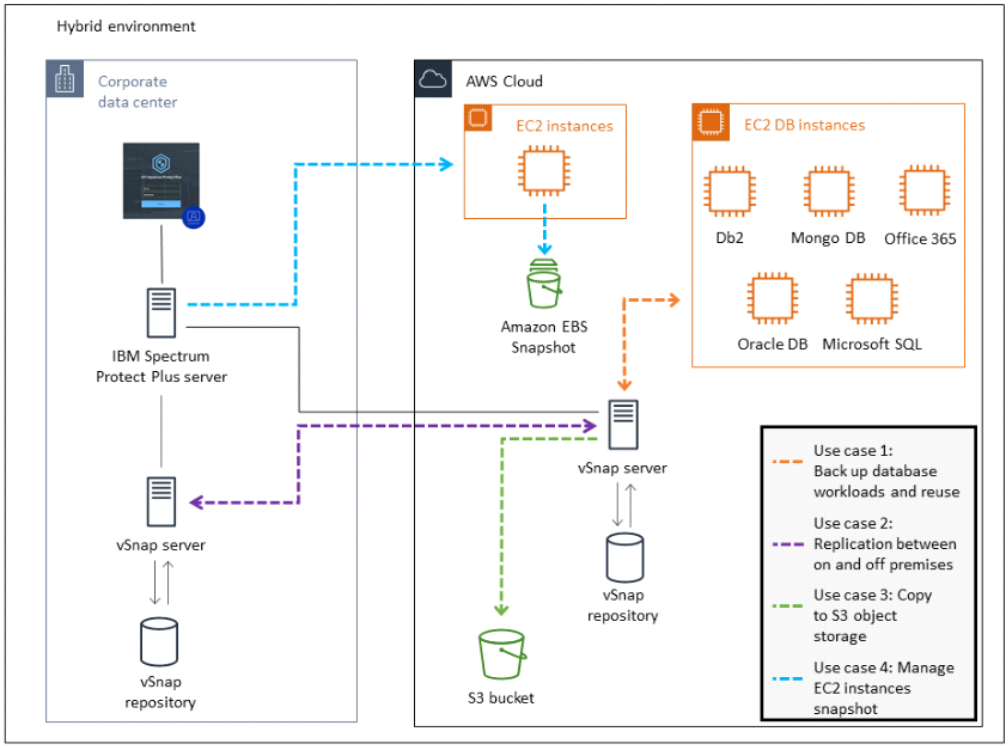

Figure 2-5 Example of a hybrid IBM Spectrum Protect Plus environment

All on cloud environment

In this configuration, the IBM Spectrum Protect Plus server and the vSnap server are deployed in IBM Cloud or on AWS on an existing or new Virtual Private Cloud (VPC). An on-premises IBM Spectrum Protect Plus server and a Hyper-V or VMware infrastructure is not required. For more information about the cloud solutions, see the following web pages:

- IBM Spectrum Protect Plus on IBM Cloud
- IBM Spectrum Protect Plus on AWS Marketplace

This option might be interesting for new IBM Spectrum Protect users who want to protect workloads on the cloud and do not currently have IBM Spectrum Protect Plus running in an on-premises environment.

Figure 2-6 shows an all-on-cloud environment in an example with AWS.

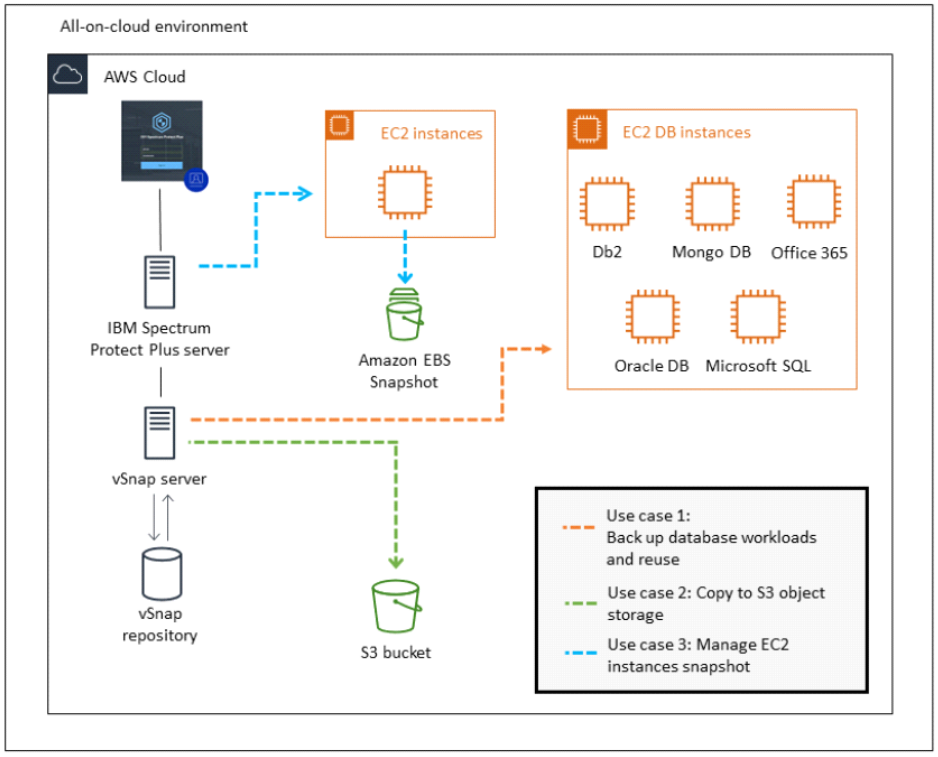

Figure 2-6 Example of an all on cloud IBM Spectrum Protect Plus environment

2.3.2 Scaling the environment

Managing the data growth in a data protection environment is a key challenge. Typical questions are about how many components are needed and how large do they need to be sized to fulfill any future needs?

IBM Spectrum Protect Plus server

A typical IBM Spectrum Protect Plus environment uses a single IBM Spectrum Protect Plus server and the default catalog sizes are suitable for most environments.

In an installation that is spread over multiple physical locations, the communication between the different components must still work correctly. Therefore, the network round-trip time must be at a level that allows the IBM Spectrum Protect Plus server to communicate with the vSnap servers and the other components, such as the VMware vCenter servers.

Also, the IBM Spectrum Protect Plus GUI must be still accessible for all required users. For example, if an installation exceeds continental boundaries, the use of multiple IBM Spectrum Protect Plus servers is preferred. However, if the network connection between an IBM Spectrum Protect Plus server in Germany and a vSnap server in Italy is not sufficient, a separate IBM Spectrum Protect Plus server can be required.

Also, in installations with multiple large VMware vCenters, it is beneficial from an organizational and scaling perspective to use multiple IBM Spectrum Protect Plus servers.

vSnap server

A typical IBM Spectrum Protect Plus environment uses multiple vSnap servers with a maximum initial capacity of 100 TB.

Because a key design aspect of an IBM Spectrum Protect Plus environment is to localize data movement, you often must have at a minimum as many vSnap servers as sites exist in your environment. Required logical separation of backup data can lead to installing even more vSnap servers.

The initial size of 100 TB per vSnap server provides the ability to scale-up this vSnap server to cope with natural data growth. In principle, it is better to start initially with more, smaller vSnap server instead of using less, but larger systems. The use of multiple vSnap servers (within a Site) for a single SLA allows for an initial distribution of the backup workload throughout the available vSnap servers.

2.3.3 Disaster Recovery and high availability with IBM Spectrum Protect Plus

The protection of the IBM Spectrum Protect Plus environment consist of the following types of protection:

- Backup data: Protection is achieved by using RAID protected storage for the vSnap pool and replicating backup data to another vSnap server. For more information about replication, see Chapter 15, “Replication and additional copies” on page 405. For more information about recovering from a vSnap server failure, see this IBM Support web page.
- vSnap server metadata: Protection is achieved by creating regular backups of the metadata information. For more information about metadata and protection, see “Protect vSnap server configuration (metadata)” on page 15.
- IBM Spectrum Protect Plus catalog: Protection is achieved by regular backups of the catalog. For more information about catalog backup, see “Data catalogs” on page 8.

Another option to protect the IBM Spectrum Protect Plus server is to run it in a high availability or vSphere Fault Tolerance configuration:

- High Availability enables the IBM Spectrum Protect Plus server to be available with a minimum amount of downtime if a VM host failure occurs. This high availability can be achieved through VMware vSphere High Availability or Microsoft failover clustering.
- Fault Tolerance enables the IBM Spectrum Protect Plus server to be continuously available if a VM host fails. This fault tolerance can be achieved through VMware vSphere Fault Tolerance (FT).

For more information about these other configuration options, see this IBM Support web page.

2.3.4 Network

The network is a crucial component for an IBM Spectrum Protect Plus implementation because the different components need to interact with each other. In one case, dedicated ports for the communication are required. Alternatively, separate networks should be used to cope with the backup and restore workload or to isolate backup traffic from application and user networks. We describe information about network configuration in Chapter 4, “Networking” on page 83.

2.3.5 Encryption

In ransomware attacks and other security breaches, your data protection solution is the last defense and is critical to protect your snapshot copies. Encryption is one solution to “harden” the backup snapshots. As encryption pertains to various layers and elements of an IBM Spectrum Protect Plus environment, there are two encryption approaches to consider: *Data at Rest* to encrypt stored snapshots and *Data in transit* to encrypt data transferred between components and sites.

Data at rest: Encryption of data at rest can be accomplished by using encryption capable storage devices, such as the IBM disk or tape storage subsystems, file systems, or applications that support encryption. Encryption capable devices implement inline transparent encryption of data as it flows onto and off the associated media. Encryption of data ensures that, if a physical loss of media (through theft or replacement) occurs, the data is unreadable and its confidentiality is maintained.

Data In transit: Data that is being transferred between applications or components is in transit. *Data in transit* is also referred as *data in flight* or *data in motion*. Throughout this publication, the term data in flight is used. With IBM Spectrum Protect Plus, SSL is the industry standard used for inter-component communication and data transfer. SSL is a set of rules governing authentication and encrypted communication between clients and servers.

SSL is widely used on the internet by an increasing number of applications. SSL is positioned as a protocol layer between the Transmission Control Protocol (TCP) layer and the application to form a secure connection between clients and servers so that they can communicate in a secure manner over a network.

Encryption approaches in IBM Spectrum Protect Plus

A user’s decision to enable or disable encryption is driven by several factors:

- Performance considerations
- Security requirements and policies
- Data center design

The schema that is shown in Figure 2-7 on page 40 shows the various hardware and software components within an IBM Spectrum Protect Plus environment. The orange-colored network paths indicate that the network traffic is encrypted when data is in flight by using the SSL protocol. Components that are displayed in green support optional encryption of data at rest.

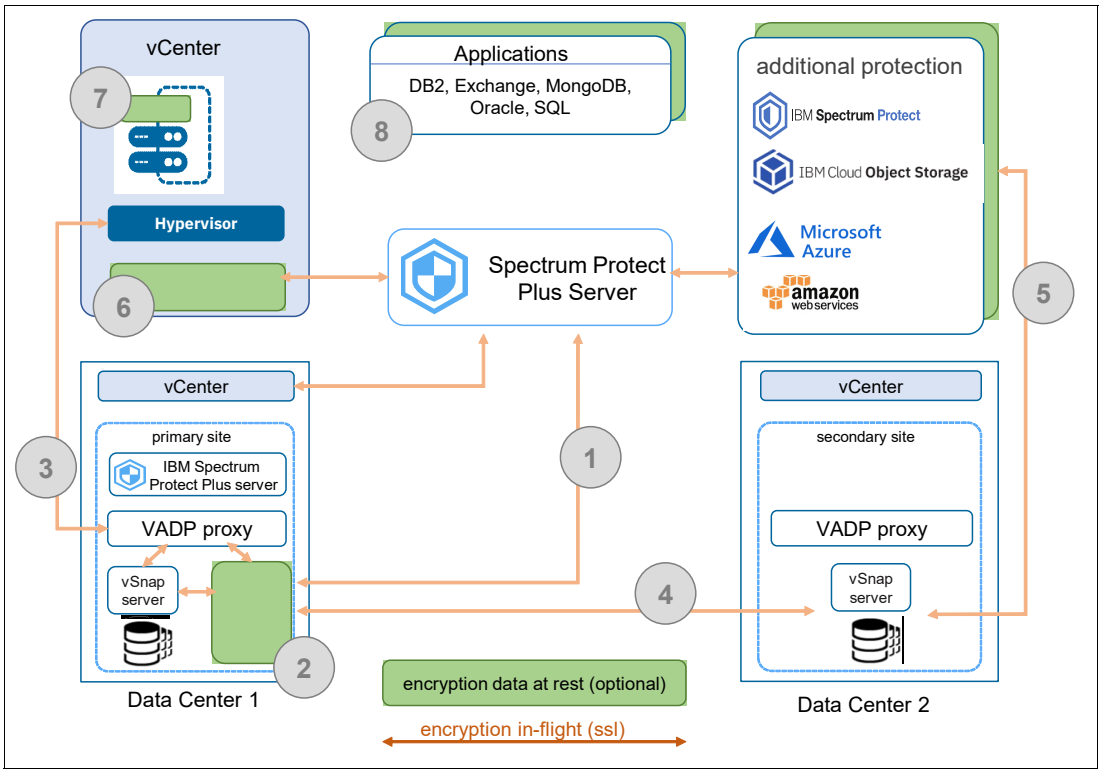

Figure 2-7 Encryption at rest and in flight in an IBM Spectrum Protect Plus environment

The following components are shown in Figure 2-7:

1. IBM Spectrum Protect Plus Server to vSnap server: Communication between the server and vSnap servers take place over the HTTPS protocol, so it is encrypted using Transport Layer Security (TLS).
2. vSnap server Data Encryption: IBM Spectrum Protect Plus provides the option to encrypt data using AES 256-bit encryption in the vSnap server at-rest. Data that is obtained during a backup or replication operation to a vSnap server can be encrypted after the data is compressed or deduplicated. The data remains encrypted in the vSnap server until it is read for a restore, data reuse, or replication operation, at which time it is decrypted.

   To use the vSnap server encryption feature, a vSnap server must be initialized with encryption enabled during creation. All disks of a vSnap pool use the same key file, which is automatically generated upon pool creation. The vSnap server is using AES 256-bit encryption. Encryption is a function of the CPUs. The sizing information from the Blueprints and the sizing tool assume that encryption is used. Encryption uses 5 - 10% extra CPU resources.

   Note: Encryption can be enabled only when the pool is created and is *not* revertible. However, compression and deduplication can be changed on a vSnap pool at any time.

3. VADP proxy to VMware hypervisors: VADP proxies in IBM Spectrum Protect Plus support the following VMware transport modes: SAN, HotAdd, NBDSSL, and NBD. Although every enterprise is unique and has different priorities in terms of size, speed, reliability, and complexity, the following general guidelines apply to the Transport Mode selection:
   - SAN transport mode should be used in a direct storage environment because this mode is fast and reliable.

- HotAdd transport mode should be used if the VADP proxy is virtualized. This mode supports all vSphere storage types.
- NBD or NBDSSL transport mode (LAN) is the fallback mode because it works in physical, virtual, and mixed environments. However, with this mode, data transfer speed might be compromised if network connections are slow. *NBDSSL* mode is similar to *Network Block device (NBD)* mode except that data transferred between the VADP proxy and the hypervisor server is encrypted using SSL.

4. vSnap server to vSnap server replication: Data that is replicated between two vSnap servers is protected by SSH encryption using Secure Sockets Layer (SSL). When replicating data, the replication process protects the data in-flight. If the data of a vSnap server was encrypted at rest, data is decrypted when reading from the source vSnap server and encrypted in-flight when data is transferred by way of an SSL connection.

   You need to ensure that the target vSnap server is also encrypted if you intend to store the replicated data in an encrypted format. There is no dependency between source and target vSnap server in terms of encryption configuration as data gets deciphered when data is read from a vSnap server.

5. vSnap server for extra protection, such as IBM Spectrum Protect, IBM Cloud Object Storage: The object agent uses HTTPS to transfer data to another data store. The communication protocol is encrypted using Transport Layer Security (TLS) and protecting the data in-flight. The user needs to configure the additional protection target system to provide at-rest encryption if the user intends to protect the offloaded data at-rest.

6. Transparent encryption on VMware datastore or encrypted storage device: Data can be encrypted at rest using the storage subsystem encryption capabilities, such as self-encrypting disks or software encryption at LUNs or pool level. Another option to encrypt VMware related data is the use of an encrypted vSAN instead of encrypting single VMs. The datastore encryption does not rely on the underlaying hardware capabilities and is not mutually exclusive. However, the administrator should not use more than one encryption solution.

7. Virtual machines: Encrypted virtual machines are supported in vSphere 6.5 environments and later. These VMs can be backed up and restored at the virtual machine-level to their original location using VMware's VMcrypt. If you are restoring to an alternative location, the encrypted virtual machine is restored without encryption, and must be encrypted manually through the vCenter Server after the restore completes.

   When a virtual machine's VMDK files are backed up, the files are retrieved decrypted from the hypervisor and are stored decrypted on the vSnap server. If the vSnap server is encrypted, the data is protected by the file system encryption.

   For more information about VMware vSAN and VMcrypt, see this VMware Knowledge base web page.

8. Application Level encryption: Many applications provide encryption capabilities at the application level. The following resources provide a brief overview and introduction to the specific application-level encryption implementations:
   - IBM Knowledge Center
   - Mongo DB
   - severalnines.com blog
   - Oracle: Transparent Data Encryption
   - Microsoft SQL Server

2.3.6 Architectural decisions template

This section provides architectural decision examples and a template to help you capture and document the rationale and justification for different decisions. Examples in Table 2-3 - Table 2-8 on page 44 show how to articulate such decisions for a typical enterprise customer.

Remember: The examples are for illustration only. The decisions can be dependent of your environment. You must evaluate how different options apply to your business requirements and base your own decisions on what fits best in your environment.

Table 2-3 Architect decisions 01

Subject area	Area of concern	Topic	Topic of interest
Architectural decision	Enable deduplication with compression.	AD ID	AD01
Issue or problem	What is the technology to use for reducing the data capacity for storing backups copies.		
Assumptions	► Compression is enabled on the vSnap server. ► Aggressive RTOs requirements such as “boot from backup” to resume production work instantly, such as starting an Oracle database directly from vSnap server, versus first migrating it back to production storage.		
Motivation	Reducing storage cost while providing rapid recovery		
Alternatives	► Enable deduplication ► Do not enable deduplication ► Only use deduplication with flash/SSD disks		
Decision			
Justification			
Implications			
Derived requirements			
Related decisions			

Table 2-4 Architect decisions 02

Subject area	Area of concern	Topic	Topic of interest
Architectural decision	Use IBM Storwize® v5000 Distributed RAID	AD ID	AD02
Issue or problem	How do we protect the vSnap server from disk failures; for example, RAID protection?		
Assumptions	There is a pre-existing investment in storage that provides storage hardware RAID such as the IBM Storwize v5000		
Motivation	Ensure stable and proven backup storage environment.		
Alternatives	1. Software RAID 2. Hardware DRAID		
Decision			
Justification			

Subject area	Area of concern	Topic	Topic of interest
Implications			
Derived requirements			
Related decisions			

Table 2-5 Architect decisions 03

Subject Area	Area of Concern	Topic	Topic of Interest
Architectural decision	Encrypt at rest method by using vSnap server native encryption and encrypt between, for example, data centers where data goes over TLS (https) connection.	AD ID	AD03
Issue or problem	How do you provide an encryption at rest and in flight solution for data stored on IBM Spectrum Protect Plus vSnap servers that aligns with technical best practices?		
Assumptions	All customers are subject to GDPR unless there is a formal letter saying that they are not.		
Motivation	The solution must ensure both the security of the data from physical device theft and minimize the risk of data loss through key loss.		
Alternatives	1. Encrypt at rest. 2. Encrypt over TLS (https) connection.		
Decision			
Justification			
Implications	N/A		
Derived requirements	N/A		
Related decisions	N/A		

Table 2-6 Architect decisions 04

Subject Area	Area of Concern	Topic	Topic of Interest
Architectural D decision	Microsoft Active Directory authentication	AD ID	AD04
Issue or problem	What authentication method to use for the administrator that manages the IBM Spectrum Protect Plus infrastructure.		
Assumptions	There is a Microsoft Active Directory infrastructure.		
Motivation	Centralized authentication management		
Alternatives	Microsoft Active Directory Local authentication		
Decision			
Justification			
Implications			

Subject Area	Area of Concern	Topic	Topic of Interest
Derived requirements			
Related decisions			

Table 2-7 Architect decisions 05

Subject Area	Area of Concern	Topic	Topic of Interest
Architectural decision	Use physical vSnap servers for large environments and virtual vSnap server for small environments.	AD ID	AD05
Issue or problem	Determine the deployment of physical or virtual vSnap server.		
Assumptions	Most deployments require both separate vSnap servers and single-server deployments of IBM Spectrum Protect Plus.		
Motivation	Using both alternatives increases efficiency.		
Alternatives	Virtualize all vSnap servers. Provide physical vSnap servers with hardware RAID.		
Decision			
Justification			
Implications			
Derived requirements			
Related decisions			

Table 2-8 Architect decisions 06

Subject Area	Area of Concern	Topic	Topic of Interest
Architectural decision	Use Syslog based alerting	AD ID	AD06
Issue or problem	What alerting tool to use		
Assumptions	IBM Spectrum Protect Plus is used in large-scale delivery operations		
Motivation			
Alternatives	► API ► Email ► Syslog		
Decision			
Justification			
Implications			
Derived requirements			

Subject Area	Area of Concern	Topic	Topic of Interest
Related decisions			

2.4 Sizing

The information about the different backup workloads as described in 2.2, “Design” on page 27 are now required to perform the sizing of the solution. The IBM Spectrum Protect Plus Blueprints include a sizing tool that is based on Microsoft Excel to help you plan the following aspects of the solution:

- Quantity and capacity of vSnap servers
- Quantity of VADP proxies
- Replication requirements

In Chapter 3, “How to Use the Sizing Tool” of *IBM Spectrum Protect Plus BluePrint* the tool is explained followed by a sizing example. The Blueprints and the sizing tool are available at this IBM Support web page.

For more information, see the IBM Spectrum Protect and IBM Spectrum Protect Plus YouTube channel.

A playlist on YouTube with multiple videos about sizing an IBM Spectrum Protect Plus solution is provided at this web page.

3

Installation and deployment

This chapter reviews the prerequisites for installing and deploying an IBM Spectrum Protect Plus environment, and describes the following essential components:

- IBM Spectrum Protect Plus server
- vSnap Backup Storage server
- VADP proxy server

This chapter is not meant to duplicate the installation procedure that is available in IBM Knowledge Center, but rather provides guidance for the initial configuration and other essential settings following the installation and deployment.

This chapter includes the following topics:

3.1 Overview

The general installation and deployment overview that is shown in Figure 3-1 provides a high-level list of the tasks that are described in this chapter.

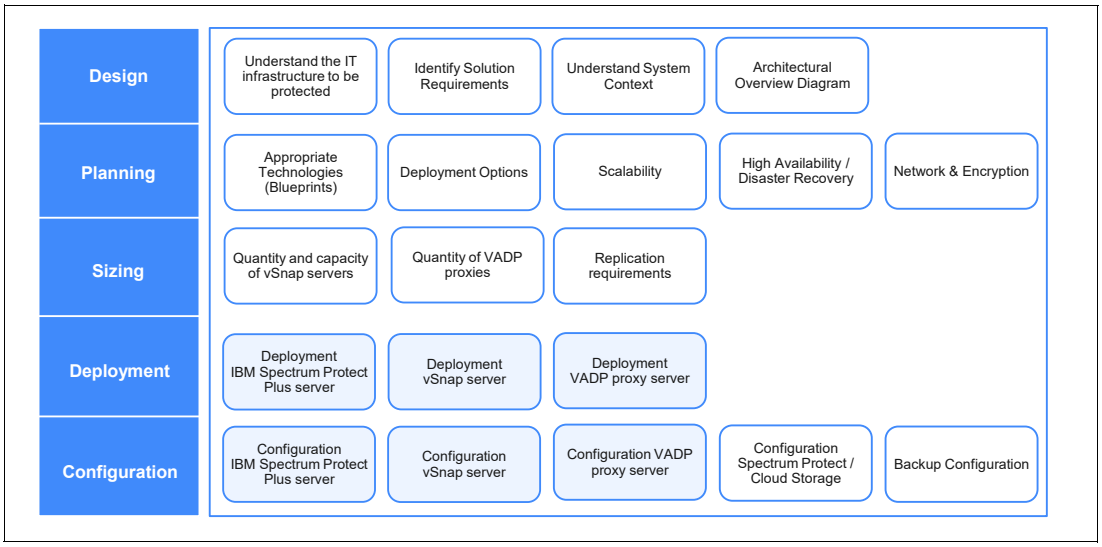

Figure 3-1 Installation and deployment overview

3.2 Prerequisites for an IBM Spectrum Protect Plus deployment

Before starting the installation and deployment of an IBM Spectrum Protect Plus environment, prerequisites and planning aspects must be reviewed and understood. For more information about planning considerations, see Chapter 2, “Solution architecture, planning, and design” on page 25.

One objective of the planning phase is to know the amount and size of vSnap Backup Storage servers and the number of VADP proxies that are required for an adequate backup infrastructure.

For a complete list of requirements, see see IBM Support’s IBM Spectrum Protect Plus - All Requirements Doc.

The document lists all requirements for the IBM Spectrum Protect Plus components and the requirements for the hypervisors, operating system platforms, Kubernetes, databases, and applications. The requirements for the IBM Spectrum Protect Plus server, vSnap server, and VADP proxy is listed under the “System requirements” section.

It is essential to also prepare and check your DNS infrastructure to include entries for your IBM Spectrum Protect Plus environment before starting the deployment. All IBM Spectrum Protect Plus system components can use DHCP, but the usage of static IP address assignments, including properly configured DNS names, is desirable for a productive IBM Spectrum Plus environment.

A common time zone and network time protocol server (NTP server) setting is recommended to be used for the IBM Spectrum Plus server, vSnap server, and VADP proxy server. For more information about configuring the time zone and NTP server, see 3.4.7, “Changing and verifying the schedules of the predefined jobs” on page 81.

3.2.1 IBM Spectrum Protect Plus server requirements

The IBM Spectrum Protect Plus server must be deployed on a VMware or Microsoft Hyper-V infrastructure. For VMware, a vCenter server is required to perform backup and restore activities on VMware virtual machines.

The download size of the installation image (`.ova` for VMware, and `.exe` for Hyper-V) is approximately 4.8 GB. The thick provisioned size of an installed IBM Spectrum Protect Plus server virtual machine (VM) is 548.0 GB. The provisioned size of the server VM is preconfigured and also suitable for large IBM Spectrum Protect Plus deployments.

Always use thick provisioned disks for the IBM Spectrum Protect Plus server VM in a production environment, and consider a deployment on SSD or Flash drives data stores to get a better VM performance. The IBM Spectrum Protect Plus server uses several internal catalogs, which benefit from a faster VM data store, especially in larger installations.

For more information about supported Hypervisor versions, CPU, and memory requirements, see IBM Support's IBM Spectrum Protect Plus - All Requirements Doc.

For more information about component requirements, see IBM Knowledge Center.

Note: The OVF package for VMware contains advanced configuration options, which might pose a security risk. Review the following advanced configuration options:

```
spp.vm.mode = appliance
disk.enableUUID = true
```

3.2.2 vSnap Backup Storage server requirements

A vSnap server is the primary backup destination for IBM Spectrum Protect Plus. It can be installed on a physical server or as a preconfigured VM in a VMware or Microsoft Hyper-V infrastructure.

Users can use a sizing tool, which is part of the IBM Spectrum Plus Blueprints. The chapter "Physical or virtual vSnap server deployment" in *IBM Spectrum Protect Plus Blueprints* should help you decide, based on your requirements, whether to use a physical or virtual vSnap server deployment. The most crucial factor is that the required sizing of the resources (memory, CPU, and so on) for the vSnap server is done correctly.

Along with the decision about a physical and virtual deployment is to decide what storage system technology to use. It can be storage out of an existing RAID-protected data store (for example in VMware) or are a simple disk storage for which you must use the software RAID features that are provided by the vSnap server to protect your backup storage against disk failures.

For vSnap server deployment on a physical server and dedicated storage system, it is preferable to use the hardware RAID feature of the storage system. If supported, use RAID 6, as recommended in *IBM Spectrum Protect Plus Blueprints*.

Also, in *IBM Spectrum Protect Plus Blueprints,* the chapter "Server and Storage Configuration" describes requirements and recommendations to build a vSnap Backup Storage server. It also includes more information about the required CPU and memory for the system.

IBM Support's IBM Spectrum Protect Plus - All Requirements Doc also lists details of supported Hypervisor and Operating System versions and the minimum CPU and memory requirements.

For more information about the requirements, see *vSnap Server Installation and User's Guide*, which is available at IBM Knowledge Center.

3.2.3 VADP proxy server requirements

The VADP proxy is required for moving the data from the VMware vSphere data stores to the vSnap Backup Storage server, for protection of VMware virtual machines. The proxy can be installed on a physical or virtual server.

You can install the vSnap server (backup storage provider) and VADP proxy on the same physical or virtual system. IBM Spectrum Protect Plus optimizes data movement by eliminating an NFS mount when these two systems are colocated.

If you choose to colocate these two components, ensure that the system is sized correctly for both workloads, taking the sum of the CPU, memory, and storage requirements of the two components. The IBM Spectrum Protect Sizing Spreadsheet helps to estimate the required CPU, memory, and storage.

For more information about supported operating system versions and the CPU and memory requirements (especially when running a combination of vSnap and VADP proxy servers), see IBM Support's IBM Spectrum Protect Plus - All Requirements Doc.

3.3 Installation and deployment

Following the planning and sizing process, the installation of the IBM Spectrum Protect Plus solution can be started. Remember the following aspects from the planning phase:

- An architectural overview of the installed solution is required:
  - VMware or Hyper-V infrastructure?
  - How many locations or data centers will be translated into site definition to achieve local data placement?
  - What type of disk storage is used for the vSnap servers as primary backup destination?
- Network setup:
  - Does the environment use the default VM network configuration?
  - Does IBM Spectrum Protect Plus/vSnap server intend to use different network configurations than the default (multi-tenant networking setup; for example, when infrastructure features a separate network for backup and production purposes. See Chapter 4, "Networking" on page 83.)
- How is the backup data protected?

  For more information, see Chapter 15, "Replication and additional copies" on page 405.
- A suitable solution sizing that uses the Blueprint Sizing tool is required:
  - How many VMs, databases, and applications should be backed up?
  - What is the expected amount of data to be backed up?
  - How long should the data be retained?

The chapter “How to use the Sizing Tool” from *IBM Spectrum Protect Plus Blueprints* provides guidance for a successful sizing.

The planning that is described in Chapter 2, “Solution architecture, planning, and design” on page 25 and the suitable chapters of the *IBM Spectrum Protect Plus Blueprints* provide the required information for installing the IBM Spectrum Protect Plus solution:

- How many sites must be configured?
- How many vSnap servers must be deployed? Virtual or physical installation?
- How many VADP proxy servers must be installed?
  - Virtual or Physical installation?
  - Will the VADP proxies run with vSnap server in the same VM or server?

The result of the planning and sizing phase in our example is that we install an IBM Spectrum Protect Plus server and three virtual vSnap Backup Storage servers with an integrated VADP proxy server in a VMware environment in three locations, as shown in Figure 3-2.

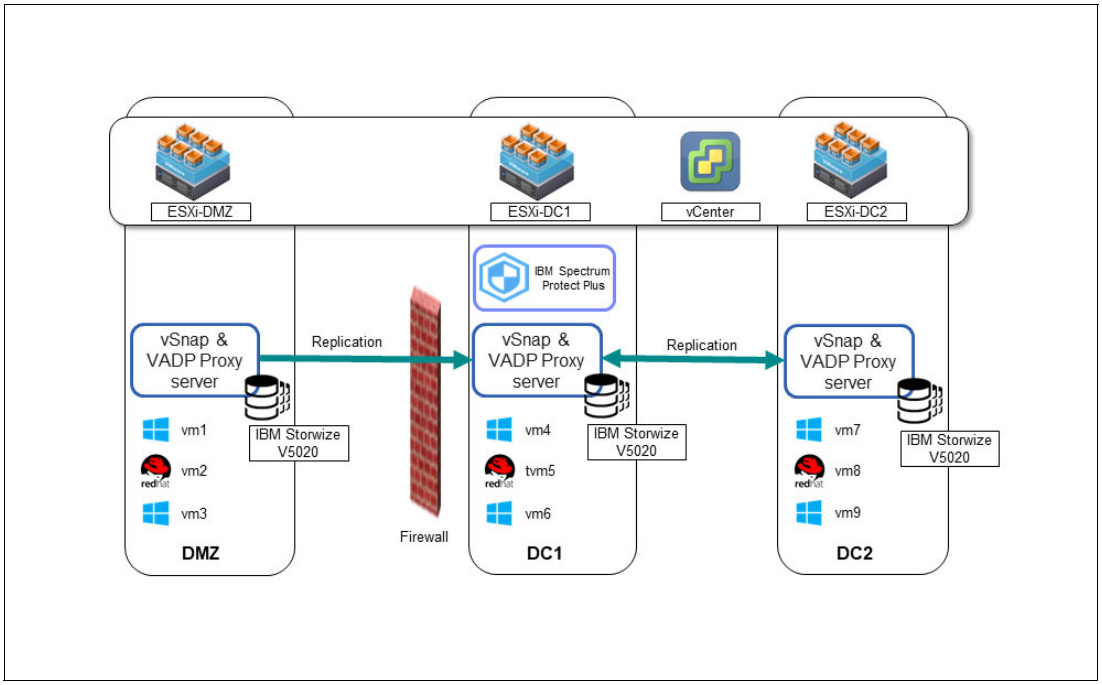

Figure 3-2 IBM Spectrum Protect Plus solution with three sites

For more information about the steps for implementing the solution that is shown in Figure 3-2, see 3.3.1, “Deploying the IBM Spectrum Protect Plus server”, and 3.3.2, “Deploying the vSnap Backup Storage server”.

For more information about the installation process, see these resources:

- IBM Knowledge Center
- IBM Spectrum Protect Plus Blueprints

After all components are deployed, see 3.4, “Configuring IBM Spectrum Protect Plus environment”.

3.3.1 Deploying the IBM Spectrum Protect Plus server

Before starting to deploy the IBM Spectrum Protect Plus server as a virtual appliance in a VMware environment, you must have the following information available:

- Virtual machine (VM) name for the IBM Spectrum Protect Plus server.
- Which Datacenter, ESX host, and data store that is to be used for the VM (Virtual disk format: Thick Provision).
- VM Network interface is recommended for the default deployment of IBM Spectrum Protect Plus in a VMware infrastructure. Select a VM Network that pointed to the default gateway and router to ensure that after the implementation of the OVA is completed, you can connect to the IBM Spectrum Protect interface and continue with the next configuration steps.
- Hostname (usually, the VM name) of the IBM Spectrum Protect Plus server.
- Network IP address for the IBM Spectrum Protect Plus server.
- Netmask (network prefix) for the IP subnet; for example, 24 for a `255.255.255.0` subnet mask.
- Default gateway.
- DNS server names.
- DNS domain name collated with the hostname as the fully qualified domain name (FQDN).

Create a worksheet or a list to gather the necessary details, as shown in Table 3-1.

Table 3-1 IBM Spectrum Protect Plus Server (ISPP) Worksheet

Item	Configuration details (IBM SPP)
Virtual machine name	
Name of the compute resource (ESXi)	
Storage(Datastore)	
Hostname (IBM Spectrum Protect Plus)	
Network IP address	
Network Prefix (Netmask)	
Default Gateway	
DNS	
Domain	

Using the VMware vSphere client to deploy an OVA template

When all information is available, deploy the IBM Spectrum Protect Plus server by using the VMware vSphere client.

Note: For more information about how to log in to the vCenter Server by using the vSphere Client, see this web page.

After logging in to the vSphere Client, start the deployment wizard by right-clicking the data center to use and selecting **Deploy OVF Template**, as shown in Figure 3-3.

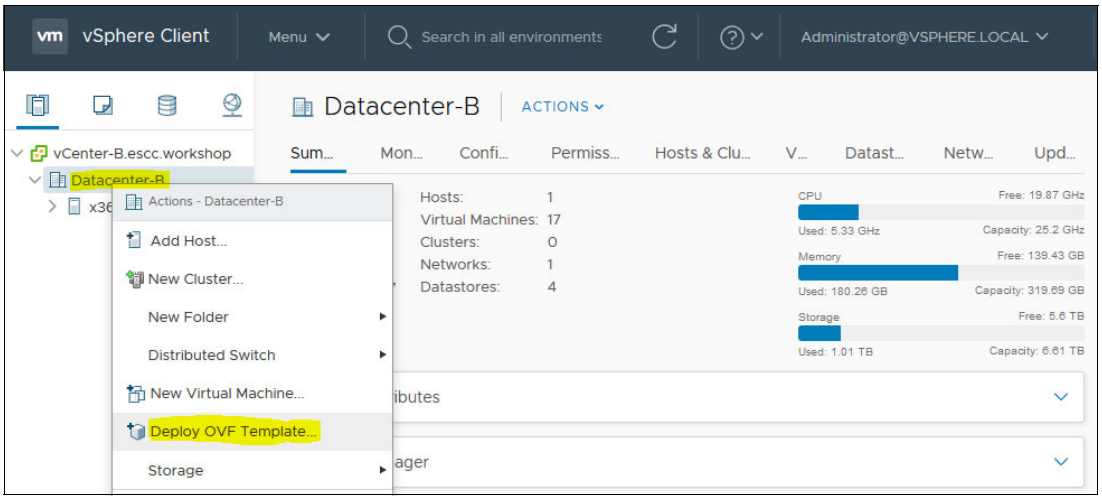

Figure 3-3 Deploy an OVF or OVA Template in the VMware vSphere Client

Note: For more information about deploying OVA templates, see this web page.

In step 6 of the deployment OVF Template configuration wizard (see Figure 3-4), select your storage configuration.

The following options are available for the virtual disk format:

- Thick Provision Eager Zero
- Thick Provision Lazy Zeroed (Default configuration)
- Thin Provision

Tip: The preferred choice is Thick Provision Lazy Zeroed because it is faster.

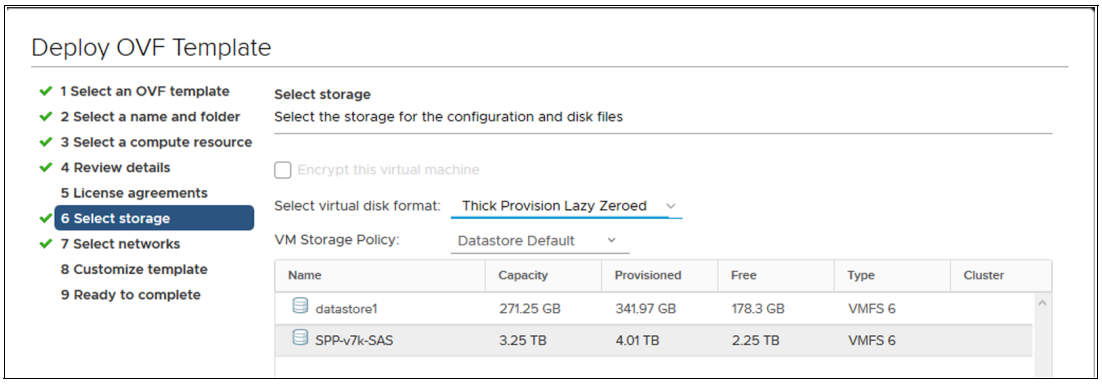

Figure 3-4 Deploy OVF Template Virtual disk format

The deployment wizard guides you through the required settings. Most settings must be entered in the Customize template section, as shown in Figure 3-5. The template details about disk requirement sizing are shown during the OVA deployment.

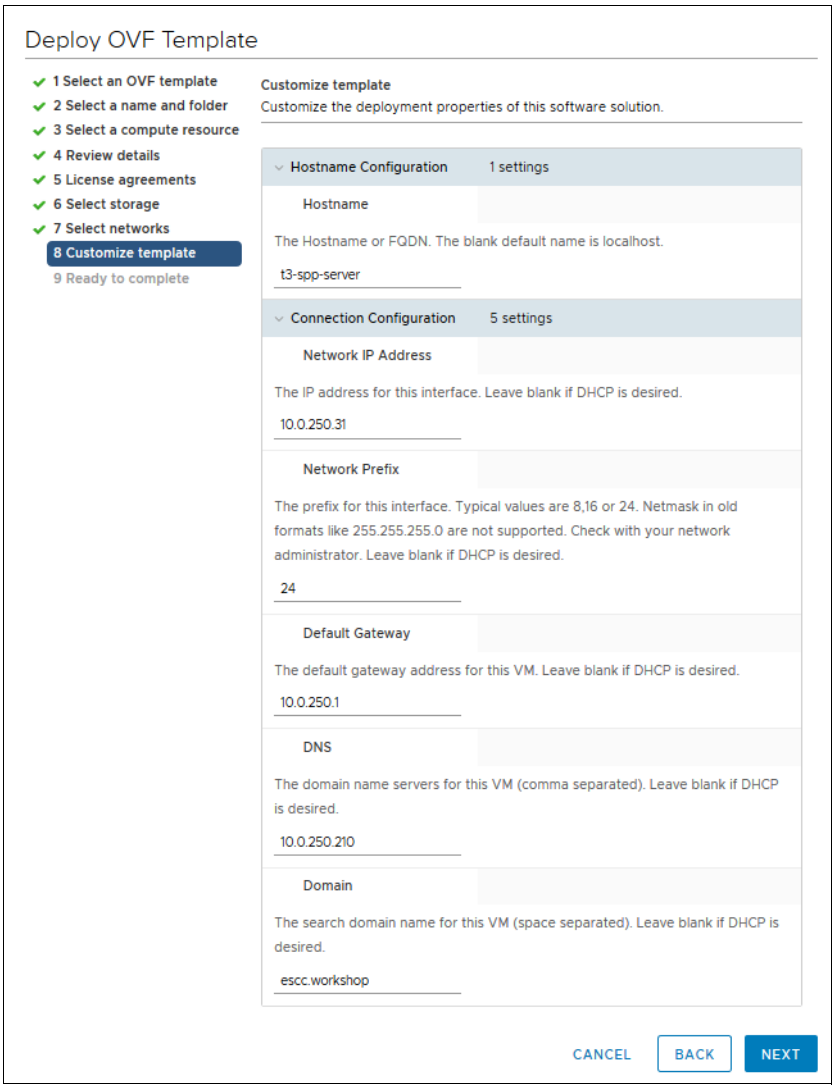

Figure 3-5 Wizard to deploy an OVF or OVA Template with options to customize a template

Finally, a summary page is displayed in which you can review all your settings before deploying the VM.

The vSphere Client shows a progress bar during the deployment, as shown in Figure 3-6.

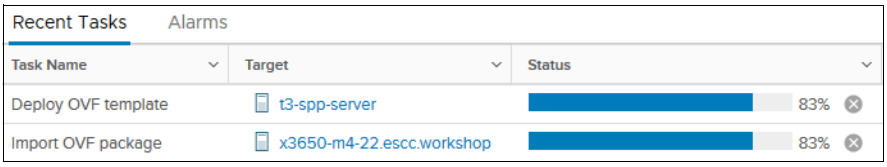

Figure 3-6 vSphere Client progress bar for the deployment of the OVF or OVA Template

Wait until the deployment completes and then, start the VM.

Important: Give the VM several minutes after starting to initialize completely, especially for the first start after the deployment.

The next step is to deploy the vSnap and VADP proxy servers. After all components are installed and deployed, go through the configuration steps that are described in 3.4, “Configuring IBM Spectrum Protect Plus environment” on page 56.

3.3.2 Deploying the vSnap Backup Storage server

In our example, we deploy a virtual vSnap Backup Storage server. For more information about the deployment of a physical vSnap server, see the *IBM Spectrum Protect Plus Blueprints*.

Before the vSnap server is deployed as a virtual appliance in an VMware environment, you must have the following information available:

- VM name for the vSnap server
- Which data center, ESX host, and data store are to be used for the VM
- Virtual disk format: Thick Provision Lazy Zero
- VM Network interface to be used from the VMware infrastructure
- IP address of the IBM Spectrum Protect Plus server
- Host name (most likely the VM name) of the vSnap server
- Network IP address for the vSnap server
- Netmask (Network Prefix) for the IP subnet; for example: 24 for a `255.255.255.0` subnet mask
- Default Gateway
- DNS server names
- DNS domain name that is used with the hostname as Full Qualified Domain Name (FQDN)

Use a worksheet (see as shown in Table 3-2) or create a list to gather the required configuration details for Deployment virtual vSnap server.

Table 3-2 IBM Spectrum Protect vSnap server worksheet

Item	vSnap configuration details
Virtual machine name	
Name of the compute resource (ESXi)	
Storage(Datastore)	
Hostname (IBM Spectrum Protect Plus)	
Network IP address	
Network Prefix(Netmask)	
Default Gateway	
DNS	
Domain	

After all information is available, deploy the vSnap server by using the VMware vSphere Client, as shown in “Using the VMware vSphere client to deploy an OVA template” on page 52.

Wait until the deployment completes and then, start the VM.

This step is repeated to deploy all three vSnap servers, as shown in Figure 3-2 on page 51.

Important: Give the VM several minutes after starting to initialize completely, especially for the first start after the deployment.

3.4 Configuring IBM Spectrum Protect Plus environment

After the IBM Spectrum Protect Plus server and the vSnap Backup Storage server are deployed, and the VADP proxy is installed from the IBM Spectrum Protect Plus UI, complete the following steps:

1. Configure the IBM Spectrum Protect Plus server.
2. Configure the vSnap Backup Storage server.
3. Configure NTP for the IBM Spectrum Protect Plus and the vSnap servers.
4. Connect the vSnap and VADP proxy servers with the IBM Spectrum Protect Plus server.
5. Add an SLA for IBM Spectrum Protect Plus catalog backup.
6. Back up the vSnap server system configuration.

Completing the configuration steps results in creating an IBM Spectrum Protect Plus environment in accordance with best practices that is ready to perform backups of VMs, file systems, and databases, as described in Chapter 6, “Backing up and restoring virtualized systems” on page 171, Chapter 8, “Backing up and restoring databases” on page 255, and Chapter 7, “Backing up and restoring Windows file system data” on page 227.

3.4.1 Configuring the IBM Spectrum Protect Plus server

As described in 3.3.1, “Deploying the IBM Spectrum Protect Plus server” on page 52, the IBM Spectrum Protect Plus server was deployed and as a last step the VM is started.

Important: Give the VM several minutes after starting to initialize completely, especially for the first start after the deployment.

The IBM Spectrum Protect Plus graphical user interface (GUI) can be accessed from a supported web browser by using https (for example, `https://spp-server/`). For more information about the GUI and its function, see “User interfaces” on page 6.

Initial log in to IBM Spectrum Protect Plus

When the IBM Spectrum Protect Plus server VM is up and running, starting the GUI displays the initial a log in dialogue that is shown in Figure 3-7 on page 57.

Note: If this is your first time logging in, the default user name is `admin` and the default password is `password`.

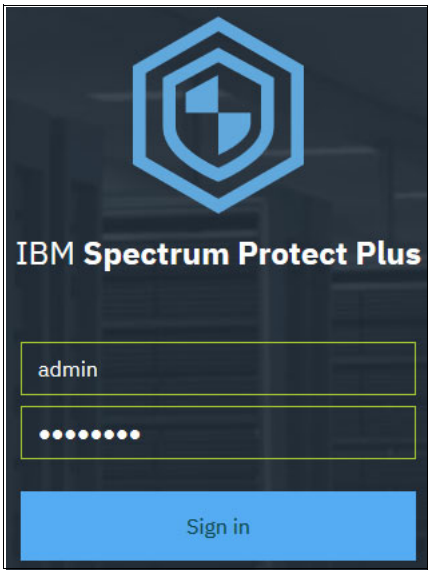

Figure 3-7 IBM Spectrum Protect Plus GUI - Sign in

Tip: If the log in dialogue is not displayed, it can mean that the I P address you specified during installation was invalid. Instead of repeating the installation, run the **`nmtui`** command to set a valid IP address. For more information about the **`nmtui`** command, see Example 4-3 on page 94.

The first time that you access IBM Spectrum Protect Plus, the following tasks must be completed before entering the Dashboard view:

- Change the user name and password of the administrative ID `admin`.
- Change the password for the local system user ID `serveradmin`**.**
- Initialize the built-in vSnap server.

Changing the user name and password of the administrative ID admin

At initial login, the default user ID `admin` must be changed to a meaningful name (see Figure 3-7). As a best practice, the recommendation is to change the ID `admin` to the name `superuser`. This recommendation is made because this initial administrative account features a dedicated user role with the same name. Even if later personalized user IDs or users from an LDAP directory are added, this user remains and cannot be deleted. If you prefer to set it to another name, that name also can work, except for `admin`, `root`, or `test`.

The following special user IDs are available:

- `admin`: (with role of superuser) manages IBM Spectrum Protect UI environment. The initial default admin user ID must be renamed. The default password must also be changed (see Figure 3-8 on page 58).
- serveradmin: (with role of the root rights) can be used for administration and installation activities; for example (VADP proxy) in IBM Spectrum Protect UI. Serveradmin can be used when accessing systems through the command line interface (CLI).

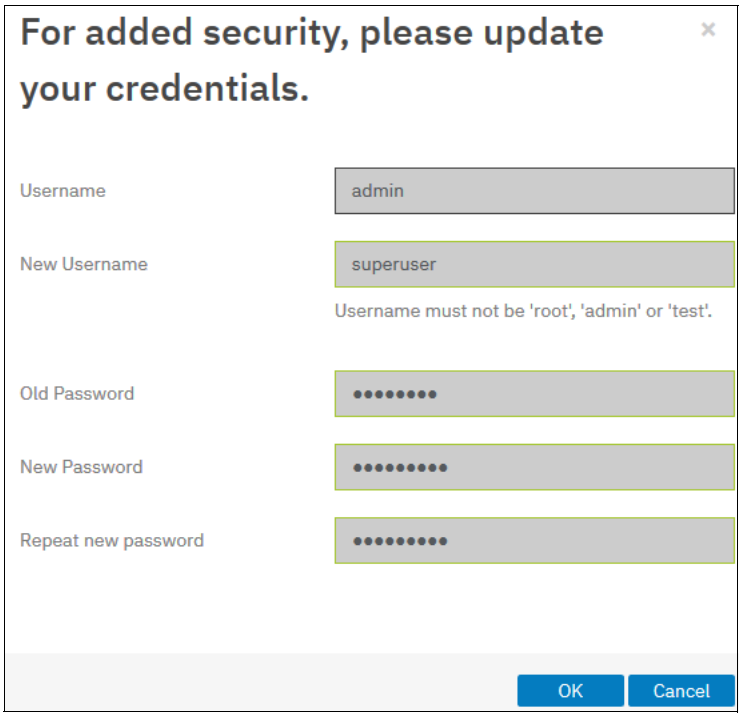

Figure 3-8 Change the initial login credentials

Note: It is not possible to rename as `admin`, `root`, or `test`. Also, the minimum password length must be at least eight characters.

Changing the password for the local system user ID serveradmin

You must also change the password of the local system user ID `serveradmin`. This user ID is used for logging in at the operating system level and can be used for logging in to the administrative console web interface. To log in to the administrative console, use IP address or fully qualified domain name and port 8090:

```
https://HOSTNAME:8090/
```

Note: In IBM Spectrum Protect Plus version 10.1.6, the initial password for the user ID `serveradmin` is `sppDP758-SysXyz`. During password change process, ensure that you follow the minimum password requirements.

Figure 3-9 on page 59 shows the change password dialogue for `serveradmin`. The new password for the `serveradmin` user ID must adhere to the following rules:

- Minimum acceptable password length is 15 characters
- Minimum of:
  - 8 characters in the new password must not be present in the old password
  - 1 numerical digit in the new password
  - 1 uppercase character in the new password
  - 1 lowercase character in the new password
  - 1 other character in the new password
- Maximum of:
  - 3 identical consecutive characters are allowed in the new password
  - 4 identical consecutive class of character are allowed in the password

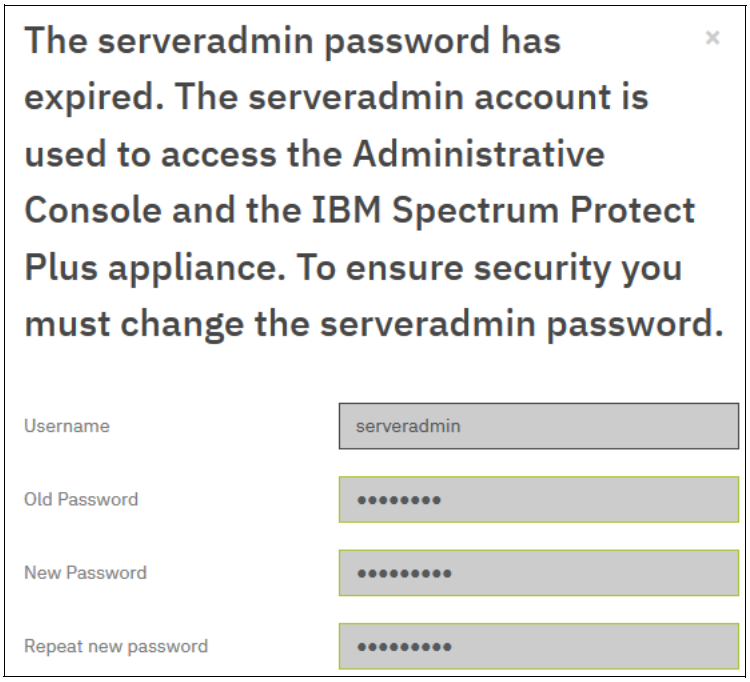

Figure 3-9 Change the serveradmin password

Initializing the built-in vSnap server

Initialize the built-in vSnap server that belongs to the site Demo. This step *cannot* be skipped. Select **Initialize without encryption**, as shown in Figure 3-10.

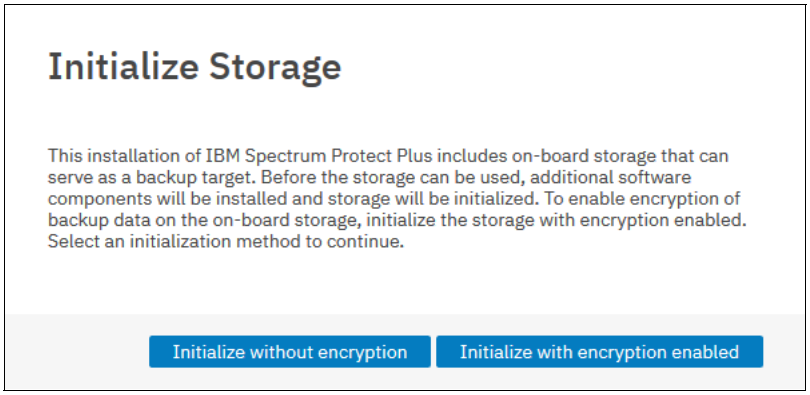

Figure 3-10 Initialize vSnap server storage

After completing these three initial login steps, you enter the IBM Spectrum Protect Plus Dashboard view, as shown and explained in “User interfaces” on page 6.

Defining and configuring required sites

The defining sites in IBM Spectrum Protect Plus is used to organize resources and to organize a local backup data placement. For more information about site principles, see 1.2.2, “Site” on page 12. Following our example in Figure 3-2 on page 51, we configure the following sites:

- DC1
- DC2
- DMZ

First, we rename the predefined **Primary** and **Secondary** site and add a third one under **System Configuration** → **Site**.

To rename a site, click the edit icon in front of the site name. After the first two sites are renamed, the third site is added by clicking **Add Site**, as shown in Figure 3-11.

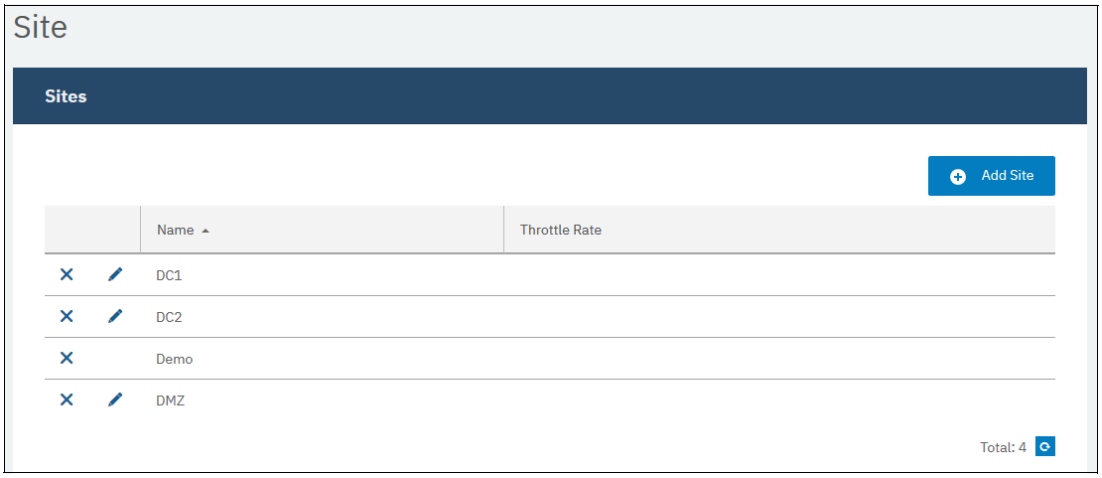

Figure 3-11 Defining the required sites in IBM Spectrum Protect Plus

Deleting the DEMO configuration

As explained in Chapter 1, “IBM Spectrum Protect Plus product architecture and components” on page 1, the initial deployment of the IBM Spectrum Protect Plus server includes a configuration of a Demo setup. This Demo setup consists of the following items:

- vSnap Backup Storage server with 100 GB capacity
- VADP proxy server
- Demo SLA configured for a daily backup with a retention of 1 month
- Demo site

The intention of this Demo configuration is for product demonstration purposes only. It should not be used for a productive backup; therefore, it can be deleted in real installations.

Deleting the local built-in vSnap server

We do not use the built-in vSnap server because of resource and performance reasons.

In the IBM Spectrum Protect Plus GUI, select **System Configuration** → **Backup Storage** → **Disk** and delete the vSnap server on the localhost in the Demo site by clicking the delete icon sign on the left side, as shown in Figure 3-12.

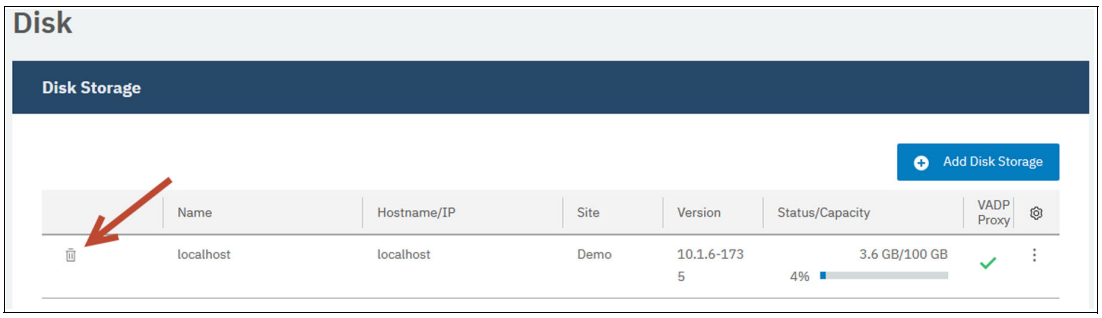

Figure 3-12 Local vSnap server disk storage for Demo site

You must confirm the deletion by entering a confirmation code and clicking **DELETE**, as shown in Figure 3-13.

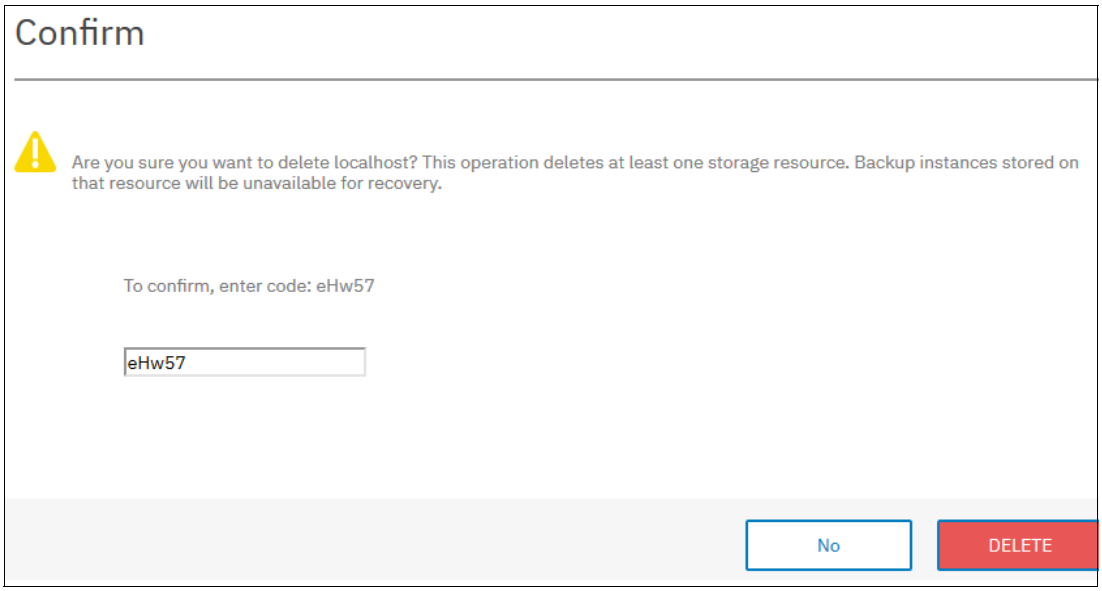

Figure 3-13 Confirm deletion

Suspending the local built-in VADP proxy

As a best practice recommendation, the built-in VADP proxy should not used and can be suspended. Because it cannot be removed completely, it must be suspended.

In the IBM Spectrum Protect Plus GUI, select **System Configuration** → **VADP Proxy**. Then, select **localhost**, click the three dots, and then, select **Suspend**, as shown in Figure 3-14.

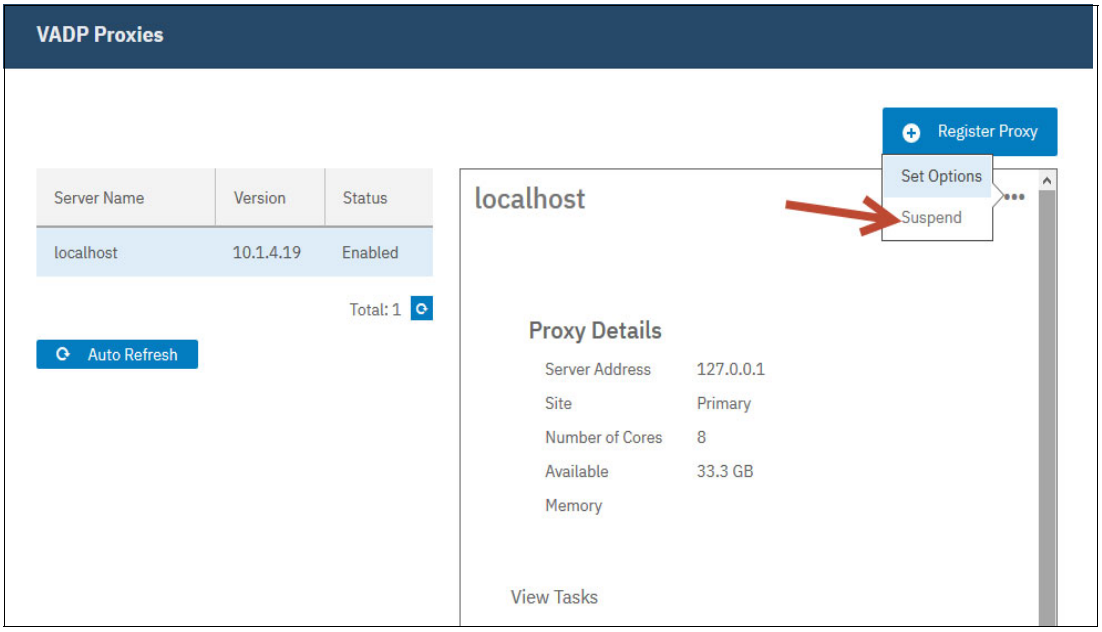

Figure 3-14 Suspend local built-in VADP proxy

Deleting DEMO SLA

The SLA Demo for demonstration purposes is not used. To delete it, choose **Manage Protection** → **Policy Overview**, select the **SLA Policy Demo**, and click the delete icon on the left side next to the edit icon . You must confirm the deletion by entering a confirmation code and clicking **OK**, as shown in Figure 3-15.

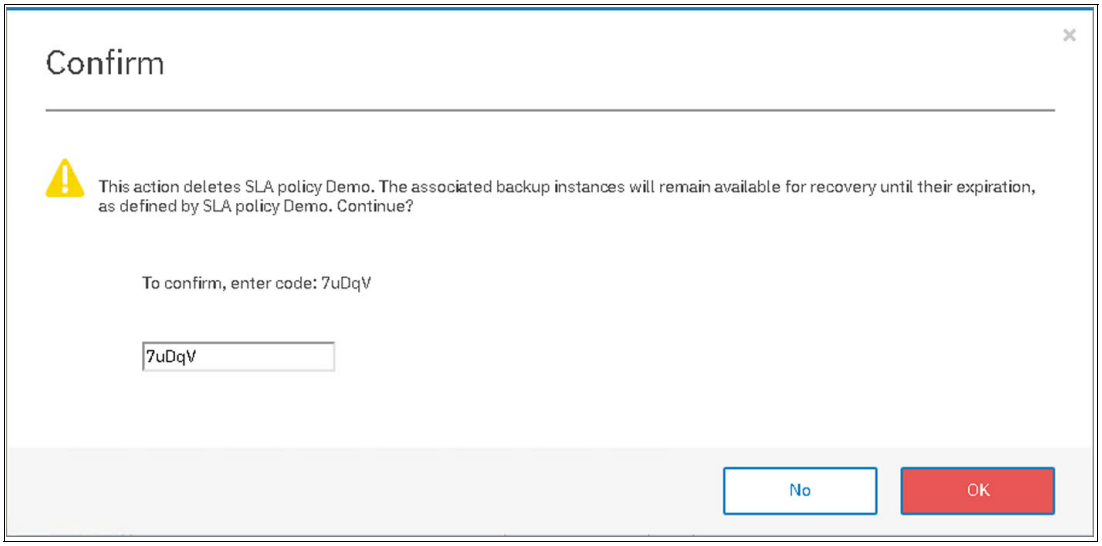

Figure 3-15 Confirm deletion

Deleting Demo Site

Finally, to clean up the IBM Spectrum Protect Plus environment from the sample configuration, it is also a best practice to delete the Demo site. Select **System configuration** → **Site**, click the **X** that is next to the Demo site, and confirm the deletion. After the Demo site is deleted, the Sites view displays the remaining configured DC1, DC2, and DMZ sites, as shown in Figure 3-16.

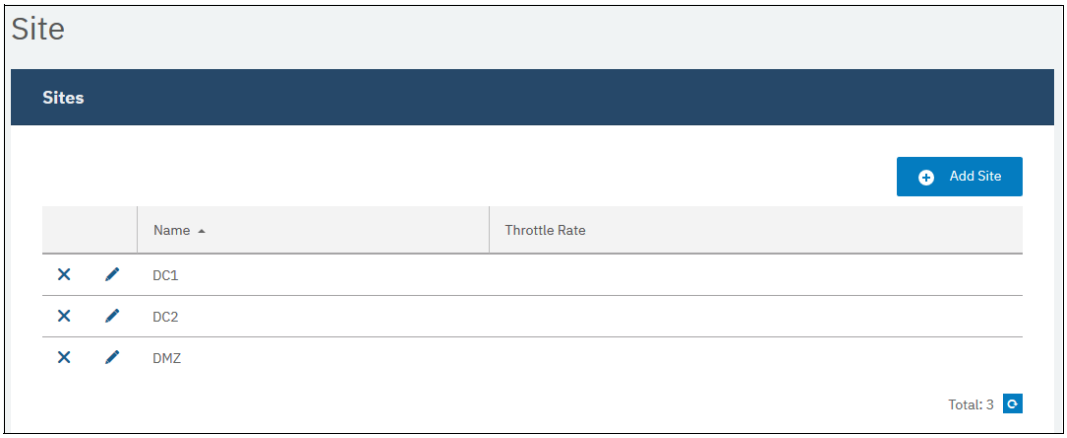

Figure 3-16 Sites - Demo site deleted

The initial configuration of the IBM Spectrum Protect Plus server is now complete.

3.4.2 Configuring the vSnap Backup Storage server

In 3.3.2, “Deploying the vSnap Backup Storage server” on page 55, the vSnap Backup Storage server is deployed and as a last step the VM is started.

Important: Give the VM several minutes after starting to initialize completely, especially for the first start after the deployment.

After the vSnap Backup Storage server is deployed, you must change the password for the local system user ID `serveradmin` before the VM for the vSnap server before can be used.

Changing serveradmin password

After a new deployment of the vSnap backup storage appliance, you *must* change the `serveradmin` password at the initial login.

Log in to the vSnap server with the SSH protocol by using a terminal or PuTTY client. Use the `serveradmin` user ID and change the password. It is possible to use the same password that was used for the `serveradmin` user at the IBM Spectrum Protect Plus server.

Note: In IBM Spectrum Protect Plus vSnap server version 10.1.6, the initial password of the user ID `serveradmin` is `sppDP758-SysXyz`.

Example 3-1 shows the change password dialogue. The new password for the `serveradmin` user ID must follow the following password rules:

- Minimum acceptable password length is 15 characters
- Minimum of:
  - 8 characters in the new password must not be present in the old password
  - 1 numerical digit in the new password
  - 1 uppercase character in the new password
  - 1 lowercase character in the new password
  - 1 other character in the new password
- Maximum of:
  - 3 identical consecutive characters are allowed in the new password
  - 4 identical consecutive class of character are allowed in the password

Example 3-1 Initial vSnap server login

```
login as: serveradmin
serveradmin@t3-spp-vsnap's password:
You are required to change your password immediately (root enforced)
------------------------------------------------------------------
Be sure to adhere to vSnap hardware and memory requirements
 as described in IBM Spectrum Protect Plus Blueprints
 accessible from the IBM Knowledge Center for IBM Spectrum Protect Plus.

------------------------------------------------------------------
WARNING: Your password has expired.
You must change your password now and login again!
Changing password for user serveradmin.
Changing password for serveradmin.
(current) UNIX password:
New password:
Retype new password:
```

After the password is changed, you must log in again by using the new password.

Creating vSnap server users (optional)

The vSnap server administrator optionally can decide to create more vSnap server admin users. To differentiate activities within an IBM Spectrum Protect Plus environment, it is better to create another user that we call `vsnapadmin` for the vSnap server with the `vsnap_admin` role by using the vSnap **`vsnap user create --username vsnapadmin --password PASSWORD`** command, as shown in Example 3-2. This user is used in a later step to connect the vSnap server and the IBM Spectrum Protect Plus server.

Example 3-2 Create vsnapadmin user

```
[serveradmin@t3-spp-vsnap ~]$ vsnap user create --username vsnapadmin --password
PASSWORD
vsnap user create --username vsnapadmin --password PASSWORD

UID: 1002
GID: 1002
NAME: vsnapadmin
ROLE: vsnap_admin

[serveradmin@t3-spp-vsnap ~]
```

Important: If a physical vSnap server is used for the role of VADP proxy, you must grant sudo privileges to the user that is created in Example 3-2. In this case, run the following commands:

```
cd /etc/sudoers.d
echo "vsnapadmin ALL=(ALL) NOPASSWD: ALL" > vsnapadmin
chmod 440 vsnapadmin
```

In our example, we initialize the vSnap servers by using the IBM Spectrum Protect Plus GUI, as described in 3.4.4, “Connect vSnap and VADP servers with IBM Spectrum Protect Plus server” on page 67. We also use the GUI to add virtual disks to the vSnap server.

An alternative is to initialize the vSnap server by using the CLI and create the vSnap pool by using the CLI. For large installations this approach is preferred.

The required CLI commands to create the vSnap server users are described in 5.4, “vSnap server CLI” on page 156. For more information, see the following resources:

- The “vSnap Server Installation and Setup” chapter of *IBM Spectrum Protect Plus Blueprints*
- IBM Knowledge Center

3.4.3 Configuring NTP for the IBM Spectrum Protect Plus and vSnap servers

To help avoid issues that can result from time zone differences, use a Network Time Protocol (NTP) server to synchronize time zones across IBM Spectrum Protect Plus resources in your environment. Such resources include the IBM Spectrum Protect Plus virtual appliance, storage arrays, hypervisors, and application servers.

If the time zones are out of sync, you might experience errors during application registration, metadata cataloging, inventory, backup or restore, or file restore jobs.

To synchronize time zones for IBM Spectrum Protect Plus server and the vSnap Backup Storage server, use the CLI on each of the servers. After the initial deployment, the time zone is set to Coordinated Universal Time.

Checking the current time zone setting

Run the **timedatectl** command to show the current settings, as shown in Example 3-3.

Example 3-3 CLI - timedatectl

```
[serveradmin@t3-spp-server ~]$ timedatectl
      Local time: Wed 2019-07-24 07:25:07 UTC
  Universal time: Wed 2019-07-24 07:25:07 UTC
        RTC time: Wed 2019-07-24 07:25:03
       Time zone: UTC (UTC, +0000)
     NTP enabled: yes
NTP synchronized: no
 RTC in local TZ: yes
      DST active: n/a

Warning: The system is configured to read the RTC time in the local time zone.
         This mode cannot be fully supported. It will create various problems
         with time zone changes and Daylight Saving Time adjustments. The RTC
         time is never updated; it relies on external facilities to maintain it.
         If possible, use RTC in UTC by calling
         'timedatectl set-local-rtc 0'.
[serveradmin@t3-spp-server ~]$ sudo timedatectl set-local-rtc 0
```

Reconfigure to read the RTC

The Real Time Clock (RTC) is not a good solution to spread time information across several servers.

Use the command **sudo timedatectl set-local-rtc 0** to disable using the RTC time.

Configure/set the time zone

As a next step the time zone need to be set using **timedatectl** command. Using this command avoids deleting /etc/localtime and creating a new link to it.

Use the **sudo timedatectl set-timezone Europe/Berlin** command to set the time zone in our example to Europe/Berlin.

Edit chrony configuration

To enable the NTP time server in your environment, change the server parameter in the **/etc/chrony.conf** chrony configuration file to a local reachable NTP server and disable the default server settings commenting out by adding the # character in front of the server name.

As shown in Example 3-4, we use the vi editor to show the activation of the time server **ntp.escc.workshop** and the deactivation of the default time servers by remarking them in the configuration file. (Example 3-4 shows the last lines of the configuration file.)

Run the **sudo vi /etc/chrony.conf** command to edit chrony configuration file.

Example 3-4 The /etc/chrony.conf file

```
...
# Serve time even if not synchronized to a time source.
#local stratum 10

# Specify file containing keys for NTP authentication.
#keyfile /etc/chrony.keys

# Specify directory for log files.
logdir /var/log/chrony

# Select which information is logged.
#log measurements statistics tracking
#server 0.north-america.pool.ntp.org
#server 1.north-america.pool.ntp.org
#server 2.north-america.pool.ntp.org
#server 3.north-america.pool.ntp.org
server ntp.escc.workshop
```

Restart the chrony daemon

After the change in the **/etc/chrony.conf** the daemon needs to be restarted.

Use the command **sudo systemctl restart chronyd** to restart the chrony daemon and check the chronyd status with the command **sudo systemctl status chronyd** as shown in Example 3-5.

Example 3-5 CLI - systemctl status chronyd

```
[serveradmin@t3-spp-server ~]$ sudo systemctl restart chronyd
[serveradmin@t3-spp-server ~]$ sudo systemctl status chronyd
? chronyd.service - NTP client/server
   Loaded: loaded (/usr/lib/systemd/system/chronyd.service; enabled; vendor
preset: enabled)
   Active: active (running) since Wed 2019-07-24 10:37:23 CEST; 5s ago
     Docs: man:chronyd(8)
           man:chrony.conf(5)
  Process: 6689 ExecStartPost=/usr/libexec/chrony-helper update-daemon
(code=exited, status=0/SUCCESS)
  Process: 6685 ExecStart=/usr/sbin/chronyd $OPTIONS (code=exited,
status=0/SUCCESS)
 Main PID: 6687 (chronyd)
   CGroup: /system.slice/chronyd.service
           ··6687 /usr/sbin/chronyd
```

```
Jul 24 10:37:23 t3-spp-server systemd[1]: Starting NTP client/server...
Jul 24 10:37:23 t3-spp-server chronyd[6687]: chronyd version 3.2 starting (+CMDMON
+NTP +REFCLOCK +RTC +PRIVDROP...EBUG)
Jul 24 10:37:23 t3-spp-server chronyd[6687]: Frequency -20.168 +/- 46.716 ppm read
from /var/lib/chrony/drift
Jul 24 10:37:23 t3-spp-server systemd[1]: Started NTP client/server.
Hint: Some lines were ellipsized, use -l to show in full.
```

Check the active time server

To check the current active time server in use, use the command **`chronyc sources`** as shown in Example 3-6.

Example 3-6 chronyc sources

```
[serveradmin@t3-spp-server ~]$ chronyc sources
210 Number of sources = 1
MS Name/IP address         Stratum Poll Reach LastRx Last sample
===============================================================================
^? ntp.escc.workshop      1   6     3     2   -20.0s[ -20.0s] +/-  10.6s
[serveradmin@t3-spp-server ~]$
```

3.4.4 Connect vSnap and VADP servers with IBM Spectrum Protect Plus server

After the deployment of the IBM Spectrum Protect Plus server, the required vSnap Backup Storage servers and the VADP proxies needs to be connected and enabled in IBM Spectrum Protect Plus.

Add identities to the environment

Credentials are required to connect a vSnap server and a VADP proxy resource in IBM Spectrum Protect Plus.

For the vSnap server, we use (optionally) the created `vsnapadmin` account (see “Creating vSnap server users (optional)” on page 64); for the registration VADP proxy, we use the local system account `serveradmin` (see “Changing serveradmin password” on page 63). Both credentials are added as identities, as described in “Identities, keys, and certificates” on page 11.

The credentials are added to IBM Spectrum Protect Plus and then are referred to as “Use existing user” when the resource is added. The Identity view under **Accounts** → **Identity** lists both identities after they are added, as shown in Figure 3-17 on page 68.

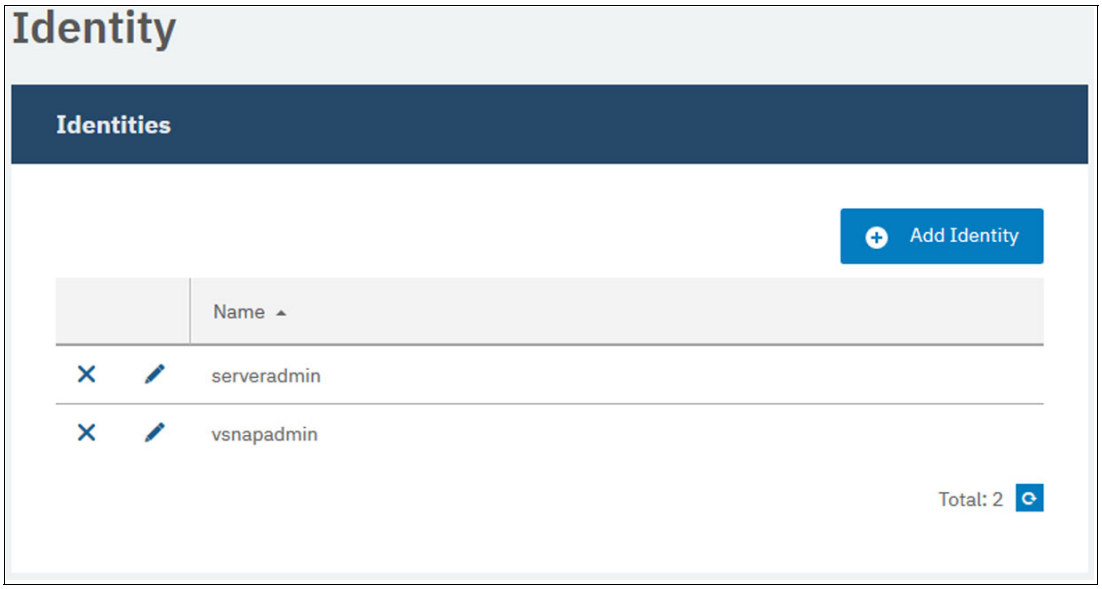

Figure 3-17 Identity view

Installing VADP proxy

The VADP proxies that are installed together with the vSnap Backup Storage servers in the IBM Spectrum Protect Plus GUI are listed under **System Configuration** → **VADP Proxy** listed. After initial deployment, you can see that only the default localhost is shown, which is at the moment in suspended state as described in “Suspending the local built-in VADP proxy” on page 61.

Register the VADP proxy by clicking **Register Proxy** (see Figure 3-18).

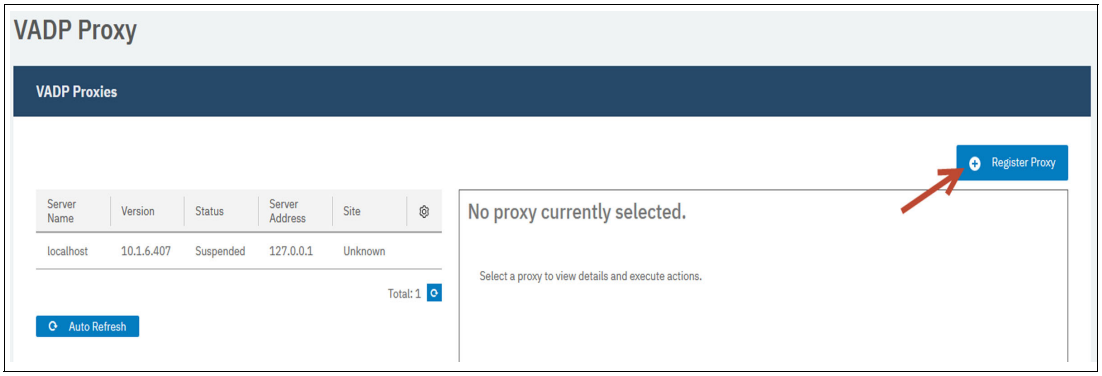

Figure 3-18 Register VADP proxy

Complete the following steps to install and add the VADP proxy to the suitable site:

1. Enter the host name or the IP address of the vSnap server where the VADP proxy is installed.
2. Select a site from the list for the installation of the VADP proxy.
3. Use the `serveradmin` admin credentials from predefined identities, as shown in Figure 3-19 on page 69.

   The second option is to enter the `serveradmin` credentials with the password.
4. Click **Install** (see Figure 3-19 on page 69).

Figure 3-19 Install VADP proxy

After the installation process is complete, an installed VADP proxy is enabled automatically, as shown in Figure 3-20.

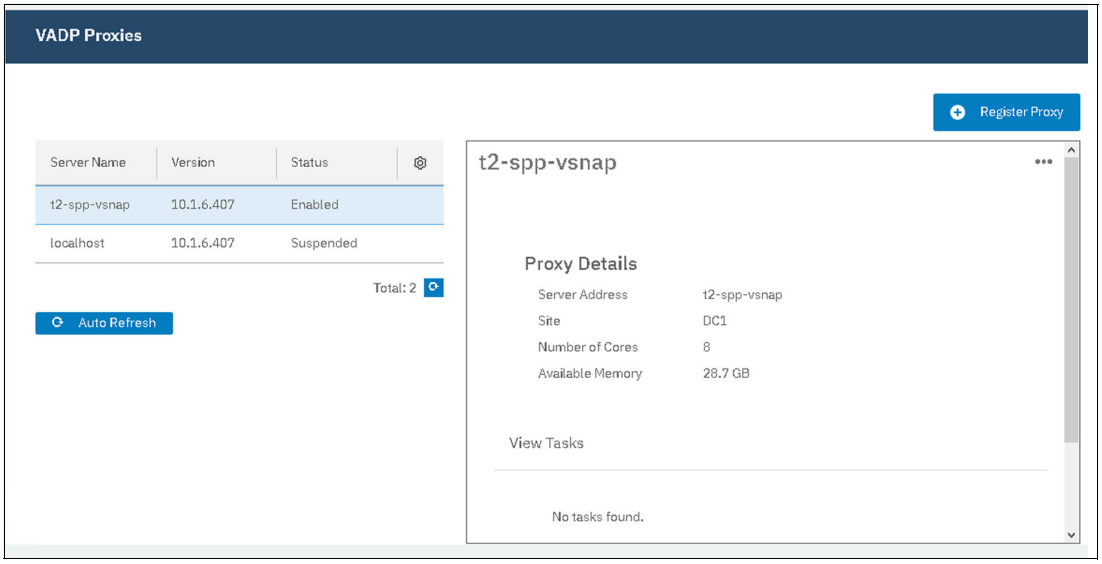

Figure 3-20 Installed VADP proxy

Configuring VADP proxy settings

After enabling all VADP proxies, you must adjust their configuration according to the output of the sizing tool and the recommendations from *IBM Spectrum Protect Plus Blueprints*.

The configuration menu of the VADP proxy (see Figure 3-21) is entered by clicking the three dots in the right corner of the VADP Proxy view (see Figure 3-20 on page 69) and selecting **Proxy Options**.

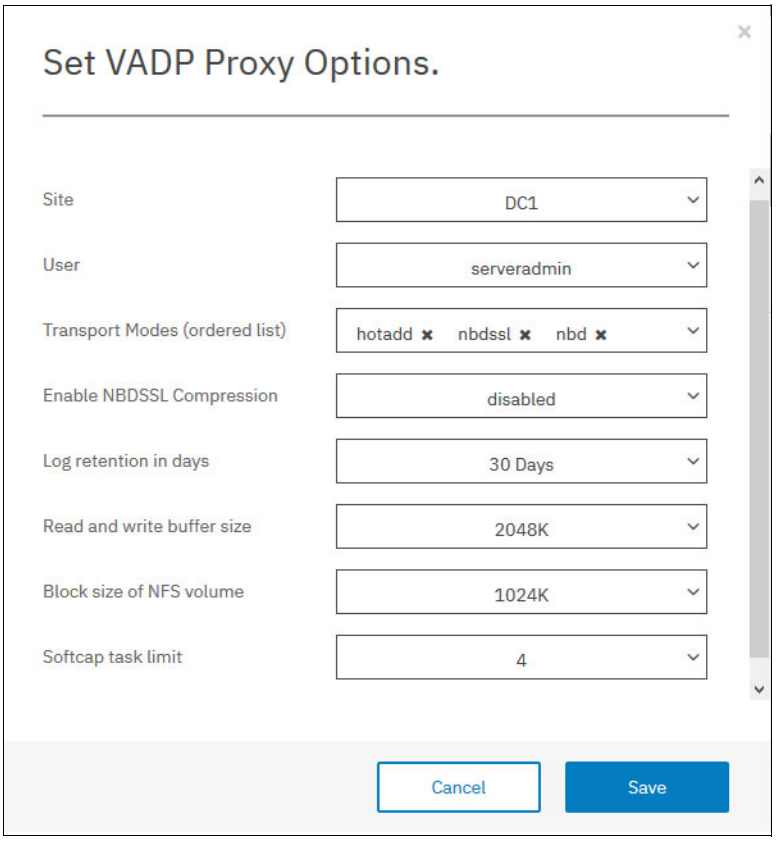

Figure 3-21 Set VADP Proxy Options

In our example that is shown in Figure 3-21, we adjusted the following options:

- Site: DC1 (the site for this VADP proxy)
- User: serveradmin
- Transport Modes: We disabled SAN transport because it is a mode that cannot be used on a virtual VADP proxy. After all initial full backups are complete, we consider disabling HotAdd as well.
- Softcap task limit: We set it to **4** as recommended in the Blueprints for a smaller virtual VADP proxy.

We kept NBDSSL Compression as **disabled** because we use a 10 Gb network and the other settings remain unchanged.

Repeat this configuration task on the remaining two VADP proxies.

Registering vSnap Backup Storage server

Next, you must register your vSnap server as disk storage by using the IBM Spectrum Protect Plus GUI.

In the IBM Spectrum Protect Plus GUI, select **System configuration** → **Backup Storage** → **Disk** and then, click **Add Disk Storage**, as shown in Figure 3-22.

Figure 3-22 Add Disk Storage

Enter the hostname of the vSnap Backup Storage server, select the site, and choose the user ID **`vsnapadmin`** to add the disk storage. Click **Save**, as shown in Figure 3-23.

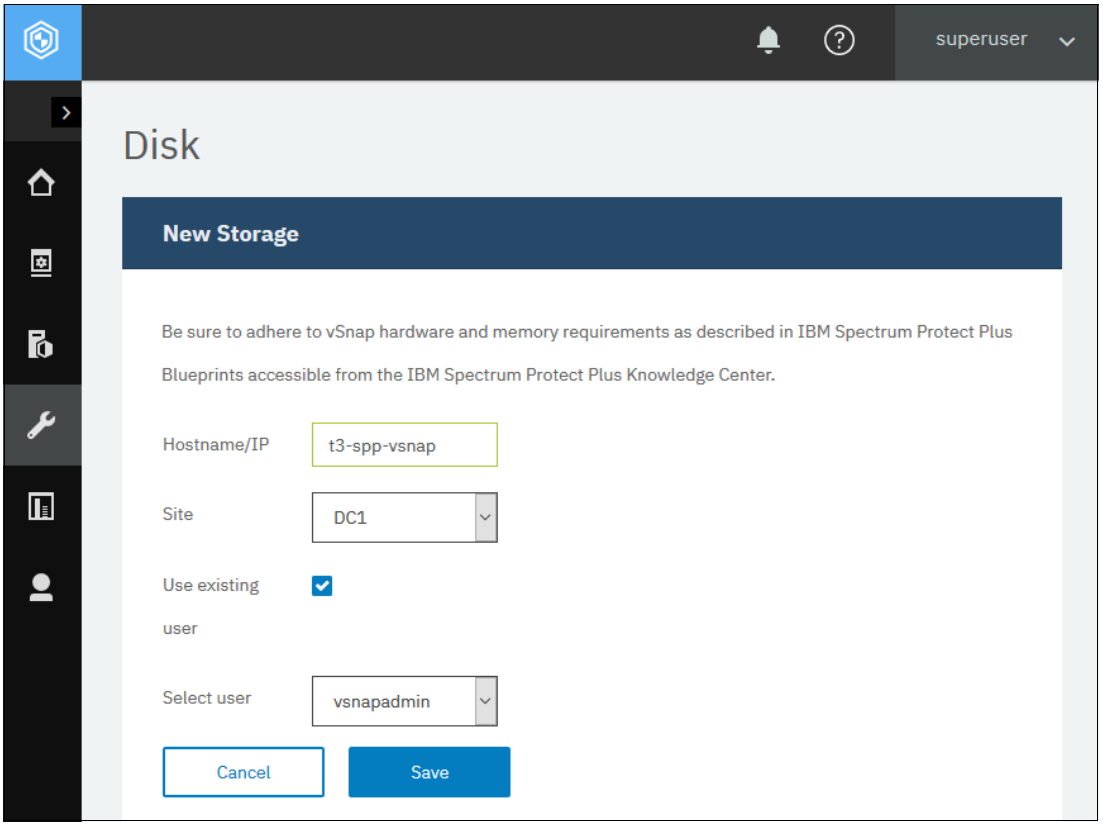

Figure 3-23 Register vSnap server

Repeat this step for the remaining two vSnap servers.

Initializing vSnap pool

After a vSnap server is deployed as a virtual appliance and added to IBM Spectrum Protect Plus, a pool must be configured. In this section, we describe how to initialize the vSnap pool.

Important: To Initialize pool is a one-time decision. The decision if encryption is used must be made before you start the initializing pool. This choice *cannot* be reversed.

To use the simple initialization that is available in the IBM Spectrum Protect Plus GUI, select **System Configuration** → **Backup Storage** → **Disk** in the GUI and click the three dots for your vSnap server. Select **Initialize** or **Initialize With Encryption** from the pull-down menu, as shown in Figure 3-24.

Figure 3-24 Initialize disk with or without encryption

For more information about encryption, see 2.3.5, “Encryption” on page 39.

After the simple initialization process is completed, the Status/Capacity column shows a utilization bar for the vSnap server, as shown in Figure 3-25.

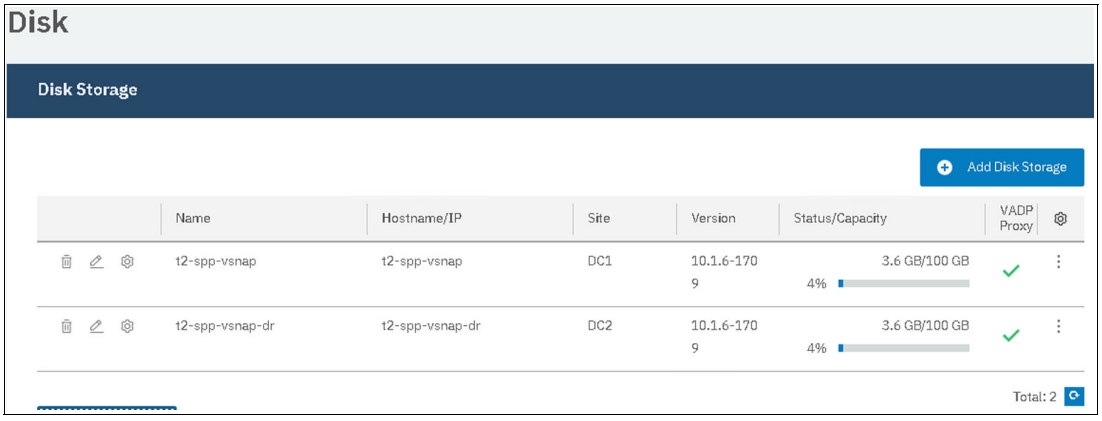

Figure 3-25 vSnap server view with all server initialized

Configuring vSnap backup storage server

In this step, adjust the configuration of the vSnap server. To enter the configuration, click the settings icon that is in front of the vSnap server hostname. The Manage Backup Storage window opens, as shown in Figure 3-26.

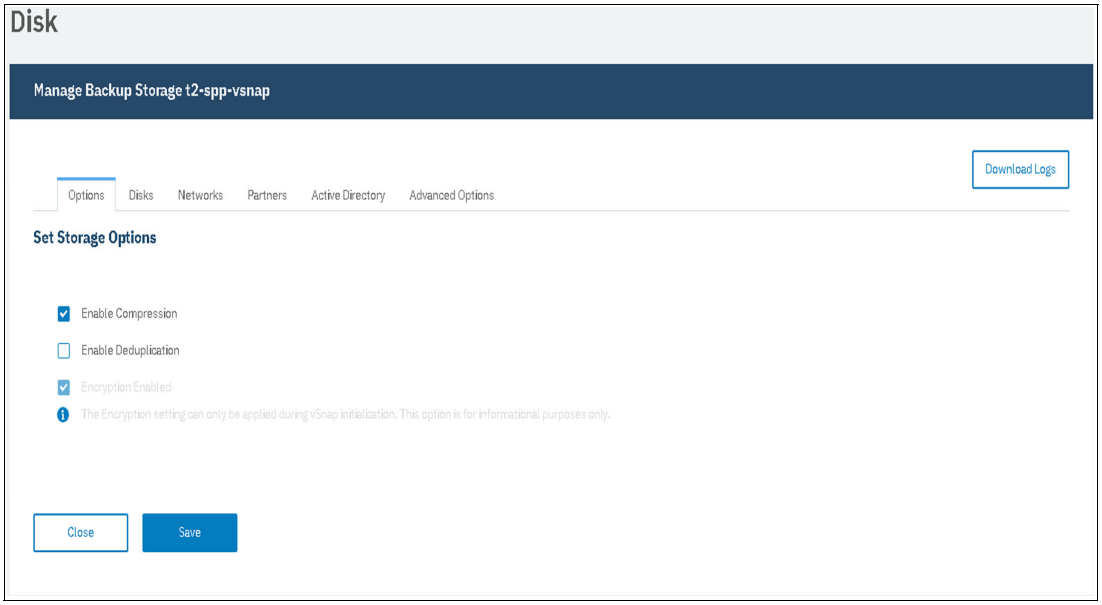

Figure 3-26 Manage Backup Storage - Configure the vSnap Backup Storage server

In the Manage Backup Storage window, the following tabs are available for configuring the details of the vSnap server:

- Set Storage Options

  By default, compression is enabled and deduplication is disabled. Enable deduplication if you planned for the deduplication from a CPU and memory perspective. Both options can be turned on and off later as well.

  If encryption was enabled during initialization, it is shown here. It is a one-time decision and cannot be turned on or off later.

- Disk Storage Options

  New disks can be added to back up storage here. For more information about how to add virtual disk (vDisks) to the vSnap servers, see IBM Knowledge Center.

  Newly attached disks to the vSnap server are available for use. Select the disk and click **Save**, as shown on Figure 3-27 on page 74.

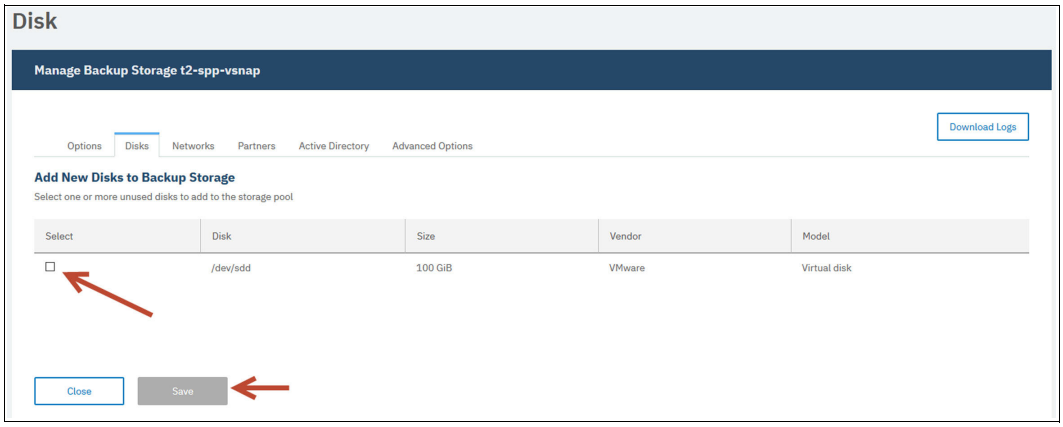

Figure 3-27 Add New Disk to the Backup Storage

- Network Interface Controllers Options

  This tab provides an overview of the network interface controllers (NICs) that are enabled and configured on the vSnap server. By default, both NICs options are selected for the selected workload:

  – Backup and Management workload
  – Replication Workload between vSnap servers

  The NICs can be enabled or disabled by selecting the Configuration button. It is recommended to use separate network interfaces for different workload purposes, as shown in Figure 3-28.

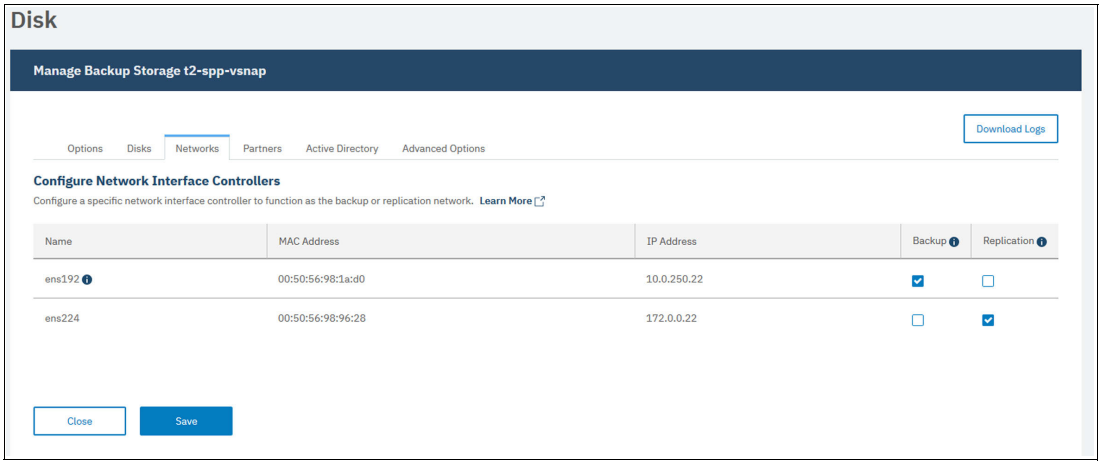

Figure 3-28 Configure Network Interface Controllers

- Configure Storage Partners

  A Storage Partner must be configured if you want to replicate backup data between two vSnap servers (primary or auxiliary storage). This configuration is discussed in 13.3, “Replication of backup data - Details” on page 211. The best practice is to set up replication protection between the vSnap server to protect primary data and the configuration for the disaster case scenario. The detailed configuration of the Storage Partners is shown in Figure 3-29 on page 75.

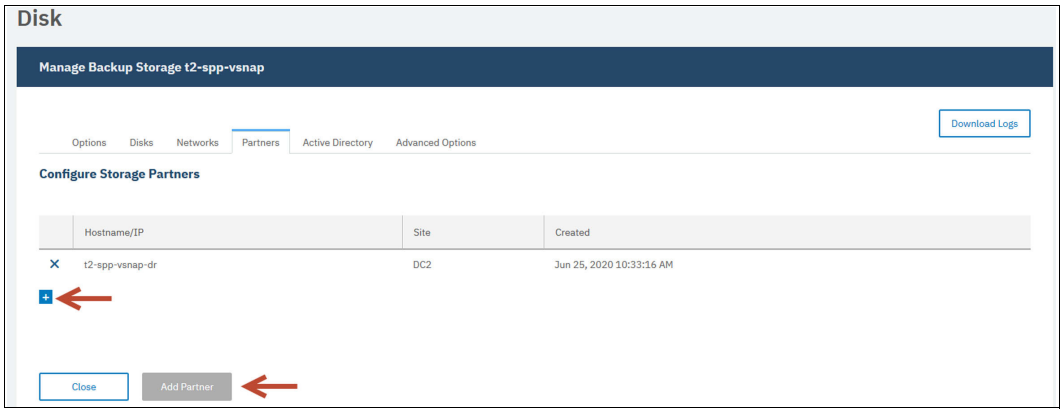

Figure 3-29 Configure (vSnap server) Storage Partners

- Active Directory Options

  The Active directory view provides an option to associate backup storage to the active directory domain. The benefit of this option is that after storage is part of the active directory domain, Microsoft SQL Server log backup jobs use domain authentication.

  The configuration also allows you to skip requirements for the local staging disk area for log backup operation.

- Advanced Configuration Options

  The number of concurrent backup streams is set to unlimited. Other options are to pause the backup stream or to limit it to a configurable number of streams, as shown in Figure 3-30.

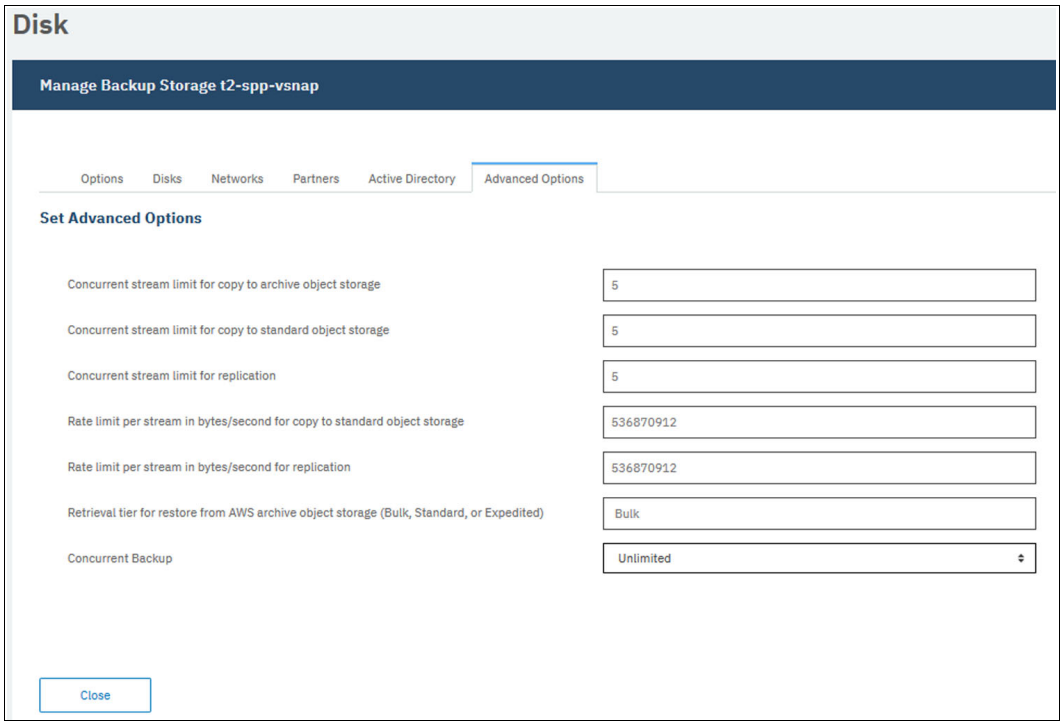

Figure 3-30 Advanced configuration options

In the right upper corner, click **Download Logs** to start downloading a vSnap server log file.

3.4.5 Adding an SLA for IBM Spectrum Protect Plus catalog backup

To recover the IBM Spectrum Protect Plus environment if a failure or disaster occurs, it is essential to secure the catalog by regularly backing up to a different location.

In the IBM Spectrum Protect Plus GUI, select **Manage Protection** → **Policy Overview** and then, click **Add SLA Policy**. Name the new SLA policy `SPP-DB-Backup`.

We recommend using the following settings, as shown in Figure 3-31:

- Retention time: 5 days
- Frequency: 12 hours to run it twice a day
- Start time: 6:00, which dictates that a schedule is started at 6 AM and 6 PM
- Target site: DC1

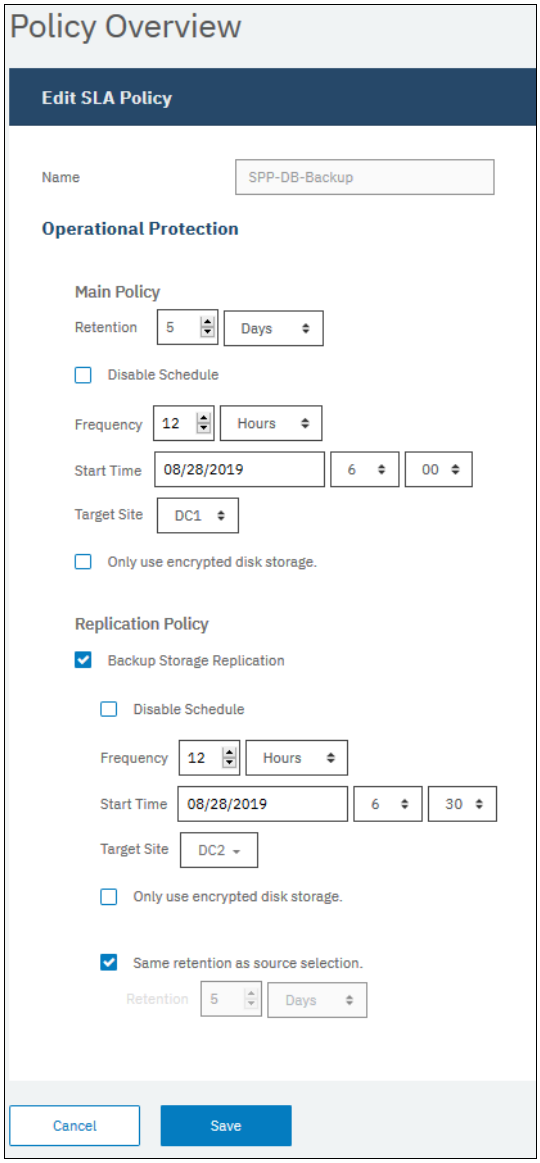

Figure 3-31 New SLA Policy for IBM Spectrum Protect Plus Catalog Backup

To protect the catalog backup, we configure replication to another vSnap server by using the site DC2, with the following settings:

- Frequency: 12 hours to run it twice a day
- Start time: 6:30, which dictates that a schedule is started at 6:30 AM and 6:30 PM
- Target Site: DC2
- Same retention is used as on source (5 days)

Note: Replication between two vSnap servers requires a two-step configuration: one step is the SLA configuration, which was just done, and a second step is to set the relationship between both vSnap servers. For more information about setting up replication, see 15.3, “Replicating backup data” on page 409.

The section Additional Protection is kept cleared. Click **Save** to create the SLA policy.

Assigning an SLA Policy for the catalog backup

To assign the SLA policy `SPP-DB-Backup` for the IBM Spectrum Protect Plus catalogs, select **Manage Protection** → **IBM Spectrum Protect Plus** → **Backup** and then, select **SPP-DB-Backup** in the SLA Policy section, as shown in Figure 3-32. Click **Save**.

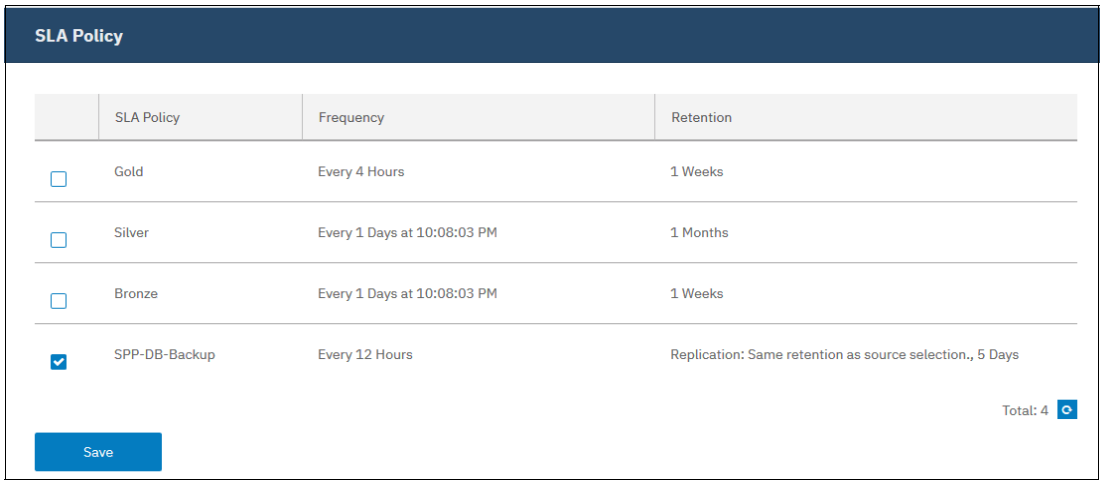

Figure 3-32 Select SLA Policy for IBM Spectrum Protect Plus Catalog Backup

3.4.6 Backing up the vSnap server system configuration

A vSnap server consists of two sets of data:

- The vSnap pool (or storage pool), which is the logical organization of disks into a pool of storage that is by the vSnap server component.
- Configuration and metadata information, which is in the `/etc/vsnap` directory of the vSnap server.

You can back up the configuration and metadata information for use cases where the vSnap pool data is intact and valid but the configuration or metadata information is lost or not available. This issue can occur in the following situations:

- The vSnap server compute environment is lost, but the storage is not. An example is a vSnap server that is running as a VM and the storage that is backing the vSnap pool is on a physical RDM (pRDM) disk. In this case, the vSnap server VM is lost, but the data on the pRDM disk is still valid.

- The vSnap server compute environment must be changed, but the storage does not. An example is a vSnap server that is running in a VM and the storage that is backing the vSnap pool is on a pRDM disk. In this case, the vSnap server must be rebuilt or re-created on a new VM, but the data on the pRDM disk is still valid.

The backup procedure is based on the **vsnap system config backup** command, which creates a compressed TAR file. The resulting file can then be securely copied to a central location, such as the IBM Spectrum Protect Plus server. To back up the vSnap server configuration, run the command that is shown in Example 3-7.

Example 3-7 Command to create a vSnap server config file backup

```
vsnap system config backup --outfile backup.tgz
```

For reliability and convenience, it is suggested to automate the backup of each vSnap server's config files and keep a set of historical backups on a centralized location, such as the IBM Spectrum Protect Plus server itself. Therefore, we demonstrate how to achieve this with some simple steps and a script that is creating backups and offloading the backup file to the SPP server.

Complete the following steps on the IBM Spectrum Protect Plus server (see Example 3-8):

1. At the Linux command line, create a user **vsnapadmin** if it does not exist.
2. Create a directory structure that stores the backup files of all vSnap servers.
3. Change permission of the new directory structure, and set the group and ownership to the `vsnapadmin` user.

Example 3-8 Commands to run on the SPP Server

```
sudo su
vsnap user create --username vsnapadmin --password pass4VsnapAdmin
mkdir /var/opt/spp/
mkdir /var/opt/spp/backup
chown vsnapadmin:vsnapadmin /var/opt/spp -R
```

Complete the following steps on each vSnap Server, as shown in Example 3-9 on page 79:

1. At the Linux command line, create a user `vsnapadmin` if it does not exist.
2. Create a directory structure that stores the backup files of the local vSnap server.
3. Change permission of the new directory structure and set the group and ownership to the `vsnapadmin` user.
4. Create an SSH key pair to enable password-less authentication and communication between the vSnap server and the IBM Spectrum Protect Plus server.
5. Transfer the keys to the IBM Spectrum Protect Plus server.
6. Create a shell script that generates the backups and transfers the files from the vSnap server to the IBM Spectrum Protect Plus server.
7. Create a crontab entry to schedule a repeating backup job.

Example 3-9 Commands to run on the vSnap server as root user

```
sudo su
vsnap user create --username vsnapadmin --password pass4VsnapAdmin
mkdir /var/opt/spp/
mkdir /var/opt/spp/backup
chown vsnapadmin:vsnapadmin /var/opt/spp -R
exit
```

8. Log in to the vSnap server as the `vsnapadmin` user and create the SSH keys, as shown in Example 3-10. The key pair is stored under the user's home directory.

Example 3-10 Create ssh keys as user vsnapadmin

```
ssh-keygen

Generating public/private rsa key pair.
Enter file in which to save the key (/home/vsnapadmin/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/vsnapadmin/.ssh/id_rsa.
Your public key has been saved in /home/vsnapadmin/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:mY4hr6ceO9XGegAXeCyjkE7koBxrOXvdhm/AJNy35cg vsnapadmin@t4-spp-vsnap
The key's randomart image is:
+---[RSA 2048]----+
|o+   o           |
|B.= = +          |
|+O + * o .       |
|..+ * * =o       |
| . ..B.ES.       |
|  .  o*++        |
|    ..o*.        |
|    .++ .        |
|   .=+ .         |
+----[SHA256]-----+
```

9. Transfer the public key of the `vsnapadmin` user to the SPP server and test the connectivity by querying the server's hostname (see Example 3-11, Example 3-12 on page 80, and Example 3-13 on page 80).

Example 3-11 The ssh-copy-id command

```
[vsnapadmin@t7-spp-vsnap ~]$ ssh-copy-id -i ~/.ssh/id_rsa.pub vsnapadmin@t4-spp-server
/usr/bin/ssh-copy-id: INFO: Source of key(s) to be installed: "/home/vsnapadmin/.ssh/id_rsa.pub"
The authenticity of host 't4-spp-server (10.0.250.41)' can't be established.
ECDSA key fingerprint is SHA256:YHzaTntje1zi9Hm3wd6JIb+7+UC6Q/+sIV/9hm3nqGA.
ECDSA key fingerprint is MD5:b6:cf:ab:0b:6O:6f:68:15:d6:62:4c:9e:78:e5:5f:44.
Are you sure you want to continue connecting (yes/no)? yes
/usr/bin/ssh-copy-id: INFO: attempting to log in with the new key(s), to filter out any that are already installed
/usr/bin/ssh-copy-id: INFO: 1 key(s) remain to be installed -- if you are prompted now it is to install the new
keys
vsnapadmin@t4-spp-server's password:

Number of key(s) added: 1

Now try logging into the machine, with:   "ssh 'vsnapadmin@t4-spp-server'" and check to make sure that only the
key(s) you wanted were added.
```

Example 3-12 Create the sample backup script /home/vsnapadmin/vsnap_backup.sh

```
#!/bin/bash

keepFilesCnt=5
bakFileName="/var/opt/spp/backup/${HOSTNAME}_backup.$(date +%Y%m%d-%H%M%S).tgz"

echo
echo "creating new vsnap: $bakFileName"
vsnap system config backup --outfile $bakFileName

bakDirSrc="/var/opt/spp/backup"
bakDirTgt="/var/opt/spp/backup"
username="vsnapadmin"
sppServerName="t4-spp-server"

echo "copying backup file to remote location"
scp $bakFileName $username@${sppServerName}:${bakDirTgt}

echo
echo "up to $keepFilesCnt newest files on server: <<< $sppServerName >>> for vsnap: <<< $HOSTNAME >>> "
echo
ssh $username@$sppServerName "ls -t ${bakDirTgt}/${HOSTNAME}* | awk 'NR<=cnt' cnt=$keepFilesCnt"

#ssh $username@$sppServerName "ls -t ${bakDirTgt}/${HOSTNAME}* | awk 'NR>cnt' cnt=$keepFilesCnt | xargs rm"


totalBackupFiles=`ssh $username@$sppServerName "ls -1 ${bakDirTgt}/*backup* | wc -l"`
echo "total backups on spp server: $totalBackupFiles"
totalBackupFilesThisVsnap=`ssh $username@$sppServerName "ls -1 ${bakDirTgt}/${HOSTNAME}* | wc -l"`
echo "total backups on spp server from vSnap $HOSTNAME: $totalBackupFilesThisVsnap"

echo
spaceUsed=`ssh $username@$sppServerName "du -h -s $bakDirTgt"`
echo "space occupied / available on server's ($sppServerName) backup directory:  $spaceUsed"
echo
ssh $username@$sppServerName "df -h -P $bakDirTgt"
```

Example 3-13 Change the script's permissions and make it executable

```
chmod 755 /home/vsnapadmin/vsnap_backup.sh
```

The script can be run manually by user `vsnapadmin` to create backups and to test the functions. Example 3-14 and Example 3-15 show the output of a manual script run.

Example 3-14 Run vsnap_backup.sh manually

```
[vsnapadmin@t4-spp-vsnap backup]$ /home/vsnapadmin/vsnap_backup.sh

creating new vsnap: /var/opt/spp/backup/t4-spp-vsnap_backup.20190820-093827.tgz
backup created: /var/opt/spp/backup/t4-spp-vsnap_backup.20190820-093827.tgz
copying backup file to remote location
t4-spp-vsnap_backup.20190820-093827.tgz 100%   14KB   9.3MB/s   00:00

up to 5 newest files on server: <<< t4-spp-server >>> for vsnap: <<< t4-spp-vsnap >>>

/var/opt/spp/backup/t4-spp-vsnap_backup.20190820-093827.tgz
/var/opt/spp/backup/t4-spp-vsnap_backup.20190820-093429.tgz
/var/opt/spp/backup/t4-spp-vsnap_backup.20190820-093104.tgz
/var/opt/spp/backup/t4-spp-vsnap_backup.20190820-093041.tgz
/var/opt/spp/backup/t4-spp-vsnap_backup.20190820-093020.tgz


total backups on spp server: 16
total backups on spp server from vsnap t4-spp-vsnap: 14

space occupied / available on server's (t4-spp-server) backup directory:  260K  /var/opt/spp/backup

Filesystem              Size  Used Avail Use% Mounted on
/dev/mapper/centos-root   26G  5.5G   20G  22% /
```

Example 3-15 Crontab entry on vSnap server to create a backup every night at 23:30 / 11:30 pm

```
crontab -e
30 23 * * * /home/vsnapadmin/vsnap_backup.sh
```

Note: We disabled a file version housekeeping, which deletes old files on the local vSnap server backup directory or on the SPP Server. The reason for this change is that the backup files have a limited size (with only a few KB) and we leave it up to the administrator to decide whether only the most recent files or a longer history of backup files are kept.

The script can be easily modified by adding or uncommenting the single line in Example 3-16 to remove all files but the last `X` files for a specific vSnap server, where `X` is specified with the variable `keepFilesCnt`.

Example 3-16 Remove older backup files from SPP server (optional)

```
ssh $username@$sppServerName "ls -t ${bakDirTgt}/${HOSTNAME}* | awk 'NR>cnt' cnt=$keepFilesCnt | xargs rm"
```

3.4.7 Changing and verifying the schedules of the predefined jobs

After the deployment, the schedules for the predefined Inventory and Maintenance jobs are set to default values, which we recommend changing to satisfy best practices.

The following three schedules must be updated:

- Change the schedule for the maintenance job to run at 13:00.
- Change the schedule for the Storage Server Inventory to 19:00.
- Change the schedule for the Hypervisor/Application Server Inventory to 19:30 after you add the resources as described in Chapter 6, “Backing up and restoring virtualized systems” on page 171, and Chapter 8, “Backing up and restoring databases” on page 255.

Changing the maintenance job schedule

In **Jobs and Operations**, choose **Schedule** from the available tabs, and select **Maintenance**, as shown in Figure 3-33.

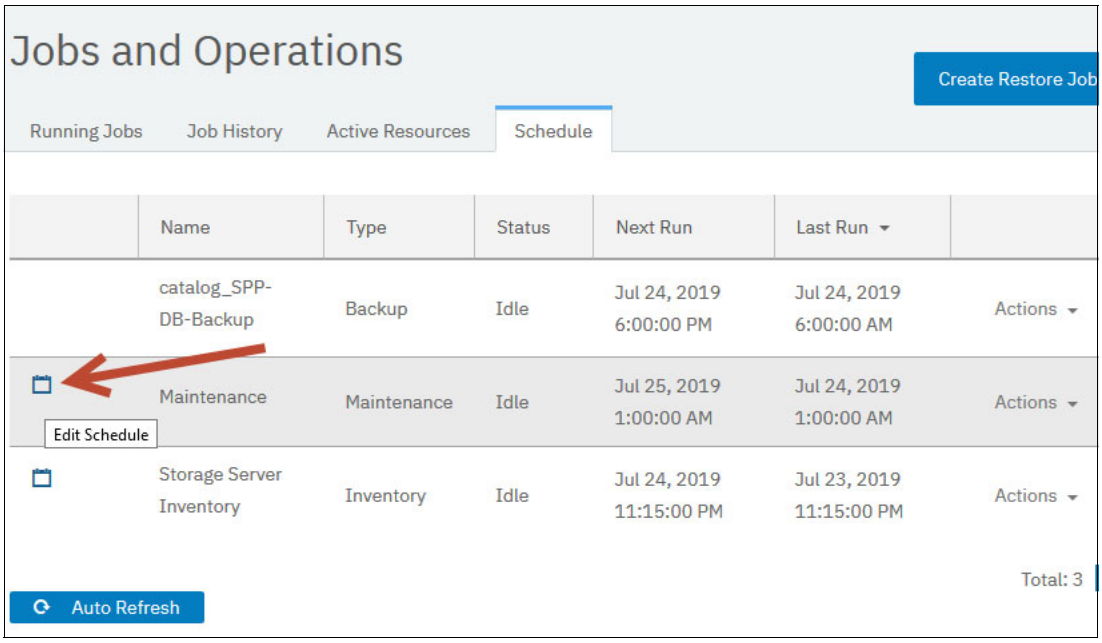

Figure 3-33 Edit Schedule for Maintenance Job

Select to run it daily at 1 PM (13:00) and click **Save**, as shown in Figure 3-34.

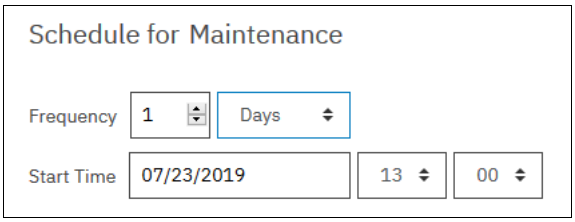

Figure 3-34 Schedule for Maintenance job

Changing the inventory job schedule

Change the schedule for Storage Server **Inventory** job to start at 7:00 PM (19:00) and to a frequency of daily, as described in “Changing the maintenance job schedule” on page 81.

Changing the Hypervisor and Application Server Inventory schedule

Change the schedule for **Hypervisor** and **Application Server Inventory** to start at 7:30 PM (19:30) and to a frequency of daily the same way, as described in “Changing the maintenance job schedule” on page 81.

4

Networking

The network is a crucial component of an IBM Spectrum Protect Plus implementation because it establishes communication paths between the different components of IBM Spectrum Protect Plus, the backup clients, and users.

In larger enterprise environments, it is common practice to implement dedicated networks for different purposes, such as to isolate backup, application, and management networks from each other. Having a dedicated network for backup and restore workloads brings the following advantages:

- Performance benefits: Data protection systems usually handle large amounts of data. Separating backup and restore traffic from application or management data ensures that the performance or the availability of an application is not impacted by running backup or restore jobs.
- Security benefits: If an enterprise is hit by a hacker or malware attack, the data protection systems are often your last line of defense where isolated network segments improves the security of the mission critical backups.

This chapter gives an introduction to Networking with IBM Spectrum Protect Plus, discusses the aspects that need to be considered in the planning phase, and shows the configuration steps based on an example.

This chapter includes the following topics:

4.1 IBM Spectrum Protect Plus networking

Depending on the design of the existing infrastructure, the number and purpose of network segments might vary. Possible configurations include:

- Separating backup traffic to a dedicated network
- Separating replication traffic to a dedicated network with a route to a remote site
- Separating vSnap servers to different purpose network segments; for example, DMZ
- Having all components in a single network (for example, for small test environments)

Note: In a heterogeneous network environment, you find a combination or mixture of these examples.

Figure 4-1 shows a logical network diagram for a two-site environment, which uses a dedicated backup network in the main location to isolate backup and restore workloads from application, user, admin, and replication traffic.

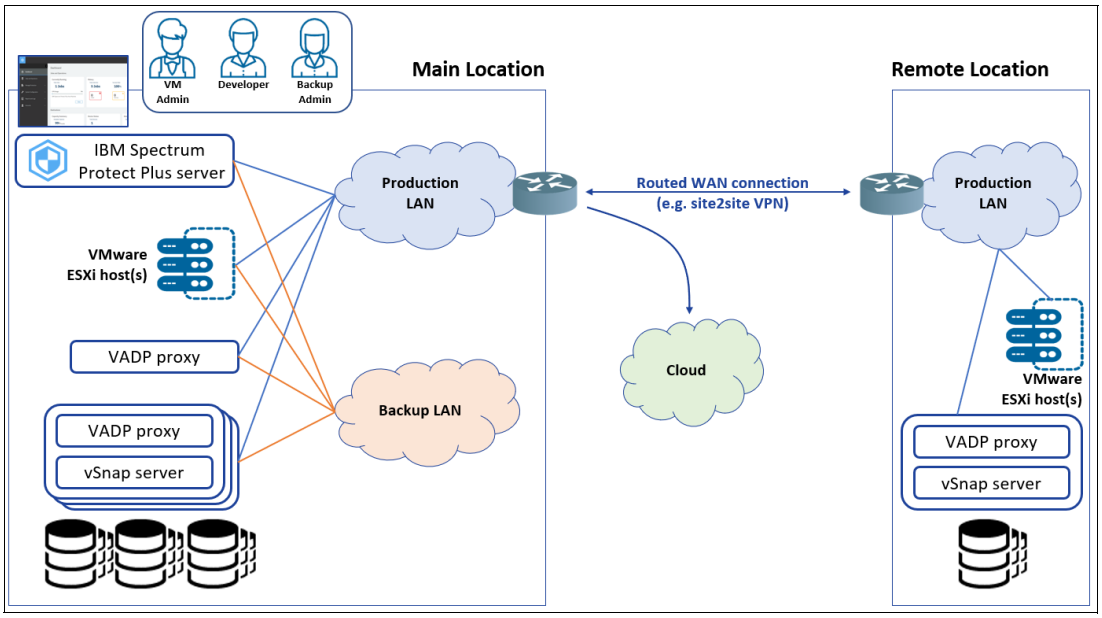

Figure 4-1 IBM Spectrum Protect Plus Networking Overview

When planning the network for an IBM Spectrum Protect Plus environment, consider the following points:

- Which protocols should be separated from each other? IBM Spectrum Protect Plus differentiates between management, backup/restore, and replication workloads.

  Although vSnap servers and VADP proxies are located close to the backup clients, admins, replication partners, and targets for extra copies (IBM Spectrum Protect or a Cloud storage provider) can be remote.

- Back up and restore operations usually require direct data flows between a backup client or hypervisor and the IBM Spectrum Protect vSnap server. If backup traffic is confined to a dedicated network, the backup client or hypervisor must be attached to this network by using a separate adapter with an assigned IP address.

- Are all components connected to a “flat” layer-2 network, or are there any routed connections in between? If components are communicating over routed connections, static routes might need to be maintained.

- Are firewalls in the communication paths? Different IBM Spectrum Protect Plus components use specific TCP ports to communicate with other components; therefore, these ports must be enabled on the firewalls before starting the implementation. These firewalls can be network components that are interconnecting different VLANs or part of an operating system (for example, on a hypervisor or application backup client).

4.2 Understanding network data flows

IBM Spectrum Protect Plus can differentiate between management, backup/restore, and replication traffic. The component that controls which VLAN is being used for which kind of workload is the vSnap server. Although management traffic is always enabled on all network interfaces, backup/restore and replication traffic can be enabled or disabled on a specific interface.

In VMware environments, another component, the VADP proxy, is used to move virtual machine (VM) backup data between the hypervisor (ESXi host) and the vSnap server storage.

The vSnap server controls which interface is used for backup traffic that flows between the VADP proxy and the vSnap server (VADP backend). However, the vSnap server cannot control which interface or VLAN is used for backup data that flows between the VADP proxy and the hypervisor (VADP front end). The path that backup data uses between the ESXi hypervisor and VADP proxy depends on various characteristics of the network, such as the Virtual Disk Transport method that is used, or the DNS name resolution.

Figure 4-2 shows the possible data flows between a VMware hypervisor, a VADP proxy, and a vSnap server in an environment with a dedicated network for backup data.

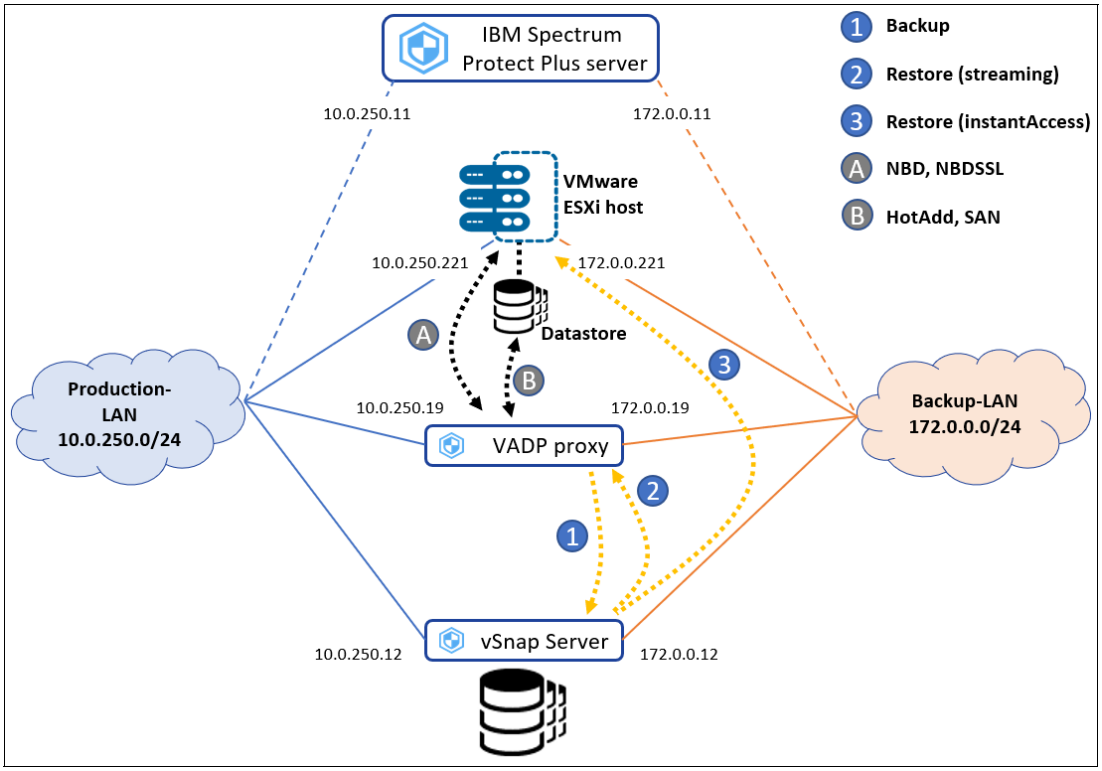

Figure 4-2 Backup data flows in a VMware environment with separate backup network

The data flows can be divided into two areas:

- Between VADP proxy and vSnap server (VADP backend)
- Between hypervisor and VADP proxy server (VADP front end)

4.2.1 VADP backend data flow

Assume that the VADP proxy and the vSnap server are connected to the Production LAN, but also to a dedicated backup LAN and the vSnap server forces backup/restore traffic to use the backup LAN interface. Consider the following points:

- VM backups traverse the backup LAN when the VADP proxy sends data to the vSnap server (Number 1 in Figure 4-2 on page 85). If the VADP proxy and the vSnap server component are on the same Operating System, no traffic appears on the backup adapter.
- VM restores traverse the backup LAN from the vSnap server to the VADP proxy when a STREAMING restore is performed, which is the default restore method (Number 2 in Figure 4-2 on page 85). If the VADP proxy and the vSnap server component are on the same Operating System, no traffic appears on the backup adapter.
- InstantAccess restores (test restores or file restores) do not support the streaming restore through the VADP proxy, but also traverse the backup LAN when the vSnap server sends data directly to the ESXi hypervisor (Number 3 in Figure 4-2 on page 85).

Important: To support direct communication between the vSnap server and the ESXi hypervisor by way of the backup LAN, the hypervisor must have a VMkernel port with an IP address that is configured in that network segment.

4.2.2 VADP front end data flow

The path that backup data uses between the hypervisor and the VADP proxy depends on the Virtual Disk Transport method and DNS name resolution. It cannot be controlled by a configuration setting in IBM Spectrum Protect Plus.

Note: The following Virtual Disk Transport methods are supported when the VADP proxy is a VM:

- NBD
- NBDSSL
- HotAdd

The following Virtual Disk Transport methods are supported when the VADP proxy is a physical server:

- NBD
- NBDSSL
- SAN

Consider the following points:

- VM backups traverse the LAN between the hypervisor and the VADP proxy, when NBD or NBDSSL Virtual Disk Transport method are being used (A in Figure 4-2 on page 85). If the hypervisor was registered to the VMware vCenter server by using a DNS name, the data flows to the IP address that is resolved to that name. If the DNS entry relates to an IP address in the Production LAN, this IP address is used to transfer the data. If the DNS entry relates to an IP address in the backup LAN, this network is used.

- VM backup data from a hypervisor does not traverse the LAN when HotAdd or SAN Virtual Disk Transport Method are used (B in Figure 4-2 on page 85).

Note: In Microsoft Hyper-V environments, the backup data flows directly between the hypervisor and the vSnap server, without a datamover (VADP proxy) in between. Therefore, the vSnap server has end-to-end control of which path is used for backup data.

4.3 Establishing connections through firewalls

Because all IBM Spectrum Protect Plus components, hypervisors, VMs, and protected applications and file systems must communicate with each other, all of the required communication ports must be opened between those components and associated services.

Note: Firewalls might exist at different locations in the communication path; for example, hardware firewalls in the network and software firewalls that are part of the backup clients operating system. Accordingly, all firewalls must be configured.

4.3.1 Communication between IBM Spectrum Protect Plus components

For more information about which network ports that must be opened in a firewall, see IBM Support's IBM Spectrum Protect Plus - All Requirements Doc.

From there, select the correct version number and browse to the System Requirements document. Here, you can review the requirements for each IBM Spectrum Protect Plus component, including a description of incoming and outgoing TCP ports. In addition, you can review a brief description of the services that are associated with these ports.

4.3.2 Communication to VMs, applications, and file systems

For the following backup use cases, it is required to connect into a backup client system:

- File indexing for Windows and UNIX VMs
- File system backup for virtual or physical Windows servers
- Backup of applications that are running in virtual or physical servers, Kubernetes clusters, or in cloud environments

Table 4-1 lists the TCP ports that must be opened between IBM Spectrum Protect Plus components and VMs or physical servers when in-guest backups must be performed.

Table 4-1 In-guest backup communication ports

Purpose	TCP Ports	Source	Target
Windows file indexing	5985	IBM Spectrum Protect Plus server, vSnap servers	Windows VM
	22	Windows VM	IBM Spectrum Protect Plus server

Purpose	TCP Ports	Source	Target
UNIX file indexing	22	IBM Spectrum Protect Plus server, vSnap servers	UNIX VM
	22	UNIX VM	IBM Spectrum Protect Plus server
Windows file system backup (for VMs and physical servers)	5985, 5986	IBM Spectrum Protect Plus server	Windows file system
	9085	File Level Restore Browser	Windows file system
	445	Windows file system	vSnap server
Windows application backup, Windows pre- and post-scripts	5985, 5986	IBM Spectrum Protect Plus server	Windows application server
	3260, 445	Windows application server	vSnap server
	443	Windows application server	IBM Spectrum Protect Plus server
UNIX application backup, UNIX pre- and post-scripts	22	IBM Spectrum Protect Plus server	UNIX application server
	111, 2049, 20048	UNIX application server	vSnap server
	443	UNIX application server (Oracle only)	IBM Spectrum Protect Plus server
Kubernetes cluster	Assigned by the NodePort service in Kubernetes	IBM Spectrum Protect Plus server	Kubernetes
	111, 443, 2049, 20048	Kubernetes	vSnap server
Office 365	22	IBM Spectrum Protect Plus server	Proxy host server
	111, 443, 2049, 20048	Proxy host server	vSnap server

Note: Consider the following points:

- The default ports that are used by IBM Spectrum Protect Plus to access the WinRM service on a Windows system (5985) and SSH for UNIX systems (22) can be changed. by selecting **System Configuration** → **Global Preferences** → **General**.
- Do not change any Global Preferences unless directed by IBM support. Changing the SSH and WinRM ports from their defaults affects all communication to UNIX or Windows systems, which can cause unforeseen side effects.

The following section presents an overview of an environment that uses a separate backup and production network. We show you how to configure the IBM Spectrum Protect Plus components to use the production network for management and replication traffic only, and to route backup traffic to the dedicated backup network.

4.4 Configuring IBM Spectrum Protect Plus to use a dedicated backup network

To establish a dedicated backup network, multiple components must be configured correctly. These components include the IBM Spectrum Protect Master server, the VADP proxies (in VMware environments), the vSnap servers, and VMs or physical servers that host an application that is backed up by way of the dedicated backup network.

For a general introduction to networking with IBM Spectrum Protect Plus, for information about data flows in different networks and for prerequisites, see 4.1, “IBM Spectrum Protect Plus networking” on page 84.

Figure 4-3 gives an overview of an environment that is used as an example to demonstrate the necessary configuration steps.

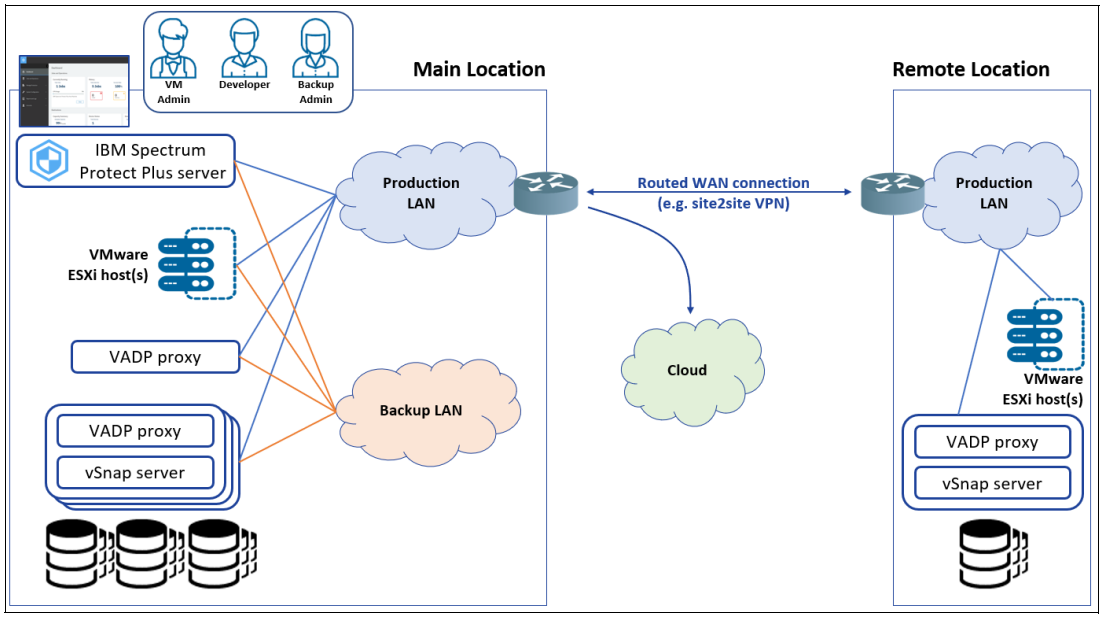

Figure 4-3 Example Network Topology

The network topology in our example consists of a main data center with a Production LAN and a backup LAN, which contains the IBM Spectrum Protect Plus virtual appliance, a vSnap server, a VADP proxy, and an ESXi hypervisor, which is managed by a vCenter.

The Production LAN contains two network routers: one router acts as the default gateway and connects all systems to the companies intranet; the second router establishes a VPN connection to a remote data center where another vSnap server with integrated VADP proxy is implemented to back up the remote VMware datacenter. In the remote data center, no dedicated backup LAN exists; therefore, the vSnap server and VADP proxy are connected to the Production LAN only.

Figure 4-4 shows the IP networks, network routers, and IP addresses that are used in the example environment.

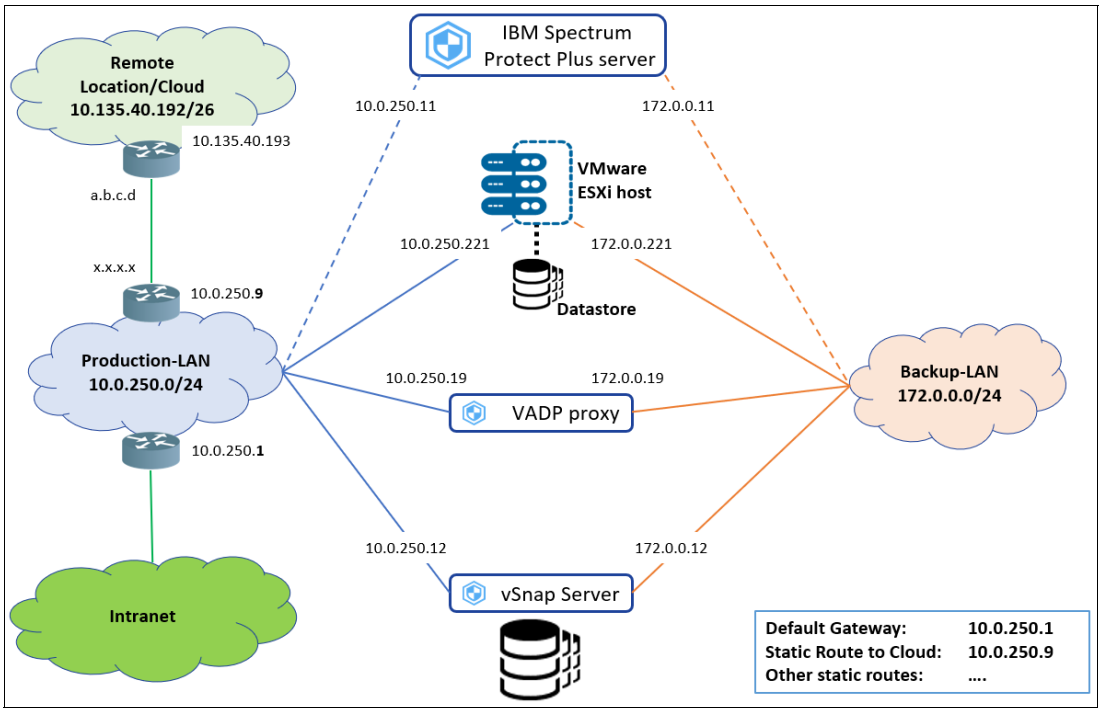

Figure 4-4 Example IP-Network Topology and Routing

In the following sections, our assumption is that the IBM Spectrum Protect Plus server, vSnap servers, and VADP proxies are deployed and configured and that they are connected to the Production LAN. To connect all components to the backup LAN, complete the following steps:

1. Physically attach the VMware ESXi or Microsoft Hyper-V Hypervisors to the backup network and specify an IP address in the corresponding IP address range.
2. Connect the IBM Spectrum Protect Plus virtual appliance to the backup network and specify an IP address in the corresponding IP address range.
3. Attach the IBM Spectrum Protect Plus vSnap servers to the backup network and configure them to use the new connection for all protocols that are related to back up workloads.
4. Connect the IBM Spectrum Protect Plus VADP proxies to the backup network (VMware only).

Configuration details are discussed next.

4.4.1 Preparing the VMware ESXi or Microsoft Hyper-V Hypervisors

The VMware ESXi or Microsoft Hyper-V hosts must be equipped with more network adapters that are connected to the backup network. Various scenarios are possible and highly dependent to the established network infrastructure; for example, the use of dedicated physical network adapters or VLAN tagging on the switchports.

Important: For more information about how to correctly connect the hypervisors to the backup network, contact your network architect or administrator.

Virtual switches

In VMware, virtual switches are used to interconnect physical network adapters to the virtual network interface cards (NICs) of the VMs. Not all VMs must be equipped with another virtual NIC and IP address in the backup LAN. This configuration is required only if a vSnap server must be mounted inside of the VM to support in-guest application backups.

Figure 4-5 shows a basic set-up for VMware that uses two virtual switches (vSwitch0 and vSwitch1) and two networks: the Production LAN (“VM Network”) and the backup LAN.

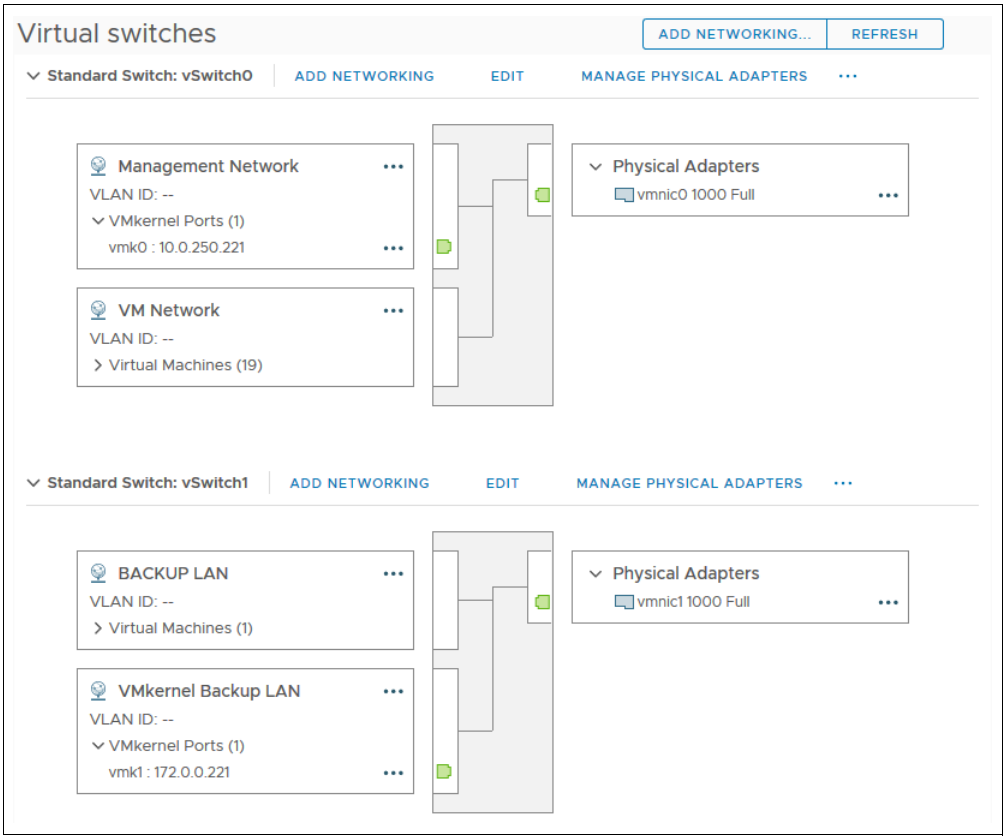

Figure 4-5 ESXi virtual switches for Production LAN and backup LAN

VMkernel adapters

In the case of VMware, a VMkernel port is required to allow the assignment of an IP address to the hypervisor inside a network segment. An example configuration is shown in Figure 4-6.

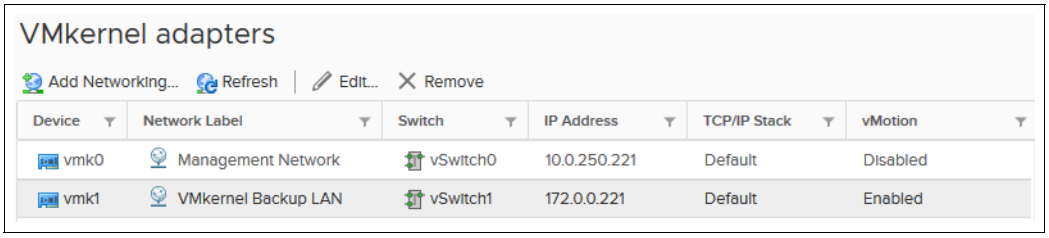

Figure 4-6 VMkernel adapters

Note: Each hypervisor must have its own IP address in the backup network to enable VADP, NFS, SMB, or iSCSI connections between the host and the vSnap server.

A Microsoft Hyper-V server cannot perform a backup or restore operations to or from a vSnap server by way of the backup LAN when it has no IP address assigned in that network.

A VMware ESXi host can perform backups and streaming restores because these services use the VADP proxy to mount resources from the vSnap server. However, InstantAccess or Test restores mount the vSnap server directly to the hypervisor by way of iSCSI or NFS and fail if the host has no IP address assigned in the backup LAN.

Configuring the hypervisor firewall

Finally, we must ensure that the hypervisor firewall policies allow to mount the vSnap server.

In our example, we edit the ESXi host security profile to allow the NFSClient service to accept outgoing connections to the backup LAN IP address range. In the vSphere GUI (vCenter server web interface), for each host that shall use the dedicated backup LAN, select **ESXi Host** → **System** → **Firewall** → **NFS Client** and add the allowed IP ranges. Figure 4-7 shows how to enable outgoing NFS connections to back up LAN IP range `172.0.0.0/24`.

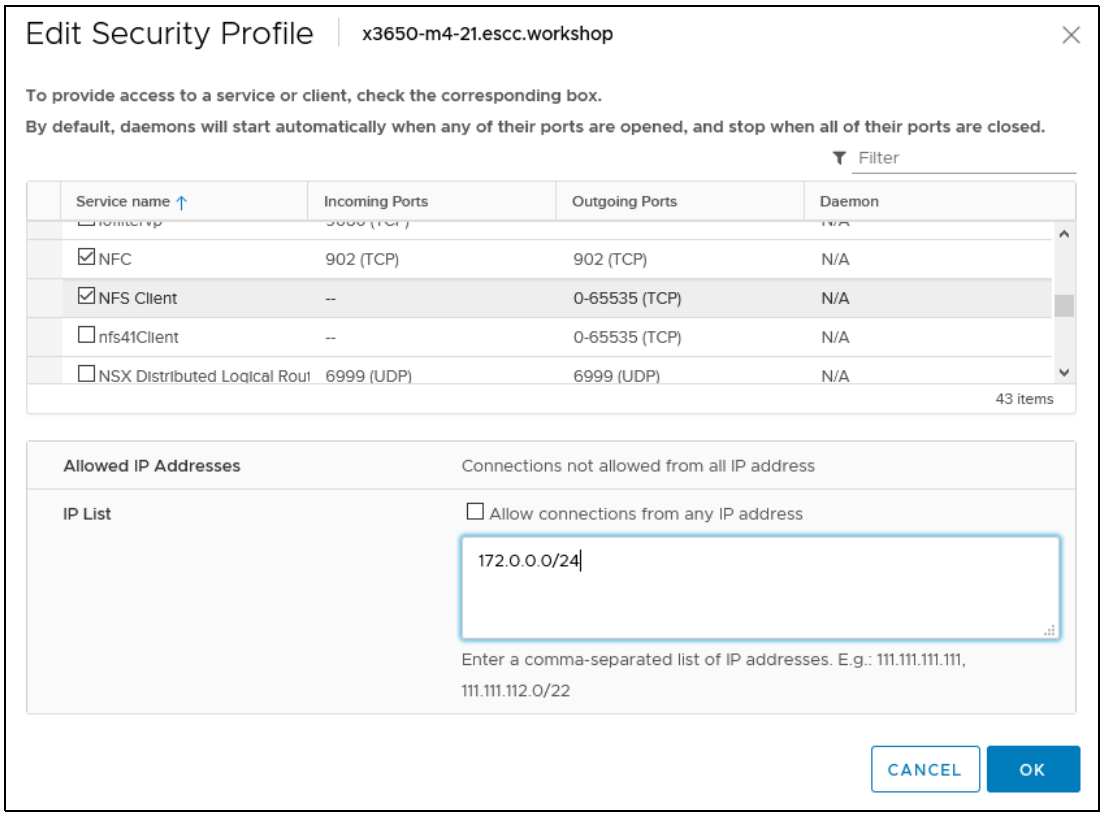

Figure 4-7 Editing the Hypervisor Security Profile

4.4.2 Preparing the IBM Spectrum Protect Plus virtual appliance

By default, the IBM Spectrum Protect Plus VM only uses one network adapter during deployment, which should be attached to the management network.

Add a virtual network interface card (NIC)

Edit the VM configuration of the IBM Spectrum Protect Plus virtual appliance and add a secondary virtual network interface. Then, attach it to the backup network.

Figure 4-8 shows the addition of a second network adapter to the IBM Spectrum Protect Plus server VM.

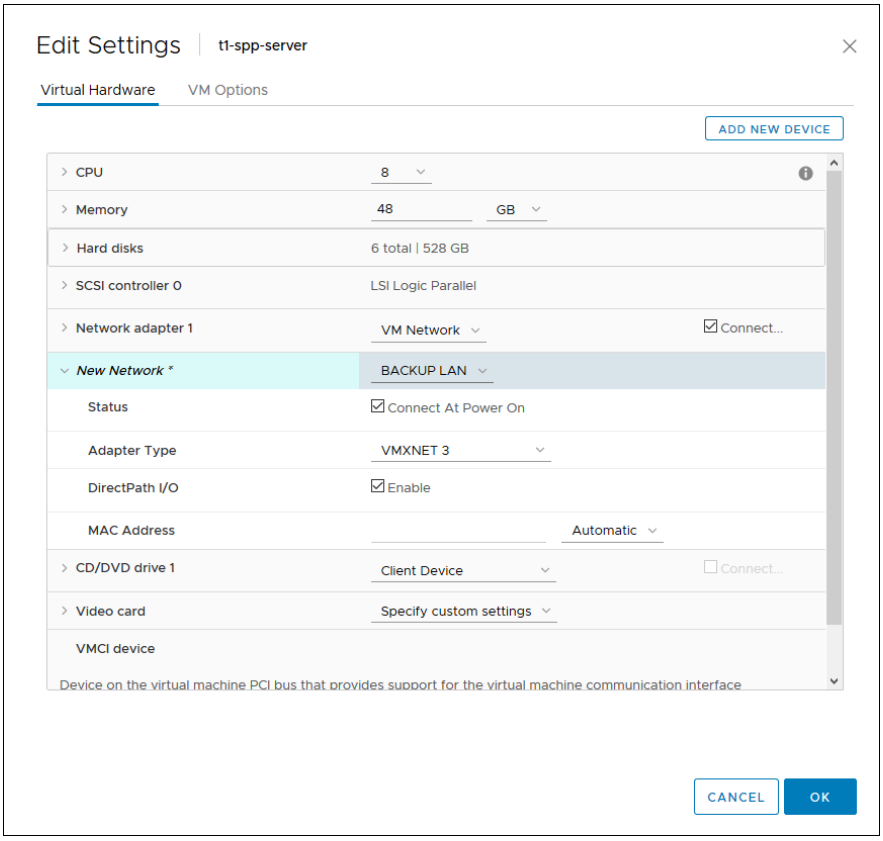

Figure 4-8 Adding a backup network adapter to the IBM Spectrum Protect Plus server

Configuring a backup LAN IP address and static routes

After the extra network adapter is attached to the IBM Spectrum Protect Plus server, it must be configured with an IP address and a static route that are needed in this example to connect to the vSnap server in the remote data center later on.

Note: The configuration of static routes is required only if multiple gateways are used to reach different network segments that contain IBM Spectrum Protect Plus components or backup clients, such as hypervisors, VMs, or physical servers with applications. Depending on the network and router design, static routes might not be required. Discuss this issue with your network administrator.

Complete the following steps:

1. Log in to the IBM Spectrum Protect Plus server by using the command line that uses the SSH protocol to configure the new network adapter by using the NetworkManager Text User Interface (`nmtui`) tool to assign an IP address that belongs to the backup network IP address range.
2. Use the `serveradmin` user ID to connect to the SPP server command line, as shown in Example 4-1.

Example 4-1 Log in to the IBM Spectrum Protect Plus server by using the ssh protocol

```
login as: serveradmin
Restricted access
serveradmin@10.0.250.11's password:
Last login: Mon Jun 15 14:53:37 2020 from 10.0.250.10
------------------------------------------------------------------
  IBM Spectrum Protect Plus  10.1.6 build [ 1972 ]
------------------------------------------------------------------
```

3. Validate that the new network adapter is visible in the operating system and note the adapter name, by running the **`ip a`** command, as shown in Example 4-2. In this case, the new backup LAN adapter is using device `ens192`.

Example 4-2 Check the SPP operating system for the new network adapter

```
[serveradmin@t1-spp-server ~]$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: ens160: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group
default qlen 1000
    link/ether 00:50:56:98:14:15 brd ff:ff:ff:ff:ff:ff
    inet 10.0.250.11/24 brd 10.0.250.255 scope global noprefixroute ens160
       valid_lft forever preferred_lft forever
    inet6 fe80::58b7:c15b:2943:bdf1/64 scope link noprefixroute
       valid_lft forever preferred_lft forever
3: ens192: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group
default qlen 1000
    link/ether 00:50:56:98:77:0e brd ff:ff:ff:ff:ff:ff
       valid_lft forever preferred_lft forever
    inet6 fe80::250:56ff:fe98:770e/64 scope link
       valid_lft forever preferred_lft forever
...
```

4. Start the `nmtui` tool to configure the backup LAN adapter with an IP address, as shown in Example 4-3.

Example 4-3 Starting the nmtui tool to configure network settings

```
------------------------------------------------------------------
  IBM Spectrum Protect Plus  10.1.6 build [ 1972 ]
------------------------------------------------------------------
[serveradmin@t1-spp-server ~]$ nmtui
```

5. Specify the IP address with the subnet mask, for the new interface (`ens192`), as shown in Figure 4-9.

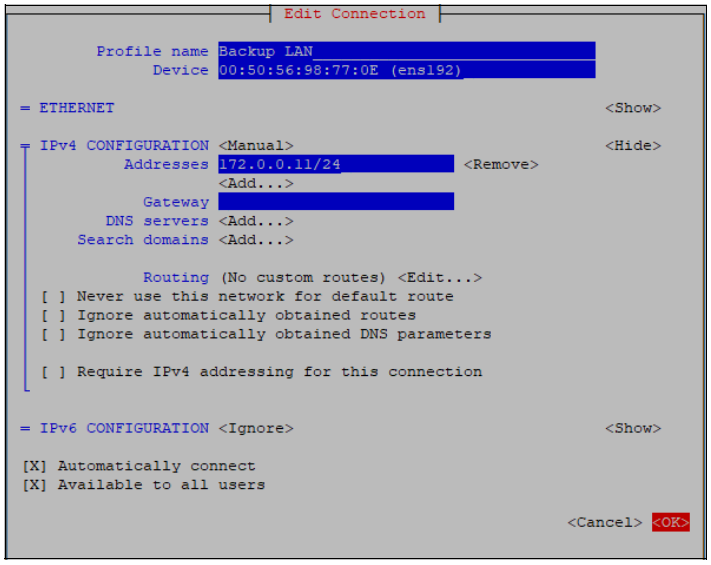

Figure 4-9 Adding an IP address to the backup network interface

In our example environment, no gateway is in the backup LAN; therefore, the fields for Gateway, DNS, and Search domains are empty. However, we must modify the network interface that is connected to the Production LAN (`ens160`) and specify a static route to the special VPN gateway that connects to the remote side. The addition of a static route is shown in Figure 4-10. It adds a route to remote network `10.135.40.192/26` by way of gateway `10.0.250.9`.

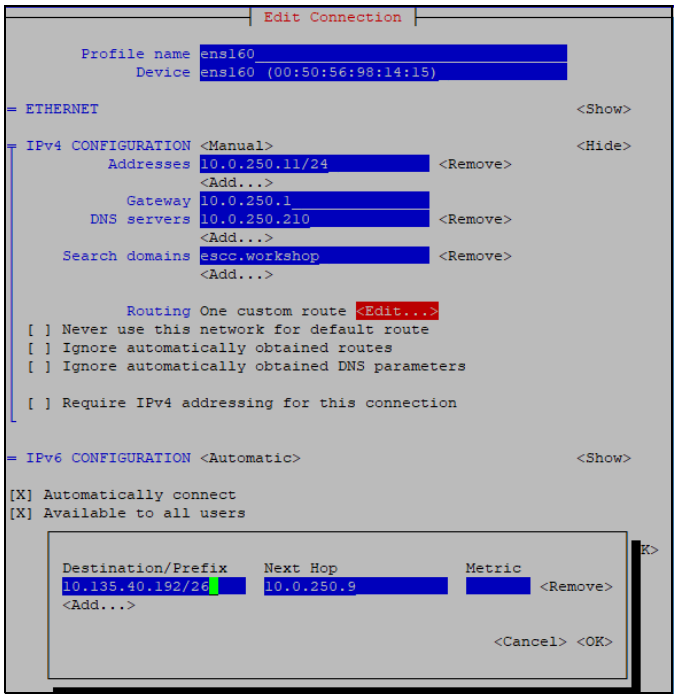

Figure 4-10 Adding a static route to the remote data center on the production network interface

6. When the configuration is complete, validate that the IP and static route configuration is correct, that the local hypervisors can be pinged by using their backup network IP address, and that the remote site is reachable as well (see Example 4-4).

Example 4-4 Validating the SPP server backup network configuration

```
---------------------------------------------------------------
  IBM Spectrum Protect Plus  10.1.6 build [ 1972 ]
---------------------------------------------------------------
[serveradmin@t1-spp-server ~]$ ip a <== Show IP configuration
...
2: ens160: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group
inet 10.0.250.11/24 brd 10.0.250.255 scope global noprefixroute ens160
 ...
3: ens192: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group
inet 172.0.0.11/24 brd 172.0.0.255 scope global noprefixroute ens192
...
[serveradmin@t1-spp-server ~]$ ip r <== Show Routing configuration
default via 10.0.250.1 dev ens160 proto static metric 100
10.135.40.192/26 via 10.0.250.9 dev ens160 proto static metric 100
172.0.0.0/24 dev ens192 proto kernel scope link src 172.0.0.11 metric 101
...
[serveradmin@t1-spp-server ~]$ ping 172.0.0.221 <== Backup LAN IP of ESXi host
PING 172.0.0.221 (172.0.0.221) 56(84) bytes of data.
64 bytes from 172.0.0.221: icmp_seq=1 ttl=64 time=0.321 ms
64 bytes from 172.0.0.221: icmp_seq=2 ttl=64 time=0.228 ms
64 bytes from 172.0.0.221: icmp_seq=3 ttl=64 time=0.237 ms
64 bytes from 172.0.0.221: icmp_seq=4 ttl=64 time=0.189 ms
...
--- 172.0.0.221 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3053ms
rtt min/avg/max/mdev = 0.189/0.243/0.321/0.051 ms
...
[serveradmin@t1-spp-server ~]$ ping 10.135.40.193 <== Router IP in remote location
PING 10.135.40.193 (10.135.40.193) 56(84) bytes of data.
64 bytes from 10.135.40.193: icmp_seq=1 ttl=63 time=2.28 ms
64 bytes from 10.135.40.193: icmp_seq=2 ttl=63 time=2.34 ms
64 bytes from 10.135.40.193: icmp_seq=3 ttl=63 time=2.21 ms
64 bytes from 10.135.40.193: icmp_seq=4 ttl=63 time=2.46 ms
...
--- 10.135.40.193 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3004ms
rtt min/avg/max/mdev = 2.217/2.327/2.464/0.102 ms
[serveradmin@t1-spp-server ~]$
```

When the IBM Spectrum Protect Plus server is correctly attached to the backup network, the same configuration must be done on the VADP proxies and the vSnap servers.

4.4.3 Preparing the IBM Spectrum Protect Plus vSnap server

By default, the IBM Spectrum Protect Plus vSnap server must be attached to the management network. To route backup traffic to a dedicated backup network, the vSnap servers, and VADP proxies also must be connected to the backup network.

If the vSnap server is a VM, use the hypervisor tools to edit the VM configuration to add a second network adapter and connect it to the new backup network. If the vSnap server is a physical machine, you must add a second physical network adapter and attach it to the backup VLAN.

The following figures show how to attach a virtual vSnap server to the backup network, and to configure it so that the backup related protocols (replication, NFS, SMC, and iSCSI) are bound to the backup network adapter.

Because the vSnap server is a VM in this scenario, the adapter is added by editing the VM settings, as shown in Figure 4-11.

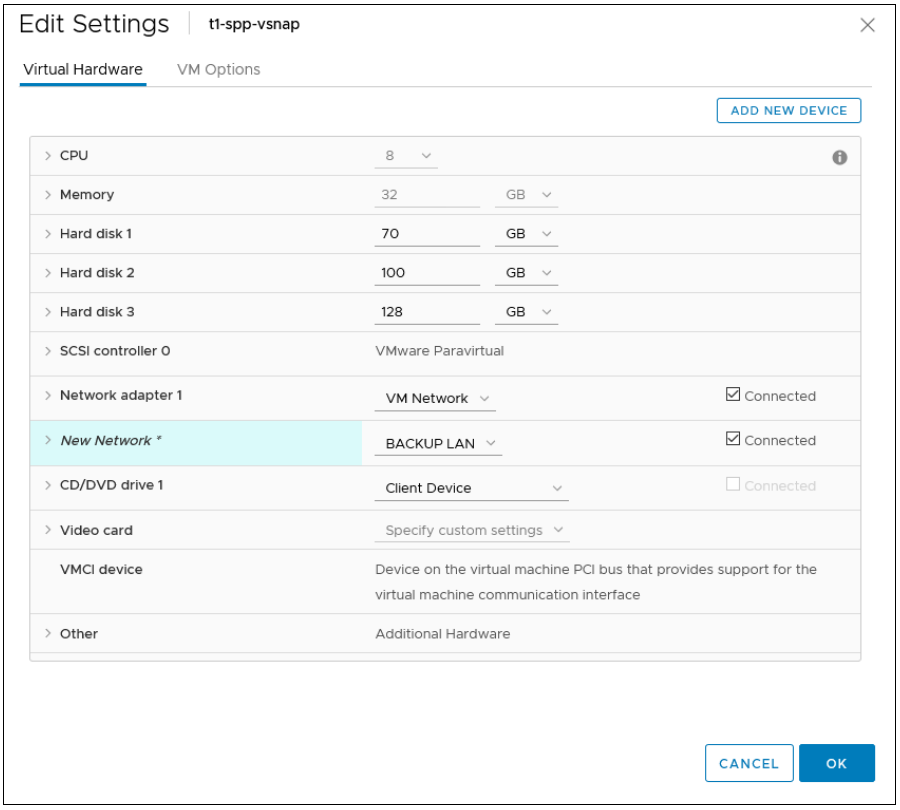

Figure 4-11 Adding a backup network interface to a virtual vSnap server

After the adapter is added to the vSnap server VM, log in to the command line interface by using the SSH protocol and use `nmtui` to specify the proper IP settings, as shown in Figure 4-12.

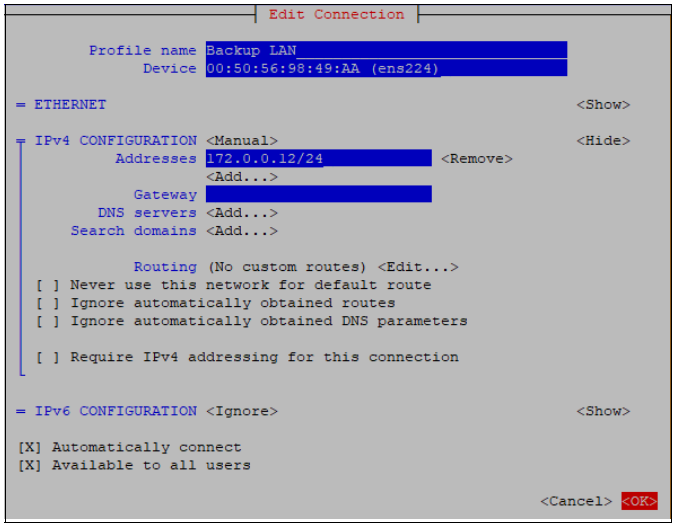

Figure 4-12 Configuring the IP address for a vSnap server backup network adapter

In the same way as it was done for the IBM Spectrum Protect Plus virtual appliance, a static route must be established for the vSnap server to enable communication to the secondary vSnap server in the remote data center (see Figure 4-13).

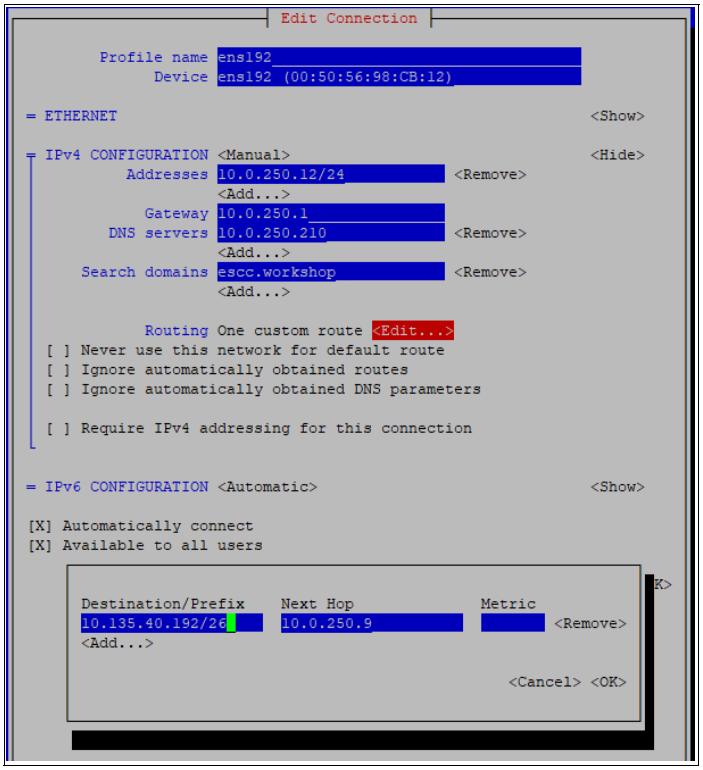

Figure 4-13 Adding a static route to the remote data center on the production network interface

When the configuration is complete, validate that the IP and static route configuration is correct, that the local hypervisors can be pinged by suing their backup network IP address, and that the remote site is reachable as well (see Example 4-5).

Example 4-5 Validating the vSnap server backup network configuration

```
login as: serveradmin
serveradmin@10.0.250.12's password:
Last login: Mon Jun 29 12:21:15 2020 from 10.5.250.141
------------------------------------------------------------------
Be sure to adhere to vSnap hardware and memory requirements
 as described in IBM Spectrum Protect Plus Blueprints
 accessible from IBM Knowledge Center.
------------------------------------------------------------------
[serveradmin@t1-spp-vsnap ~]$ ip a <== Show IP configuration
2: ens192: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group inet
10.0.250.12/24 brd 10.0.250.255 scope global noprefixroute ens192
...
3: ens224: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group inet
172.0.0.12/24 brd 172.0.0.255 scope global noprefixroute ens224
...
[serveradmin@t1-spp-vsnap ~]$ ip r <== Show Routing configuration
default via 10.0.250.1 dev ens192 proto static metric 100
10.0.250.0/24 dev ens192 proto kernel scope link src 10.0.250.12 metric 100
10.135.40.192/26 via 10.0.250.9 dev ens192 proto static metric 100
172.0.0.0/24 dev ens224 proto kernel scope link src 172.0.0.12 metric 101
[serveradmin@t1-spp-vsnap ~]$ ping 172.0.0.221 <== Backup LAN IP of ESXi host
PING 172.0.0.221 (172.0.0.221) 56(84) bytes of data.
64 bytes from 172.0.0.221: icmp_seq=1 ttl=64 time=0.198 ms
64 bytes from 172.0.0.221: icmp_seq=2 ttl=64 time=0.264 ms
64 bytes from 172.0.0.221: icmp_seq=3 ttl=64 time=0.243 ms
64 bytes from 172.0.0.221: icmp_seq=4 ttl=64 time=0.240 ms
--- 172.0.0.221 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3102ms
rtt min/avg/max/mdev = 0.198/0.236/0.264/0.026 ms
...
[serveradmin@t1-spp-vsnap ~]$ ping 10.135.40.193 <== Router IP in remote location
PING 10.135.40.193 (10.135.40.193) 56(84) bytes of data.
64 bytes from 10.135.40.193: icmp_seq=1 ttl=63 time=2.27 ms
64 bytes from 10.135.40.193: icmp_seq=2 ttl=63 time=2.52 ms
64 bytes from 10.135.40.193: icmp_seq=3 ttl=63 time=2.25 ms
64 bytes from 10.135.40.193: icmp_seq=4 ttl=63 time=2.48 ms
--- 10.135.40.193 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3002ms
rtt min/avg/max/mdev = 2.255/2.383/2.523/0.120 ms
[serveradmin@t1-spp-vsnap ~]$
```

4.4.4 Preparing the IBM Spectrum Protect Plus VADP proxy (VMware only)

It is a common practice to have the VADP proxy component installed on the vSnap server. If the VADP proxy is installed on a separate machine, it also must have access to both networks (management and backup).

If the VADP proxy is a VM, use the hypervisor tools to edit the VM config to add a second network adapter and connect it to the new backup network. If the VADP proxy is a physical machine, you must add a second physical network adapter and attach it to the backup VLAN.

After this process is done, you must set an IP address and static route to remote vSnap servers by using the `nmtui` tool. (Because this procedure is the same as for the IBM Spectrum Protect Plus virtual appliance or the vSnap server, we do not show this process in more detail here.)

4.4.5 Enabling or disabling specific protocols on a network interface

Backup/restore or replication traffic can be enabled or disabled on a network interface by using one of the following methods:

- GUI
- CLI

Managing vSnap server network interfaces by using the GUI

To configure a NIC for backup and replication operations, complete the following steps:

1. In the navigation pane, click **System Configuration** → **Backup Storage** → **Disk** and select the management icon of the vSnap server that you want to configure.
2. On the **Networks** tab, choose the configuration that you want for your listed NICs:
   - To configure a NIC for transfers of data for backup and restore operations only, select **Backup**. This selection means that during backup and restore operations, connections are made to the vSnap server by using the IP address of this NIC. When multiple NICs are specified for backup, the first one that connects successfully is used.
   - To configure an NIC for transfers of data for replication purposes only, select **Replication**. This selection means that during incoming replication operations to a vSnap server, connections are made by using the IP address of this NIC on the target vSnap server. When multiple NICs specify Replication on the target vSnap server, the first target IP address that connects successfully from the source vSnap server is used.
   - To configure a NIC for both replication, and backup and restore data transfers, select both **Backup** *and* **Replication**.

In our example, we use two vSnap servers, one in the main location (site “Primary”) and one in the remote data center (site “IBM-Cloud”), as shown in Figure 4-14 on page 101.

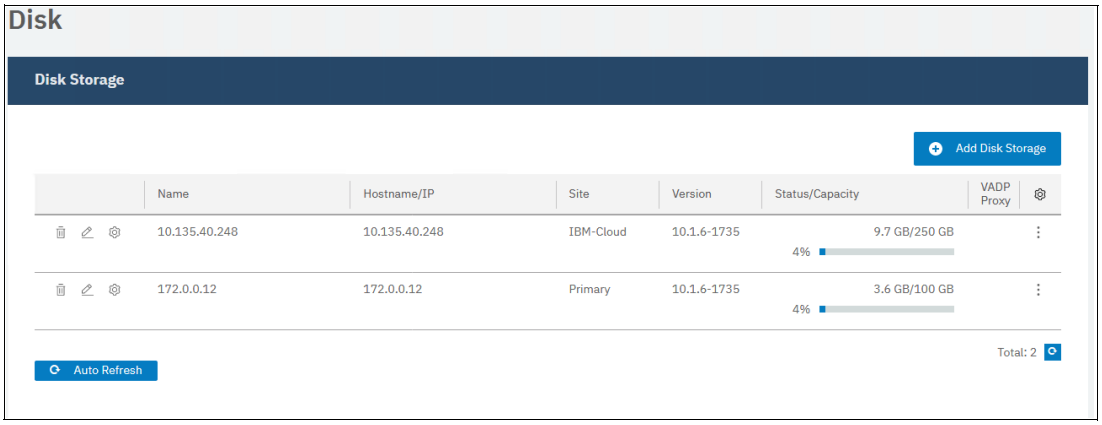

Figure 4-14 Disk storage in different sites

Figure 4-15 shows the configuration in the main location where we have two NICs: one for backup and one for replication traffic.

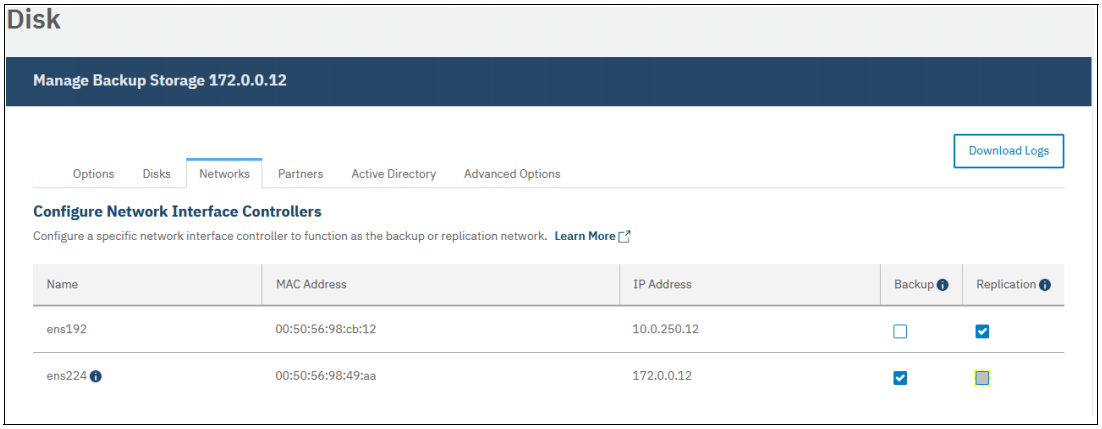

Figure 4-15 Assignment of protocols to network interfaces - main location

Figure 4-16 shows the configuration in the remote data center where we have a vSnap server that is connected to the production LAN only.

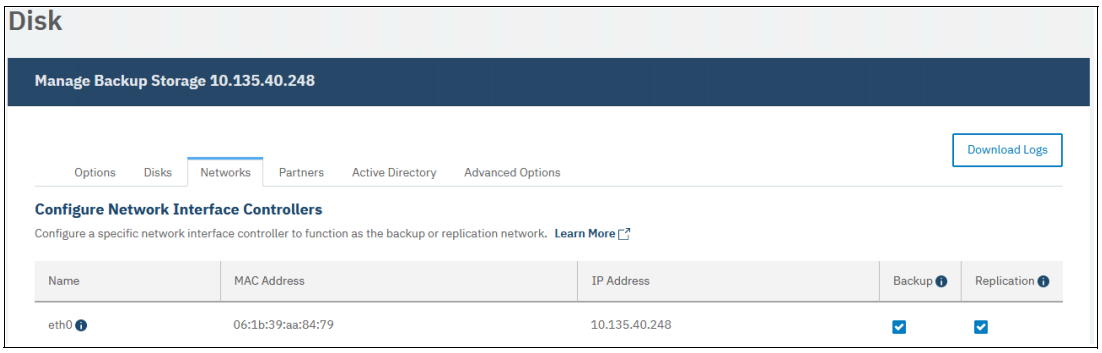

Figure 4-16 Assignment of protocols to network interfaces - remote location

Note: The icon that is next to the interface name (“i” sign) shows which NIC is being used for management traffic. Usually, this interface is used to register the vSnap server to the IBM Spectrum Protect Plus server.

Manage vSnap server Network interfaces by using the CLI

You can run the **vsnap network** commands to view or modify the protocols that can be used on a NIC.

Run the **vsnap network show** command to identify which protocols are enabled on the available network adapters, as shown in Example 4-6.

Example 4-6 The vsnap network show command

```
[serveradmin@t1-spp-vsnap ~]$ vsnap network show
ID           | NAME   | MAC ADDR          | IPV4 ADDR   | SERVICES
-------------------------------------------------------------------------------
00505698cb12 | ens192 | 00:50:56:98:cb:12 | 10.0.250.12 | mgmt, repl, nfs, smb, iscsi
0050569849aa | ens224 | 00:50:56:98:49:aa | 172.0.0.12  | mgmt, repl, nfs, smb, iscsi

[serveradmin@t1-spp-vsnap ~]$
```

By default, *all* protocols are enabled on an adapter. In our example, we want to limit backup workloads to the backup adapter in IP network `172.0.0.0/24` and replication traffic to the production LAN interface in IP network `10.0.250.0/24`.

Run the **vsnap network update** command, as shown in Example 4-7.

Example 4-7 The vsnap network update command

```
[serveradmin@t1-spp-vsnap ~]$  vsnap network update --id 00505698cb12 --services mgmt,repl

ID: 00505698cb12
NAME: ens192
MAC ADDRESS: 00:50:56:98:cb:12
IPV4 ADDRESSES:
    10.0.250.12

IPV6 ADDRESSES:
    fe80::19f2:a215:68c8:8498

SERVICES:
    mgmt
    repl

SERVICE TYPES:
    repl

[serveradmin@t1-spp-vsnap ~]$ vsnap network update --id 0050569849aa --services
mgmt,nfs,smb,iscsi

ID: 0050569849aa
NAME: ens224
MAC ADDRESS: 00:50:56:98:49:aa
IPV4 ADDRESSES:
    172.0.0.12

IPV6 ADDRESSES:
    fe80::250:56ff:fe98:49aa

SERVICES:
    mgmt
    nfs
    smb
    iscsi
```

```
SERVICE TYPES:
    data

[serveradmin@t1-spp-vsnap ~]$ vsnap network show
ID           | NAME   | MAC ADDR          | IPV4 ADDR   | SERVICES
-------------------------------------------------------------------------------
00505698cb12 | ens192 | 00:50:56:98:cb:12 | 10.0.250.12 | mgmt, repl
0050569849aa | ens224 | 00:50:56:98:49:aa | 172.0.0.12  | mgmt, nfs, smb, iscsi

[serveradmin@t1-spp-vsnap ~]$
```

This set up ensures that replication workloads are routed through NIC `ens192` and backup data is routed to the vSnap server through NIC `ens224`.

Important: Ensure that the `mgmt` service remains enabled on the interface that was used to register the vSnap server in IBM Spectrum Protect Plus. To prevent issues in the device management, enable the **mgmt** service on all available adapters.

4.4.6 Special Configuration: Forcing VADP front end traffic to use a dedicated interface

As described in Chapter 4.2, “Understanding network data flows” on page 85, IBM Spectrum Protect Plus cannot control which interface is being used for data that flows between the VADP proxy and the hypervisor (VADP front end). The selection of the interface is determined most often by DNS name resolution.

The effect of DNS is shown in the following example:

- IBM Spectrum Protect Plus is configured to backup a VMware environment.
- The VMs to be protected are on one or more ESXi hypervisors, which are managed through a vCenter server. The vCenter is registered to the IBM Spectrum Protect virtual appliance, not the hypervisors themselves.
- The hypervisors was registered to the vCenter server using DNS names.

Important: During a backup or streaming restore operation, the VADP proxy uses the IP addresses that resolve to the hypervisor DNS names. If these IP addresses belong to the production LAN, VADP front end traffic traverses through this network, independent of the fact which network is being used between VADP proxy and vSnap server.

Figure 4-17 shows a VM backup operation in the IBM Spectrum Protect Plus job log.

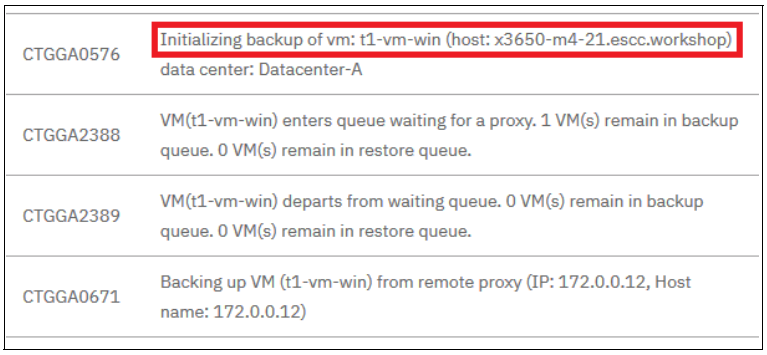

CTGGA0576	Initializing backup of vm: t1-vm-win (host: x3650-m4-21.escc.workshop) data center: Datacenter-A
CTGGA2388	VM(t1-vm-win) enters queue waiting for a proxy. 1 VM(s) remain in backup queue. 0 VM(s) remain in restore queue.
CTGGA2389	VM(t1-vm-win) departs from waiting queue. 0 VM(s) remain in backup queue. 0 VM(s) remain in restore queue.
CTGGA0671	Backing up VM (t1-vm-win) from remote proxy (IP: 172.0.0.12, Host name: 172.0.0.12)

Figure 4-17 SPP job log with DNS name of the ESXi hypervisor

In our example, the host name `x3650-m4-21.escc.workshop` resolves to the IP address in the Production LAN (`10.0.250.221`), which causes all VADP front end traffic to flow through this network. The intention was to use a dedicated backup network for all backup and restore workloads; therefore, we must change the name resolution.

To force the VADP proxy to use the backup LAN for front end traffic, we must override the name resolution on that server by specifying a corresponding entry in the `/etc/hosts` file (see Example 4-8).

Example 4-8 DNS override on the VADP proxy to force front end traffic to backup LAN

```
[serveradmin@t1-spp-vadp ~]$ cat /etc/hosts
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
...
172.0.0.221 x3650-m4-21 x3650-m4-21.escc.workshop
...
[serveradmin@t1-spp-vadp ~]$
```

As a result, the VADP proxy communicates with the ESXi hypervisor through the backup LAN.

4.4.7 Editing firewall ports

Although the IBM Spectrum Protect Plus “Master server” and the vSnap server can be deployed from VM images that have all required internal firewall rules in place, some IBM Spectrum Protect Plus components can be installed manually onto a virtual or physical Linux server; for example, the VADP proxy or the vSnap server component. In this case, the internal firewall rules of the underlying operating system must be configured to allow the required communication.

For more information about how to edit firewall rules in various Linux operating systems, see IBM Knowledge Center.

For more information about the TCP ports that must be opened in network firewalls to allow communication between the different IBM Spectrum Protect Plus components, the hypervisors, VMs, and the protected applications, see Chapter 4.3, “Establishing connections through firewalls” on page 87.

4.4.8 Testing network connectivity

When all internal and external firewall configurations are complete, the IBM Spectrum Protect Plus Service tool can be used to validate that a connection can be established to a specific component and port.

For more information about how to use this tool, see 5.3.9, “Testing network connectivity” on page 152.

5

Daily operations and maintenance

In this chapter, daily operations and maintenance tasks are described.

This chapter includes the following topics:

5.1 Role-based access control overview

Through role-based access control, IBM Spectrum Protect Plus can be tailored for individual users, giving them access to only the features and resources that they require. This feature enables you to create self-service models for users.

For example, a database administrator can start an environment with consistent copies of the database without knowing server login credentials or storage credentials. This approach saves time for the server and storage team.

Role-based access control consists of the following elements:

- Users: A user account associates a resource group with a role. To enable a user to log in to IBM Spectrum Protect Plus and use its functions, you must first add the user as an individual user or as part of an imported group of LDAP users, and then assign resource groups and roles to the user account.
- Roles: Roles define the actions that can be performed on the resources that are defined in a resource group. While a resource group defines the resources that will be made available to a user account, a role sets the permissions to interact with the resources defined in the resource group. For example, if a resource group is created that includes backup and restore jobs, the role determines how a user can interact with the jobs.
- Resources: A resource group defines the resources that are available to a user. Every resource that is added to IBM Spectrum Protect Plus can be included in a resource group, along with individual IBM Spectrum Protect Plus functions and screens. For example, a resource group can include an individual hypervisor, with access to only backup and reporting functionality. When the resource group is associated with a role and a user, the user sees only the screens that are associated with backup and reporting for the assigned hypervisor, such as resources are hypervisors, databases, and individual screens.

Figure 5-1 shows the relationship between users, roles, and resources and their purpose.

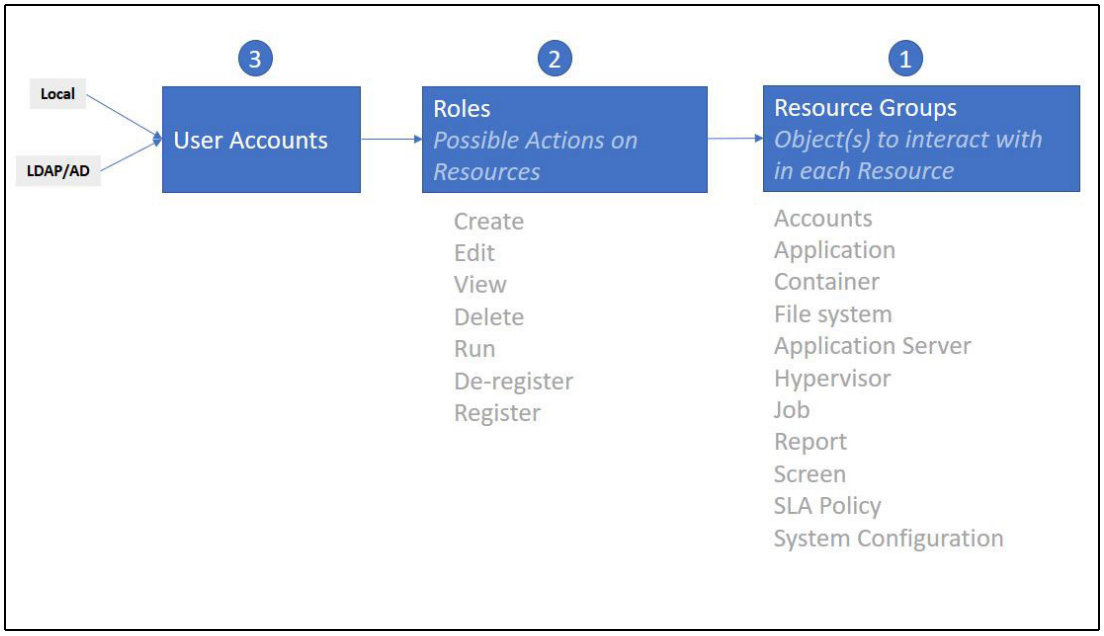

Figure 5-1 Role Based Access overview

The suggested order of creation when you define a new set of Role based Access rules is also shown in Figure 5-1 on page 108, starting with the Resource Groups, then the Role, and finally, assign these two to a User.

For more information about role-based access control (RBAC), see IBM Knowledge Center.

To illustrate RBAC, how it works, and how to use it, we use a practical use case in this chapter.

Assume that we are running IBM Spectrum Protect Plus in a context where the following requirements must be met:

- The IBM Spectrum Protect Plus environment is configured for two different environments that are isolated from each other in terms of management.
- BackupOnly and RestoreOnly profiles are created to satisfy some sort of separation of duties.
- Separate teams are managing VMware, Oracle, and SQL resources, and each team sees only what is under their responsibility.
- Two Oracle databases cannot be restored by the same team. One team cannot restore a database that is not under its responsibility.

Figure 5-2 shows our example the scenario. Next, we see how to meet our requirements by using IBM Spectrum Protect Plus RBAC.

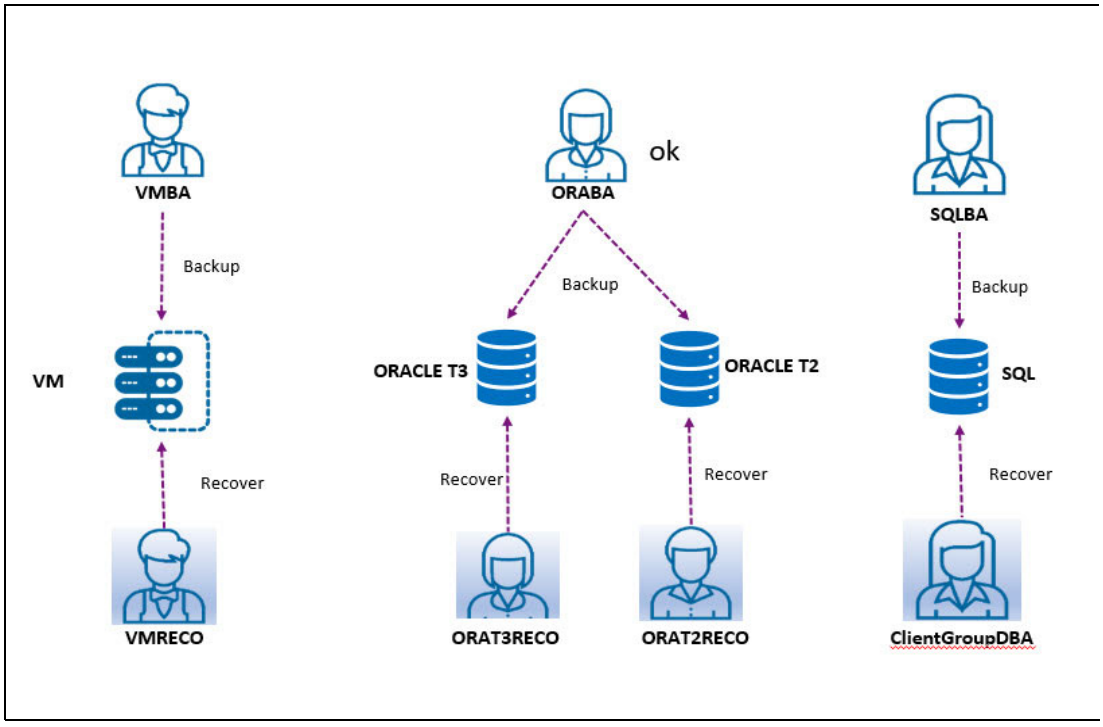

Figure 5-2 RBAC scenario of possible actions and users

5.1.1 Planning user, roles, and resource groups

Consider that you are planning to deploy IBM Spectrum Protect Plus and must provide access to IBM Spectrum Protect Plus for specific users for tasks. Because the RBAC structure that is composed by the three components user, roles, and resource groups, it is better to start configuring the resource group, then the roles, and eventually define the user, as shown in Figure 5-1 on page 108.

To meet our requirements, we create specific resource groups for administering virtual machines (VMs), Oracle DBA, and SQL DBA. We also create dedicated roles for these administrators with the BackupOnly or the RestoreOnly permissions and only resources they are allowed to access.

5.1.2 Creating resource groups

Resource groups define the resources with which a user can work. Every resource that is added to IBM Spectrum Protect Plus can be included in a resource group, along with individual IBM Spectrum Protect Plus functions and screens.

For example, a resource group can include a VMware hypervisor, with access to only backup and reporting functionality. When the resource group is associated with a role and a user, the user sees only the screens that are associated with backup and reporting for the assigned hypervisor.

In our example, we create four resource groups: one for VMware, two for Oracle (instance segregation) and one for SQL:

- `resourceORA_t2`:
  - Resource: Application → Oracle → IBM Spectrum Protect Plus on t2-vm-lx machine
  - Resource: Application Server → Oracle → t2-vm-lx
  - Resource: Screen → All
  - Resource: SLA Policy → All
  - Resource: Job → All
  - Resource: Accounts → Identity → All

  To create a resource group, click **Accounts** → **Resource group** → **Create Resource Group** in the IBM Spectrum Protect Plus web interface.

  Figure 5-3 on page 111 shows the resource group `resourceORA_t2` that was created for the purpose of our example. When selecting a resource, always click the lower-left blue **Add Resources** button. Otherwise, the resource is not added to the resource group.

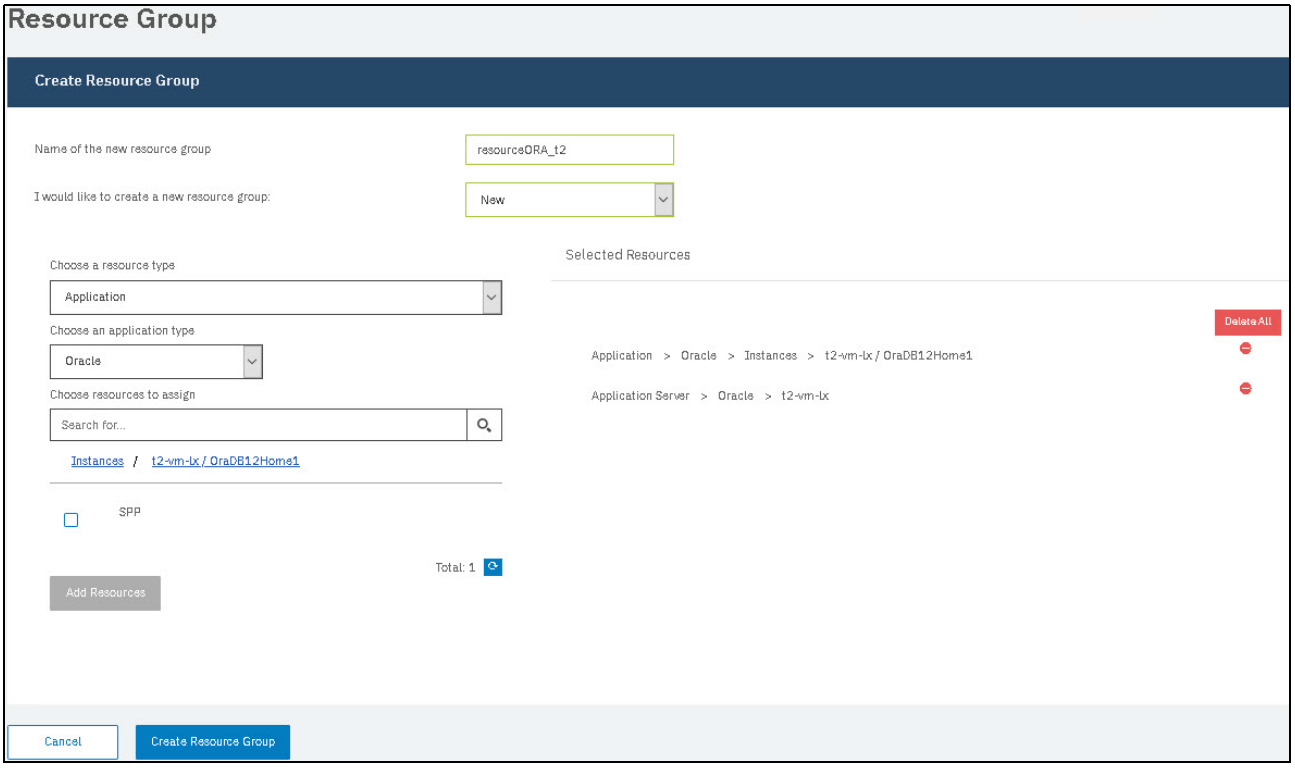

Figure 5-3 Building the resource group resourceORA_t2 for the purpose of our example.

- `resourceVMW`:
  - Resource: Hypervisor → VMware → vCenters → All VMware
  - Resource: Screen → All
  - Resource: SLA Policy → All
  - Resource: Job → All
  - Resource: Report → VM Environment → All
  - Resource: System Configuration → Site → Primary → System Configuration → Site → Secondary
- `resourceORA_t3`:
  - Resource: Application → Oracle → IBM Spectrum Protect Plus on t3-vm-lx machine
  - Resource: Application Server → Oracle → t2-vm-lx
  - Resource: Screen → All
  - Resource: SLA Policy → All
  - Resource: Job → All*
  - Resource: Accounts → Identity → All
- `resourceSQL`:
  - Resource: Application → SQL → All SQL
  - Resource: Screen → All
  - Resource: SLA Policy → All
  - Resource: Job → All
  - Resource: Accounts → Identity → All

The Job resource is important because by using it, the user can see the jobs log from the IBM Spectrum Protect Plus interface, and submit new jobs from its own operation.

After it is completed, the resource groups looks like the example that is shown in Figure 5-4

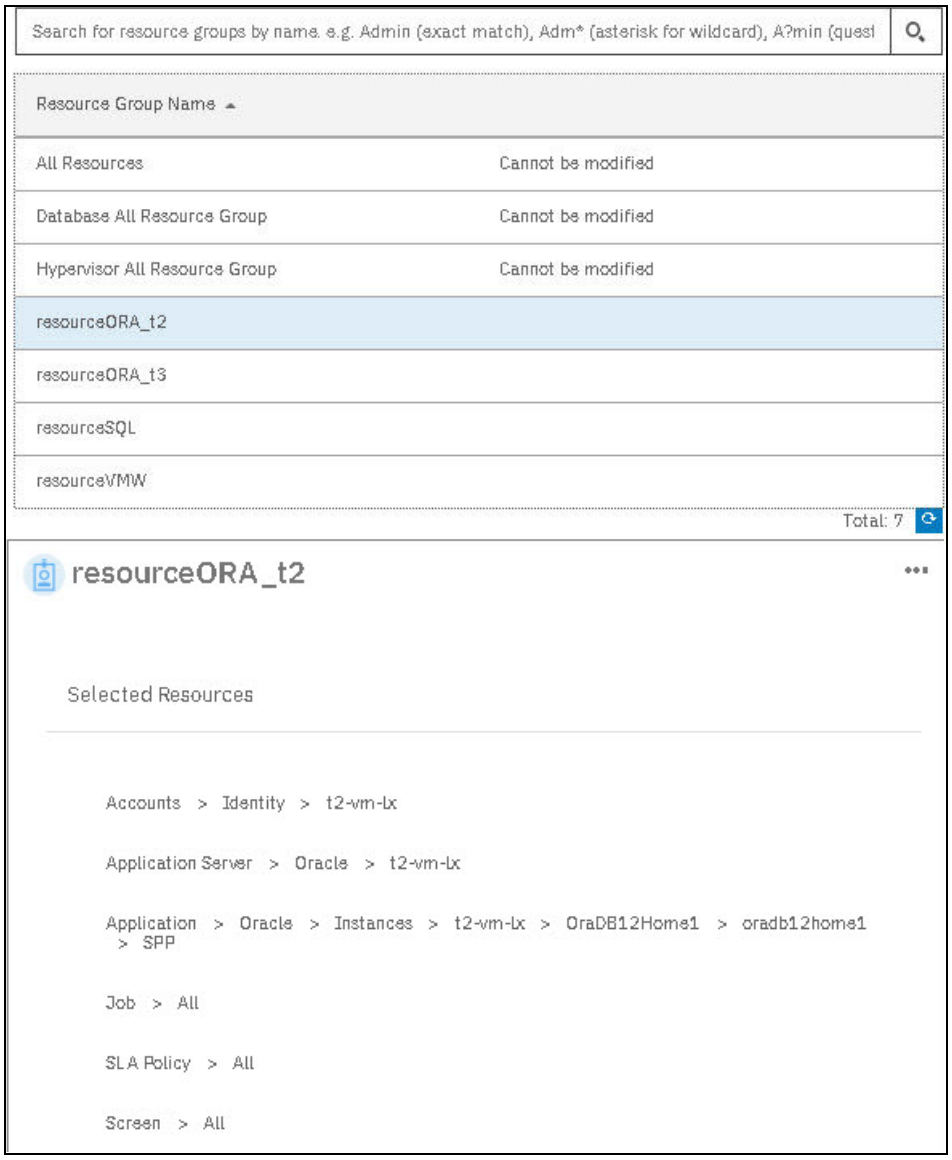

Figure 5-4 Resource group detailed view when clicking the resource group name

In our example, we assigned SLA Policy Resource to **All**, but you can restrict it to a specific SLA policy if you have a dedicated SLA Policy per workload type. Likewise, we assigned Accounts resources to **All**, but you can restrict it to specific accounts that are required for the profile you are creating.

Consider the following points when you create a resource group:

- The resource types Application and Application Server do not provide the same level of granularity. The Application allows you to select specific objects down to the database name, whereas the Application Server allows a selection at the server level only.
- Resource group allows you to restrict access to specific IBM Spectrum Protect Plus objects being backed up as we described with `resourceORA_t3` and `resourceORA_t2`. You can set specific permissions down to a specific resource name (such as a database Instance) within the same Resource type.

In our example, we restricted the resourceORA resource group to only one Oracle database. This type of segregation can be done in a multi-client environment. You can have one resource group per client per resource type.

Note: The resource *Screen* allows assigned users to open the IBM Spectrum Protect Plus web interface. Always assign the user to the relevant screens; otherwise, the user cannot use the assigned resources in IBM Spectrum Protect Plus.

5.1.3 Creating roles

Roles define the actions that can be performed on the resources within a resource group. To create a custom role, click **Accounts** → **Role** in the IBM Spectrum Protect Plus web interface.

Note: The SUPERUSER role cannot be assigned to any individual user account; it is dedicated to the sppadmin user.

To meet the requirements listed in 5.1, “Role-based access control overview” on page 108, we use predefined roles and create specific roles when we want more restriction:

- `RestoreOnly_DBA`: A role that we created that is be used by the DBA in charge of recoveries:
  - Built upon the predefined model Restore Only
  - We removed the permission on Hypervisor
- `Application Admin`: Pre-defined role in IBM Spectrum Protect Plus. We use it for the DBA in charge of configuring and performing backup.
- `RestoreOnly_VM`: A role that we created that is to be used by VM administrator who is in charge of VM restore:
  - Built upon the predefined model Restore Only
  - We removed the permission on Applications
- `BackupOnly_VM`: A role that we created that is to be used by VM administrator that is in charge of VM backup:
  - Built upon the predefined model Backup Only
  - Be removed the permission on Applications

Figure 5-5 on page 114 shows the creation of the `Role RestoreOnly_VM` that is based on the Restore Only template.

By selecting the role, the assigned permissions are displayed (as shown on the right side of Figure 5-5 on page 114).

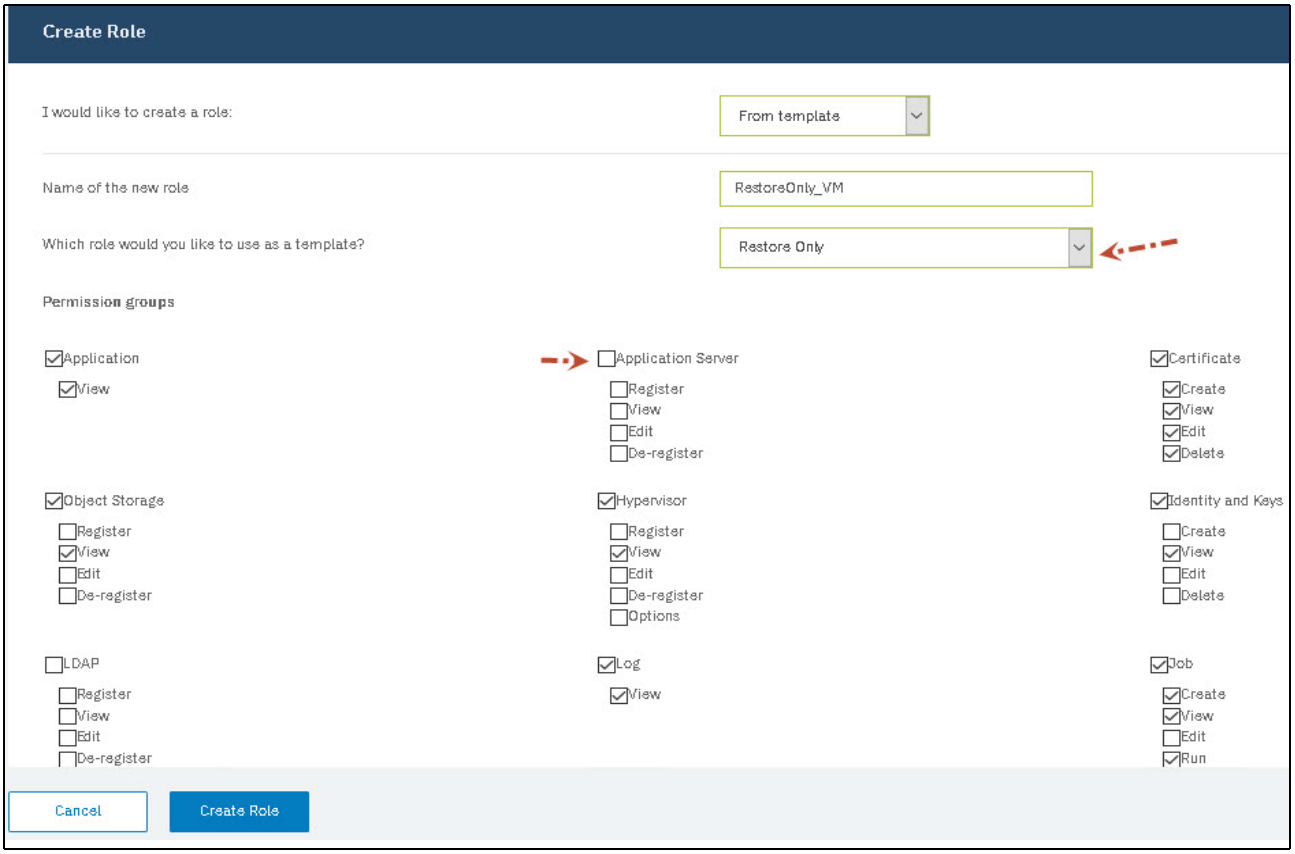

Figure 5-5 RestoreOnly_VM role creation based on Restore Only t*emplate*

To modify or delete a role, click the three dots (**...**), as shown in Figure 5-6.

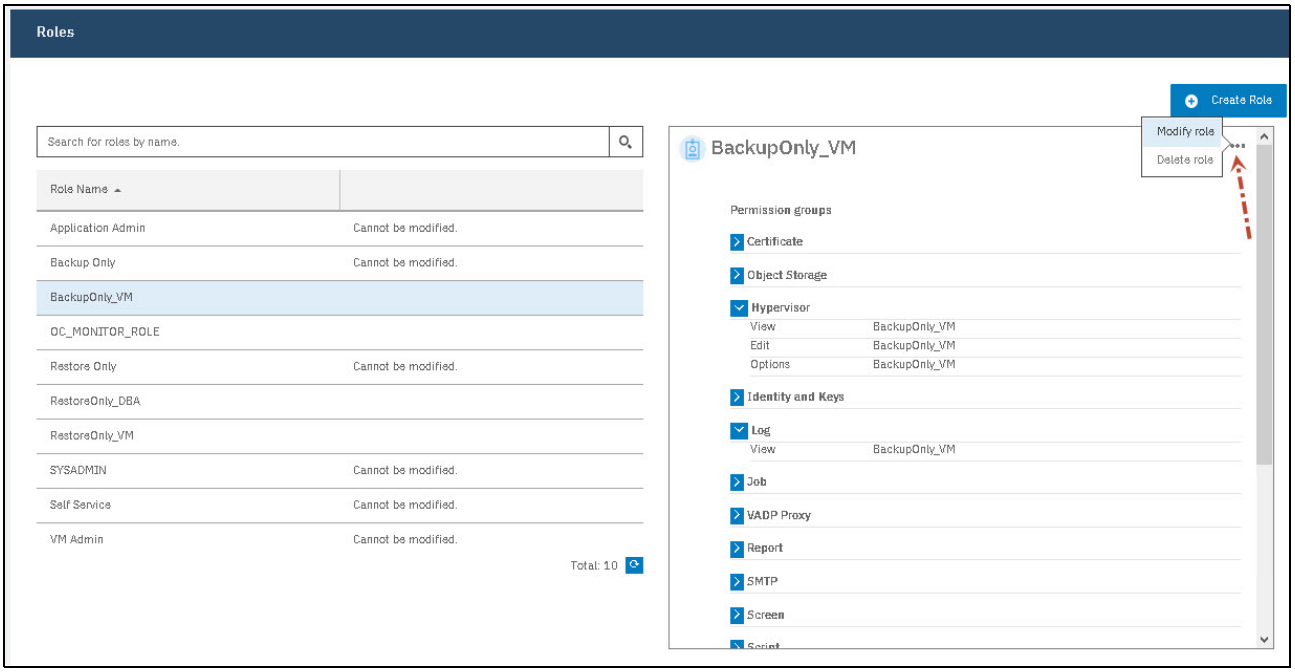

*Figure 5-6 Edit Role by selecting the three dots and s*electing *the Modify Role option*

5.1.4 Creating users

In our example, we are creating the seven users that are listed in Table 5-1, as they were presented in the use case description in Figure 5-2 on page 109.

These users are defined as local IBM Spectrum Protect Plus users, except for SQLBA and SQLRECO, which are defined as LDAP groups to show the integration of IBM Spectrum Protect Plus with LDAP. Before creating a user that is based on an LDAP groups, see how to configure the connectivity between LDAP and IBM Spectrum Protect Plus as described in 5.3.4, “Configuring LDAP and SMTP” on page 137

Table 5-1 List of User created for our practical example

User name	Resource group	Role
VMBA	resrouceVMW	BackupOnly_VM
VMRECO	resourceVMW	RestoreOnly_VM
ORABA	resourceORA_t2,resourceORA_t3	Backup Only (pre defined)
ORAT3RECO	resrouceORA_t3	RestoreOnly_DBA
ORAT2RECO	resourceORA_t2	RestoreOnly_DBA
SQLBA	resourceSQL	Backup Only (pre defined)
ClientGroupDBA	resrouceSQL	RestoreOnly_DBA

To create a user, select **Accounts** → **User** → **Add User.**

Figure 5-7 shows the example of ORABA User creation, as a local user.

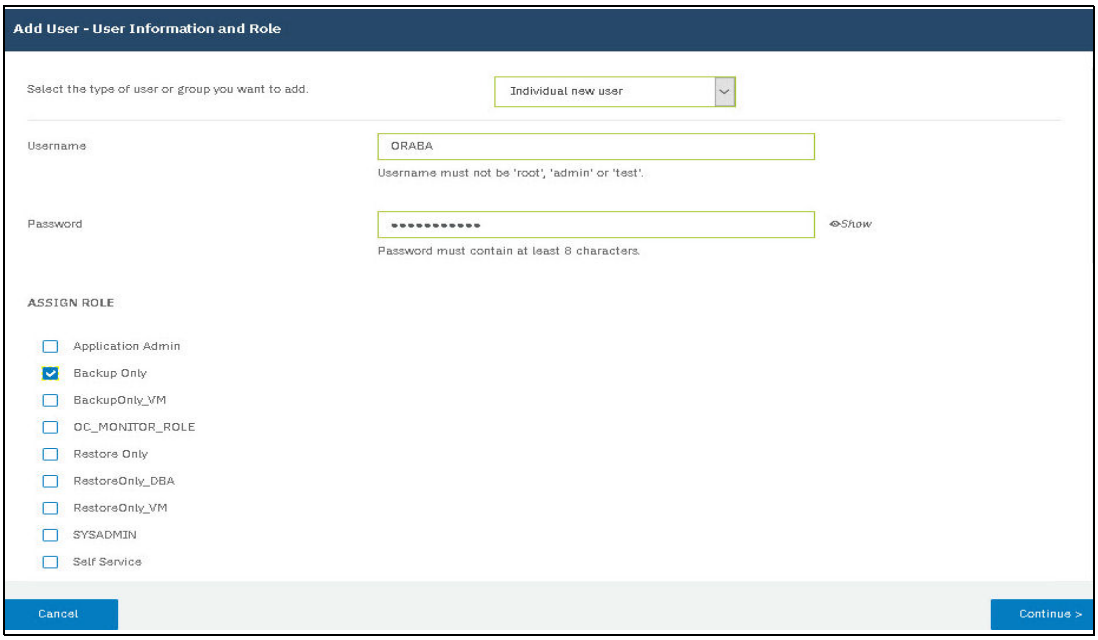

Figure 5-7 User Creation and role assignment

The next window in the User creation wizard (see shown Figure 5-8) is where we assigned the resource groups with which the user accesses only two specific databases environments.

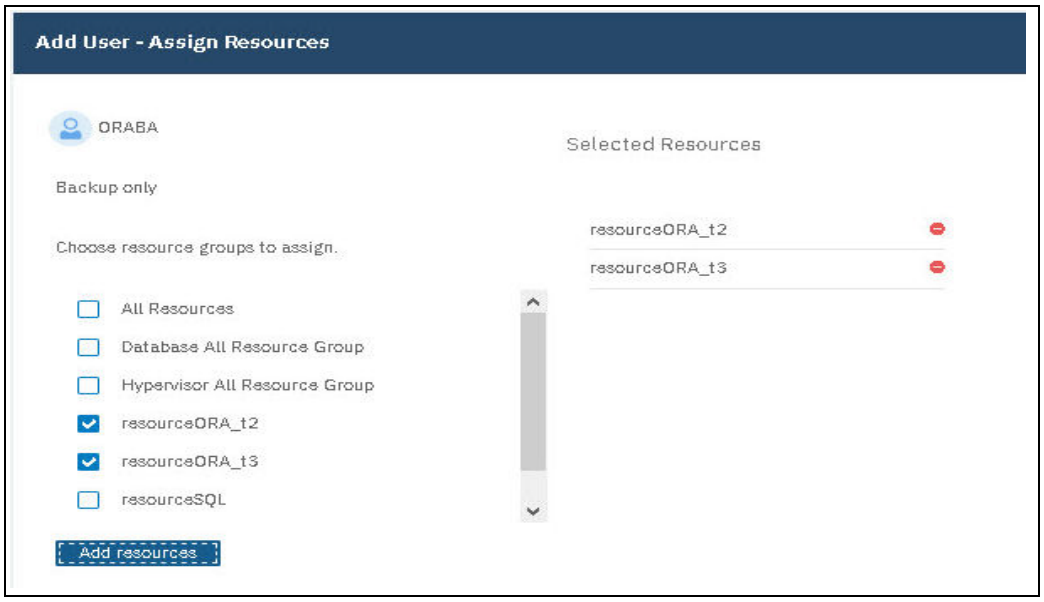

Figure 5-8 User creation resource group assignment window

Figure 5-9 shows the same step of user creation, but uses the LDAP Group instead of a specific user.

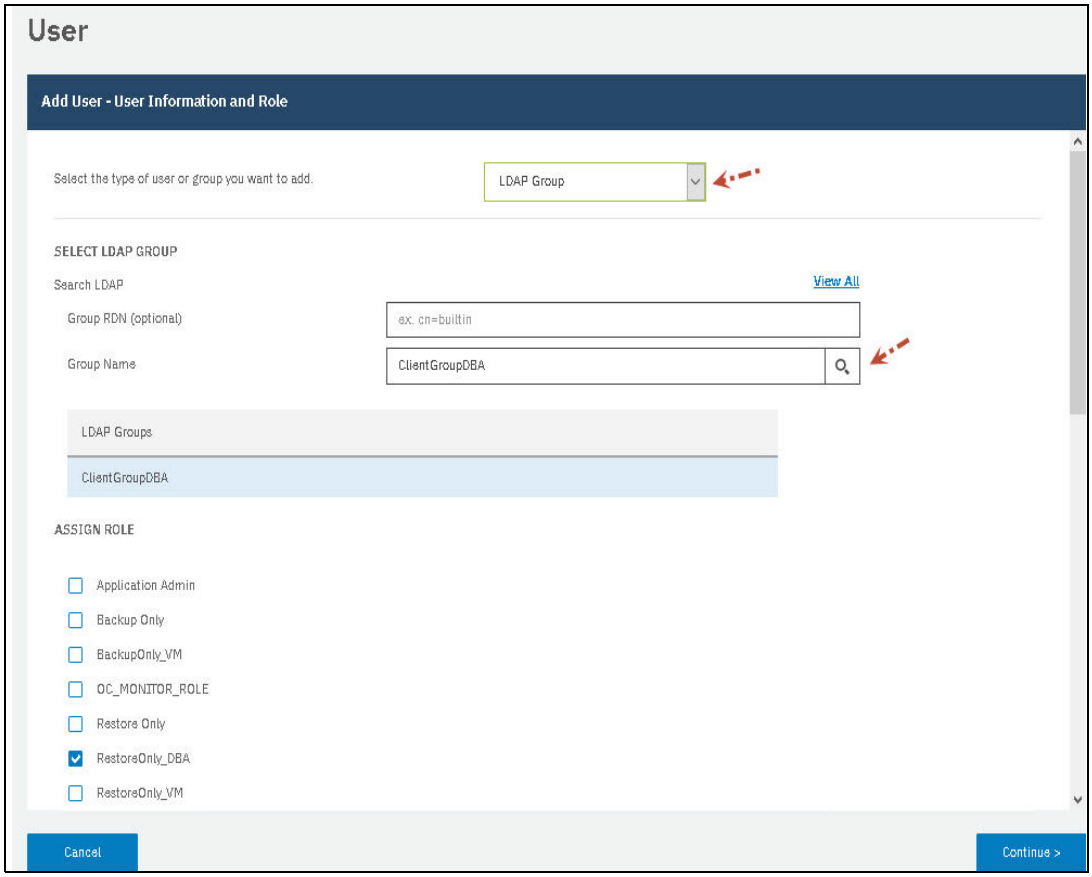

Figure 5-9 User group creation based on LDAP Group

Therefore, each user that is a member of the group ClientGroupDBA must have the Role Restoreonly_DBA and the resource group in our example because these properties are assigned to the group, as shown in Figure 5-10.

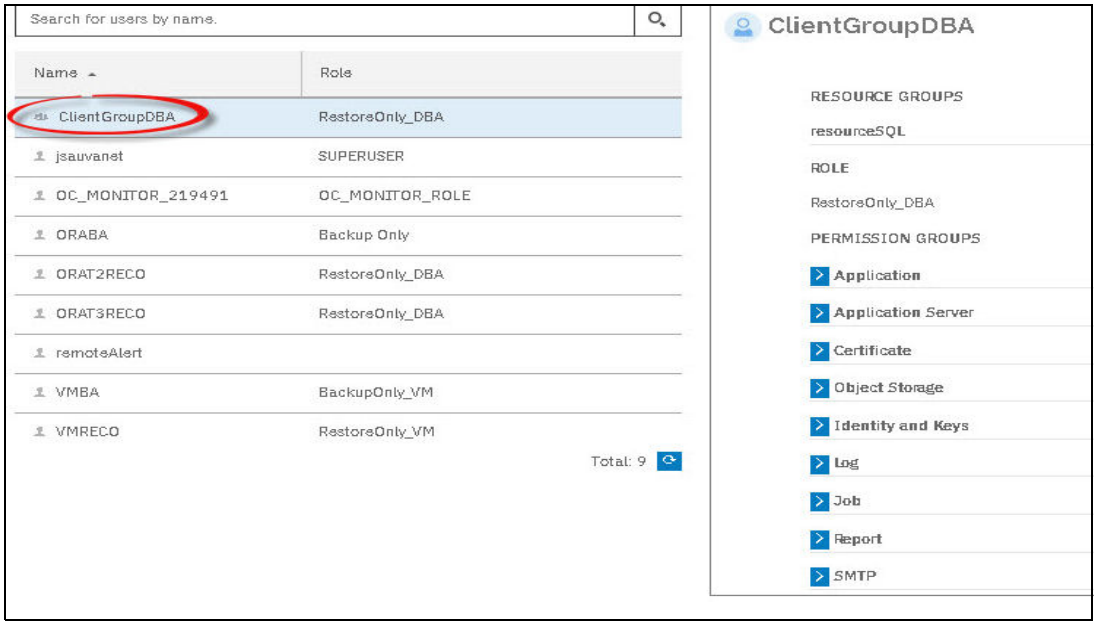

Figure 5-10 ClientGroupDBA which is an LDAP group being defined as Spectrum Protect User

5.2 Daily operations

This section describes the features that are available in IBM Spectrum Protect Plus for the daily operations. In addition to the main dashboard that shows an overview of the current and recent activity, the following features were introduced in version 10.1.6 that can help you in your daily activity:

- IBM Spectrum Protect Plus integrated to Spectrum Protect Operations Center
- Enable logging IBM Spectrum Protect Plus alerts to the system log, which can be configured from Global Preferences.

The following features are available from previous releases:

- Built in and configurable HTML or CSV reports
- Command line operations (see 5.4, “vSnap server CLI” on page 156)

5.2.1 IBM Spectrum Protect Plus in Spectrum Protect Operations Center

Starting with IBM Specturm Protect Plus 10.1.6, you can link IBM Spectrum Protect Plus to IBM Spectrum Protect Operation Center to get a dashboard view in the Operation Center.

Complete the following steps:

1. Start IBM Spectrum Protect Operations Center.
2. Select **Overview** → **Protect Plus** → **Add Server**.
3. Specify your IBM Spectrum Protect Plus server address and port (default is 443).
4. Accept the Certificate information.

5. Enter the log in credentials for a user account that can create custom user roles and user accounts on the IBM Spectrum Protect Plus Server, as shown in Figure 5-11.

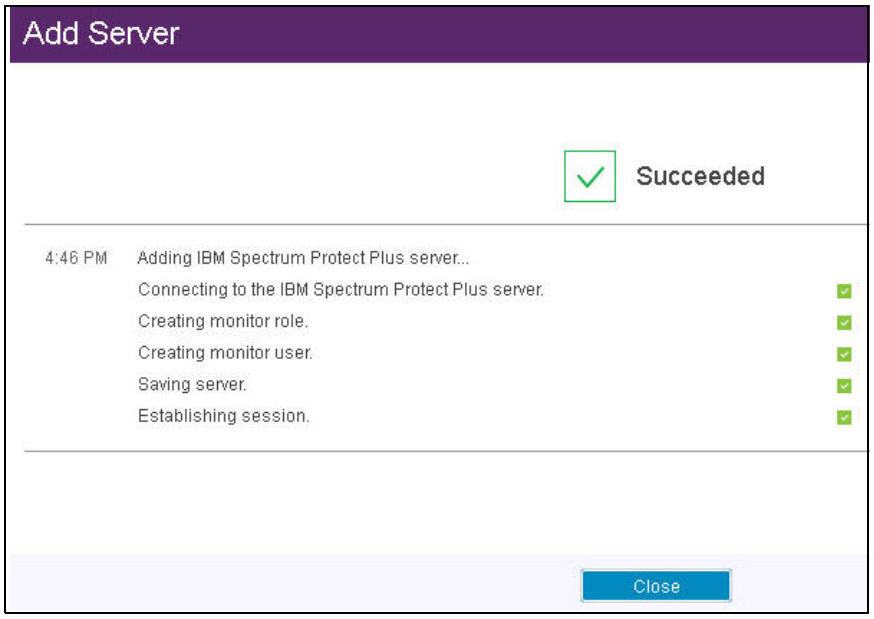

Figure 5-11 Operations Center end of IBM Spectrum Protect Plus server registration

This operation creates a user and role dedicated to operation center (for information exchange) in the IBM Spectrum Protect Plus environment, as shown in Figure 5-12

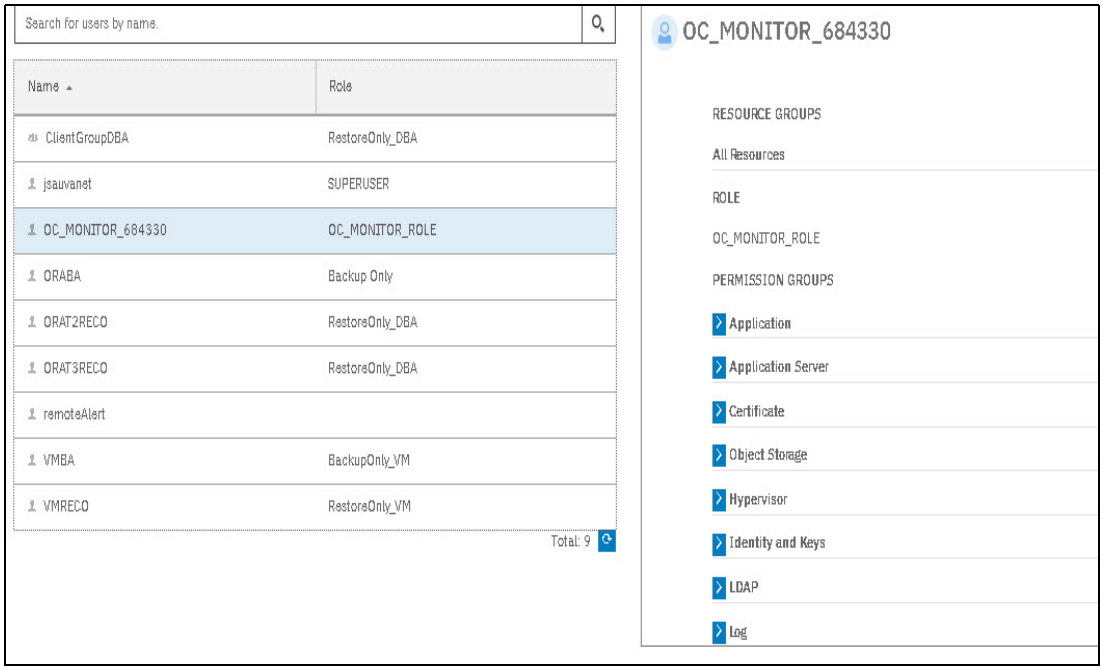

Figure 5-12 Operations Center user in IBM Spectrum Protect Plus with its created role

After the IBM Spectrum Protect Plus server is registered in the Operations Center, a new dashboard that shows a summary of IBM Spectrum Protect Plus key indicators is available. This dashboard has the same look and feel than the traditional Operations Center dashboard, as shown in Figure 5-13. You can interact with Jobs and Operations gauges to get more information about the Backup, Restore, Inventory, and Maintenance activities.

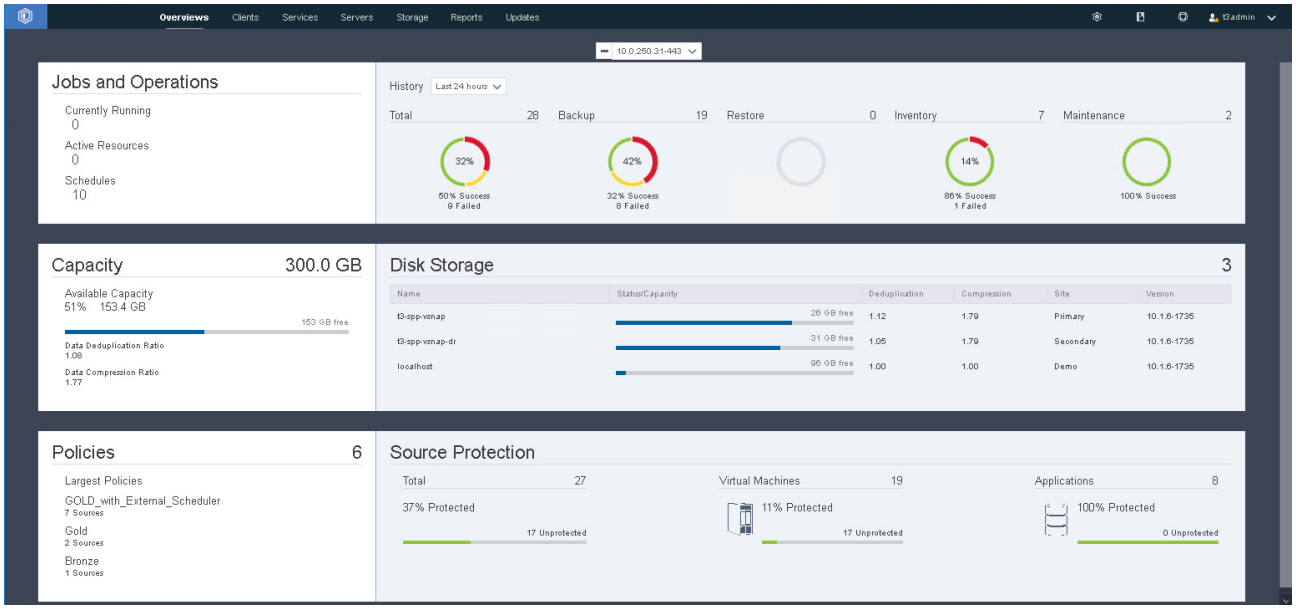

Figure 5-13 Operation Center's IBM Spectrum Protect Plus dashboard

Also, by setting the IBM Spectrum Protect Operations Center URL in **System Configuration** → **Global Preferences,** a new icon (see Figure 5-14) appears in the top right corner of the IBM Spectrum Protect Plus interface. This icon is a link to the Operations Center.

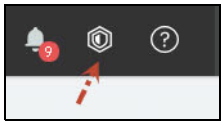

Figure 5-14 Operations Center link from IBM Spectrum Protect Plus interface

By clicking this icon, a pop-up window appears that includes the Operations Center log in window. You then must log in by using the Operations Center credentials to browse through the IBM Spectrum Protect Plus dashboard.

5.2.2 Built-in and custom reports

IBM Spectrum Protect Plus reports are available by selecting **Reports and Logs** → **Reports**.

A set of predefined reports are available to run immediately, or you can define a schedule that automatically generates the report and sends results by email to any specified address (for more information about how to configure SMTP, see 5.3.4, "Configuring LDAP and SMTP" on page 137).

Figure 5-15 shows the Reports Schedule wizard. By clicking the “**+**” sign, you can create a customized report from the predefined report. By using this option, you can filter out information to build different reports for different audiences or different entities if you are running in a multi-tenant environment.

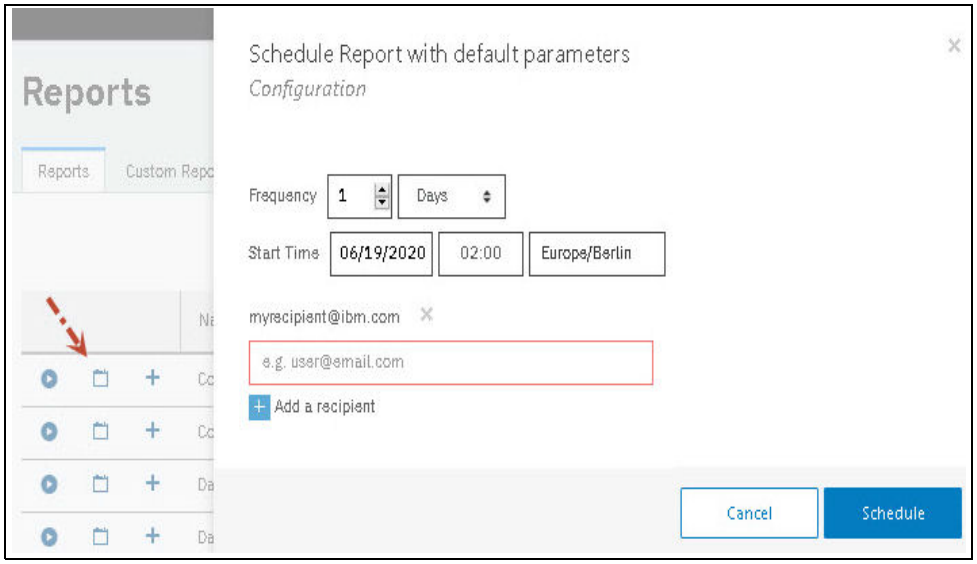

Figure 5-15 Schedule report wizard

Any custom report that you created can be found in **Reports and Logs** → **Reports** → **Custom Reports** tab.

The following report categories are available:

- *Backup Storage utilization* reports the vSnap server storage utilization or VM backup storage consumption.
- *Protection* reports the various workload backup status.

  Specific reports that include “SLA Policy RPO compliance” in their name are giving a clear view of whether your backups are compliant with the SLA Policy. Figure 5-16 on page 121 shows an example of the Database SLA Policy RPO Compliance where we can see that half of our selected database is in compliance. In addition to the overview section of these reports, detailed information is provided to better understand what is or what is not in compliance.
- *System* reports IBM Spectrum Protect Plus/vSnap server configuration, Jobs, and License information:
  - Configuration
  - Job reports the Job status for IBM Spectrum Protect Plus housekeeping activities and backup/restore related activities
  - License provides information about the license information that gives the front-end capacity per workload types (VM, Physical Machine, Office365, Container Persistent Volume, and so on)
- *VM Environment* reports information about the VMware environment. You can view information about datastore space and usage, VM snapshots sprawl, VM LUNs, and VM storage.

An example of a Database SLA Policy RPO compliance report is shown in Figure 5-16.

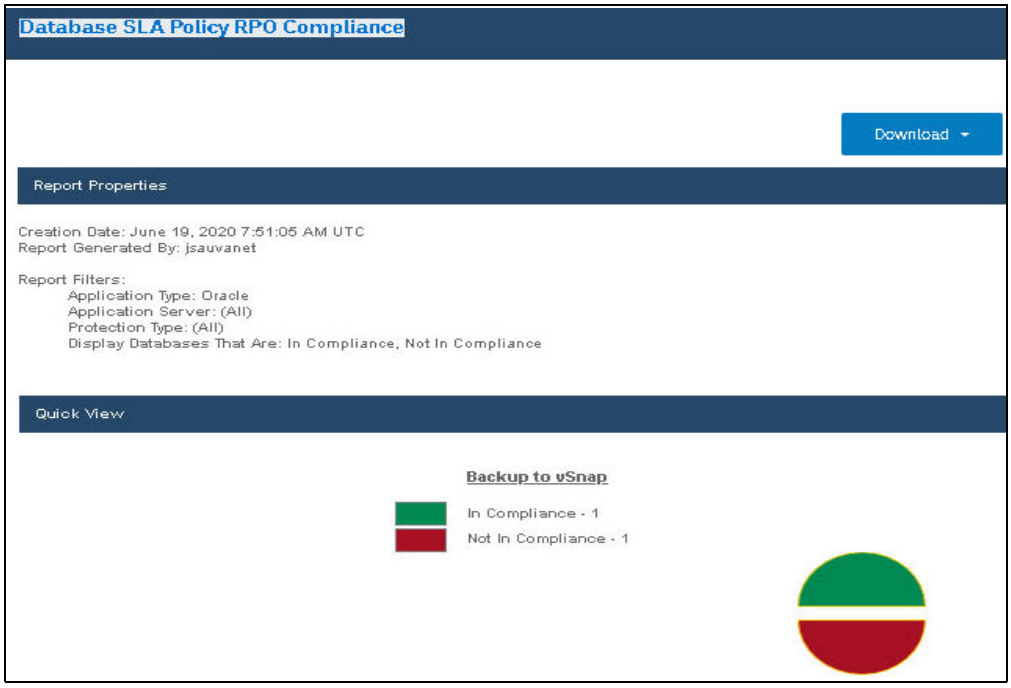

Figure 5-16 Database SLA Policy RPO compliance report

5.2.3 Morning Healthcheck routine

As part of the daily job, you must ensure that your IBM Spectrum Protect Plus environment is healthy. You can define and schedule the following suggested reports as your daily routine:

- By using the System report category, you can create custom report for IBM Spectrum Protect Plus Housekeeping, as shown in Figure 5-17.

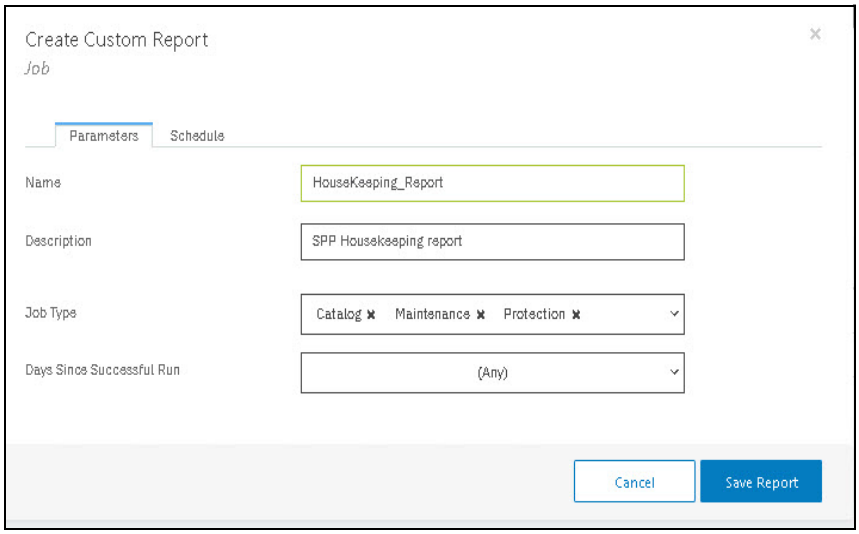

Figure 5-17 IBM Spectrum Protect Plus Housekeeping custom report creation

- Use the System report category to create customer report for IBM Spectrum Protect Plus Backup and Recovery results, as shown in Figure 5-18.

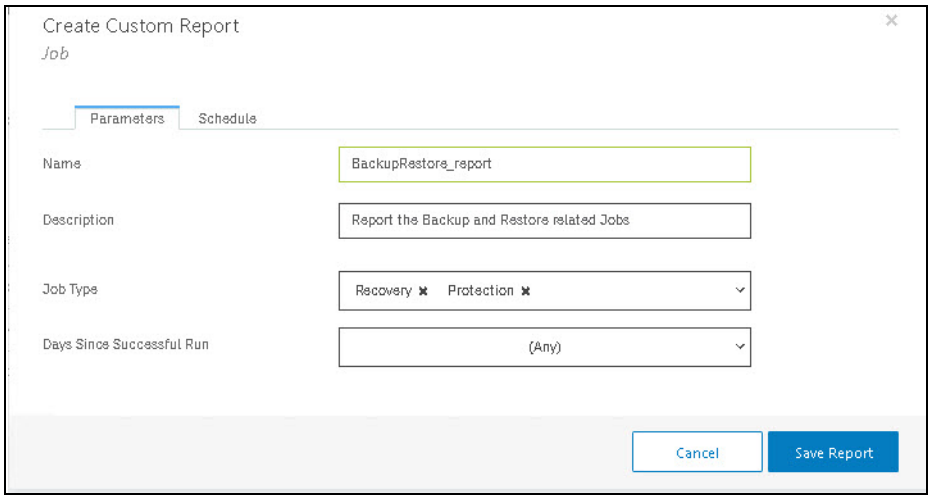

Figure 5-18 IBM Spectrum Protect Plus Backup and REcovery operation report creation

- By using the Protection report category, you can create three custom reports to view the SLA Policy RPO compliance for Database, File System, and VM. Figure 5-19 on page 123 shows an example of SLA Policy RPO compliance report for Applications.

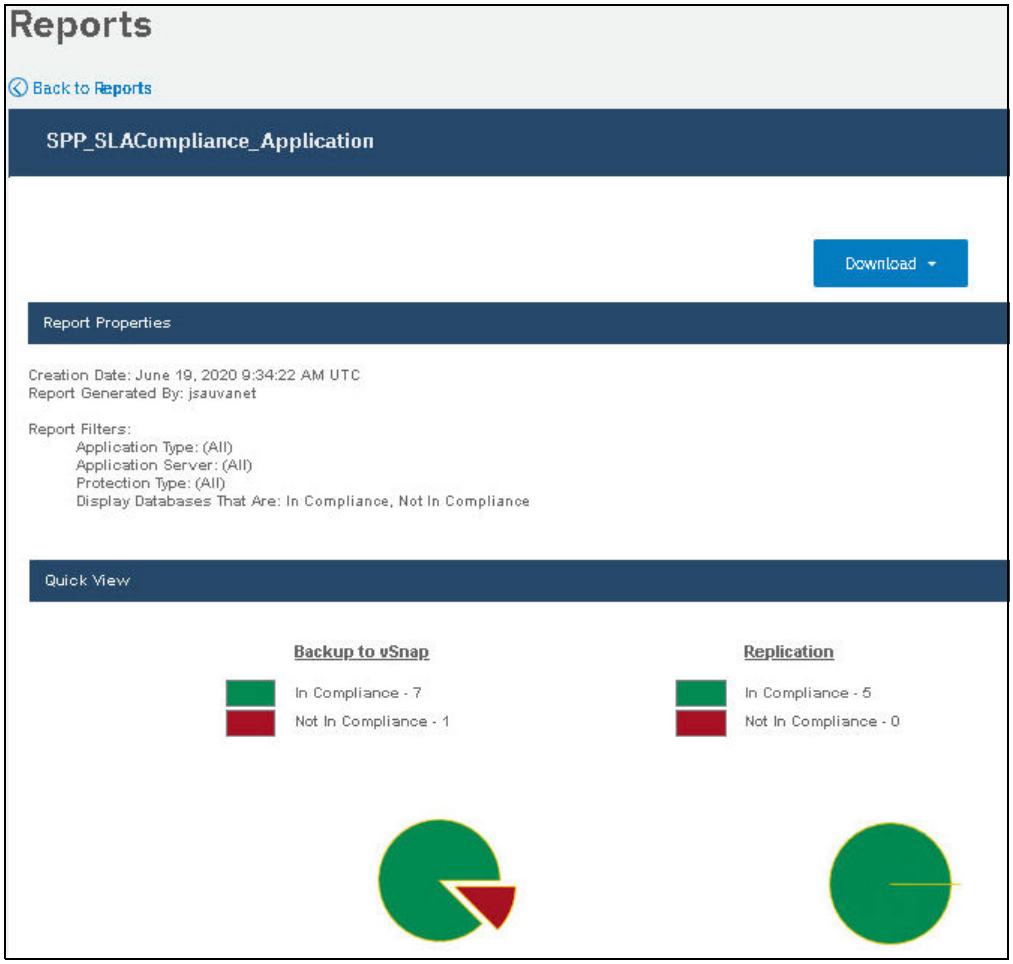

Figure 5-19 IBM Spectrum Protect Plus SLA Policy RPO compliance report for Applications

Note: The SLA Policy RPO compliance report can be used to review the vSnap server replication status. As shown in Figure 5-19, whenever the SLA Policy is configured to replicate the vSnap server and the replication did not occur or is not up to date, the SLA Policy RPO report flags the specific resource (here, the VM) as non-compliant. Whenever the compliance is not met, more information is provided in the report, as shown in Figure 5-20 on page 124.

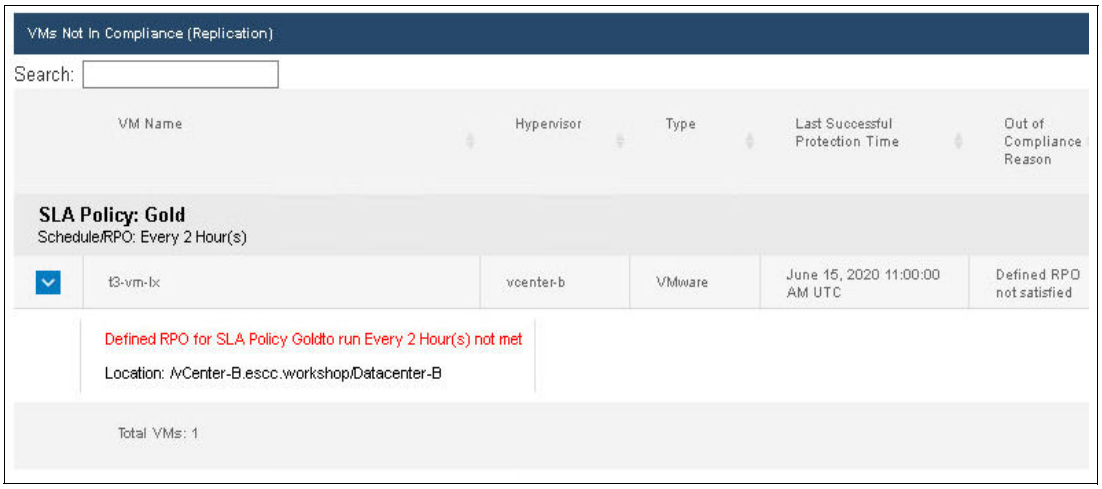

Figure 5-20 SLA Policy RPO report, non-compliance detailed information

Note: The content of the report is limited to what the Resource Group is allowing for the user who is defining or running that report.

5.3 Maintenance

This section provides guidance about how IBM Spectrum Protect Plus users can maintain the product. We discuss different approaches to handle daily operations by using the Graphical User Interface (GUI) and the Command Line Interface (CLI).

5.3.1 Update IBM Spectrum Protect Plus components

If you have an IBM Spectrum Protect Plus environment running, you might need to upgrade to a newer version. However, you must consider several issues before you upgrade your systems.

You can update the IBM Spectrum Protect Plus virtual appliance, vSnap servers, and the VADP proxy servers to get the latest features and enhancements. Software patches and updates are installed by using the IBM Spectrum Protect Plus administrative console or CLI for these components.

You must update the IBM Spectrum Protect Plus server first. After the IBM Spectrum Protect Plus update completes, you must update any external vSnap and VADP proxy servers in your environment.

Updating IBM Spectrum Protect Plus virtual appliance

Use the IBM Spectrum Protect Plus Administrative Console to update the virtual appliance. The IBM Spectrum Protect Plus server update can be run offline or online if you have internet access. Offline means that you must provide a downloaded upgrade package; the online update process downloads the required packages from the internet:

- For online updates, access to the FTP site `public.dhe.ibm.com`. The administrator console checks for available updates automatically and displays available updates.

Note: If you want to run online updates but can see only the offline mode, check your internet connectivity and reattempt access to the FTP site (`public.dhe.ibm.com`).

- For offline updates, download the prerequisite IBM Spectrum Protect Plus update file from the IBM Fix Central website to a directory on the computer that is running the browser for the Administrative Console. For more information about files for IBM Spectrum Protect version 10.1.6.x and how to obtain it from Fix Central, see this web page.
- Ensure that your IBM Spectrum Protect Plus environment is backed up before you run updates. For more information about backing up your environment, see IBM Knowledge Center.
- Ensure that no jobs are running during the update procedure. Pause the schedule for any jobs that have a status of IDLE or COMPLETED. You can perform this task from IBM Spectrum Protect Plus GUI by selecting **Jobs and Operations** → **Schedule** and then, click **Pause Schedule** for each of the jobs.

Complete the following steps to update the IBM Spectrum Protect Plus server:

1. From a supported web browser, access the Administrative Console at `https://hostname:8090/`, where *`hostname`* is the IP address of the VM where the application is deployed.
2. In the log-in window, select **Authentication Type** → **System**.
3. Enter the password for the userID serveradmin.

Note. Starting with version 10.1.5, the initial password of the user ID serveradmin is `sppDP758-SysXyz`.

4. Select **Updates and hotfix management**.
5. Click **Browse** to browse for the update (`.iso`) file to upload to the appliance and then, click **Upload Update Image (or) Hotfix**, as shown in Figure 5-21 on page 126.

Note: You can select only one update file at a time.

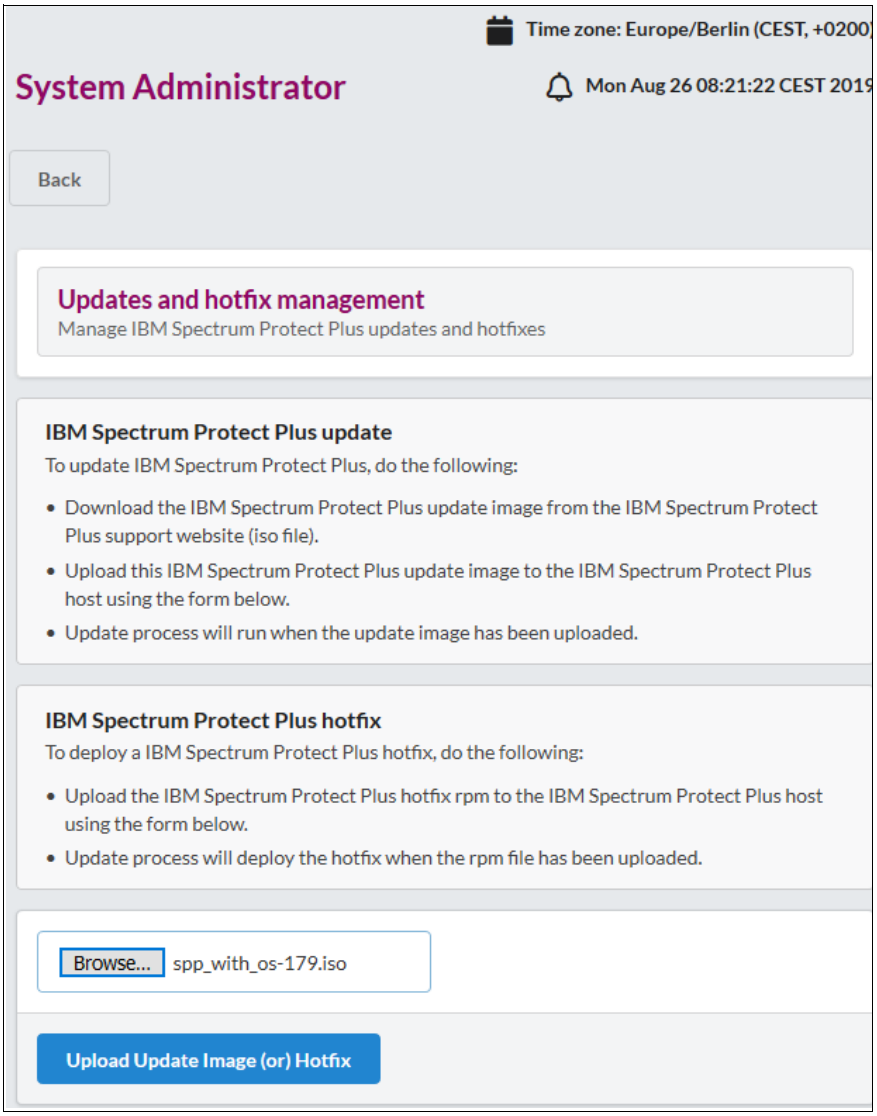

Figure 5-21 Administrative Console: Update and hotfix management

6. Acknowledge the warning, as shown in Figure 5-22.

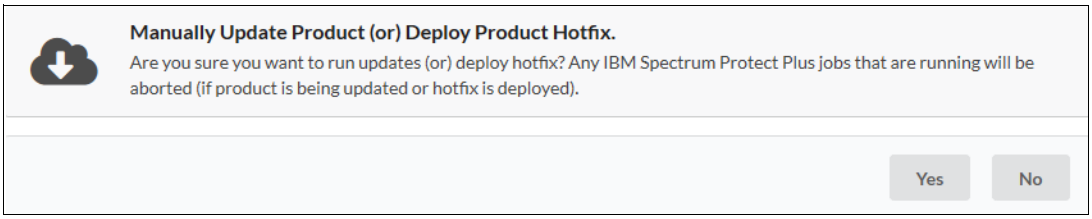

Figure 5-22 Run updates (or) deploy hotfix warning

Note: The update process begins after the update image is uploaded to the appliance.

7. When the update is complete, the VM where the application is deployed automatically restarts.

Note: HTML content from previous versions of IBM Spectrum Protect Plus might be stored in your browser cache. Clear the cache before logging in to an updated version of IBM Spectrum Protect Plus to ensure that you are viewing the latest content changes.

8. Do not forget to release the paused jobs, if any such jobs exist. Log in to IBM Spectrum Protect Plus GUI and select **Jobs and Operations** → **Schedule** and find the jobs that you paused. Click **Actions** → **Release Schedule**.

Updating vSnap servers

The built-in vSnap server is updated with the IBM Spectrum Protect Plus server. You must update extra vSnap servers that are installed on virtual or physical appliances separately.

Test restore jobs must complete before an update to vSnap server is started. Jobs that are not completed or canceled when an upgrade is started are not visible when the update completes. If jobs are not visible when the update completes, rerun the test restore jobs.

You also might be required to update the operating system for the vSnap servers before updating the servers. For more information about operating system requirements, see Component requirements at IBM Knowledge Center.

Checking the version and operating system

To check the version and operating system for your vSnap servers, log on to the vSnap server by using SSH as the serveradmin user. Use the vSnap CLI to run the **`vsnap system info`** command, as shown in Example 5-1.

Example 5-1 The vSnap system info

```
[serveradmin@t3-spp-vsnap SPP]$ vsnap system info
ID: 530149c9d9f24263a28c4f76dcd6ee6e
INIT STATUS: Ready
MULTIPOOL SUPPORT: No
LOCAL REPL SUPPORT: No
HOSTNAME: t3-spp-vsnap
FQDN: t3-spp-vsnap.escc.workshop
OS NAME: CentOS Linux
OS VERSION: 7.7.1908
API VERSION: 1.0
VSNAP VERSION: 10.1.6-1735
NGINX VERSION: 1.12.2-3.el7
UWSGI VERSION: 2.0.17.1-2.el7
NFS VERSION: 1.3.0-0.61.el7
SAMBA VERSION: 4.9.1-10.el7_7
ZFS VERSION: 0.8.3.10.1.6-20200511.el7
CRYPTSETUP VERSION: 2.0.3-5.el7
ADS DOMAIN: N/A
ADS SERVER: N/A
ADS WORKGROUP: N/A
MAINTENANCE MODE: No
```

Updating a vSnap server

Before you begin the update, ensure that you backed up your IBM Spectrum Protect Plus environment, as described in IBM Knowledge Center.

Download the vSnap server `.run` update file and copy it to a temporary location on the vSnap server.

To update a vSnap server, complete the following steps:

1. Log on to the vSnap server by using SSH as the serveradmin user.
2. From the directory where the <*updatefilename*> run file is stored, make the file executable and run the installer by running the following commands:

   ```
   chmod +x <updatefilename>.run
   sudo ./<updatefilename>.run
   ```

   You are prompted to agree to the License conditions. After confirmation, the vSnap server packages are installed.

   Accept the reboot question to commit the changes.
3. Do not forget to release the paused jobs, if any exist. Log in to IBM Spectrum Protect Plus GUI and select **Jobs and Operations** → **Schedule** and find the paused jobs. Select **Actions** → **Release Schedule**.

Updating the operating system for a physical vSnap server

If you have installed the vSnap server on a machine that is running Red Hat Enterprise Linux, you must update the operating system to version 7.7 before you update the vSnap server. For instructions about how to update the operating system, see the Red Hat Enterprise Linux documentation.

Updating the operating system for a virtual vSnap server

If the operating system is CentOS Linux version 7.4 or earlier, you must update the operating system before you update the vSnap server. To update the operating system, follow the instructions in Updating vSnap servers to version 10.1.2. The version 10.1.2 installation includes CentOS Linux version 7.5.

Updating VADP proxies

Updating the IBM Spectrum Protect Plus virtual appliance automatically updates all the VADP proxies that are associated with the virtual appliance. In rare scenarios, such as loss of network connectivity, you must update the VADP proxy manually.

If a VADP proxy update is available for external proxies during a restart of the IBM Spectrum Protect Plus virtual appliance, the update will be automatically applied to any VADP proxy associated with an identity. To associate a VADP proxy with an identity, navigate to **System Configuration** → **VADP Proxy**. Click the **options icon** and select **Set Options**. Through the User setting, select a previously entered username and password for the VADP proxy server.

To update a VADP proxy manually, complete the following steps:

1. Navigate to the **System Configuration** → **VADP Proxy** page in IBM Spectrum Protect Plus GUI.
2. The VADP Proxy page displays each proxy server. If a newer version of the VADP proxy software is available, an update icon displays in the **Status** field.
3. Ensure that there are no active jobs that use the proxy, and then click the update icon.

   The proxy server enters a suspended state and installs the latest update. When the update completes, the VADP proxy server automatically resumes and enters an enabled state.

If you are attempting to update as a non-root user, special instructions need to be followed in order to push-install or push-update a VADP proxy:

1. Create a file in the ***/etc/sudoers.d/*** directory.

   ```
   sudo cd /etc/sudoers.d/
   ```

2. Write the text to the file and save it by pressing **CTRL+D** on the keyboard when done.

   ```
   sudo cat > 99-vadpuser
            Defaults !requiretty
            vadpuser ALL=NOPASSWD: /tmp/cdm_guestapps_vadpuser/runcommand.sh
       <<Press CTRL+D>>
   ```

3. Set the appropriate permissions on the file.

   ```
   sudo chmod 0440 99-vadpuser
   ```

Applying early availability updates (efix)

Early availability updates provide fixes for authorized program analysis reports (APARs) and minor issues between IBM Spectrum Protect Plus releases. These updates are available in bundles from this IBM Support web page.

Early availability updates might not contain fixes for all IBM Spectrum Protect Plus components. For more information about how to obtain and install interim fixes, see the download information that is published when the fixes are available.

5.3.2 IBM Spectrum Protect Plus troubleshooting (log files)

For each job that 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:

- Click **Download .zip** in the **Jobs and Operations** menu to download the log collection 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.

Collecting log files for troubleshooting

To troubleshoot the IBM Spectrum Protect Plus application, you can download an archive of log files that are generated by IBM Spectrum Protect Plus.

To collect log files for troubleshooting, complete the following steps:

1. Click the user menu, and then, click **Download System Logs**. (The download process can take some time to complete.)
2. Open or save the file log `.zip` file, which contains individual log files for different IBM Spectrum Protect Plus components.

Log location

The log location IBM Spectrum Protect Plus is available at this IBM Support web page.

To view the real-time log files for the different components in IBM Spectrum Protect Plus, SSH to the IBM Spectrum Protect Plus server. The log files can be found in the locations that are listed in Table 5-2 on page 130.

Table 5-2 IBM Spectrum Protect Plus log locations

Log	Location
Virgo logs	`/opt/virgo/serviceability/logs`
vSnap log	`/opt/vsnap/log`
VMDKbackupproxy	`/data/log/vmdkbackupproxy`
RabbitMQ	`/data/log/rabbitmq`
nodejscdmservice	`/data/log/node-cdm-service/`
Mongo	`/data/log/mongo`
Guestdeployer	`/data/log/guestdeployer`
adminconsole	`/data/log/adminconsole/`

5.3.3 Managing the vSnap server

The following sections describe aspects of managing the vSnap server.

For more information about vSnap server management, see the documentation that is available at this IBM Support web page.

Managing disks

A vSnap server creates a storage pool by using disks that are provisioned to the vSnap server. In the case of virtual deployments, the disks can be RDM or virtual disks provisioned from data stores on any backing storage. In the case of physical deployments, the disks can be local or SAN storage attached to the physical server. The local disks might have external redundancy enabled by way of a hardware RAID controller, but if not, a vSnap server can also create RAID-based storage pools for internal redundancy.

Disks that are attached to vSnap servers must be thick provisioned. If disks are thin provisioned, the vSnap server will not have an accurate view of free space in the storage pool, which might lead to data corruption if the underlying data store runs out of space.

If a vSnap server was deployed as part of a virtual appliance, it includes a 100 GB starter virtual disk that can be used to create a pool. You can add disks before or after creating a pool and use them to create a larger pool or expand an existing pool. I

If job logs report that a vSnap server is reaching its storage capacity, more disks can be added to the vSnap pool. Alternatively, creating SLA policies force backups to use an alternative vSnap server.

It is essential to protect against vSnap server file system corruption, which might be caused by a VMware data store on a vSnap server reaching its capacity. Create a stable environment for virtual vSnap servers that do not use RAID configurations by using thick provisioned VMDKs. Replicating to external vSnap servers provides more protection.

A vSnap server becomes invalidated if the vSnap pool is deleted or if a vSnap server disk is deleted in a non-redundant RAID configuration. All data on the vSnap server is lost. If your vSnap server becomes invalidated, you must unregister the vSnap server by using the IBM Spectrum Protect Plus interface and then, run the maintenance job. When complete, the vSnap server can be registered again.

Detecting disks

If you add disks to a vSnap server, use the command line or the IBM Spectrum Protect Plus user interface to detect the newly attached disks.

Command line: Run the **`vsnap disk rescan`** command.

User interface: Click **System Configuration** → **Backup Storage** → **Disk** in the navigation pane, and then click the **Actions menu** next to the relevant vSnap server and select **Rescan**.

Creating a storage pool

If you completed the simple initialization procedure described in Chapter 3, “Installation and deployment” on page 47, a storage pool was created automatically, and the information in this section is not applicable.

To complete an advanced initialization, use the **`vsnap pool create`** command to create a storage pool manually. Before you run the command, ensure that one or more unused disks are available as described in Showing disks. For information about available options, pass the --help flag for any command or subcommand.

Specify a user-friendly display name for the pool and a list of one or more disks. If no disks are specified, all available unused disks are used. You can choose to enable compression and deduplication for the pool during creation. You can also update the compression or deduplication settings at a later time by using the **`vsnap pool update`** command.

The pool type that you specify during the creation of the storage pool dictates the redundancy of the pool:

- raid0

  This is the default option when no pool type is specified. In this case vSnap assumes your disks have external redundancy, for example, if you use virtual disks on a data store backed by redundant storage. In this case, the storage pool will have no internal redundancy.

  After a disk is added to a raid0 pool, it cannot be removed. Disconnecting the disk results in the pool becoming unavailable, which can be resolved only by destroying and recreating the pool.

- raid5

  When you select this option, the pool is composed of one or more RAID5 groups each consisting of three or more disks. The number of RAID5 groups and the number of disks in each group depends on the total number of disks you specify during pool creation. Based on the number of available disks, vSnap chooses values that maximize total capacity while also ensuring optimal redundancy of vital metadata.

- raid6

  When you select this option, the pool is composed of one or more RAID6 groups each consisting of four or more disks. The number of RAID6 groups and the number of disks in each group depends on the total number of disks that you specify during pool creation. Based on the number of available disks, vSnap chooses values that maximize total capacity while also ensuring optimal redundancy of vital metadata.

Initializing vSnap server pool

After a vSnap server is installed physically or deployed as a virtual appliance, a pool must be configured. In Example 5-2 on page 132, we show with the **`vsnap pool show`** and the **`vsnap volume show`** commands on the command line interface (CLI) that no pool or volume is configured yet.

Example 5-2 The vSnap command line interface (CLI)

```
[serveradmin@t3-spp-vsnap ~]$ vsnap pool show
TOTAL: 0
[serveradmin@t3-spp-vsnap ~]$ vsnap volume show
ERROR: VolumeInfoError: Failed to collect volume information
```

Before the vSnap pool can be used, it must be initialized. The simple initialization method that is available within the IBM Spectrum Protect Plus GUI was described in “Initializing vSnap pool” on page 72.

For servers that are deployed in a physical environment, the vSnap server console offers more options for initializing the server, including the ability to create a storage pool by using advanced redundancy options and a specific list of disks.

To initialize a vSnap server by using the vSnap server console, complete the following steps:

1. Log in to the vSnap server console with the user ID serveradmin. You can also use a user ID that has vSnap admin privileges that you create by using the **`vsnap user create`** command.

   Note: In Spectrum Protect Plus version 10.1.4, the initial password of the user ID `serveradmin` is `sppDP758`. Starting with version 10.1.5, the password was changed to `sppDP758-SysXyz`.

2. Run the **`vsnap system init --skip_pool`** command. The command requires no further interaction and completes all initialization tasks except for the creation of a storage pool. The process might take 5 - 10 minutes to complete.

After the initialization process is completed, the **Status/Capacity** column shows a utilization bar for your vSnap server. Let’s check now what we see on the vSnap CLI when we enter the show command for the vSnap pool again, as shown in Example 5-3.

Example 5-3 The vsnap pool show command list details of a vSnap pool

```
[serveradmin@t3-spp-vsnap ~]$ vsnap pool show
TOTAL: 1

ID: 1
NAME: primary
POOL TYPE: raid0
STATUS: ONLINE
HEALTH: 100
COMPRESSION: Yes
COMPRESSION RATIO: 1.01
DEDUPLICATION: No
DEDUPLICATION RATIO: 1.00
ENCRYPTION:
    ENABLED: Yes
    TYPE: disk

TOTAL SPACE: 100.00GB
FREE SPACE: 96.39GB
USED SPACE: 3.61GB
DATA SIZE BEFORE DEDUPLICATION: 116.50KB
DATA SIZE BEFORE COMPRESSION: 47.50KB
```

```
CREATED: 2019-06-18 13:03:50 UTC
UPDATED: 2019-06-18 13:03:50 UTC
DISKS PER RAID GROUP: 1
DISKS IN POOL:
    RAID0:
        /dev/dm-3
```

The command **vsnap pool show** in Example 5-3 on page 132 lists the details of our vSnap pool. We can see a pool of 100 GB capacity with one disk where compression and encryption is enabled and deduplication is disabled.

In Example 5-4, the **vsnap volume show** command lists one volume in the pool (ID 1). This volume with the ID 1 refers to a cloud cache area that is created with initialization of the vSnap pool. It is not used for backup data. For more information, see “Preparing the disk cache area” on page 420.

Example 5-4 The vsnap pool show command list the cloud cache volume after initialization

```
[serveradmin@t3-spp-vsnap ~]$ vsnap volume show
ID | TYPE       | POOL | IS CLONE | TOTAL    | FREE    | USED    | NAME                | TAGS
-------------------------------------------------------------------------------------------
1  | filesystem | 1    | No       | 100.00GB | 96.39GB | 24.00KB | vsnap_metadata_cloud | N/A
```

If we now use the **vsnap snapshot show** command, it does not list any snapshot, which is at this time expected, because no backup was done so far.

Expanding vSnap pool

If IBM Spectrum Protect Plus reports that a vSnap server is reaching its storage capacity, the vSnap pool must be expanded. To expand a vSnap pool, you must first add virtual or physical disks on the vSnap server by adding virtual disks to the vSnap server VM or adding physical disks to the vSnap physical server. For more information about creating virtual disks, see the vSphere documentation.

To expand a vSnap pool, complete the following steps:

1. In the navigation pane, click **System Configuration** → **Backup Storage** → **Disk**.
2. Select **Actions** → **Rescan** for the vSnap server that you want to rescan.
3. Click the manage icon the manage icon that is associated with the vSnap server, and then expand the **Add New Disks to Backup Storage** section.
4. Add and save the selected disks. The vSnap pool expands by the size of the disks that are added.

Use the command line to expand a storage pool:

Command line: Run the **vsnap pool expand** command, as shown in Example 5-5. For information about available options, pass the **--help** flag for any command or subcommand.

Example 5-5 The vsnap pool disk expand --help command

```
[serveradmin@t3-spp-vsnap ~]$ vsnap pool expand --help
Usage: vsnap pool expand [OPTIONS]

  Expand a storage pool by adding disks.

Options:
```

```
  --id TEXT         ID of the pool to expand.  [required]
  --disk_list TEXT  Comma separated list of disks (name or UUID). If not
                    specified, all available disks are used.
  --force           Skip checking if disks are unused. Use this option if you
                    want to use a partition instead of an entire disk.
  --help            Show this message and exit.
[serveradmin@t3-spp-vsnap ~]$
```

The following example shows an expansion of the `vsnap_pool` with another disk (100.00 GB).

1. List the current pool by running the **`vsnap pool show`** command. It shows only one disk is assigned to this pool (see last line `DISKS IN POOL:`). Note that the pool ID is 1 and the TOTAL SPACE is 100.00 GB, as shown in Example 5-6.

Example 5-6 The vsnap pool show command before adding a disk

```
[serveradmin@t3-spp-vsnap ~]$ vsnap pool show
TOTAL: 1

ID: 1
NAME: primary
POOL TYPE: raid0
STATUS: ONLINE
HEALTH: 100
COMPRESSION: Yes
COMPRESSION RATIO: 1.36
DEDUPLICATION: Yes
DEDUPLICATION RATIO: 1.01
ENCRYPTION:
    ENABLED: No

TOTAL SPACE: 100.00GB
FREE SPACE: 71.16GB
USED SPACE: 28.84GB
DATA SIZE BEFORE DEDUPLICATION: 25.51GB
DATA SIZE BEFORE COMPRESSION: 34.77GB
CREATED: 2019-08-28 20:41:19 UTC
UPDATED: 2019-09-02 18:34:03 UTC
DISKS PER RAID GROUP: 1
DISKS IN POOL:
    RAID0:
        /dev/sdb1

[serveradmin@t3-spp-vsnap ~]$
```

2. To list all the available disks, run the **`vsnap pool show`** command, as shown in Example 5-7.

Example 5-7 The vsnap disk show command

```
[serveradmin@t3-spp-vsnap ~]$ vsnap disk show
UUID                             | TYPE | VENDOR | MODEL        | SIZE     | USED AS     | NAME
-------------------------------------------------------------------------------------------------
36000c294686b281a370e7612e55e153b | SCSI | VMware | Virtual disk | 50.00GB  | xfs         | /dev/sda
36000c29ac6130630aec2f02c2aa865a8 | SCSI | VMware | Virtual disk | 100.00GB | vsnap_pool  | /dev/sdb
36000c2960aa178c16a010c527c249df4 | SCSI | VMware | Virtual disk | 128.00GB | LVM2_member | /dev/sdc
```

```
[serveradmin@t3-spp-vsnap ~]$
```

3. Use the vCenter to add a new disk to your vSnap VM.
4. After a new disk is added to your vSnap server, check if a new disk is available by running the **vsnap pool show** command. Review the USED AS column, which shows the newly added disk as unused, as shown in Example 5-8.

Example 5-8 The vSnap disk show command after adding a new disk

```
[serveradmin@t3-spp-vsnap ~]$ vsnap disk show
UUID                             | TYPE | VENDOR | MODEL        | SIZE     | USED AS     | NAME
-------------------------------------------------------------------------------------------------
36000c294686b281a370e7612e55e153b | SCSI | VMware | Virtual disk | 50.00GB  | xfs         | /dev/sda
36000c29ac6130630aec2f02c2aa865a8 | SCSI | VMware | Virtual disk | 100.00GB | vsnap_pool  | /dev/sdb
36000c2960aa178c16a010c527c249df4 | SCSI | VMware | Virtual disk | 128.00GB | LVM2_member | /dev/sdc
36000c291396b40e6b66f62265ac7cf04 | SCSI | VMware | Virtual disk | 100.00GB | unused      | /dev/sdd

[serveradmin@t3-spp-vsnap ~]$
```

5. To add the new disk to the vsnap_pool run the **vsnap pool expand --id TEXT --disk_list TEXT** command, as shown in Example 5-9. After completion of the command, it shows two disks are assigned to this pool (see the last line DISKS IN POOL:), and the TOTAL SPACE expanded to 200.00 GB.

Example 5-9 The vsnap pool expand command

```
[serveradmin@t3-spp-vsnap ~]$ vsnap pool expand --id 1 --disk_list /dev/sdd

ID: 1
NAME: primary
POOL TYPE: raid0
STATUS: ONLINE
HEALTH: 100
COMPRESSION: Yes
COMPRESSION RATIO: 1.36
DEDUPLICATION: Yes
DEDUPLICATION RATIO: 1.01
ENCRYPTION:
    ENABLED: No

TOTAL SPACE: 200.00GB
FREE SPACE: 167.55GB
USED SPACE: 32.45GB
DATA SIZE BEFORE DEDUPLICATION: 25.51GB
DATA SIZE BEFORE COMPRESSION: 34.77GB
CREATED: 2019-08-28 20:41:19 UTC
UPDATED: 2019-09-03 10:55:52 UTC
DISKS PER RAID GROUP: 1
DISKS IN POOL:
    RAID0:
        /dev/sdb1
        /dev/sdd1

[serveradmin@t3-spp-vsnap ~]$
```

6. Check if the new disk is part of the pool by running the **vsnap disk show** command. Review the USED AS column, which shows the newly added disk as part of the vsnap_pool, as shown in Example 5-10.

Example 5-10 The vsnap disk show command after expansion

```
[serveradmin@t3-spp-vsnap ~]$ vsnap disk show
UUID                              | TYPE | VENDOR | MODEL        | SIZE     | USED AS     | NAME
----------------------------------------------------------------------------------------------------
36000c294686b281a370e7612e55e153b | SCSI | VMware | Virtual disk | 50.00GB  | xfs         | /dev/sda
36000c29ac6130630aec2f02c2aa865a8 | SCSI | VMware | Virtual disk | 100.00GB | vsnap_pool  | /dev/sdb
36000c2960aa178c16a010c527c249df4 | SCSI | VMware | Virtual disk | 128.00GB | LVM2_member | /dev/sdc
36000c291396b40e6b66f62265ac7cf04 | SCSI | VMware | Virtual disk | 100.00GB | vsnap_pool  | /dev/sdd

[serveradmin@t3-spp-vsnap ~]$
```

Reducing vSnap pool

To reduce the vSnap pool, disks that are similar to expanding disks can be removed, as shown in Example 5-11.

Example 5-11 The vsnap pool disk remove --help command

```
[serveradmin@t3-spp-vsnap ~]$ vsnap pool disk remove --help
Usage: vsnap pool disk remove [OPTIONS]
Remove a failed disk in a pool.
Options:
  --id TEXT    ID of the pool to remove disk from.  [required]
  --disk TEXT  Disk (name or UUID) to remove.  [required]
  --help       Show this message and exit.
```

Uninstall vSnap server

You can remove a vSnap server from your IBM Spectrum Protect Plus environment.

Ensure that no jobs use SLA policies that define the vSnap server as a backup target. To view the SLA policies that are associated with jobs, see the Backup page for the hypervisor or application that is scheduled for backup. For example, for VMware backup jobs, click **Manage Protection** → **Hypervisors** → **VMware**.

Complete the following steps:

1. Log on to the vSnap server console by using the serveradmin user ID.

   Note: Starting with version 10.1.5, the password of the serveradmin user ID is sppDP758-SysXyz.

   You can also use a user ID that has vSnap server administrator privileges that you create by running the **vsnap user create** command.

2. Run the following commands:

   ```
   systemctl stop vsnap
   yum remove vsnap
   ```

After a vSnap server is uninstalled, the configuration is retained in the /etc/vsnap directory. The configuration is reused if the vSnap server is reinstalled. The configuration is removed if you ran the optional commands to remove the configuration data.

3. Optional: If you do not plan to reinstall the vSnap server after it is uninstalled, remove the data and configuration by running the following commands:

```
rm -rf /etc/vsnap
rm -rf /etc/nginx
rm -rf /etc/uwsgi.d
rm -f /etc/uwsgi.ini
```

5.3.4 Configuring LDAP and SMTP

You can add a Lightweight Directory Access Protocol (LDAP) and Simple Mail Transfer Protocol (SMTP) server in the IBM Spectrum Protect Plus environment for user account and report features.

Adding an LDAP server

You must add an LDAP server to create IBM Spectrum Protect Plus user accounts by using an LDAP group. These accounts allows users to access IBM Spectrum Protect Plus by using LDAP user names and passwords.

Note: Only *one* LDAP server can be associated with an instance of Spectrum Protect Plus.

You can add a Microsoft Active Directory or OpenLDAP server. OpenLDAP does not support the `sAMAaccountName` user filter that is commonly used with Active Directory. Also, the `memberOf` option must be enabled on the OpenLDAP server.

To register an LDAP server, complete the following steps:

1. In the navigation pane, click **System Configuration** → **LDAP/SMTP**.
2. In the **LDAP Servers** pane, click **Add LDAP Server**.
3. Populate the fields in the **LDAP Servers** pane, as described next.

Host Address

The IP address of the host or logical name of the LDAP server.

Port

The port on which the LDAP server is listening. The typical default port is 389 for non-SSL connections or 636 for SSL connections.

SSL

Enable the SSL option to establish a secure connection to the LDAP server.

Use existing user

Enable to select a previously entered user name and password for the LDAP server.

Bind Name

The bind distinguished name that is used for authenticating the connection to the LDAP server. IBM Spectrum Protect Plus supports simple bind.

Password

The password that is associated with the Bind Distinguished Name.

Base DN

The location where users and groups can be found.

User Filter

A filter to select only those users in the Base DN that match certain criteria. An example of a valid default user filter is cn={0}.

Tips: Consider the following points:

- To enable authentication by using the sAMAccountName Windows user naming attribute, set the filter to samaccountname={0}. When this filter is set, users log in to IBM Spectrum Protect Plus by using only a user name. A domain is not included.
- To enable authentication using the user principal name (UPN) naming attribute, set the filter to userprincipalname={0}. When this filter is set, users log in to IBM Spectrum Protect Plus by using the username@domain format.
- To enable authentication by using an email address that is associated with LDAP, set the filter to mail={0}.

The **User Filter** setting also controls the type of user name that appears in the IBM Spectrum Protect Plus display of users.

User RDN

The relative distinguished path for the user. Specify the path where user records can be found. An example of a valid default RDN is cn=Users.

Group RDN

The relative distinguished path for the group. If the group is at a different level than the user path, specify the path where group records can be found.

4. Click **Save.**

After the SMTP server is added, the **Add LDAP Server** button is no longer available.

If your IBM Spectrum Protect Plus server communicates by using Secure Sockets Layer (SSL) to the LDAP environment, register LDAP with SSL authentication at this web page.

Adding an SMTP server

Adding an SMTP server enables email communications to be sent from IBM Spectrum Protect Plus server.

Note: Only **one** SMTP server can be associated with IBM Spectrum Protect Plus.

To configure a SMTP sever, complete the following steps:

1. In the navigation pane, click **System Configuration** → **LDAP/SMTP**.
2. In the **SMTP Servers** pane, click **Add SMTP Server**.
3. Enter the fields in the **SMTP Servers** pane, as described next.

Host Address

The IP address of the host, or the path and host name of the SMTP server.

Port

The communications port of the server that you are adding. The typical default port is 25 for non-SSL connections or 443 for SSL connections.

Username

The name that is used to access the SMTP server.

Password

The password that is associated with the user name.254 IBM Spectrum Protect Plus: Installation and User's Guide.

Timeout

The email timeout value in milliseconds.

From Address

The address that is associated with email communications from IBM Spectrum Protect Plus.

Subject Prefix

The prefix to add to the email subject lines sent from IBM Spectrum Protect Plus.

4. Click **Save**.

To test the SMTP connection, click the **Test SMTP Server** button, then enter an email address. Click **Send**. A test email message is sent to the email address to verify the connection. After the SMTP server is added, the **Add SMTP Server** button is no longer available.

Editing settings for an LDAP or SMTP server

Edit the settings for an LDAP or SMTP server to reflect changes in your IBM Spectrum Protect Plus environment.

To edit the settings for an LDAP or SMTP server, complete the following steps:

1. From the navigation menu, click **System Configuration** → **LDAP/SMTP**.
1. Click the **edit icon** that is associated with the server. The edit pane is displayed.
1. Revise the settings for the server, and then, click **Save**.

Deleting an LDAP or SMTP server

Delete an LDAP or SMTP server when it becomes obsolete. Ensure that the server is not in use by IBM Spectrum Protect Plus before deleting the server.

To delete an LDAP or SMTP server, complete the following steps:

1. From the navigation menu, click **System Configuration** → **LDAP/SMTP**.
2. Click the **delete icon** that is associated with the server.
3. Click **Yes** to delete server.

5.3.5 Administrative Console

IBM Spectrum Protect Plus provides an additional GUI, the Administrative Console.

Using the administrative console, you can complete the following tasks:

- Get details about the installed product versions.
- Manage and install the licenses.
- Manage and install certificates, e.g. Active Directory LDAP certificates.
- Apply and install Spectrum Protect Plus Software Updates.
- Perform System Actions, such as Start/Stop the server, restart the VM, and configure the time zone.

Logging in to the Administrative Console

Access the Administrative Console GUI by using the address that is shown in Example 5-12.

Example 5-12 Administrative Console login

```
https://spp-server-hostname:8090
```

In the login window shown in Figure 5-23, select one of the following authentication types shown in the Table 5-3 Authentication Type list.

Table 5-3 Administrative Console - Authentication Type

Authentication Type	Login information
IBM Spectrum Protect Plus	To log in as an IBM Spectrum Protect Plus user with SYSADMIN privileges, enter your administrator user name and password. If you log in by using the admin user account, you are prompted to reset the user name and password. You cannot reset the user name to admin, root, or test.
System (recommended)	To log in as a system user, enter the server admin password. In Spectrum Protect Plus version 10.1.4, the initial password of the user ID **serveradmin** is `sppDP758`. Starting with version 10.1.5, the password was changed to `sppDP758-SysXyz`. You are prompted to change this password during the first login.

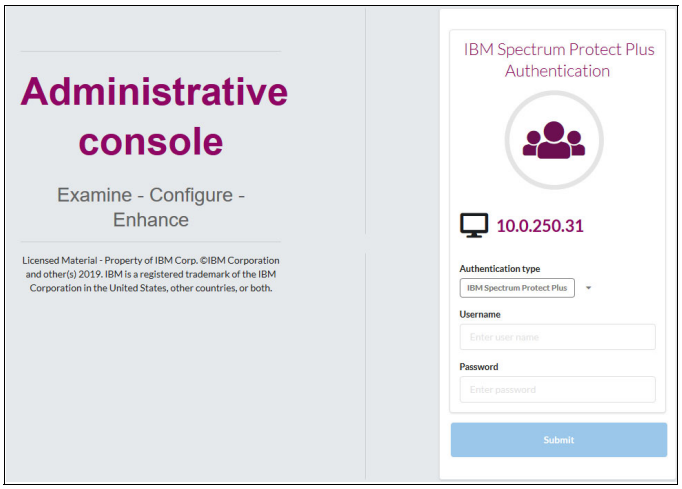

Figure 5-23 Administrative Console - Login

After logging on to the Administrative Console, you see the available options, as shown in Figure 5-24.

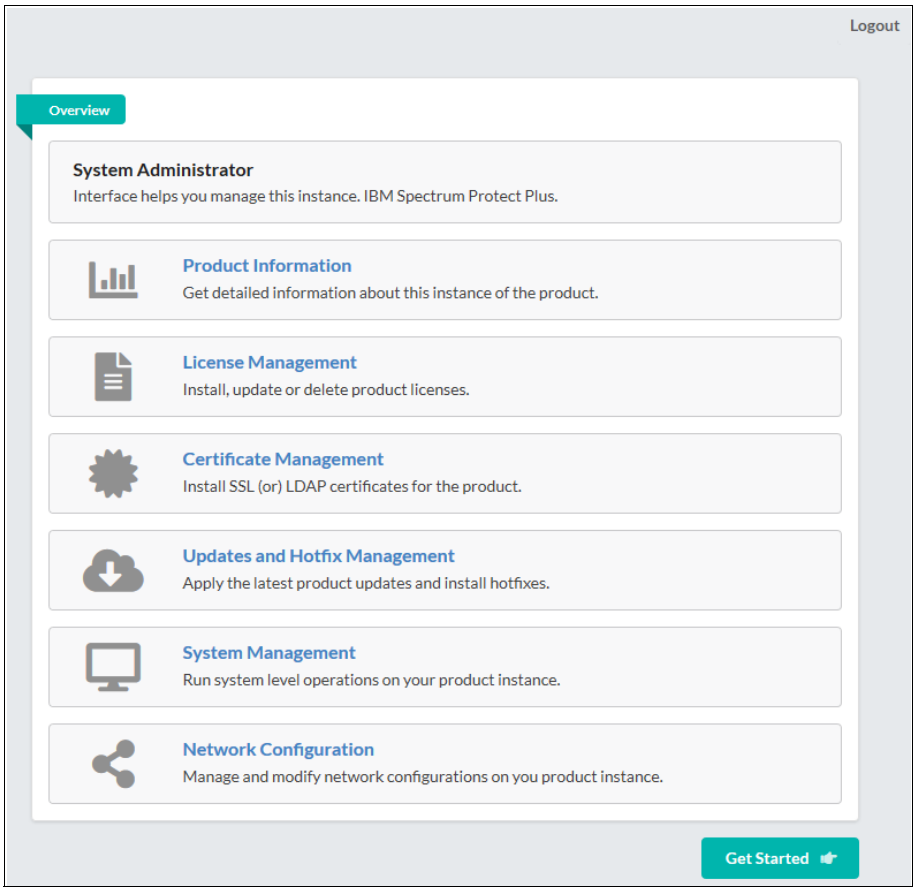

Figure 5-24 Administrative Console - System Administrator pane

System Management

Select the **System Management** option to manage your instance of the IBM Spectrum Protect Plus server, as shown in Figure 5-25 on page 142.

By using this panel, you can either and start Spectrum Protect Plus applications and components, or restart the Spectrum Protect Plus appliance. This second action is stopping and starting the VM along with all Spectrum Protect Plus services.

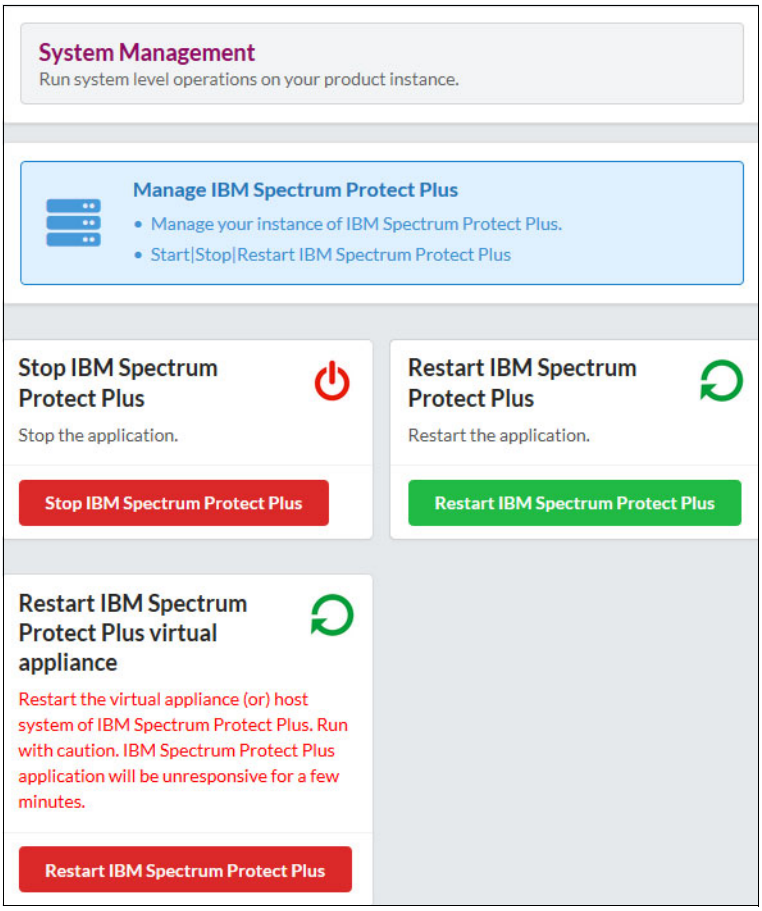

Figure 5-25 Administrative Console - System Management

Using Subject Alternative Name attribute in SSL certificates

To establish secure connections in IBM Spectrum Protect Plus, you can upload an SSL certificate by using the administrative console.

Technote 739663 provides information about the use of an HTTPS certificate that is issued by Microsoft Certificate Authority.

This section provides more information about how to create certificate signing requests that include so-called Subject Alternative Name (SAN):

> *A SAN or subject alternative name is a structured way to indicate all of the domain names and IP addresses that are secured by the certificate*[1].

If the SAN are not defined, the following browsers report warnings, even if the certificate is correctly signed:

- Mozilla Firefox posts a warning: “Potential security Alert” if the user is accessing the server by way of the IP address instead of host name
- Google Chrome reports an error NET::ERR_CERT_COMMON_NAME_INVALID when accessing the GUIs by host name
- Google Chrome posts an alert NET::ERR_CERT_COMMON_NAME_INVALID when accessing the GUI by IP

[1] Source Entrrust DataCard, see: https://www.entrustdatacard.com/blog/2019/march/what-is-a-san-and-how-is-it-used

- Microsoft Edge posts a certificate error DLG_FLAGS_SEC_CERT_CN_INVALID when accessing the GUI by IP

Note: The instructions that we provide next do not include more information about SSL and signed certificates or the principles of trusted certificate chains. For more information about these topics, see this web page.

Generating the certificate

Complete the following steps to generate SSL certificates with Subject Alternative Name:

1. Create the `CertRequest.conf` file.

   Create an OpenSSL configuration file that provides SSL defaults for items, such as the server's Distinguished Name (DN), organizational unit (OU), and country.

   It is recommended to add the alternative names DNS.1 and IP.1 as shown in Example 5-13.

Example 5-13 Certificate configuration

```
> vi /etc/ssl/certs/CertRequest.conf
[req]
distinguished_name = req_distinguished_name
req_extensions = v3_req
prompt = no

[req_distinguished_name]
C  = DE
ST = HE
L  = Kelsterbach
O  = IBM
OU = ESCC
CN = spp-server-itso.sle.kelsterbach.de.ibm.com

[v3_req]
keyUsage = keyEncipherment, dataEncipherment
extendedKeyUsage = serverAuth
subjectAltName = @alt_names

[alt_names]
DNS.1 = spp-server-itso.sle.kelsterbach.de.ibm.com
DNS.2 = spp-server-itso
IP.1  = 9.1yy.122.2xx
```

2. Create the SPP server private key:

   ```
   cd /etc/ssl/certs
   openssl genrsa -out spp-server-itso.key 2048
   ```

3. Create a certificate signing request (CSR):

   ```
   openssl req -new -out spp-server-itso.csr -key spp-server-itso.key \
    -config /etc/ssl/certs/CertRequest.conf
   ```

4. Sign the certificate signing request by CA.

   As the CA administrator to sign the CSR (for example, for a Windows AD CA) this can be achieved by using the following command:

   ```
   certreq -attrib "CertificateTemplate:webserver" -submit spp-server-itso.csr
   ```

Transferring the certificate back to the SPP server

Complete the following steps:\ to determine the format of the signed certificate:

1. Display PEM encoded certificate:

   openssl x509 -in spp-server-itso.cer -text -noout

   ```
   Certificate:
       Data:
           Version: 3 (0x2)
           Serial Number:
               14:00:00:00:23:a7:ea:66:bb:db:84:9e:85:00:01:00:00:00:23
       Signature Algorithm: sha256WithRSAEncryption
           Issuer: DC=com, DC=ibm, DC=de, DC=kelsterbach, DC=sle, CN=CA-SLE
           Validity
               Not Before: Aug 12 18:30:28 2020 GMT
               Not After : Aug 12 18:30:28 2022 GMT
           Subject: C=DE, ST=HE, L=Kelsterbach, O=IBM, OU=ESCC,
   CN=spp-server-itso.sle.kelsterbach.de.ibm.com
   …
   …
   ```

2. If the certificate is in PEM format copy the file:

   ```
   cp spp-server-itso.cer spp-server-itso.pem
   ```

3. Display the DER encoded certificate:

   openssl x509 -in spp-server-itso.cer -inform der -text -noout

   ```
   unable to load certificate
   140514786793360:error:0D0680A8:asn1 encoding routines:ASN1_CHECK_TLEN:wrong
   tag:tasn_dec.c:1220:
   140514786793360:error:0D07803A:asn1 encoding routines:ASN1_ITEM_EX_D2I:nested
   asn1 error:tasn_dec.c:386:Type=X509
   ```

   If you receive an error the reads: “Unable to load certificate”, the format does not match the queried format.

   In this example, the certificate is a PEM certificate, *not* a DER encoded certificate.

   IBM Spectrum Protect Plus supports only Privacy Enhanced Mail (PEM) encoded certificates for HTTPS. A PEM certificate is a Base64-encoded DER certificate. To convert the DER encoded certificate to PEM format, run the following command:

   ```
   openssl x509 -inform DER -outform PEM -in cert_name.cer -out cert_name.pem
   ```

4. Merge the private key and the PEM certificate.

   The `.key` and `.crt` files must be merged such that the content of the `.key` file must be on top, followed by the content of the `.crt` file. The signer public certificate also must be added to the file in the descending order (ICA after the signed certificate, then the CA, and so on):

   – Example 1: **cat cert_name.key cert_name.cer > cert_name.crt**

   – Example 2: **cat spp-server-itso.key spp-server-itso.pem ca-sle.cer > spp-server-itso.crt**

The following command lists the different certificates:

```
[root@spp-server-itso certs]# ll /etc/ssl/certs
```

```
-rw-r--r--. 1 root root  436 Aug 12 18:31 CertRequest.conf
-rw-rw-r--. 1 root root 1460 Aug 12 18:42 ca-sle.cer
-rw-rw-r--. 1 root root 2218 Aug 12 19:13 spp-server-itso.cer
-rw-r--r--. 1 root root 3897 Aug 12 19:16 spp-server-itso.crt
-rw-r--r--. 1 root root 1216 Aug 12 19:10 spp-server-itso.csr
-rw-r--r--. 1 root root 1679 Aug 12 19:07 spp-server-itso.key
-rw-r--r--. 1 root root 2218 Aug 12 19:15 spp-server-itso.pem
```

The following example is of a server certificate with private key, signed certificate, and signer certificated (abstract):

```
>cat spp-server-itso.crt
```

```
-----BEGIN RSA PRIVATE KEY-----
MIIEpQIBAAKCAQEA97bBLCXmvgBv6say6mGMWpC9DLA+rqnOt5A+Jpi41EsXxptn
9zPZO9IRWN1jFUyTvwVhV7tdyqctj+AkhwNGSu891y5U991Mz2GptdJtO+X3okF=
…
…
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
MIIGIDCCBQigAwIBAgITFAAAACSxBPOb8Vj/5QABAAAAJDANBgkqhkiG9w0BAQsF
ADCBgTETMBEGCgmSJomT8ixkARkWA2NvbTETMBEGCgmSJomT8ixkARkWA2libTES
…
…
-----END CERTIFICATE-----
```

Download this certificate from the SPP server to your workstation and upload the certificate in the SPP administrative console (as shown in Figure 5-26) and then, restart the restart the SPP appliance.

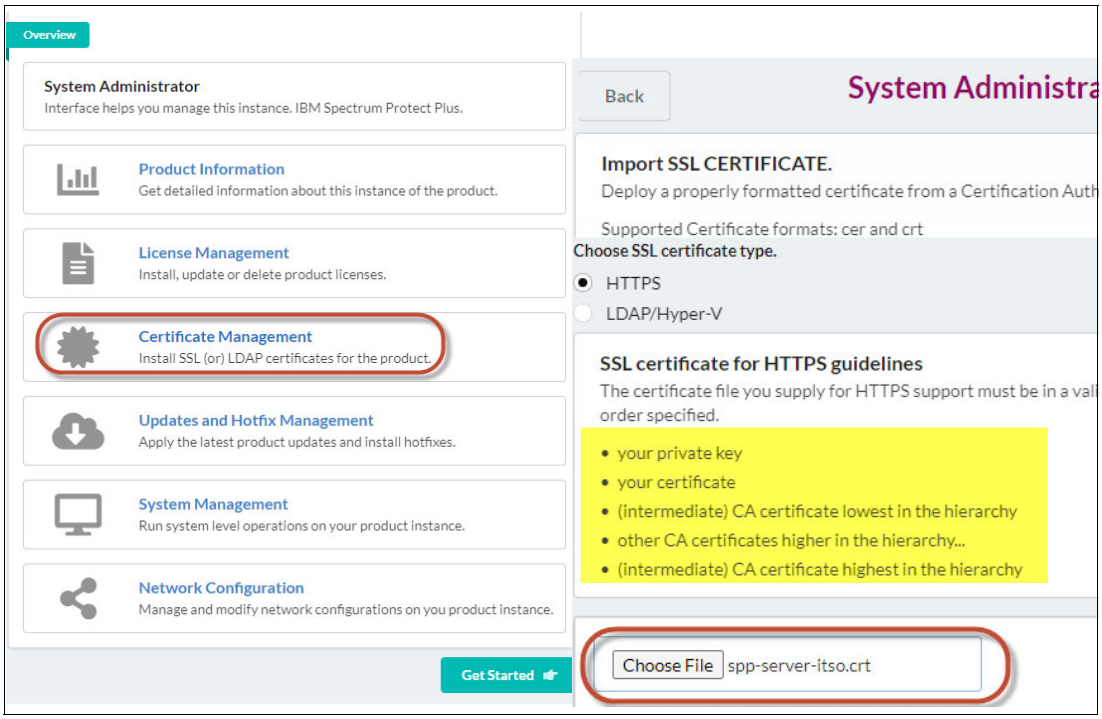

Figure 5-26 Import SSL certificate

After restarting the SPP appliance, the server's HTTPs connection should be evaluated as secure, regardless if it was started by hostname or IP, as shown in Figure 5-27.

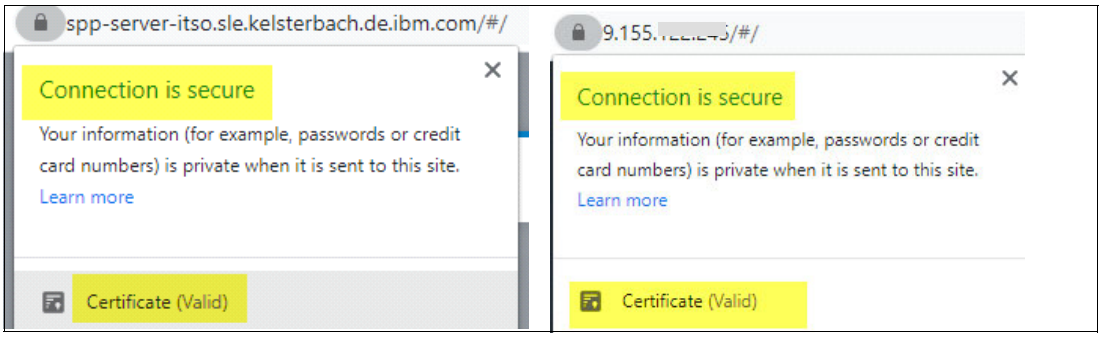

Figure 5-27 Connection secure

Some CAs might remove the SAN parameters or set their own values. Also, ensure that the management station from which you start the SPP server GUI holds your company's CA signer certificates in the workstation's keystore. The Firefox browser might need a configuration adjustment to use the workstation's default trust store instead of the predefine trusted CAs in Firefox.

5.3.6 Managing global preferences

The user administrator can manage preferences that apply to all IBM Spectrum Protect Plus operations in the Global Preferences section.

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

The Global Preferences page contains default values for parameters that apply to all IBM Spectrum Protect Plus operations. The preferences are organized into the following categories:

- Application
- General
- Job
- Logging
- Protection
- Security

Note: On the Global Preferences page, the **Integration with other storage products** section is where you can specify the URL of your Spectrum Protect Operation Center.

The default values for the Application global preferences are listed in Table 5-4.

Table 5-4 Default values for Application global preferences

Preference	Default value	Unit (if applicable)
Enable SQL Server databases restored in test mode eligible for backup	Uncheck	
Maximum volume size for backup target LUNs on Windows (TB)	256	TB

Preference	Default value	Unit (if applicable)
Maximum backup retries(k8s)	3	
Maximum concurrent servers running backups		
Allow SQL database backup when transaction log backup chain is broken	Uncheck	
Rename SQL data and log files when database is restored in production mode with new name	Uncheck	

The default values for the General global preferences are listed in Table 5-5.

Table 5-5 Default values for General global preferences

Preference	Default value	Unit (if applicable)
Access log retention (days)	30	days
Tools working folder on Linux guest	/tmp	
Tools working folder on Windows guest	c:\ProgramData	
Linux/IBM AIX® Clients Port (SSH) used for application and file indexing	22	
Windows Clients Port (WinRM) used for application and file indexing	5985	
IBM Spectrum Protect Plus Server IP Address	SPP server IP address	

The default values for the Job global preferences are listed in Table 5-6.

Table 5-6 Default values for Job global preferences

Preference	Default value	Unit (if applicable)
Job log retention (days)	60	days
Job notification status	failed	

The default value for the Logging global preference is shown in Table 5-7.

Table 5-7 Default values for Logging global preferences

Preference	Default value
Enable logging IBM Spectrum Protect Plus alerts to the system log	Uncheck

When this Logging option is enabled, the Spectrum Protect Plus messages appear in the system log, as shown in Example 5-14.

Example 5-14 SPPalert output

```
[root@t3-spp-server log]# grep SPPAlert /var/log/messages
Aug 26 23:53:03 t3-spp-server virgo[993]: ERROR [SPPAlert]
vmware_Offload_Protect:1548460344227 Job vmware_Offload_Protect (id=1006,
session=1,548,460,344,227) failed.
```

Note: Consider the general format of the alerts that are logged by IBM Spectrum Protect Plus:

- `<Type of Message>`: Can be ERROR, WARNING, and so on.
- `<Source of Message>`: Indicates the source of the message. `SPPAlert` indicates IBM Spectrum Protect Plus messages.
- `<Message>`: This contains the IBM Spectrum Protect Plus alert details.

The default values for the Protection global preferences are listed in Table 5-8.

Table 5-8 Default values for Protection global preferences

Preference	Default value	Unit (if applicable)
Number of seconds to wait before checking connection	1000	seconds
Number of times to check for valid connection	0	
Temporary folder for file index zip files	/data2/filecatalog	
Temporary folder for file indexing on Windows server	empty	
Group VMs by	Count	
Number of VMs in group	1	
Force the removal of replication relationship for last remaining snapshot	Unchecked	
Target free space error (percentage)	20	
Target free space warning (percentage)	30	
Catalog object update count	60	
VM backup status update interval (seconds)	300	
VADP proxy uses only HotAdd transport mode	Unchecked	
VM group size (GB)	5120	GB

Preference	Default value	Unit (if applicable)
vSnap auto disable deduplication when DDT size reaches resource limit	Checked	
vSnap DDT size limit as percentage of total memory cache	80	percentage (%)
vSnap DDT size limit in GB	50	GB
Used space threshold on datastore or a volume before backup cannot take snapshots of a VM (percentage)	95	percentage (%)
Backup wait timeout (seconds)	600	seconds
VMware communication timeout (seconds)	300	seconds

The default value for the Stitchery global preferences is listed in Table 5-8.

Table 5-9 Default value for Stitchery global preferences

Preference	Default value
Set Minimum Password Length (characters)	8

For more information about these preferences and their purpose, see IBM Knowledge Center.

5.3.7 Managing the IBM Spectrum Protect Plus catalog

The following sections describe aspects of managing the catalog.

Restoring the IBM Spectrum Protect Plus application

Restore IBM Spectrum Protect Plus configuration settings, restore points, search data, and job information that were backed up to the vSnap server. The data can be restored to the same location or an alternate IBM Spectrum Protect Plus location.

Attention: An IBM Spectrum Protect Plus restore operation overwrites all data that is in the IBM Spectrum Protect Plus virtual appliance or alternative virtual appliance location. All IBM Spectrum Protect Plus operations stop while the data is being restored. The user interface is not accessible, and all jobs that are running are canceled. Any snapshots that are created between the backup and restore operations are not saved.

If restoring from another copy, the cloud resource or repository server must be registered on the alternative IBM Spectrum Protect Plus location.

To restore IBM Spectrum Protect Plus data, complete the following steps:

1. In the navigation pane, click **Manage Protection** → **IBM Spectrum Protect Plus** → **Restore**.
2. Select a vSnap server, cloud resource, or repository server.

Data can be restored to the same location, or an alternative location in Disaster Recovery scenarios. Available snapshots for the server are displayed.

3. Click **Restore** for the catalog snapshot that you want to restore.
4. Select one of the following restore modes:
   - Restore the catalog and suspend all scheduled jobs

     The catalog is restored and all scheduled jobs are left in a suspended state. No scheduled jobs are started, which allows for the validation and testing of catalog entries and the creation of new jobs. Typically, this option is used in DevOps use cases.
   - Restore the catalog

     The catalog is restored and all scheduled jobs continue to run as captured in the catalog backup. Typically, this option is used in disaster recovery.
5. Click **Restore**.
6. To run the restore job, in the dialog box, click **Yes**.

Deleting IBM Spectrum Protect Plus resources from the catalog

You can use the Virtual Machines/Databases tab in the Restore Point Retention pane to expire catalog metadata that is associated with a resource from the IBM Spectrum Protect Plus catalog. Resources are added to the catalog through inventory jobs. Expiring a resource removes the metadata that is associated with a restore point from the catalog, which frees up space in the catalog and removes the restore point from recovery screens.

Expiring a resource from the catalog does not remove associated snapshots from a vSnap server or secondary backup storage.

Complete the following steps to expire a resource from the catalog:

1. In the navigation pane, click **Manage Protection** → **IBM Spectrum Protect Plus** → **Restore Point Retention**.
2. Click the **Virtual Machines/Databases** tab.
3. Use the filter to search by resource type, then enter a search string to search for a resource by name. For more information about using the search function, see Search guidelines.
4. Click the search icon.
5. Click the delete icon that is associated with a resource.
6. To confirm the expiration, in the dialog box, click **Yes**.

As the result, the catalog metadata associated with the resource is removed from the catalog.

Managing IBM Spectrum Protect Plus restore points

You can use the **Restore Point Retention** pane to search for restore points in the IBM Spectrum Protect Plus catalog by backup job name, view their creation and expiration dates, and override the assigned retention.

Expiring a job session does not remove a snapshot and related recovery point if the snapshot is locked by a replication or copy/archive relationship. Run the replication or copy/archive enabled job to set the lock to a later snapshot. The snapshot and recovery point are removed during the next run of the maintenance job.

To set a job session to expire, complete the following steps:

1. In the navigation pane, click **Manage Protection** → **IBM Spectrum Protect Plus** → **Restore Point Retention**.
2. In the Backup Sessions tab, search for the desired job session or restore point. For more information about using the search function, see Search guidelines. Alternatively, in the Virtual Machines / Databases tab, select either Applications or Hypervisors to search for the desired catalog entry by entering the name. Names can be searched by entering partial text, using the asterisk (*) as a wildcard character, or using the question mark (?) for pattern matching.
3. Optional: If searching from the Backup Sessions tab, use filters to fine-tune your search across job types and date range when the associated backup job started.
4. Click the search icon the search icon.
5. Select the job sessions you want to expire.
6. From the **Actions** list, select one of the following options:
   - **Expire** is used to expire a single job session.
   - **Expire All Job Sessions** is used to expire all unexpired job sessions for the selected job.
7. To confirm the expiration, in the dialog box, click **Yes**.

As a result, the job session is removed during the next run of the maintenance job.

5.3.8 Search guidelines

Use filters to search for an entity, such as a file or a restore point. You can enter a character string to find objects with a name that exactly matches the character string. For example, searching for the term string.txt returns the exact match, string.txt.

Regular expression search entries are also supported. For more information, see Search Text with Regular Expressions on the Microsoft SQL Docs web page.

You can also include special characters (+, -, &, |, !, (,), {, }, [,], ^, ", ~, *, ?, :, and \) in the search. You *must* use a backslash (\) escape character before any of the special characters.

For example, to search for the file `string[2].txt`, enter: `string\[2\].txt`.

Searching with wildcard

You can position a wildcard at the beginning, middle, or end of a string, and combine them within a string.

Matching a character string with an asterisk

The following examples show search text with an asterisk:

```
string* searches for terms like string, strings, or stringency
str*ing searches for terms like string, straying, or straightening
*string searches for terms like string or shoestring
```

You can use multiple asterisk wildcard in a single text string, but multiple wildcards might considerably slow down a large search.

Matching a single character with a question mark:

The following examples show search text with a question mark:

```
string? searches for terms like strings, stringy, or string1
st??ring searches for terms like starring or steering
???string searches for terms like hamstring or bowstring
```

5.3.9 Testing network connectivity

The IBM Spectrum Protect Plus Service Tool runs tests on host addresses and ports to determine if a connection can be established. You can use the Service Tool to verify whether a connection can be established between IBM Spectrum Protect Plus and a node

You can run the Service Tool from the IBM Spectrum Protect Plus command line or remotely by using a .jar file. If a connection can be established, the tool returns a green check mark. If a connection cannot be established, the error condition is displayed, along with possible causes and actions.

The tool provides guidance for the following error conditions:

- Timeout
- Connection refused
- Unknown host
- No route

Figure 5-28 and Figure 5-29 on page 153 show the look of the tool.

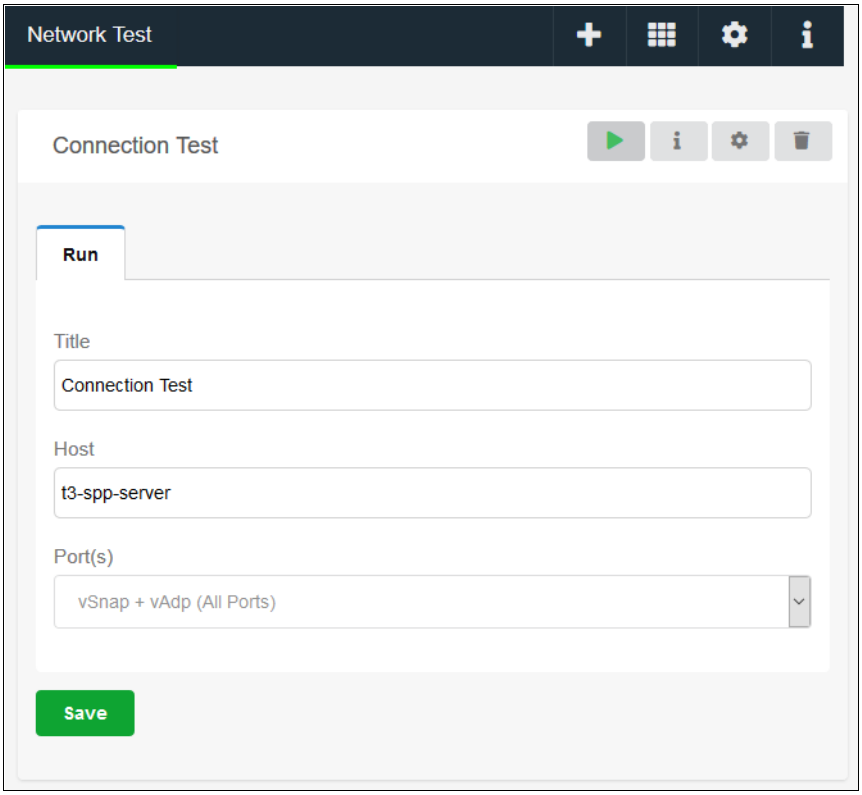

Figure 5-28 Network Connection Test Tool - Settings

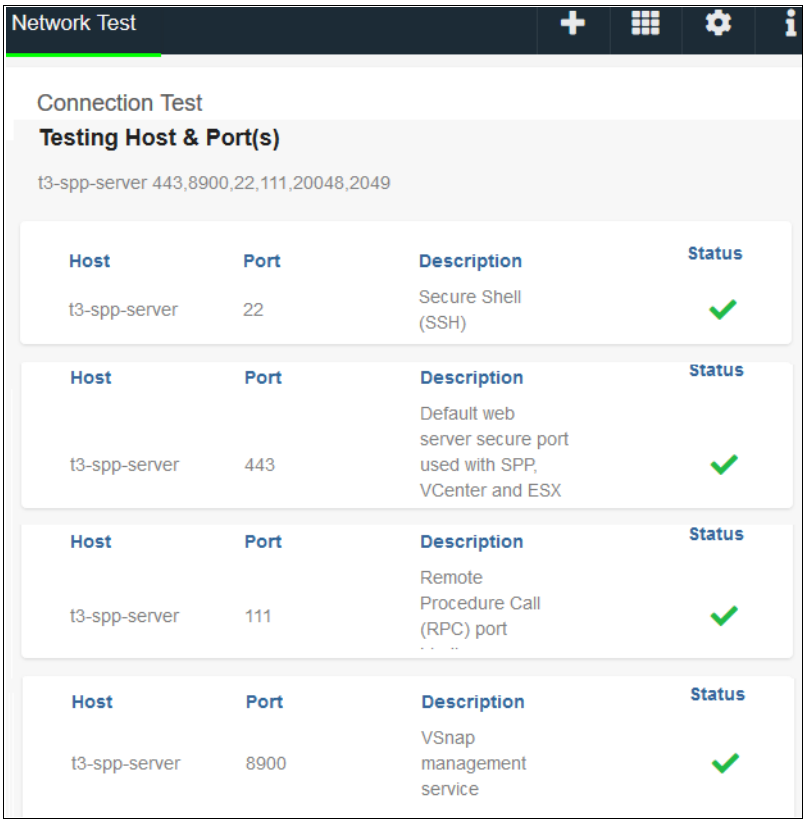

Figure 5-29 Network Connection Test Tool - Results

Running the Service Tool from a command-line interface

You can start the Service Tool from the IBM Spectrum Protect Plus virtual appliance command-line interface and run the tool in a web browser. Then, you can use the Service Tool to verify network connectivity between IBM Spectrum Protect Plus and a node. To use the tool follow these steps:

1. Log in to the IBM Spectrum Protect Plus virtual appliance by using the serveradmin user ID and access the command prompt. Issue the following command:

   ```
   # sudo bash
   ```

2. Open port 9000 on the firewall by issuing the following command:

   ```
   # firewall-cmd --add-port=9000/tcp
   ```

3. Run the tool by issuing the following command:

   ```
   # java -Dserver.port=9000 -jar /opt/ECX/spp/public/assets/tool/ngxdd.jar
   ```

4. To connect to the tool, enter the following URL in a browser:

   ```
   http://hostname:9000
   ```

   Where hostname specifies the IP address of the VM where the application is deployed.

5. To specify the node to test, populate the following fields:

   – Host: The host name or IP address of the node that you want to test.
   – Port: The connection port to test.

6. Click **Save**.

7. To run the tool, hover the cursor over the tool, and then click the green **Run** button.

If a connection cannot be established, the error condition is displayed, along with possible causes and actions.

8. Stop the tool by running the following command on the command line:

   ```
   CTRL-C
   ```

9. Protect your storage environment by resetting the firewall. Run the following commands:

   ```
   # firewall-cmd --zone=public --remove-port=9000/tcp
   # firewall-cmd --runtime-to-permanent
   # firewall-cmd --reload
   ```

Note: If the **`firewall-cmd`** command is not available on your system, edit the firewall manually to add necessary ports and restart the firewall by using iptables. For more information about editing firewall rules, see *Firewall configuration using iptables* of IBM Knowledge Center.

Running the Service Tool remotely

You can download the Service Tool as a .jar file from the IBM Spectrum Protect Plus user interface. Then, you can use the Service Tool to remotely test connectivity between IBM Spectrum Protect Plus and a node.

To use the tool remotely follow these steps:

1. In the IBM Spectrum Protect Plus user interface, click the user menu on the upper right to drop down a list and then, click **Download Test Tool**.

   A .jar file is downloaded to your workstation.

2. Launch the tool from a command-line interface. Java is only required on the system where the tool will be launched. Endpoints or target systems that are tested by the tool do not require Java.

   The following command launches the tool in a Linux environment:

   ```
   # java -jar -Dserver.port=9000 /<tool path >/ngxdd.jar
   ```

3. Run the tool by issuing the following command:

   ```
   # java -Dserver.port=9000 -jar /opt/ECX/spp/public/assets/tool/ngxdd.jar
   ```

4. To connect to the tool, enter the following URL in a browser:

   ```
   http://hostname:9000
   ```

   where hostname specifies the IP address of the VM where the application is deployed.

5. To specify the node to test, populate the following fields:

   – Host: The host name or IP address of the node that you want to test.
   – Port: The connection port to test.

6. Click **Save**.

7. To run the tool, hover the cursor over the tool, and then, click the green **Run** button.

   If a connection cannot be established, the error condition is displayed, along with possible causes and actions.

8. Stop the tool by issuing the following command on the command line:

   ```
   CTRL-C
   ```

5.3.10 Messages

IBM Spectrum Protect Plus components send messages with prefixes that help to identify which component they come from. Use the search option to find a particular message by using its unique identifier.

Messages consist of the following elements:

- A five-letter prefix.
- A number to identify the message.
- Message text that is displayed on screen and written to message logs.

Tip: Use your browser's search capability by using Ctrl+F to find the message code you are looking for.

The following example contains the Db2 agent prefix. When you click **More**, more information that explains the reason for the message are shown (see Example 5-15).

Example 5-15 Messages Example - Warning

```
Warning
Apr 16, 2019
9:14:37 AM
CTGGH0098
[myserver1.myplace.irl.ibm.com]
Database AC7 will not be backed up as it is ineligible for the backup operation.
More
```

IBM Spectrum Protect Plus message prefixes

Messages have different prefixes to help you to identify the component that issues the message.

Table 5-10 lists the prefix that is associated with each component.

Table 5-10 Messages prefixes by component

Prefix	Component
CTGGA	IBM Spectrum Protect Plus
CTGGE	IBM Spectrum Protect Plus for Microsoft SQL Server
CTGGF	IBM Spectrum Protect Plus for Oracle
CTGGG	IBM Spectrum Protect Plus for Microsoft Exchange Server
CTGGH	IBM Spectrum Protect Plus for IBM Db2
CTGGI	IBM Spectrum Protect Plus for MongoDB
CTGGK	IBM Spectrum Protect Plus for Containers
CTGGL	IBM Spectrum Protect Plus for Amazon EC2
CTGGR	IBM Spectrum Protect Plus for Microsoft Office 365
CTGGT	IBM Spectrum Protect Plus for for file systems

For more information about these messages, see IBM Knowledge Center.

5.4 vSnap server CLI

The vSnap server has two user interfaces to administer the server. Graphical access to the vSnap server configuration is provided as part of web-based IBM Spectrum Protect Plus Day-to-day GUI. Under **System Configuration** → **Backup Storage** → **Disk**, all available vSnap servers with their current status and capacity are listed.

The second user interface is a command-line interface (CLI), which can be entered by way of SSH by using the default `serveradmin` user ID. The sections “User interfaces” on page 15 and 3.4.2, “Configuring the vSnap Backup Storage server” on page 63 provide details about how to use and configure the access to these user interfaces.

In Example 5-16 we show how to log in to the CLI user interface using the `serveradmin` user ID and by running the command **`vsnap`** to list all options for the vSnap CLI.

Example 5-16 The vSnap Backup Storage CLI

```
Using username "serveradmin".
serveradmin@spp-vsnap-demo's password:
Last login: Mon Aug  5 13:18:46 2019 from 9.123.45.67
---------------------------------------------------------------
Be sure to adhere to vSnap hardware and memory requirements
 as described in IBM Spectrum Protect Plus Blueprints
 accessible from the IBM Spectrum Protect Plus Knowledge Center.

---------------------------------------------------------------
[serveradmin@spp-vsnap-demo ~]$ vsnap
Usage: vsnap [OPTIONS] COMMAND [ARGS]...

Options:
  --json     Show output in JSON format.
  --summary  Show output in summary (tabular) format.
  --detail   Show output in detail (multiline) format.
  --help     Show this message and exit.

Commands:
  archive       Manage archive resources.
  cloud         Manage cloud resources.
  disk          Manage disks.
  host          Manage volume host mappings.
  maint         Manage maintenance sessions.
  network       Manage network interfaces.
  partner       Manage partner servers.
  pool          Manage storage pools.
  relationship  Manage replication relationships.
  session       Manage replication sessions.
  share         Manage volume shares.
  snapshot      Manage volume snapshots.
  system        Manage vSnap system.
  target        Manage storage targets.
  throttle      Manage throttling events.
  user          Manage vSnap users.
  volume        Manage storage volumes.
[serveradmin@spp-vsnap-demo ~]$
```

5.5 vSnap server initialization

After the deployment or installation on a physical server, the vSnap server must be initialized, which can be done by using the GUI or CLI interface. The use of the GUI is described as part of the installation and deployment process in “Initializing vSnap pool” on page 72; and the CLI initialization is described in “Initializing vSnap server pool” on page 131.

5.6 Checking vSnap server status

To check the status of vSnap server core components, run the **vsnap_status** command, as shown in Example 5-17.

When all of the services that are checked by the **vsnap_status** command are in active state, the vSnap server is considered healthy.

Example 5-17 vsnap_status command output

```
[serveradmin@t3-spp-vsnap ~]$ vsnap_status
? vsnap.service - vSnap Services Launcher
   Loaded: loaded (/usr/lib/systemd/system/vsnap.service; enabled; vendor preset:
disabled)
   Active: active (exited) since Wed 2020-06-03 08:30:35 UTC; 3 weeks 4 days ago
 Main PID: 2440 (code=exited, status=0/SUCCESS)
   CGroup: /system.slice/vsnap.service

? vsnap-config.service - vSnap Configuration Database
   Loaded: loaded (/usr/lib/systemd/system/vsnap-config.service; disabled; vendor
preset: disabled)
   Active: active (running) since Wed 2020-06-03 08:30:37 UTC; 3 weeks 4 days ago
 Main PID: 5416 (mongod)
   CGroup: /system.slice/vsnap-config.service
           ··5416 /opt/vsnap/mongo/bin/mongod --quiet -f
/opt/vsnap/config/mongo/mongo.conf run

? vsnap-data.service - vSnap Data Services
   Loaded: loaded (/usr/lib/systemd/system/vsnap-data.service; disabled; vendor
preset: disabled)
   Active: active (exited) since Wed 2020-06-03 08:30:43 UTC; 3 weeks 4 days ago

? vsnap-api.service - vSnap API Services
   Loaded: loaded (/usr/lib/systemd/system/vsnap-api.service; disabled; vendor
preset: disabled)
   Active: active (exited) since Wed 2020-06-03 08:30:44 UTC; 3 weeks 4 days ago
 Main PID: 8390 (code=exited, status=0/SUCCESS)
   CGroup: /system.slice/vsnap-api.service

? vsnap-repl.service - vSnap Replication Services
   Loaded: loaded (/usr/lib/systemd/system/vsnap-repl.service; disabled; vendor
preset: disabled)
   Active: active (running) since Wed 2020-06-03 08:30:43 UTC; 3 weeks 4 days ago
 Main PID: 8392 (python3)
   CGroup: /system.slice/vsnap-repl.service
           ·· 8392 /opt/vsnap/venv/bin/python3 /opt/vsnap/lib/vsnap/service/repld
           ··27325 /opt/vsnap/venv/bin/python3 /opt/vsnap/lib/vsnap/service/repld
```

```
? vsnap-maint.service - vSnap Maintenance Services
   Loaded: loaded (/usr/lib/systemd/system/vsnap-maint.service; disabled; vendor
preset: disabled)
   Active: active (running) since Wed 2020-06-03 08:30:43 UTC; 3 weeks 4 days ago
 Main PID: 8394 (python3)
   CGroup: /system.slice/vsnap-maint.service
           ··8394 /opt/vsnap/venv/bin/python3
/opt/vsnap/lib/vsnap/service/maintenance/maint

? vsnap-work.service - vSnap Worker Services
   Loaded: loaded (/usr/lib/systemd/system/vsnap-work.service; disabled; vendor
preset: disabled)
   Active: active (running) since Wed 2020-06-03 08:30:43 UTC; 3 weeks 4 days ago
 Main PID: 8391 (python3)
   CGroup: /system.slice/vsnap-work.service
           ··8391 /opt/vsnap/venv/bin/python3
/opt/vsnap/lib/vsnap/service/worker/work
```

If a VADP proxy is configured on the vSnap server, run the **service remote-vadp.service status** command to check the status of that VADP proxy service, as shown in Example 5-18.

Example 5-18 Checking the status of the VADP proxy service

```
[serveradmin@t3-spp-vsnap ~]$ service remote-vadp.service status
Redirecting to /bin/systemctl status remote-vadp.service
? remote-vadp.service - Spring Cloud remote VADP service
   Loaded: loaded (/usr/lib/systemd/system/remote-vadp.service; enabled; vendor
preset: disabled)
   Active: active (running) since Wed 2020-06-03 08:56:13 UTC; 3 weeks 4 days ago
     Docs: https://www.ibm.com
  Process: 13262 ExecStartPre=/bin/bash -c /usr/bin/systemctl set-environment
JAR=$(ls -1t $INSTALL_DIR/bin/com.catalogic.vmdkbackup-*.jar| head -1)
(code=exited, status=0/SUCCESS)
  Process: 13257 ExecStartPre=/bin/bash -c /usr/bin/systemctl set-environment
SERVER_HOST=$(hostname) (code=exited, status=0/SUCCESS)
  Process: 13245 ExecStartPre=/bin/bash -c /usr/bin/systemctl set-environment
UUID=$(${INSTALL_DIR}/bin/generate_uuid.sh) (code=exited, status=0/SUCCESS)
  Process: 13243 ExecStartPre=/bin/mkdir -p ${INSTALL_DIR}/core (code=exited,
status=0/SUCCESS)
 Main PID: 13268 (java)
   CGroup: /system.slice/remote-vadp.service
           ··13268 /opt/IBM/SPP/jre/bin/java -Djdk.tls.ephemeralDHKeySize=2048
-Dspring.config.location=/opt/IBM/SPP/bin/application.yml -DINSTAL...
```

5.7 vSnap server preferences

The preferences of the vSnap server can be listed by running the **vsnap system pref get** command, as shown in Example 5-19.

Example 5-19 List vSnap server preferences

```
[serveradmin@t3-spp-vsnap ~]$ vsnap system pref get
NAME                              | DEFAULT VALUE | USER VALUE | TYPE
-------------------------------------------------------------------------
archiveCompressMinSize            | 10485760      | N/A        | integer
archiveDefaultRetrievalTier       | Bulk          | N/A        | string
archiveListObjectsMax             | 1000          | N/A        | integer
archiveLunGranular                | true          | N/A        | boolean
archiveMaxParallelClone           | 2             | N/A        | integer
archiveMaxParallelDelete          | 2             | N/A        | integer
archiveMaxStreams                 | 5             | N/A        | integer
archiveObjectSize                 | 1048576000000 | N/A        | integer
archiveOffloadMaxAttempts         | 5             | N/A        | integer
archivePartSize                   | 104857600     | N/A        | integer
archiveRestoreDays                | 3             | N/A        | integer
archiveRestoreMaxAttempts         | 50            | N/A        | integer
archiveRestoreSleepSecs           | 30            | N/A        | integer
archiveSimulatorEnabled           | false         | N/A        | boolean
archiveSimulatorRestoreSeconds    | 30            | N/A        | integer
archiveStatusUpdateFrequency      | 300           | N/A        | integer
archiveThreads                    | 4             | N/A        | integer
archiveThrottleRate               | 536870912     | N/A        | integer
archiveThrottleRatePoll           | 120           | N/A        | integer
archiveTransferRetryable          | true          | N/A        | boolean
archiveUnmountMaxAttempts         | 30            | N/A        | integer
archiveWorkers                    | 4             | N/A        | integer
asyncSnapshotWait                 | 15            | N/A        | integer
cancelPendingSessionsOnRestart    | true          | N/A        | boolean
cloudBlockSize                    | 4096          | N/A        | integer
cloudBlockmapOperationTimeout     | 86400         | N/A        | integer
cloudChannelSize                  | 536870912     | N/A        | integer
cloudCondenseChunkSize            | 2097152       | N/A        | integer
cloudCondenseInterval             | 3             | N/A        | integer
cloudConnectionTimeout            | 600           | N/A        | integer
cloudCrashOnExportTimeout         | false         | N/A        | boolean
cloudDbBlockmapDownload           | false         | N/A        | boolean
cloudDbBlockmapUpload             | true          | N/A        | boolean
cloudDeviceQueueDepth             | 16            | N/A        | integer
cloudDownloadThreads              | 16            | N/A        | integer
cloudDynamicThrottle              | true          | N/A        | boolean
cloudErrorCleanupTimeout          | 120           | N/A        | integer
cloudIOReadAttempts               | 3             | N/A        | integer
cloudIOReadIntervalSecs           | 5             | N/A        | integer
cloudIOReadTimeout                | 60            | N/A        | integer
cloudIOTimeout                    | 300           | N/A        | integer
cloudIOWriteAttempts              | 5             | N/A        | integer
cloudIOWriteIntervalSecs          | 30            | N/A        | integer
cloudKillOnExportTimeout          | true          | N/A        | boolean
cloudLocalBlockmapRetention       | 3             | N/A        | integer
```

cloudLogDebugWithRequestErrors	false	N/A	boolean
cloudLogMessageBufferSize	24	N/A	integer
cloudLogMessageFlushIntervalSecs	180	N/A	integer
cloudMaxAttempts	10	N/A	integer
cloudMaxParallelClone	2	N/A	integer
cloudMaxParallelDelete	2	N/A	integer
cloudMaxStreams	5	N/A	integer
cloudObjectSize	16777216	N/A	integer
cloudOffloadCacheSize	8589934592	N/A	integer
cloudOffloadChunkSize	131072	N/A	integer
cloudOffloadLogPeriod	300	N/A	integer
cloudOffloadRate	536870912	N/A	integer
cloudOffloadThreads	16	N/A	integer
cloudPendingUploadTimeout	2700	N/A	integer
cloudPoolExportAttemptUnforced	true	N/A	boolean
cloudPoolExportIntervalSecs	30	N/A	integer
cloudPoolExportLockTimeout	3600	N/A	integer
cloudPoolExportNumAttempts	5	N/A	integer
cloudPoolImportLockTimeout	3600	N/A	integer
cloudPoolImportTimeout	900	N/A	integer
cloudPoolSyncTimeout	3600	N/A	integer
cloudPreserveOffloadFailure	false	N/A	boolean
cloudPropagateReadErrors	true	N/A	boolean
cloudReadTimeout	600	N/A	integer
cloudRecordsCache	128	N/A	integer
cloudRecordsFlush	64	N/A	integer
cloudRequireDataDisk	true	N/A	boolean
cloudRestoreCacheSize	6553600000	N/A	integer
cloudRestoreChunkSize	1048576	N/A	integer
cloudRestoreMemcacheRetention	30	N/A	integer
cloudRestoreMemcacheSize	16777216	N/A	integer
cloudRestoreMinChunkSize	524288	N/A	integer
cloudScsiCommandTimeout	0	N/A	integer
cloudSpacemapBinaryFormat	true	N/A	boolean
cloudSpacemapTrimEnabled	false	N/A	boolean
cloudStallCount	20	N/A	integer
cloudStallInterval	30	N/A	integer
cloudSuccessCleanupTimeout	600	N/A	integer
cloudTcmuHwMaxSectors	2048	N/A	integer
cloudTcmuRingBufferSize	128	N/A	integer
cloudThrottleDirtyData	268435456	N/A	integer
cloudThrottleMemPercent	90	N/A	integer
cloudThrottleObjectsPoll	5	N/A	integer
cloudThrottleObjectsRatio	0.9	N/A	float
cloudThrottleRatePoll	120	N/A	integer
cloudTotalTempSpace	26843545600	N/A	integer
cloudTransferTimeout	604800	N/A	integer
cloudTrimMonitorIntervalSecs	10	N/A	integer
cloudTrimMonitorNumAttempts	6	N/A	integer
cloudUseL2ARC	false	N/A	boolean
cloudUseSLOG	false	N/A	boolean
cloudVerifyBlockmapOnOffload	false	N/A	boolean
cloudWaitBeforeForceExport	20	N/A	integer
createFileBasedLuns	true	N/A	boolean
ddtMassEvictOnNextReboot	false	N/A	boolean

ddtPreloadTimeout	900	N/A	integer
ddtUncompressedEntrySizeInCore	448	N/A	integer
ddtUniqueMaxAddToExistingPool	1048576	N/A	integer
ddtUniqueMaxHeadroom	209715200	N/A	integer
ddtUniqueMaxTotalForNewPool	1073741824	N/A	integer
debugOffloadToFile	false	N/A	boolean
dedupeVolsWithSmallBlocks	false	N/A	boolean
excludeAllowedHostsPrefix	N/A	N/A	string
gatewayDeviceWait	10	N/A	integer
initInstallZFS	true	N/A	boolean
lunTargetQueueDepth	32	N/A	integer
lunTargetTPU	0	N/A	integer
maintCloudObjectsCount	5000	N/A	integer
maintCloudObjectsRetention	7	N/A	integer
maintServiceCleanupFrequency	21600	N/A	integer
maintServiceFrequency	300	N/A	integer
mvrSkipInputPipeWait	true	N/A	boolean
mvrSkipOutputPipeWait	true	N/A	boolean
passwordMinLength	8	N/A	integer
poolCreateRecordSize	131072	N/A	integer
poolListTimeout	300	N/A	integer
poolManagementTimeout	3600	N/A	integer
repairAutoIncrement	1000	N/A	integer
repairMaxParallel	5	N/A	integer
replDynamicThrottle	true	N/A	boolean
replFallbackToMgmtAddr	true	N/A	boolean
replMaxRetries	3	N/A	integer
replMaxStreams	5	N/A	integer
replMgmtTimeout	840	N/A	integer
replRetryInterval	600	N/A	integer
replServerAliveCountMax	5	N/A	integer
replServerAliveInterval	120	N/A	integer
replServiceFrequency	30	N/A	integer
replSnapshotTimeout	3600	N/A	integer
replStallCount	20	N/A	integer
replStallInterval	30	N/A	integer
replStreamCompress	true	N/A	boolean
replStreamDedupe	false	N/A	boolean
replStreamResume	true	N/A	boolean
replTargetSpaceLimitRatio	0.95	N/A	float
replTransferRate	536870912	N/A	integer
replTransferTimeout	604800	N/A	integer
sessionHistoryRetentionHours	48	N/A	integer
shareIgnoreAllowedHosts	false	N/A	boolean
skipOffloadCleanup	false	N/A	boolean
smbReloadWait	15	N/A	integer
snapshotCreateTimeout	600	N/A	integer
snapshotDeleteSync	false	N/A	boolean
snapshotDestroyTimeout	7200	N/A	integer
systemNotifyEnabled	false	N/A	boolean
taskIntervalAlert	30	N/A	integer
taskParallelAlert	5	N/A	integer
taskParallelArchive	5	N/A	integer
taskParallelCloud	5	N/A	integer
taskPollingFrequency	30	N/A	integer

```
taskRetentionAlert              | 0            | N/A        | integer
taskRetentionArchive            | 86400        | N/A        | integer
taskRetentionCloud              | 86400        | N/A        | integer
volumeCreateRecordSize          | 131072       | N/A        | integer
volumeCreateTimeout             | 600          | N/A        | integer
volumeDeleteSync                | false        | N/A        | boolean
volumeDestroyTimeout            | 7200         | N/A        | integer
```

Specific values can be filtered by using the **grep** option (as shown in Example 5-20) where we list the maximum number of streams for archive, cloud, and replication operations.

Example 5-20 Filter the preference list for specific values using grep

```
[serveradmin@t3-spp-vsnap ~]$ vsnap system pref get | grep -i MaxStreams
archiveMaxStreams               | 5            | N/A        | integer
cloudMaxStreams                 | 5            | N/A        | integer
replMaxStreams                  | 5            | N/A        | integer
```

5.7.1 Changing replication streams and timeouts

The preferences of the replication process between two vSnap servers can be adjusted if required. By default, five streams are used to replicate data from the source vSnap server to the target vSnap server. If enough bandwidth, CPU and memory resources on both vSnap servers are available, the number of streams could be increased to 10, as shown in Example 5-21.

Example 5-21 Increase replication streams

```
[serveradmin@spp-vsnap-demo ~]$ vsnap system pref set --name replMaxStreams
--value 10

NAME: replMaxStreams
DEFAULT VALUE: 5
USER VALUE: 10
TYPE: integer
```

Another option is the timeout setting of a replication task. By default the timeout its set to 608400 seconds, which translate to 7 days. A large timeout can be required especially for initial replication tasks. To change this value to a timeout of 24 hours you can set it to 86400 seconds, as shown in Example 5-22.

Example 5-22 Change timeout values for the replication

```
[serveradmin@spp-vsnap-demo ~]$ vsnap system pref set --name replTransferTimeout
--value 86400

NAME: replTransferTimeout
DEFAULT VALUE: 604800
USER VALUE: 86400
TYPE: integer
```

5.8 vSnap server volumes and snapshots

In this section, we describe volumes and snapshots that are used on the vSnap server to manage the backup data. In the first step, we review what occurs when a backup is performed; in the second step, we describe what occurs when backup data is replicated.

Before a backup could be taken, a resource need to assigned to an SLA. Chapter 6.1.3, “Assigning an SLA policy” on page 174 explains this as example for VMs and chapter 1.3, “SLA backup policies” on page 19 speaks about SLA policies in general.

Replication must be configured before it can be used. In 15.3.1, “Configuring vSnap replication” on page 410 all required steps are explained.

5.8.1 Volumes and snapshots for backup data

When a SLA for backing up data is executed for the first time, immediately a new volume is created while the backup job is running. In Example 5-23 the command **`vsnap volume show`** list the newly created volume with the **ID 2**. The volume with the ID 1 refers to a cloud cache area which is created with initialization of the vSnap pool. For more information, see “Preparing the disk cache area” on page 420.

Example 5-23 Execution of a SLA to create a new volume in the vSnap server

```
[serveradmin@t2-spp-vsnap ~]$ vsnap volume show
ID | TYPE       | POOL | IS CLONE | TOTAL    | FREE    | USED    | NAME                                    | TAGS
---------------------------------------------------------------------------------------------------------------
1  | filesystem | 1    | No       | 100.00GB | 93.62GB | 24.00KB | vsnap_metadata_cloud                    | N/A
2  | filesystem | 1    | No       | 100.00GB | 93.62GB | 2.77GB  | spp_1004_2002_16b6aeb7936__group0_12_ | N/A
```

The first number referenced in the volume name (in our example 1004) is the **Job ID** of the job using this volume. In our example we have associated a VMware backup with the Bronze SLA, which created the job `vmware_Bronze` that is listed under the **Jobs and Operations** in the IBM Spectrum Protect Plus GUI. The Job `ID:1004` of `vmware_Bronze` could be figured out when looking into the first line of the job log, as shown in Figure 5-30.

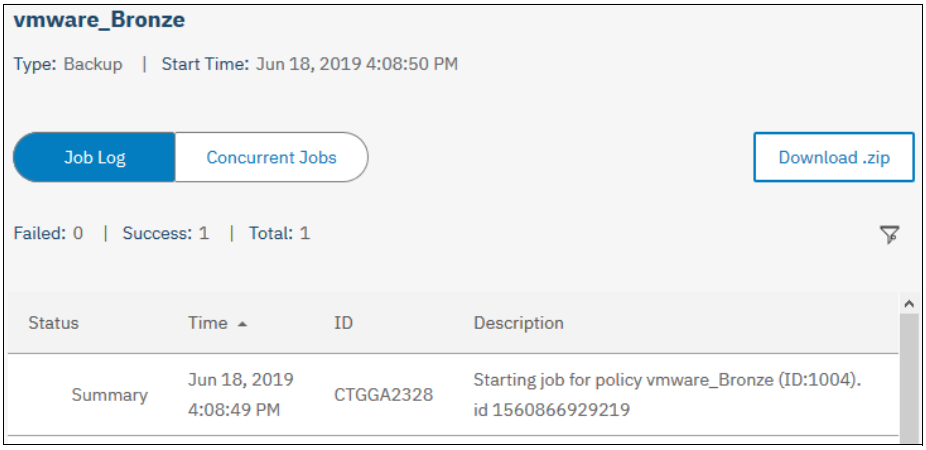

Figure 5-30 vmware_Bronze Job Log

The volume is based on a file system and the Operating System **df -h** command lists all file system of the vSnap server as shown in Example 5-24. The vSnap server file system belonging to the volume could be identified by using the volume ID, which is part of the file system. In our example it is volume ID 2 corresponding to file system /vsnap/vpool1/fs2.

Example 5-24 The df -h command list the vSnap file systems

```
[serveradmin@t2-spp-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  153M   16G   1% /run
tmpfs                                16G     0   16G   0% /sys/fs/cgroup
/dev/mapper/centos-root              34G  2.4G   31G   8% /
/dev/mapper/vsnapdata-vsnapdatalv   126G   33M  126G   1% /opt/vsnap-data
/dev/sda1                          1014M  191M  824M  19% /boot
vpool1                               85G     0   85G   0% /vsnap/vpool1
vpool1/fs1                           85G     0   85G   0% /vsnap/vpool1/fs1
vpool1/fs2                           97G   13G   85G  13% /vsnap/vpool1/fs2
tmpfs                               3.1G     0  3.1G   0% /run/user/1001
```

When running operating system commands, such as **du** or **find** as shown in Example 5-25, you can also explore the content of the file system to determine which VM is backed up in which file system.

Example 5-25 Explore the content of a vSnap server volume

```
[serveradmin@t2-spp-vsnap ~]$ du /vsnap/vpool1/fs2
12774239        /vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356
12774239        /vsnap/vpool1/fs2/b97956fe
12774240        /vsnap/vpool1/fs2

[serveradmin@t2-spp-vsnap ~]$ find /vsnap/vpool1/fs2
/vsnap/vpool1/fs2
/vsnap/vpool1/fs2/b97956fe
/vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356
/vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356/t2-vm-win.nvram
/vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356/backupMetadata.json
/vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356/t2-vm-win.vmx
/vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356/t2-vm-win.vmxf
/vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356/t2-vm-win-flat.vmdk
/vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356/backupMetadataDisks.json
/vsnap/vpool1/fs2/b97956fe/t2-vm-win.vm-356/t2-vm-win.vmdk
```

As long as the backup job is running we could see a share on the file system which is associated with the SLA. The vSnap CLI command **vsnap share show** as shown in Example 5-26 is listing the active share, where the **Volume ID 2** and the file system name **/vsnap/vpool1/fs2** could be identified.

Example 5-26 Active share

```
[serveradmin@t2-spp-vsnap ~]$ vsnap share show
ID | TYPE | PARENT VOL | PARTNER ID | NAME
------------------------------------------------------
1  | nfs  | 2          | N/A        | /vsnap/vpool1/fs2

[serveradmin@t2-spp-vsnap ~]$ vsnap share show --id 1
```

```
ID: 1
NAME: /vsnap/vpool1/fs2
SHARE TYPE: nfs
VOLUME ID: 2
PARTNER ID: N/A
CREATED: 2019-06-18 14:09:12 UTC
UPDATED: 2019-06-18 14:09:33 UTC
SHARE OPTIONS:
    ALLOWED HOSTS:
        10.0.250.27
    READ ONLY: No
```

The share is used to transfer the backup data from the hypervisor, application or database to the vSnap server.

Note: For the backup of VMware VMs a VADP proxy is required, therefore the backup data is transferred from the hypervisor through the VADP proxy to the vSnap server. For more information about the VADP proxy, see 1.2.4, “VADP proxy server” on page 16.

After the backup is completed, a Snapshot on the file system is created, which is the backup entity or the backup version. In Example 5-27 we use the command **`vsnap snapshot show`** list all snapshots.

Example 5-27 List available snapshots on the vSnap pool

```
[serveradmin@t2-spp-vsnap ~]$ vsnap snapshot show
ID | PARENT ID | CREATED                 | NAME
-------------------------------------------------------------------------
1  | 2         | 2019-06-18 14:11:53 UTC | spp_1004_2002_2_16b6aee0745

[serveradmin@t2-spp-vsnap ~]$ vsnap snapshot show --id 1

ID: 1
NAME: spp_1004_2002_2_16b6aee0745
PARENT ID: 2
PARENT NAME: spp_1004_2002_16b6aeb7936__group0_12_
POOL ID: 1
POOL NAME: primary
HAS CLONES: No
USED SPACE: 0.00KB
VERSION ID: 12686594730219022292
CREATED: 2019-06-18 14:11:53 UTC
UPDATED: 2019-06-18 14:11:53 UTC
```

The volume where the snapshot is taken from could be identified by the **Parent ID** and the **Parent Name**.

5.8.2 Volumes and snapshots for replication data

In this chapter explore what happens on a vSnap server when data is replicated between two vSnap servers. After the configuration for the vSnap server replication is completed (refer to 15.3.1, “Configuring vSnap replication”) a replication job could be started.

In our example a vSnap server partnership is created between **t2-spp-vsnap** and **t2-spp-vsnap-dr**. When a replication task of an SLA policy is started, the following things happen on the source and target vSnap server:

1. A new replication target volume is created on the target vSnap server.
2. A relationship for the source and target volume is created on the source and target vSnap server.
3. A new and empty snapshot on the target volume is created.
4. A replication session is created on the source vSnap server.
5. The target snapshot is updated on the target vSnap server after the replication session is completed.

In the following replication example we show the previous steps in detail.

As shown in Example 5-28 the replication job creates a target volume (ID 2) on the target vSnap server **t2-spp-vsnap-dr**.

Example 5-28 Replication target volume is created on the target vSnap server

```
[serveradmin@t2-spp-vsnap-dr ~]$ vsnap volume show
ID | TYPE       | POOL | IS CLONE | TOTAL    | FREE    | USED    | NAME                                                                                                | TAGS
---------------------------------------------------------------------------------------------------------------------------------------------------------------
1  | filesystem | 1    | No       | 100.00GB | 84.21GB | 24.00KB | vsnap_metadata_cloud                                                                                | N/A
2  | filesystem | 1    | No       | 100.00GB | 84.21GB | 12.18GB | spp_16b6b15c8d1_spp_1004_2002_16b6aeb7936__group0_12_ | N/A
```

Next the replication relationship between both volumes is configure on both vSnap servers. The vSnap CLI command **`vsnap relationship show`** is listing these relationships as shown in Example 5-29 for the source vSnap server `t2-spp-vsnap` and in Example 5-30 on page 167 for the target vSnap server `t2-spp-vsnap-dr`.

Example 5-29 vSnap relationship on the source server created

```
[serveradmin@t2-spp-vsnap ~]$ vsnap relationship show
ID                                   | PARTNER ADDR    | PARTNER TYPE | LOCAL ROLE | LAST SYNC | LOCAL VOL                          | REMOTE VOL
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------
7ce7be815c5f4fc9147c52a55bd6530d | t2-spp-vsnap-dr | vsnap        | primary    | N/A       | spp_1004_2002_16b6aeb7936__group0_12_ | spp_16b6b15c8d1_spp_1004_2002_16b6aeb7936__group0_12_
```

```
[serveradmin@t2-spp-vsnap ~]$ vsnap relationship show --id 7ce7be815c5f4fc9147c52a55bd6530d

ID: 7ce7be815c5f4fc9147c52a55bd6530d
PARTNER ID: 5345ec9e347b40d5ae63b2397f64c8da
PARTNER TYPE: vsnap
PARTNER ADDR: t2-spp-vsnap-dr
LOCAL ROLE: primary
LOCAL POOL ID: 1
LOCAL VOL ID: 2
LOCAL VOL NAME: spp_1004_2002_16b6aeb7936__group0_12_
REMOTE POOL ID: 1
REMOTE VOL ID: 2
REMOTE VOL NAME: spp_16b6b15c8d1_spp_1004_2002_16b6aeb7936__group0_12_
LAST SYNC STATUS: N/A
LAST SYCNED SNAP ID: N/A
LAST ATTEMPT SNAP ID: 1
CREATED: 2019-06-18 14:55:19 UTC
UPDATED: 2019-06-18 14:55:19 UTC
```

Example 5-30 vSnap relationship on the target server created

```
[serveradmin@t2-spp-vsnap-dr ~]$ vsnap relationship show
ID                               | PARTNER ADDR | PARTNER TYPE | LOCAL ROLE | LAST SYNC | LOCAL VOL                                             | REMOTE VOL
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------
7ce7be815c5f4fc9147c52a55bd6530d | t2-spp-vsnap | vsnap        | replica    | N/A       | spp_16b6b15c8d1_spp_1004_2002_16b6aeb7936__group0_12_ | spp_1004_2002_16b6aeb7936__group0_12_

[serveradmin@t2-spp-vsnap-dr ~]$ vsnap relationship show --id 7ce7be815c5f4fc9147c52a55bd6530d

ID: 7ce7be815c5f4fc9147c52a55bd6530d
PARTNER ID: bc1c6c4052a744ad86bdc25625e2c516
PARTNER TYPE: vsnap
PARTNER ADDR: t2-spp-vsnap
LOCAL ROLE: replica
LOCAL POOL ID: 1
LOCAL VOL ID: 2
LOCAL VOL NAME: spp_16b6b15c8d1_spp_1004_2002_16b6aeb7936__group0_12_
REMOTE POOL ID: 1
REMOTE VOL ID: 2
REMOTE VOL NAME: spp_1004_2002_16b6aeb7936__group0_12_
LAST SYNC STATUS: N/A
LAST SYCNED SNAP ID: N/A
LAST ATTEMPT SNAP ID: N/A
CREATED: 2019-06-18 14:55:19 UTC
UPDATED: 2019-06-18 14:55:19 UTC
```

The fields **Partner Address**, **Local Volume ID**, **Local Volume Name**, **Remote Volume ID** and **Remote Volume Name** are used to identify which vSnap servers and volumes are used for this relationship.

A empty Snapshot is created as shown by running the **`vsnap snapshot show`** command, as shown in Example 5-31. The details of the snapshot are not yet available by running **`vsnap snapshot show --id 1`** the command.

Example 5-31 Empty snapshot is created

```
[serveradmin@t2-spp-vsnap-dr ~]$ vsnap snapshot show
ID | PARENT ID | CREATED                 | NAME
--------------------------------------------------------------------
1  | 2         | 2019-06-18 14:55:20 UTC | spp_1004_2101_16b6b15c808

[serveradmin@t2-spp-vsnap-dr ~]$ vsnap snapshot show --id 1
ERROR: SnapshotInfoError: Failed to collect snapshot details for 1
```

When the replication is running a session is created for the replication on the source vSnap server. In Example 5-32 the command **`vsnap session show`** is used to list the replication session on the source server t2-spp-vsnap. The status field indicates if it is an **Active** session or a Completed session.

Example 5-32 List replication session using the vSnap CLI

```
[serveradmin@t2-spp-vsnap ~]$ vsnap session show
ID | RELATIONSHIP                     | PARTNER TYPE | LOCAL SNAP | REMOTE SNAP | STATUS | SENT   | STARTED                 | ENDED
-----------------------------------------------------------------------------------------------------------------------------------
1  | 7ce7be815c5f4fc9147c52a55bd6530d | vsnap        | 1          | 1           | ACTIVE | 9.23GB | 2019-06-18 14:55:26 UTC | N/A

[serveradmin@t2-spp-vsnap ~]$ vsnap session show --id 1

ID: 1
```

```
RELATIONSHIP ID: 7ce7be815c5f4fc9147c52a55bd6530d
PARTNER TYPE: vsnap
REPL ADDRESS: 10.0.250.23
LOCAL SNAP ID: 1
LOCAL SNAP NAME: spp_1004_2002_2_16b6aee0745
REMOTE SNAP ID: 1
REMOTE SNAP NAME: spp_1004_2101_16b6b15c808
PRIORITY: 50
STATUS: ACTIVE
SENT: 9.42GB
QUEUED: 2019-06-18 14:55:20 UTC
STARTED: 2019-06-18 14:55:26 UTC
ENDED: N/A
MESSAGE: Data transfer in progress

[serveradmin@t2-spp-vsnap ~]$ vsnap session show --id 1

ID: 1
RELATIONSHIP ID: 7ce7be815c5f4fc9147c52a55bd6530d
PARTNER TYPE: vsnap
REPL ADDRESS: 10.0.250.23
LOCAL SNAP ID: 1
LOCAL SNAP NAME: spp_1004_2002_2_16b6aee0745
REMOTE SNAP ID: 1
REMOTE SNAP NAME: spp_1004_2101_16b6b15c808
PRIORITY: 50
STATUS: COMPLETED
SENT: 12.22GB
QUEUED: 2019-06-18 14:55:20 UTC
STARTED: 2019-06-18 14:55:26 UTC
ENDED: 2019-06-18 14:57:37 UTC
MESSAGE: Completed
```

After the replication session is completed the snapshot on the target volume is updated and details of the snapshot are available which is listed in Example 5-33 by running the vSnap CLI **vsnap snapshot show** command.

Example 5-33 Snapshot on the replication target is created

```
[serveradmin@t2-spp-vsnap-dr ~]$ vsnap snapshot show
ID | PARENT ID | CREATED                 | NAME
-------------------------------------------------------------------
1  | 2         | 2019-06-18 14:57:37 UTC | spp_1004_2101_16b6b15c808

[serveradmin@t2-spp-vsnap-dr ~]$ vsnap snapshot show --id 1

ID: 1
NAME: spp_1004_2101_16b6b15c808
PARENT ID: 2
PARENT NAME: spp_16b6b15c8d1_spp_1004_2002_16b6aeb7936__group0_12_
POOL ID: 1
POOL NAME: primary
HAS CLONES: No
USED SPACE: 0.00KB
VERSION ID: 12686594730219022292
CREATED: 2019-06-18 14:57:37 UTC
```

```
UPDATED: 2019-06-18 14:57:37 UTC
```

The relationship of the volumes on the source and the target server is updated as shown in Example 5-34 for the source vSnap server and in Example 5-35 for the target vSnap server.

Example 5-34 vSnap relationship on the source server completed

```
[serveradmin@t2-spp-vsnap ~]$ vsnap relationship show --id 7ce7be815c5f4fc9147c52a55bd6530d

ID: 7ce7be815c5f4fc9147c52a55bd6530d
PARTNER ID: 5345ec9e347b40d5ae63b2397f64c8da
PARTNER TYPE: vsnap
PARTNER ADDR: t2-spp-vsnap-dr
LOCAL ROLE: primary
LOCAL POOL ID: 1
LOCAL VOL ID: 2
LOCAL VOL NAME: spp_1004_2002_16b6aeb7936__group0_12_
REMOTE POOL ID: 1
REMOTE VOL ID: 2
REMOTE VOL NAME: spp_16b6b15c8d1_spp_1004_2002_16b6aeb7936__group0_12_
LAST SYNC STATUS: COMPLETED
LAST SYCNED SNAP ID: 1
LAST ATTEMPT SNAP ID: 1
CREATED: 2019-06-18 14:55:19 UTC
UPDATED: 2019-06-18 14:55:19 UTC
```

Example 5-35 vSnap relationship on the target server completed

```
[serveradmin@t2-spp-vsnap-dr ~]$ vsnap relationship show --id 7ce7be815c5f4fc9147c52a55bd6530d

ID: 7ce7be815c5f4fc9147c52a55bd6530d
PARTNER ID: bc1c6c4052a744ad86bdc25625e2c516
PARTNER TYPE: vsnap
PARTNER ADDR: t2-spp-vsnap
LOCAL ROLE: replica
LOCAL POOL ID: 1
LOCAL VOL ID: 2
LOCAL VOL NAME: spp_16b6b15c8d1_spp_1004_2002_16b6aeb7936__group0_12_
REMOTE POOL ID: 1
REMOTE VOL ID: 2
REMOTE VOL NAME: spp_1004_2002_16b6aeb7936__group0_12_
LAST SYNC STATUS: COMPLETED
LAST SYCNED SNAP ID: 1
LAST ATTEMPT SNAP ID: 1
CREATED: 2019-06-18 14:55:19 UTC
UPDATED: 2019-06-18 14:55:19 UTC
```

The fields **LAST SYNC STATUS**, **LAST SYNCED SNAP ID**, **LAST ATTEMPT SNAP ID** in both detailed outputs provide details about the completion status and the snapshots used for this relationship. The timestamps in the CREATED and UPDATED field are not adjusted; instead, they remain on the values when the relationship was initially created.

6

Backing up and restoring virtualized systems

In this chapter, we discuss backing up, restoring, and reusing a virtualized system's data. IBM Spectrum Protect Plus supports VMware, Microsoft Hyper-V hypervisors, and Amazon Elastic Compute Cloud (EC2).

For VMware and Microsoft Hyper-V hypervisors, backups are snapshots that are based in a block-level incremental forever model: one base full backup is followed by many always incremental backups. Each backup is a *Synthetic Full* backup, which means that only one restore activity must be started whether it refers base or incremental data backups.

In addition to the virtual machine (VM) backup, the files on the VM can be indexed for single file restore out of a VM backup without requiring you to restore the full VM or dedicated virtual disks (VDisks) first. The VM Restore Options Test Mode (Instant Access), Clone Mode (Instant Recover), and Production Mode (Instant Recover) support different VM restore scenarios along with data reuse options.

EC2 data is stored in Amazon Web Services (AWS) Elastic Block Store (EBS) snapshots rather than in the vSnap server. IBM Spectrum Protect Plus manages these snapshots for backup and restore operations.

Note: For more information about the latest Hypervisor (Microsoft Hyper-V and VMware) and cloud instance (Amazon EC2) backup and restore requirements, see *IBM Spectrum Protect Plus - All Requirements Doc*.

This chapter includes the following topics:

6.1 VM backup configuration basics

You must register the Virtualized Systems that you want to protect in IBM Spectrum Protect Plus. Then, you create jobs to back up and restore the VMs and resources that are associated with the virtualized systems.

We explain how to add Virtualized Systems and assign a VM to an SLA (backup schedules). We also show how to run backup jobs based on schedules or manual and how ad hoc backups of single or multiple VMs can be performed. Finally, we show how backups can be distributed on different vSnap servers and explain the backup options that can be configured for VMs.

6.1.1 Create an identity

To access a virtualized system, resource credentials are required. It is recommended to create and maintain them before adding a resource. For more information about how to create an identity for resources, see “Identities, keys, and certificates” on page 11.

6.1.2 Add a virtualized system resource

To register virtualized systems in the IBM Spectrum Protect Plus GUI, select **Manage Protection** → **Virtualized Systems** → ***your virtualized systems type***:

- For VMware, choose **Manage vCenter** and then, click **Add vCenter** to register a VMware vCenter server.
- For Hyper-V, choose **Manage Hyper-V Server** and then, click **Add Hyper-V Server** to register a Microsoft Hyper-V server.
- For Amazon Elastic Compute Cloud (EC2), you need an AWS Identity and Access Management (IAM) user to complete the task. IAM users must have access keys and required permissions. You can use the AWS Management Console to create an IAM user.

  For more information about creating an IAM user, see this AWS web page.

For a VMware example, you must enter information, as shown in Figure 6-1, to add a VMware vCenter server.

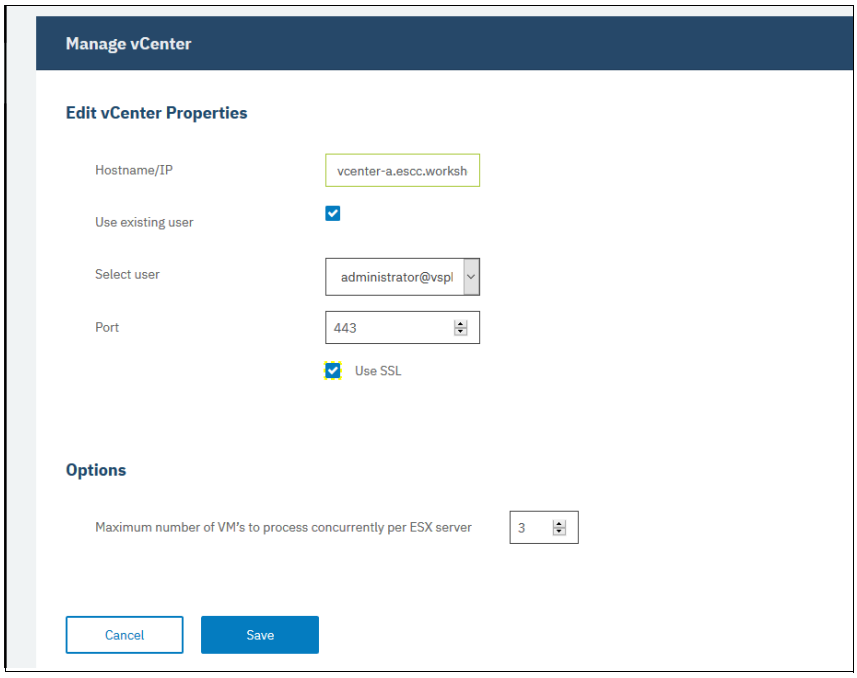

Figure 6-1 Add a VMware vCenter server as virtualized system resource

In the Manage vCenter window, specify the following information:

- Virtualized system DNS name or IP address.
- A user to access the virtualized system with suitable privileges. Choose a configured identity (for more information, see “Identities, keys, and certificates” on page 11) by selecting the **Use existing user** option or specifying the user name and password in the dialog box.
- A port to connect and whether a SSL connection is used to connect to the virtualized system.
- As another option, the maximum number of VMs to process concurrently per virtualized system server also can be configured. The default is 3, which is set per ESX server in a VMware vCenter environment.

For more information about adding a virtualized system resources and the required privileges of the user, see IBM Knowledge Center.

After a Hypervisor is added, a system-defined Hypervisor Inventory job is created and run to collect all configured VMs and settings of the Hypervisor. The job runs by default once a day, and the schedule and job log can be adjusted or viewed under **Jobs and Operations** in the IBM Spectrum Protect Plus GUI.

The last inventory job run time is shown in the VMware, Hyper-V, or Amazon EC2 Backup menu. From there, a new inventory can be started manually by clicking **Run Inventory** (as shown in Figure 6-2) for VMware.

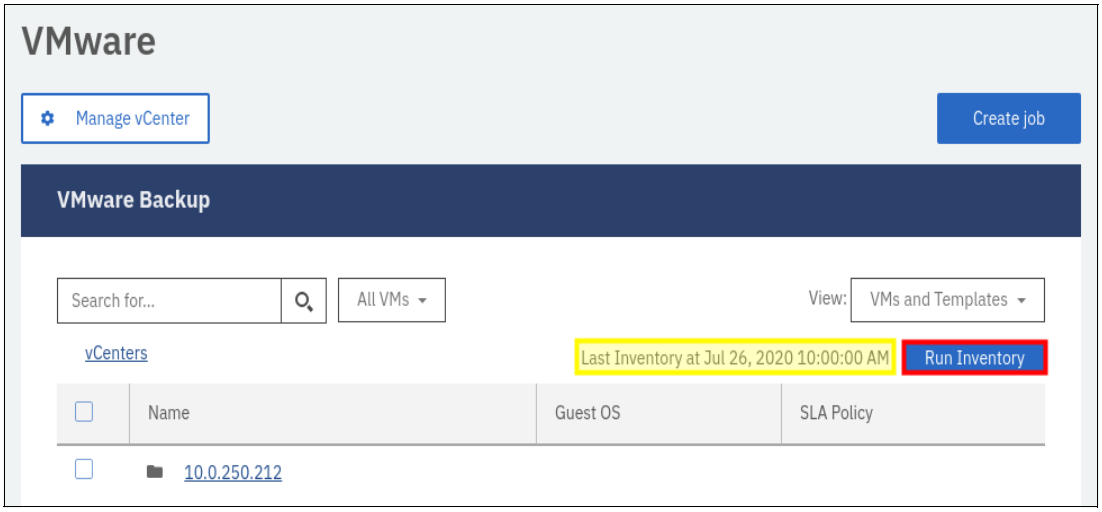

Figure 6-2 VMware Backup menu with information about the inventory job

6.1.3 Assigning an SLA policy

After a Virtualized Systems has been defined in IBM Spectrum Protect Plus, assign one or more SLA policies to the VMs that should be backed up.

An SLA must be defined before it can be assigned to a VM. For more information about how to create and configure an SLA, see 1.3, “SLA backup policies” on page 19.

Use the search function (and filter) to search for available resources, or click through the provided Virtualized Systems tree to select your VMs. Filtering options can be set for VMs and Templates, VMs, Datastore, Tags and Categories, or Hosts and Clusters. Single VMs, Folders, Tags, Datastores, Hosts and Clusters, Data centers, and vCenters can be assigned to an SLA.

In our example for VMware, as shown in Figure 6-3, we searched for `vm-win`, selected both VMs from the search result, and added them to the Bronze SLA by clicking **Select SLA Policy** and checking **Bronze**.

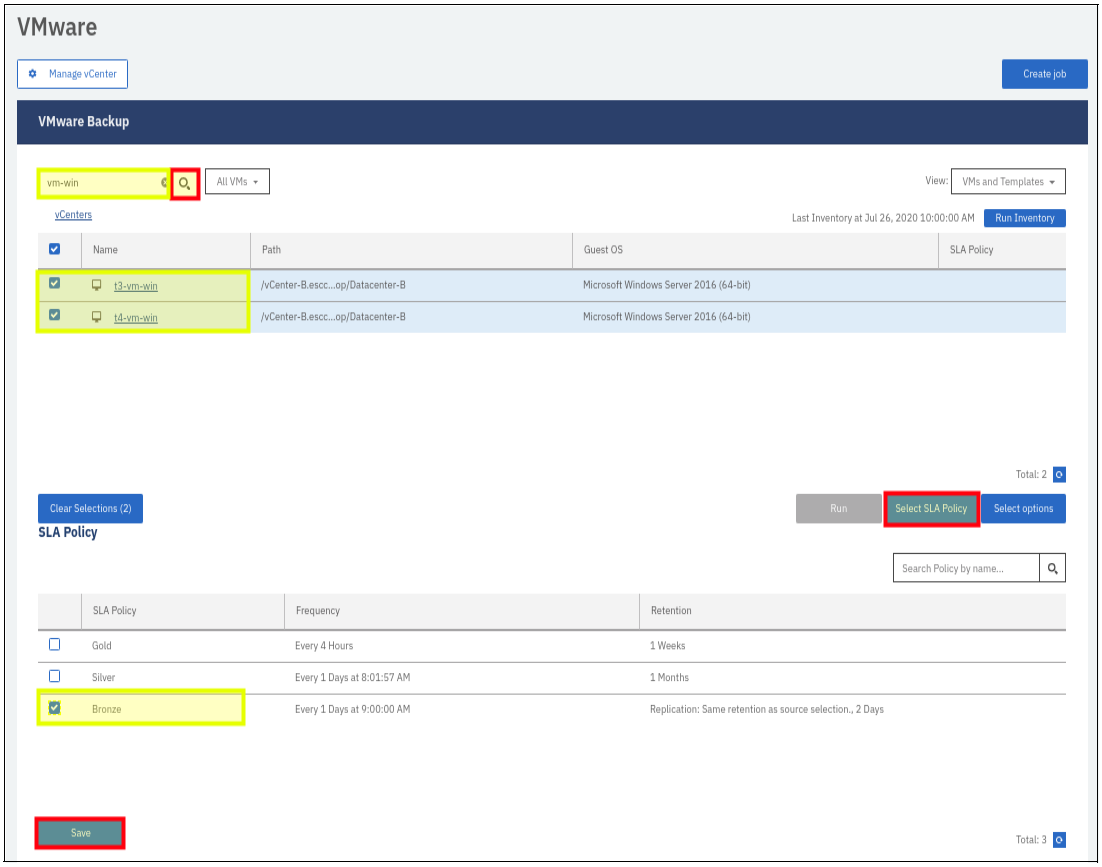

Figure 6-3 Adding VMs to a SLA

After the Save button is clicked, a new backup job Bronze is created. Scroll down to the SLA Policy Status section to review the job schedules and check the backup job log output, as shown in Figure 6-4.

Figure 6-4 SLA Backup Job listed SLA Policy Status at the VMware Backup page

Alternatively, switch to the **Jobs and Operations** menu to check the schedule and job log output. In the **Jobs and Operations** → **Schedule** view, the backup job Bronze is listed with the prefix `vmware` (`vmware_Bronze`) to indicate that it is a VMware backup job that is associated with the Bronze SLA.

Excluding single VDisks from the SLA policy for a VMware job

After you save a backup job definition, you can exclude individual VDisks (VMDK) of a VMware VM from the assigned SLA policy.

To exclude a VDisk from the SLA policy go to the same view, as shown in Figure 6-3 on page 175 and search for the VM with the disk to exclude. Click the VM to open the VMDK view and select the disk that you want to exclude, as shown in Figure 6-5.

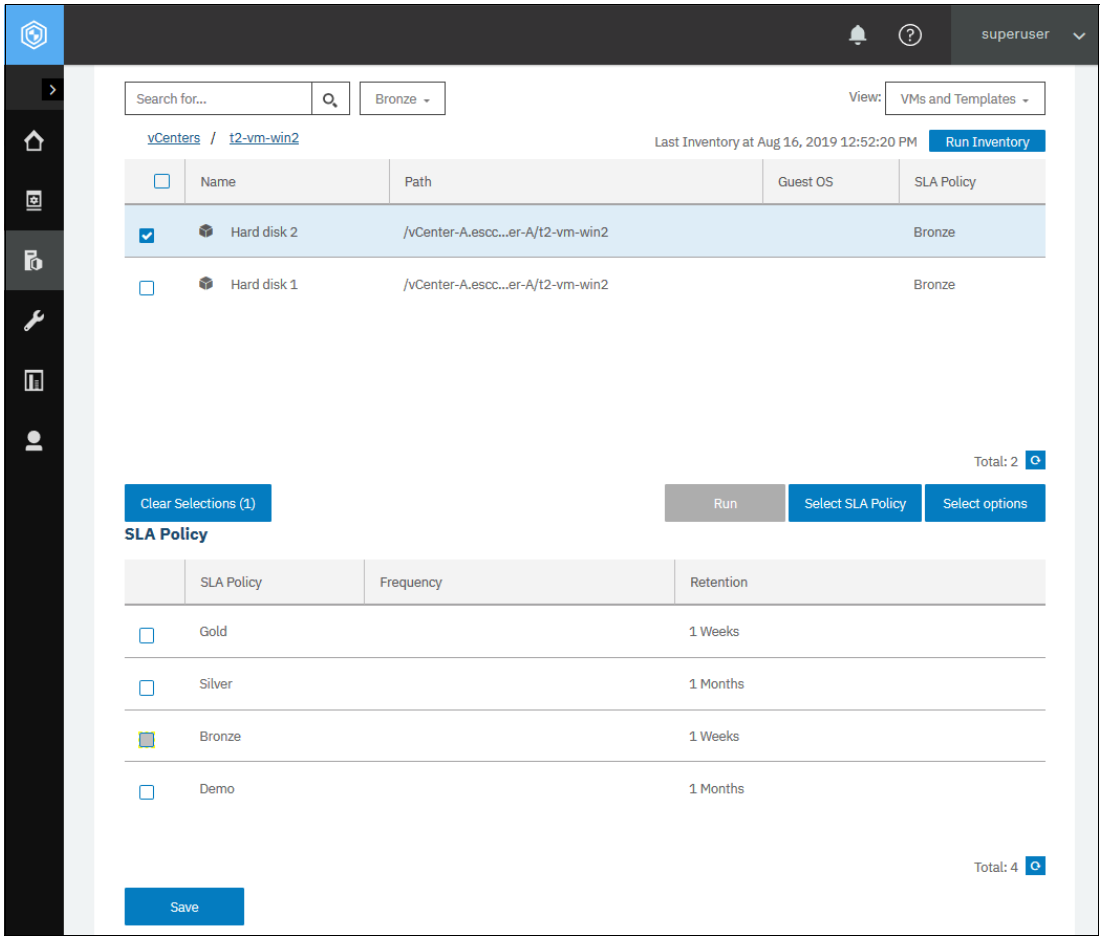

Figure 6-5 Exclude a VMDK disk from a SLA Policy

After you select the disk (in our example: Hard disk 2) click **Select SLA Policy**, clear (in our example) the **Bronze** SLA option for the disk, and click the **Save** button. The next backup for the VM will now run without backing up the `Hard Disk 2` of the VM `t2-vm-win2`.

Note: Excluding a single VDisks from an SLA policy is available only for VMs in a VMware environment.

6.1.4 Running a backup job for an SLA

The backup that is defined in an SLA is run at the scheduled time frame or started manually by selecting **Actions** → **Start**, as shown in Figure 6-4 on page 175.

When the backup is started, you can follow the backup process by monitoring the job log by expanding the Bronze section in the same view, as shown in Figure 6-6.

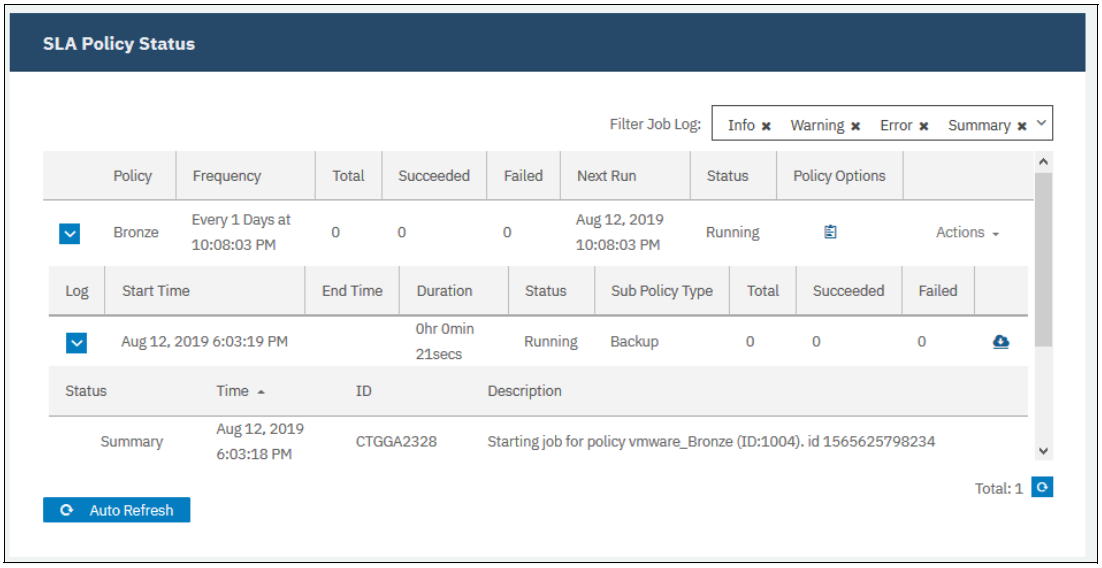

Figure 6-6 SLA Backup Job Log listed directly in the SLA Policy Status at the VMware Backup page

The following process is used as part of the backup workflow:

1. The backup job (SLA Policy) is run manually or scheduled.
2. The vSnap server is determined, which is used as target for the snapshot backup.
3. A VMware or Hyper-V snapshot of the source VM is taken.
4. A new target volume is created on the vSnap server for the initial base backup, or an existing value is used for the incremental forever backup.
5. The volume is mounted in one of the following conditions:
   - On the VADP proxy server as NFS datastore. If the VADP proxy runs on the same machine as the vSnap server, no NFS mount is required and the data transfer takes place internally on the machine.
   - On the source Hyper-V hosts by using iSCSI.
6. The backup is performed on the mounted volume (by using VADP transport method for a VMware backup).
7. The VM snapshot is deleted.
8. A vSnap snapshot of the target volume is taken on the vSnap server (that is, the “backup version”).
9. The vSnap server volume is unmounted.
10. The Spectrum Protect Plus catalog captures the details of the backup (source, vSnap server, retention, schedule, and so on).

Note: For more information about VMware snapshot behavior, see IBM Knowledge Center.

6.1.5 Running a backup for a single VM

You might want to run a backup of a single VM instead of backing up all VMs assigned to an SLA for the following reasons:

- The backup of a specific VM failed, and you want to rerun the backup.
- You are planning maintenance activities for a VM and want to have a backup of this VM first before starting the maintenance.

Rerunning a failed backup

If you find a warning or error in the Dashboard view or see a Warning or Failed SLA backup job under Jobs and Operations, analyze the log for the details of this status. If the warning or error was caused by an issue in backing up a VM and you fixed the problem, you can rerun the failed VMs.

As shown in Figure 6-7 in the Job History tab under the Jobs and Operations view, the `vmware_Bronze` job log is selected. To start a backup of the failed VMs, can click **Actions** → **Rerun failed**.

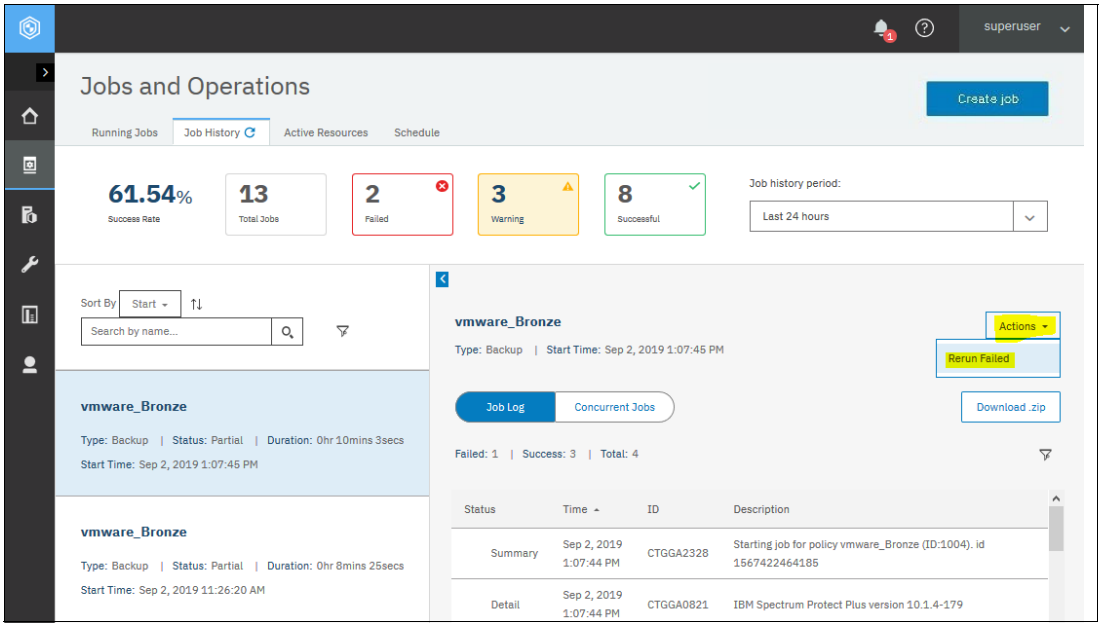

Figure 6-7 Job History with job details and the rerun failed button

Note: The job log is used to capture activities of the rerun. New status messages are added that include timestamps at the end of the job log.

Starting the SLA for a single VM (ad-hoc backup)

Situations exist in which you must run an ad-hoc backup of one or more VMs. One reason might be that a maintenance activity (for example, a software update or system reconfiguration) is planned, and you require a current backup of the VMs that are part of this maintenance activity. In such a scenario, you can use the ad-hoc backup capability.

In our VMware example, we want to start an ad-hoc backup for `t3-vm-win` VM, which we added to the Bronze SLA. Select **Manage Protection** → **Virtualized Systems** → **VMware** to view all of the VMs that are backed up with the Bronze SLA, as shown in Figure 6-8 on page 179.

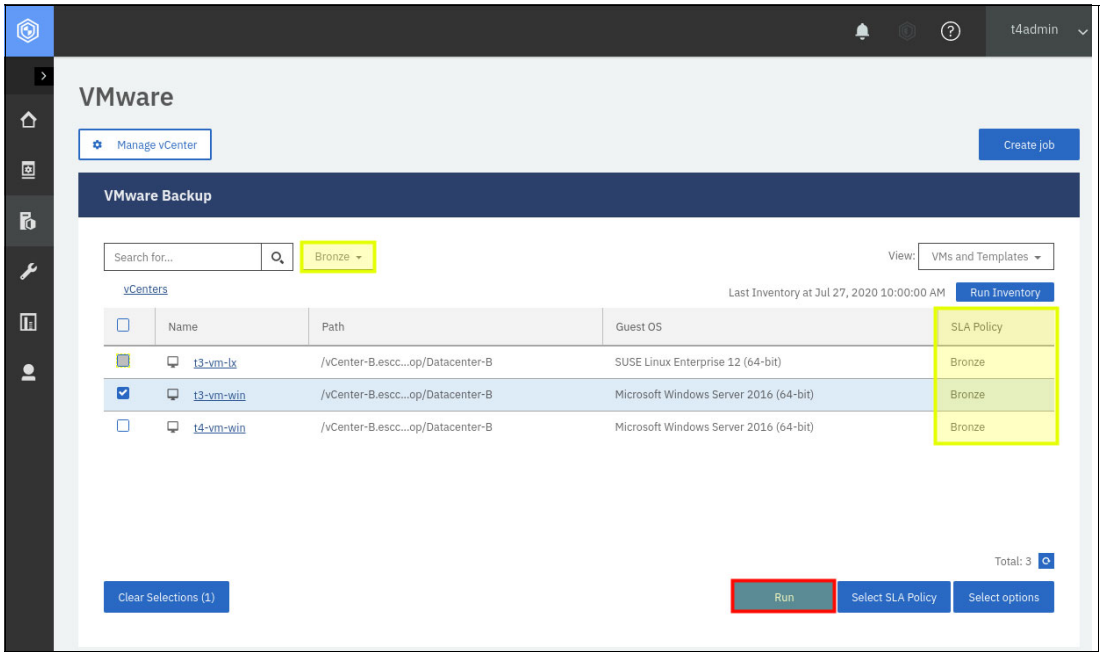

Figure 6-8 List all VMs backed up with the Bronze SLA to run an ad-hoc backup for t2-vm-win2

Now, select the `t3-vm-win` and click **Run** to start an ad-hoc backup for this VM. The backup process can be followed by scrolling to the SLA Policy Status and expanding the log view for the Bronze SLA.

Alternatively, you can select Create Job (see Figure 6-8), which starts the Job Wizard in which you can select **Ad hoc backup**, as shown in Figure 6-9.

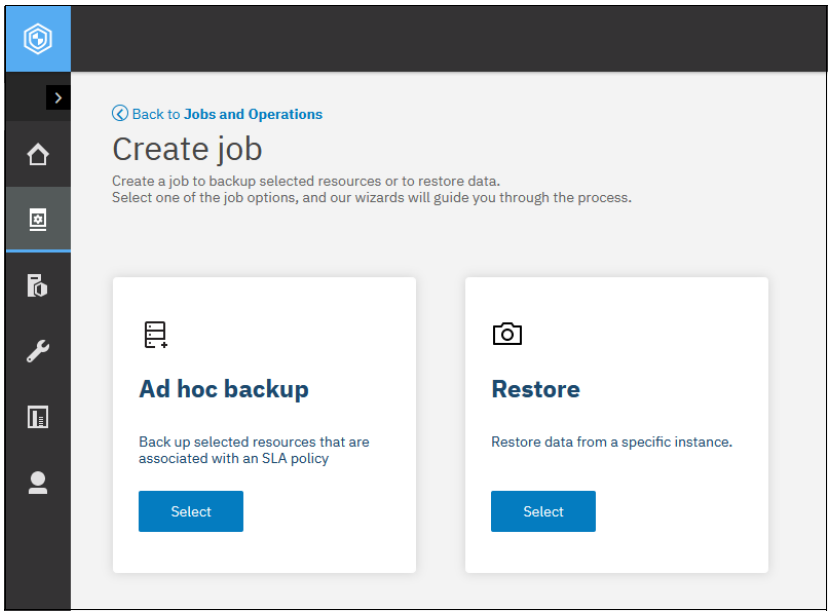

Figure 6-9 Start page of the Create job Wizard

The wizard guides you through the process of selecting the VM you want to back up and starting a backup job for those VMs. The wizard is also used to create restore jobs (for more information, see 6.3, “VM restore and data reuse” on page 189).

6.1.6 Distributing VM backups to multiple vSnap servers

When multiple vSnap servers exist for one site, it is beneficial to distribute the available VMs to be backed up by using one of the following methods in IBM Spectrum Protect Plus:

- Create at a minimum as many SLAs for the same site as the number of vSnap servers that are available in this site. Distribute the VMs over the available SLAs afterwards. Each SLA chooses one of the available vSnap servers as a backup target. Therefore, all vSnap servers are used and the VMs are distributed.
- In the IBM Spectrum Protect Plus configuration under **System Configuration** → **Global Preferences** Virtualized Systems section, you can configure how VMs are grouped. Grouping is done within an SLA. The default grouping is by size and the default size is configured to 5120 GB. This setting results in adding a specific number of VMs from an SLA to a group until their source capacity exceeds a total size of 5120 GB. Then, a new group is created within the same SLA. Each group is using another available vSnap server. The configuration can also be changed to VM counts instead of size. A group is created whenever the number of configured VMs reaches that count number.

In Example 6-1, we changed the Global Preferences setting to a VM count of two. The backup job log of the `vmware_Bronze` SLA policy includes the Example 6-1 lines, which indicate that the four VMs that are backed up are distributed over the two available vSnap servers.

Example 6-1 The vmware_Bronze backup log showing VM distribution over the available vSnap servers

```
Selected vm(s) count: 4
Using storage volume spp_1004_2002_16cbd7eb496__group1_53_ on controller t2-spp-vsnap2 for backup
Using storage volume spp_1004_2002_16c86921d3a__group0_70_ on controller t2-spp-vsnap for backup
Backing up VM (t2-vm-lx) from remote proxy (IP: 10.0.250.28, Host name: t2-spp-vsnap2)
Backing up VM (t2-vm-win2) from remote proxy (IP: 10.0.250.27, Host name: t2-spp-vsnap)
Backing up VM (t2-vm-win) from remote proxy (IP: 10.0.250.27, Host name: t2-spp-vsnap)
Backing up VM (t1-vm-lx) from remote proxy (IP: 10.0.250.28, Host name: t2-spp-vsnap2)
```

6.1.7 Backup options

IBM Spectrum Protect Plus offers backup options to be configured at an SLA or VM level. This section discusses both options and their settings.

Backup options at SLA level (backup job level)

The Backup Options at SLA Level can be accessed by clicking the **Policy Options** icon from the **SLA Policy Status** under **Manage Protection** → **Virtualized Systems** → **VMware or Hyper-V or Amazon EC2**. A configuration window appears in which the following options can be configured:

- Specify **Pre- and Post-Scripts,** which IBM Spectrum Protect Plus runs before or after running the backup job. To continue running the job if the script fails, select **Continue job/task on script error**. Scripts and script servers are configured and maintained by selecting **System Configuration** → **Script**.
- Select **Run inventory before backup** to first run an inventory job before a backup is started.
- Select **Exclude Resources** to exclude specific resources from the backup job by using single or multiple exclusion patterns. Resources can be excluded by specifying an exact match or by using wildcards. Multiple filters need to be separated by a semicolon.

- To force a new full backup for specific VMs or databases, select the **Force Base Resources** option to create a full backup for the specified resource with the next backup job execution. Multiple resources must be separated by a semicolon.

Backup options at VM level

The backup options at VM level are accessed from the same window as shown in Figure 6-3 on page 175. Search for specific VMs or select all VMs from a specific SLA for which you want to change the backup options. After the VMs are selected, choose **Select Options** and a new window appears in which the following settings can be determined:

- Skip Read-only datastores: Skips datastores that are mounted as read-only.
- Skip temporary datastores mounted for Instant Access: Skip datastores that are mounted by an Instant Access Restore.
- VADP proxy: Select the Proxy or a pool of Proxies (Site) to be used for backup.

Note: If you want to specify VMs to use specific proxies, use the VM option and specify By Proxy. It is best to let Spectrum Protect Plus decide rather than manually controlling the load balancing.

- Make VM snapshot application/file system consistent and VM Snapshot retry attempt (default 2): All VSS-compliant applications, such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL server, and the system state are quiesced.
- Fall back to unquiesced snapshot if quiesced snapshot fails Enable to fall back to a non-application or non-file-system consistent snapshot if the application consistent snapshot fails: Selecting this option ensures that an unquiesced snapshot is taken if environmental issues prohibit the capture of an application or file-system consistent snapshot.
- Truncate SQL logs: Truncates Microsoft SQL server logs during backup (see 11.3.3, “SQL database backups logs” on page 317).
- Catalog file metadata: Not selected by default, this option stores metadata about mounted file systems. The metadata is required for Single File Restore functionality.
- Exclude Files: Files that are excluded from the catalog file metadata.
- User: Guest Operating System user information. This user is also called IBM Spectrum Protect Plus agent user, which is required for the following operations:
  - File metadata cataloging (see 6.2, “Catalog file metadata for single file restore”).
  - Microsoft SQL server log truncation.
  - Specific restore operations, such as IP address reconfiguration during restore (see 6.3.10, “Restoring a VM and changing static IP address on one NIC” on page 204).

6.2 Catalog file metadata for single file restore

IBM Spectrum Protect Plus allows for single file restore out of the VM snapshot backups. This capability is helpful and efficient when only a single file or a couple of files must be restored instead of restoring the entire VM.

The file metadata is discovered when configured as part of the VM backup job locally on the VM, transferred to the IBM Spectrum Protect Plus, and stored there in a global catalog. This enables you to search for files and directories globally from the central IBM Spectrum Protect Plus GUI.

6.2.1 Configure requirements

To allow single file restore cataloging, the file metadata must be enabled as part of the backup job. To allow IBM Spectrum Protect Plus to discover the file metadata, some operating system-specific prerequisites must be fulfilled.

Note: For more information about requirements, see this web page.

In addition to the supported operating systems, file systems, software, and connectivity requirements to access a VM, a key requirement is the Authentication and Privilege Requirements.

For Windows, the credentials that are specified for the VM must include a user with the following privileges:

- The user must have “Log on as a service” rights. For more information about setting up these rights, see this web page.
- The user must have the permissions of the local administrator.

For Linux, the credentials that are specified for the VM must specify a user that has the following sudo privileges:

- The `sudoers` configuration must allow the user to run commands without a password.
- The `!requiretty` setting must be set.

The suggested approach is to create a dedicated IBM Spectrum Protect Plus agent user, which is used for the file metadata discovery. This user is centrally managed and configured in IBM Spectrum Protect Plus by creating an Identity, as described in “Identities, keys, and certificates” on page 11.

In a Windows environment, this agent user can be configured as a standard Domain User within Microsoft Active Directory and added by way of a Group Policy to the local administrator group, as shown in Figure 6-10 on page 183.

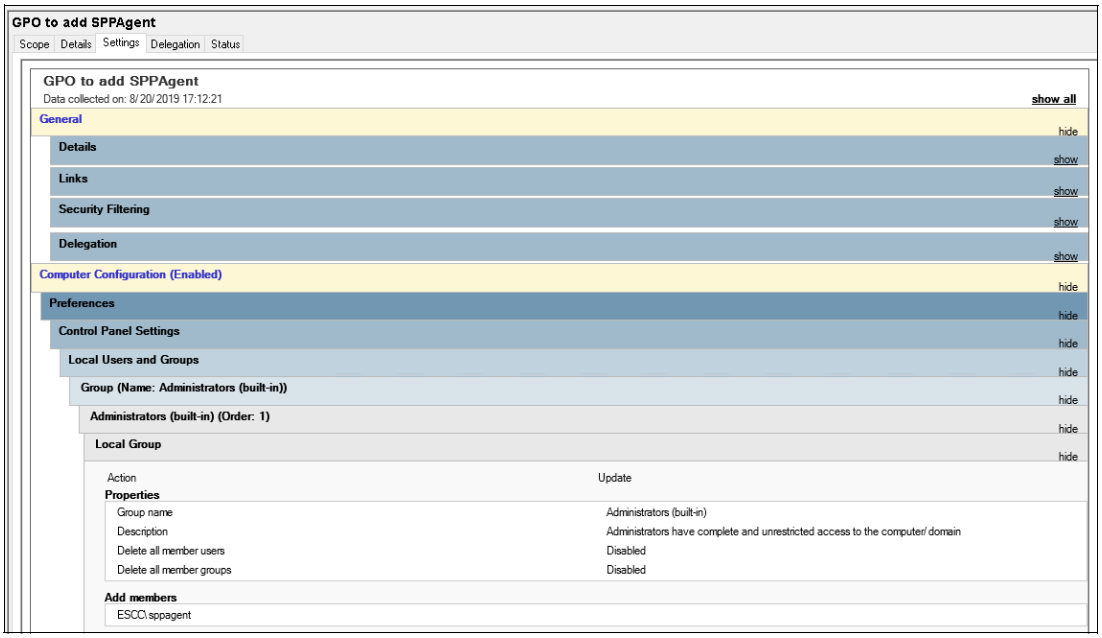

Figure 6-10 Group Policy Object (GPO) to add Domain User SPPAGENT to local Administrator Group

By adding the agent user to the local administrator group, the user automatically has the “Log on as a service” right assigned.

Note: For information about the Windows agent user, see IBM Knowledge Center.

If Linux VMs are configured with Microsoft Active Directory authentication, the same agent user account also can be used in Linux. Otherwise, a local agent user account with the privileges that are shown in Example 6-2, must be created.

Example 6-2 sudoers configuration for IBM Spectrum Protect Plus agent user

```
t2-vm-lx:~ # useradd -m sppagent
t2-vm-lx:~ # passwd sppagent
New password:
Retype new password:
passwd: password updated successfully
t2-vm-lx:~ # vi /etc/sudoers
Defaults: sppagent !requiretty
sppagent ALL=(root) NOPASSWD:ALL
```

Place the lines that are shown in Example 6-2 at the end of your `sudoers` configuration file (typically, `/etc/sudoers`). If your `sudoers` file is configured to import configurations from another directory (for example, `/etc/sudoers.d`), you can also place the lines in a new file in that directory.

6.2.2 Configuring file metadata discovery

The file metadata discovery is performed as part of the backup job and enabled as a backup option, which is described in “Backup options at VM level” on page 181.

The following example shows the configuration for the two Windows VMs that are added as shown in Figure 6-3 on page 175 to the Bronze Backup Policy.

Before we configure the backup options, we add the IBM Spectrum Protect Plus agent user as an Identity under **Accounts** → **Identity** to IBM Spectrum Protect Plus, as shown in Figure 6-11.

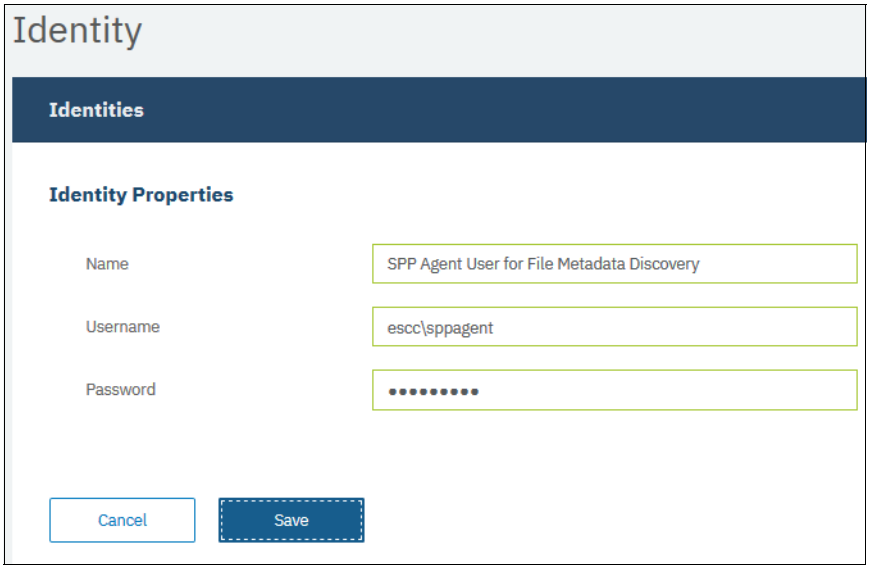

Figure 6-11 Add SPP agent User as Identity to IBM Spectrum Protect Plus

The backup options at the VM level can be accessed in the same window as where a VM is added to an SLA, as shown in the Figure 6-3 on page 175. Search for the two Windows VMs or select the Bronze SLA by using the search filter drop-down menu. Then, select the two VMs and click **Select Options**.

The window is expanded to configure the file metadata cataloging option (file discovery), as shown Figure 6-12.

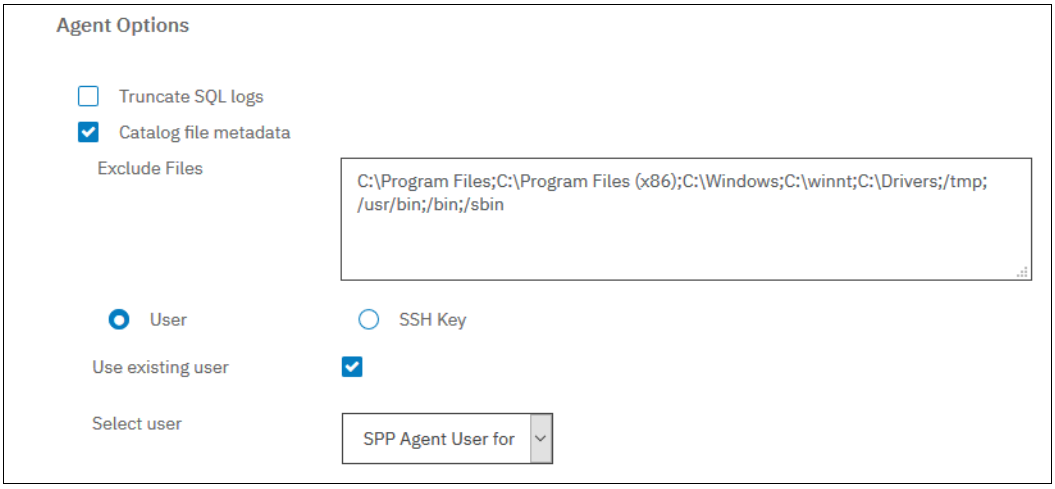

Figure 6-12 Enable Catalog file metadata

You must select **Catalog file metadata** to enable file metadata cataloging. In the Exclude Files section, you specify directories that can excluded from the catalog process. Select **Use existing user** to select the previously configured IBM Spectrum Protect Plus agent user as a cataloging user.

Otherwise, you can specify the credentials to be used in the username and password field. These credentials are then automatically added as an Identity. A third option is to select an SSH Key, which is configured in **System Configuration** → **Keys and Certificates** in the SSH Keys section by clicking **Add SSH Key**.

To store the configuration, click **Save**.

When the Bronze SLA backup `vmware_Bronze` is run again, the log file under Jobs and Operations in the tab Job History includes entries about performing file discovery, as shown in Figure 6-13.

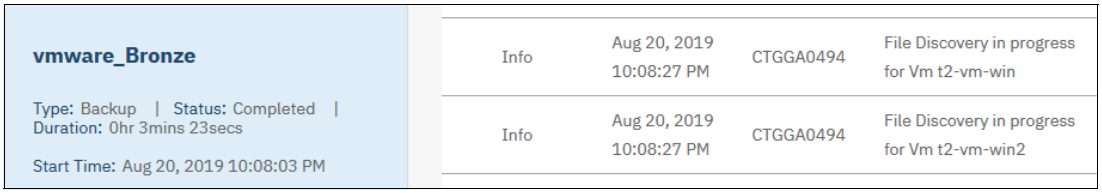

Figure 6-13 File Discovery log entries in the vmware_Bronze job log

The discovered file metadata information is stored temporarily on the client and permanently on the IBM Spectrum Protect Plus server. Metadata collection starts on the client after a successful snapshot and a `*.txt` file is created for each drive or volume. All the `*.txt` files are compressed and sent to the server. In Windows, the data is stored temporarily in the `c:\ProgramData\SPP\temp\output\` directory and in Linux in the `/tmp` directory.

The metadata is deleted from the client after it is sent to the IBM Spectrum Protect Plus server. When the `.zip` files reach the server repository on `/data2/filecatalog/` (the default location that is configured under **System Configuration** → **Global Preferences** → **Protection**), the Lucene indexing starts and runs in the background, and the file metadata is stored as an index process in `/data3/lucene`. After a backup image expires, the file index that is associated with that image also is deleted.

Configuring test file metadata

IBM Spectrum Protect Plus offers a helpful test function to verify the file metadata cataloging after configuration. If you select only a single VM to be configured, a Test button appears next to the Save button, as shown in Figure 6-14.

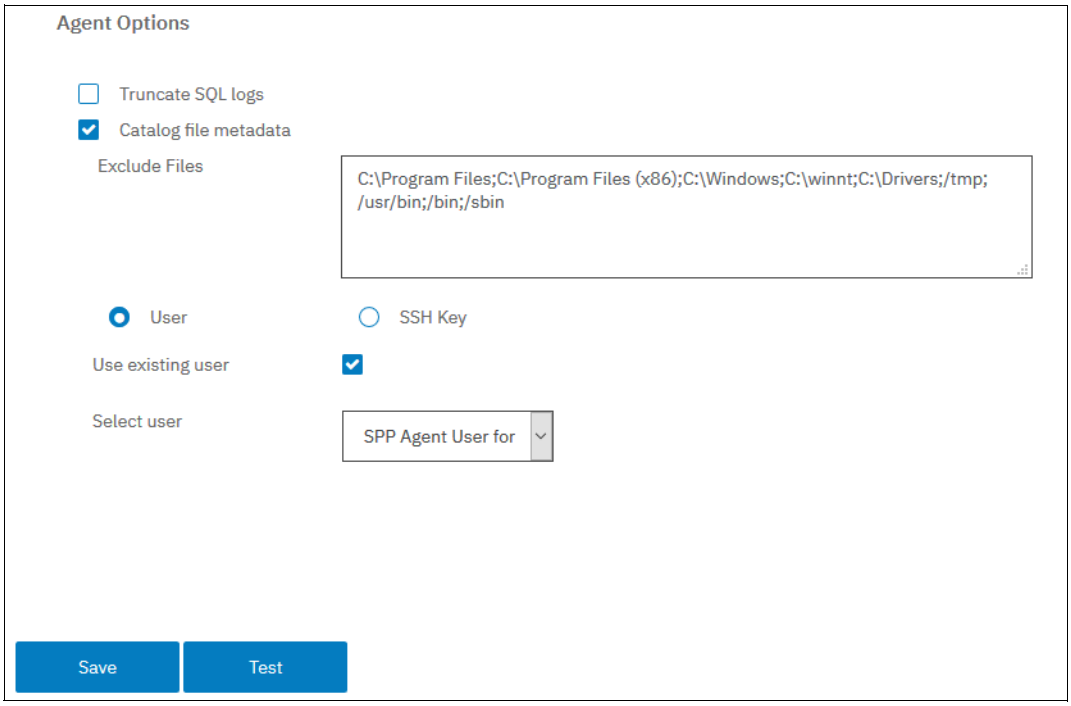

Figure 6-14 Test Button to verify file metadata configuration

If you click **Test**, IBM Spectrum Protect Plus checks if all requirements are fulfilled. You can follow this checking process in the status window that opens automatically. The tests are divided into the following steps:

1. Virtual: Basic Hypervisor test for Hypervisor Tools and DNS configuration
2. Remote: Remote executor test for session creation and remote agent deployment
3. Operating System pre-requisites:
   - Basic Windows prerequisites for file and volume operations
   - Basic Linux prerequisites for file and volume operations

If one or more of the tests are unsuccessful, these errors must be resolved first before the file metadata cataloging can occur. The test can be repeated as often as required to verify that all requirements are fulfilled.

6.2.3 Restoring single files and directories

A restore of single files and directories is possible after the file metadata cataloging is configured and the file discovery process during the VM backup job is completed successfully.

To search for files and directories, select **Manage Protection** → **File Restore**. A new view opens. A search form is displayed in the first part of the File Restore window. From here, you can complete the following actions:

1. Search for files.
2. Select files for restore.
3. Start the restore process.
4. Monitor the restore process.

You can complete these steps as one process flow without accessing any other menu or view. These steps are described next.

Step 1: Searching for files

In our example, we search for a file that is named `backup text test file.txt` in the VM `t2-vm-win2`. Therefore, we use the search pattern `*test*.txt` in the VM `t2-vm-win2` within folder `e:\*test*`, which results in two files found, as shown in Figure 6-15.

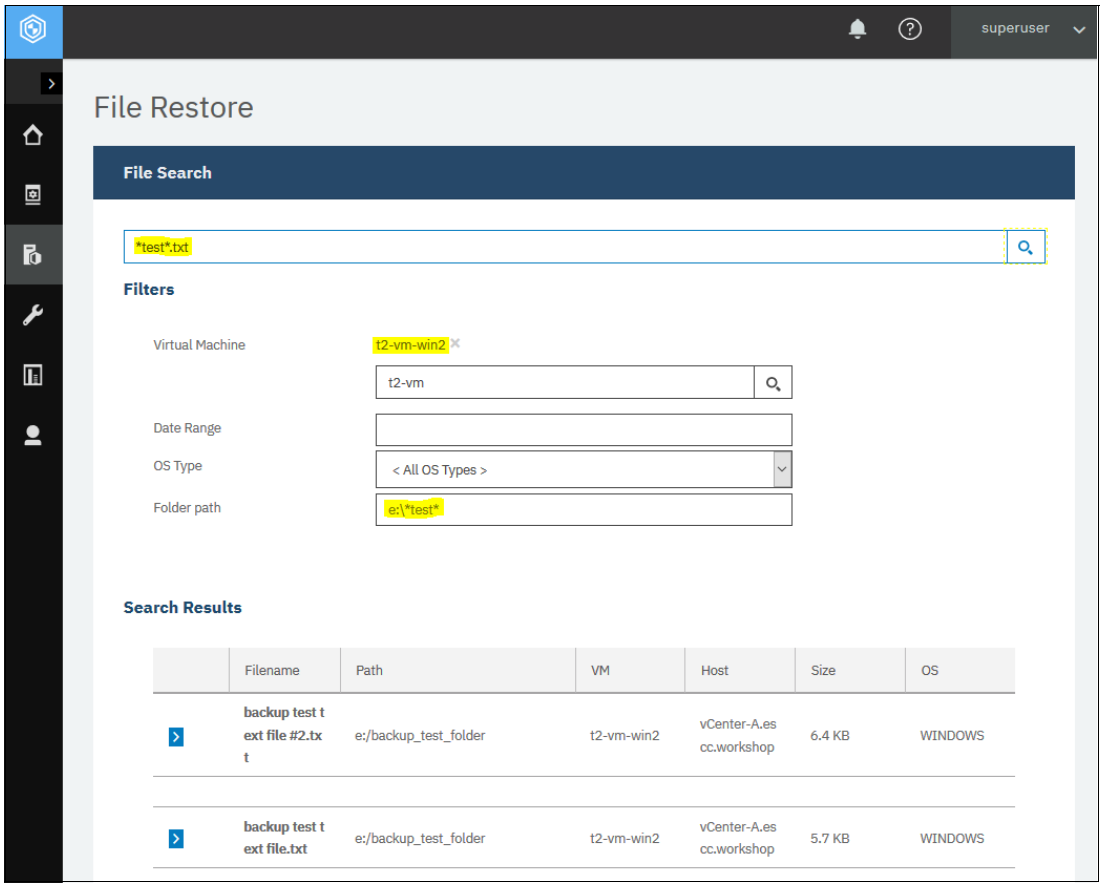

Figure 6-15 File Restore view with File Search

In addition to the search pattern for file, folder names, and specific VMs, a date range and operating system type can be specified for a more granular search.

Step 2: Selecting files for restore

From the search results (see Figure 6-15), the results can be expanded to see the available restore points for the files and the suitable version can be selected for restoring, as shown in Figure 6-16 on page 188.

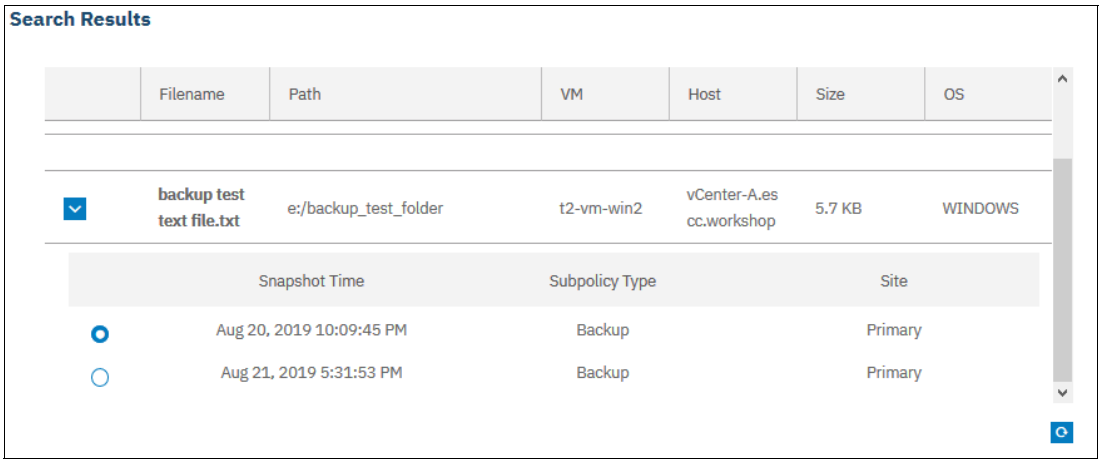

Figure 6-16 File Restore Search Results

After the wanted file for the restore is selected, scroll down to specify the restore options. In our example (see Figure 6-17), we restore the file to another VM called `t2-vm-win` into the folder `c:\restore`.

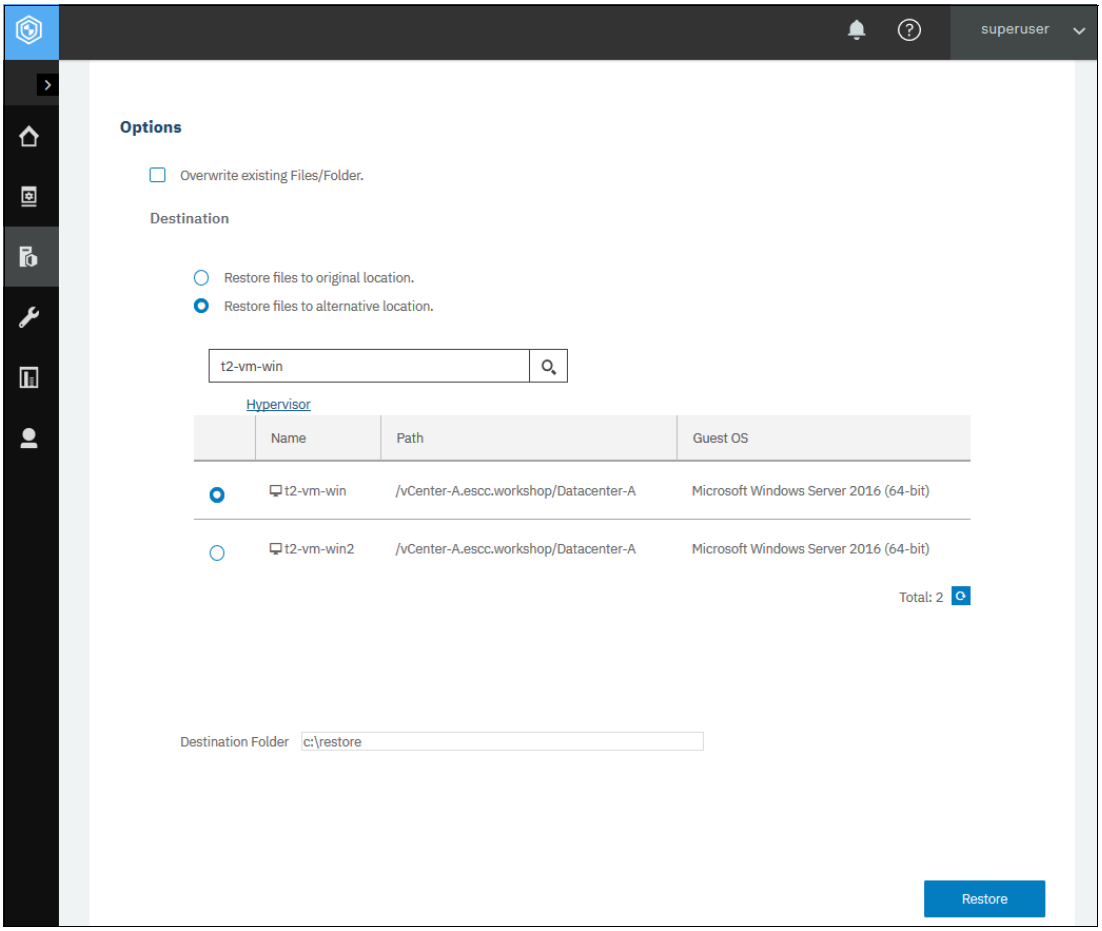

Figure 6-17 File Restore Options

In addition to the settings we use in our example, the file can be restored to the original location with the option to overwrite an existing version.

Step 3: Starting the restore

The restore is started by clicking **Restore** (see Figure 6-17 on page 188).

Step 4: Monitoring the restore process

After the restore job is started, it can be monitored in the same view by scrolling down to the Job Sessions section and expanding the log, as shown in Figure 6-18.

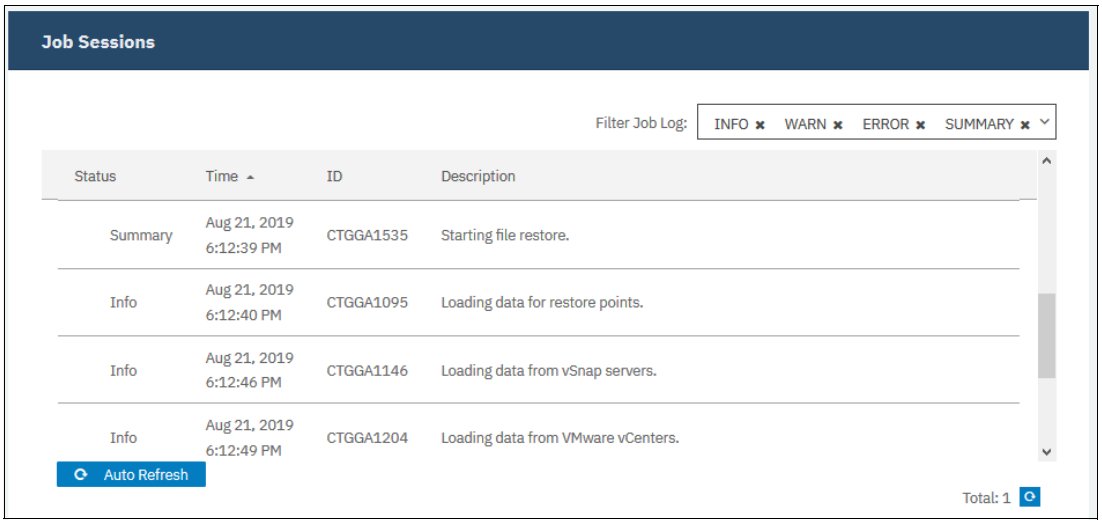

Figure 6-18 Monitor the file restore job

6.3 VM restore and data reuse

In this section, we explain the function of restoring VDisks (VMDK or VHDX) or complete VMs for recovery purposes or data reuse cases. Multiple VMs or VDisks can be part of one restore activity.

IBM Spectrum Protect Plus includes a restore wizard that simplifies restoring virtualized systems, databases, file systems, cloud-managed applications, and containers. The wizard guides you through the configuration of restore types and parameters, and optionally schedules a job that performs the restore.

During a restore, you can change the following parameters of a VM by using the wizard:

- Storage location where the VM is located (for example, use a different datastore)
- Compute resource where the VM is running on (for example, use a different ESX server)
- Display name of the VM
- VM network that the VM use (for example, restore the VM into an isolated VM network to not interfere with a still running source version of the VM)
- MAC address of a network adapter when a Test- or Clone-Mode Restore is used
- IP configuration (for example, set a different static IP or change from a static IP to DHCP)

Changing the IP configuration for VM requires a IBM Spectrum Protect Plus agent user for logging in to the VM and performing the change. For more information about setting up an IBM Spectrum Protect Plus agent user, see “Backup options at VM level” on page 181, and 6.2, “Catalog file metadata for single file restore” on page 182.

For changing the IP configuration in the case of a Linux guest operating system, the IBM Spectrum Protect Plus agent user must be root, which is not the case for file-metadata recovery.

To change the IP configuration of a Windows guest operating system, ensure that the Microsoft Visual C++ 2008 Service Pack 1 Redistributable Package MFC Security Update is installed before backing up the VMs. This update is necessary because it ensures that the packages are available if a restore with a required IP reconfiguration is needed. You can download VC 2008 SP1 Redistributable (9.0.30729.6161) at this web page.

Next, we described the restore wizard and the options that are available during a restore. Beginning with 6.3.7, “Restoring a VDisk” on page 196) we describe different restore examples to demonstrate the restore capabilities.

6.3.1 Restore wizard

The wizard can be accessed from the Jobs and Operations view by selecting **Create job** in the right corner (as shown in Figure 6-19) or from the different views under the Manage Protection menu.

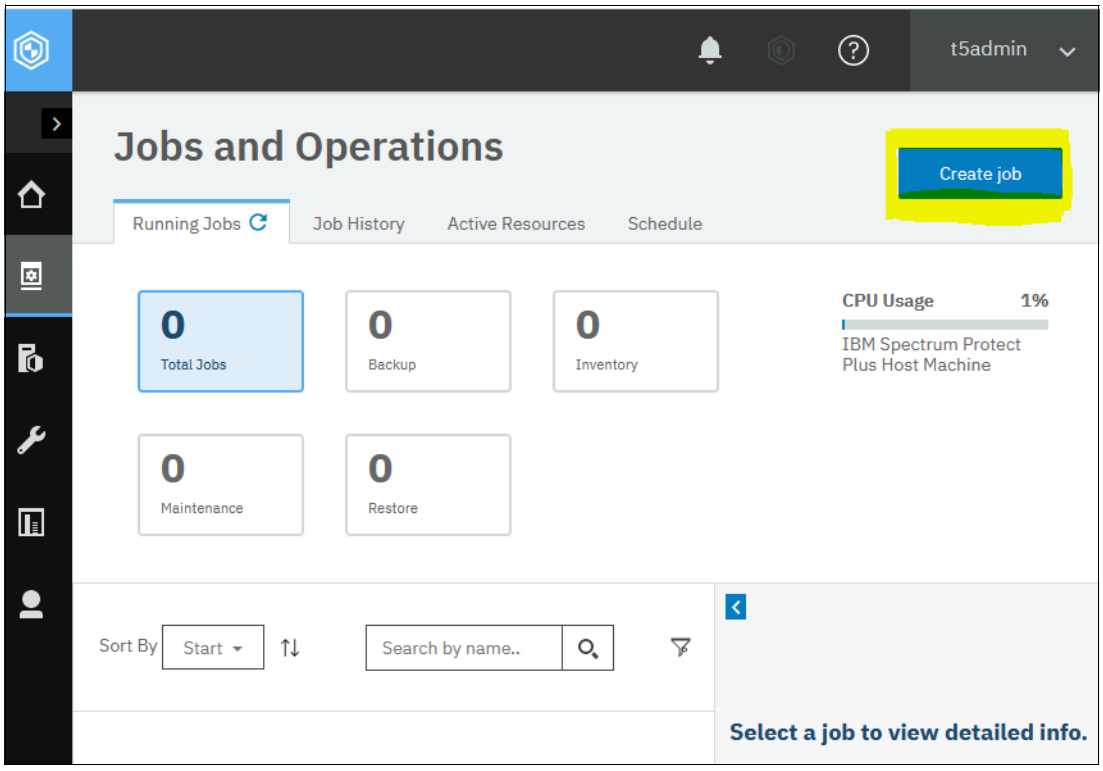

Figure 6-19 Create job button in the Jobs and Operations view

When the job wizard is started, you must decide if you want to perform an Ad hoc backup or a Restore, as shown in Figure 6-20 on page 191. For more information about ad hoc backups, see **6.1.5, “Running a backup for a single VM” on page 178**.

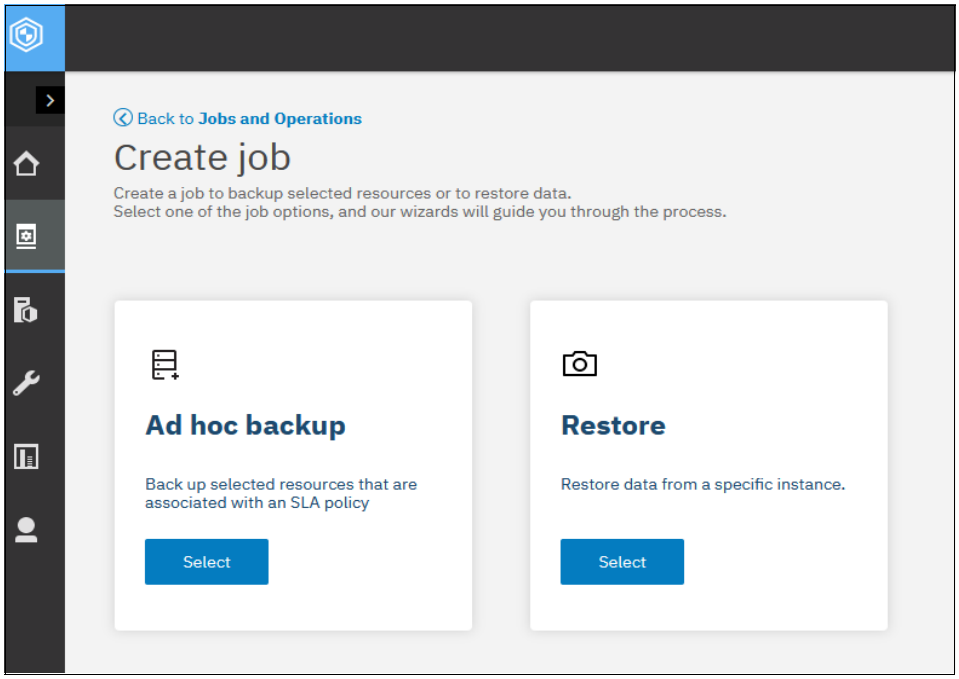

Figure 6-20

After you select **Restore**, select the Source type a restore that is to be performed, as shown in Figure 6-21.

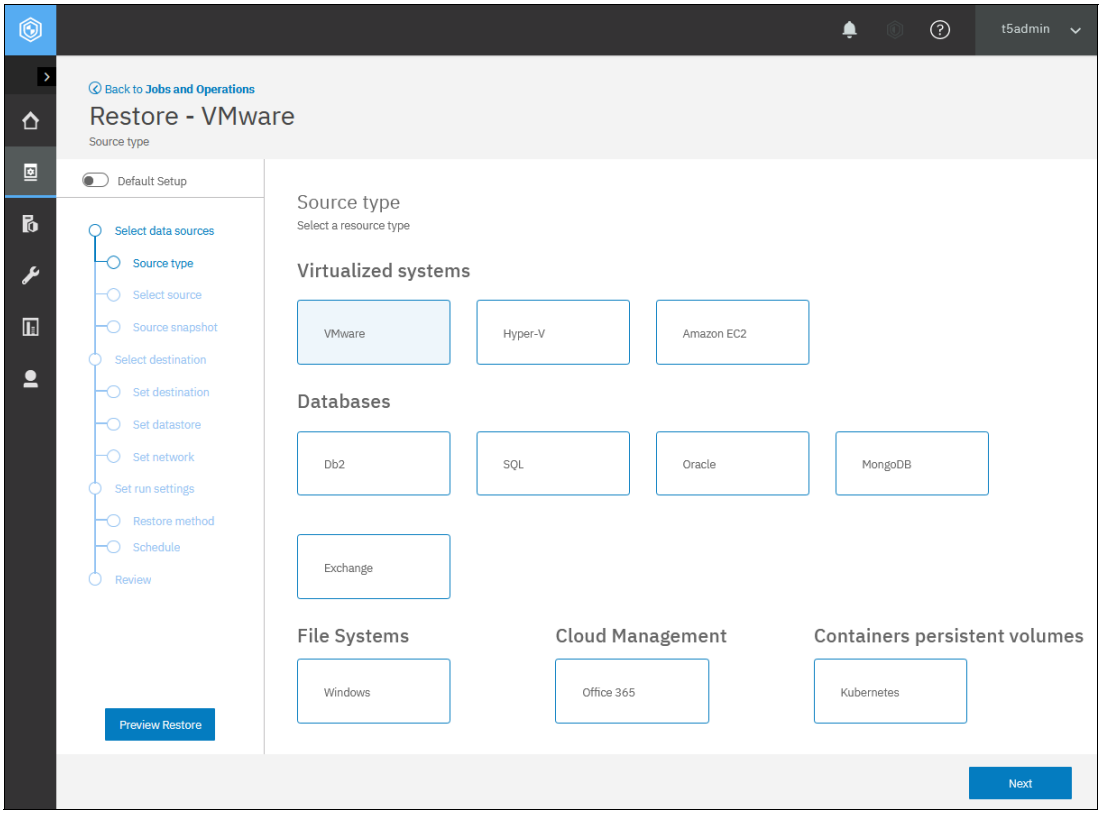

Figure 6-21 Select the source type for a restore in the Restore Wizard

If you open the wizard (for example from the **Virtualized Systems** → **VMware** specific view), you skipped the first step in which you must decide for which Source type (Virtualized systems, Databases, File Systems, Cloud Management, Containers persistent volumes) you want to create a restore job.

In our example, we show a restore for VMware. After the VMware is selected, the VMs that we want to restore also must be selected, as shown in Figure 6-22.

Complete the following steps (the number in Figure 6-22 correspond to the steps):

1. Use the Search and View Filter to find the VM that must be restored.
2. Select the VM by clicking the **+** sign that is next to the VM name.
3. The VM appears in the selected item view.

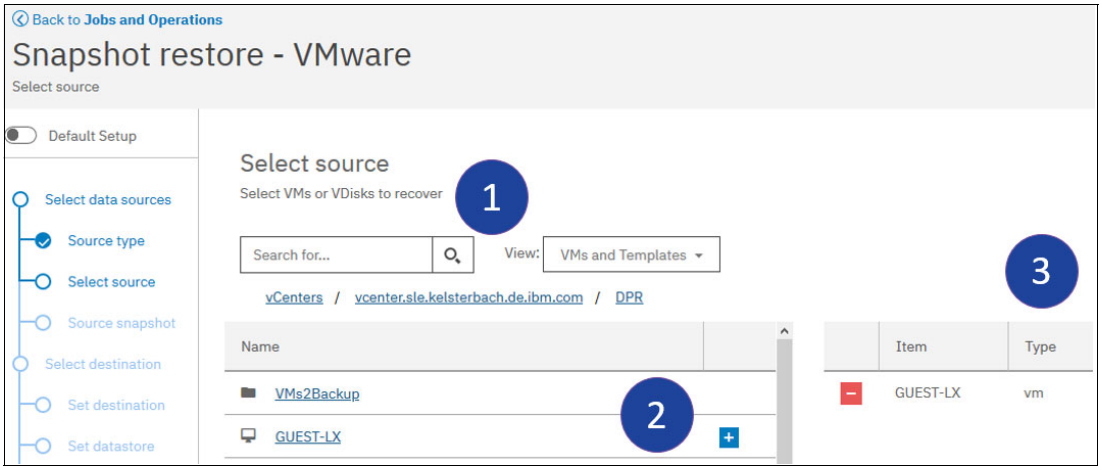

Figure 6-22 Select VMs for a restore job

If you select VMDK (VMware) or VHDX (Hyper-V) VDisks, the process flow is similar as shown in Figure 6-23. Complete the following steps (the number in Figure 6-23 correspond to the steps):

1. Search for the VM and then click the VM name.
2. Select the disk to be restored by clicking the **+** sign.
3. The VDisk appears in the selected item view.

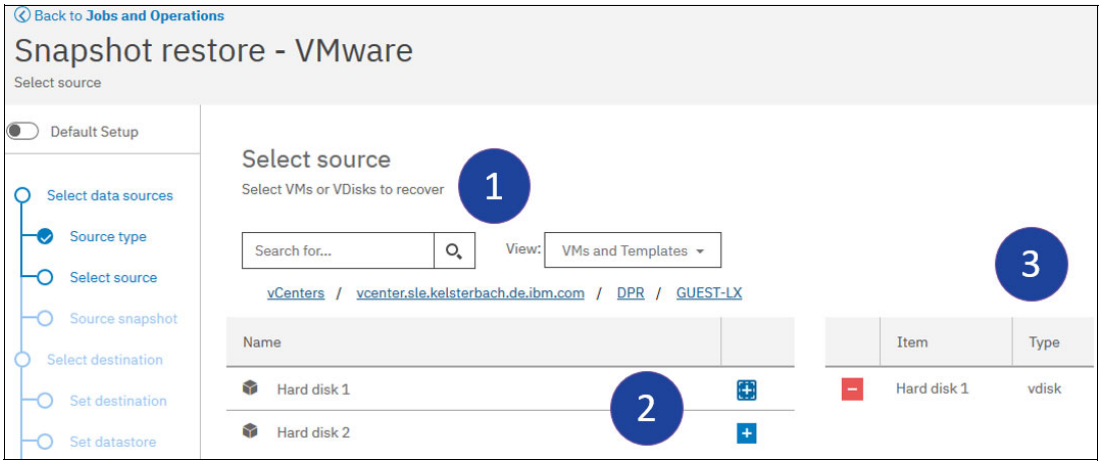

Figure 6-23 Select vDisks for a restore job

6.3.2 Restore schedule

After you select the VMs and vDisks for the restore job, you can create an On-Demand (ad-hoc) restore job by, which is run immediately after the restore wizard is completed. Alternatively, you can create a Recurring restore for which you can define a schedule and name for the job.

Figure 6-24 shows as an example of the step in the wizard in which you decide if you want to perform an On-Demand (as selected in the example) or a Recurring restore job by using the pull-down menu.

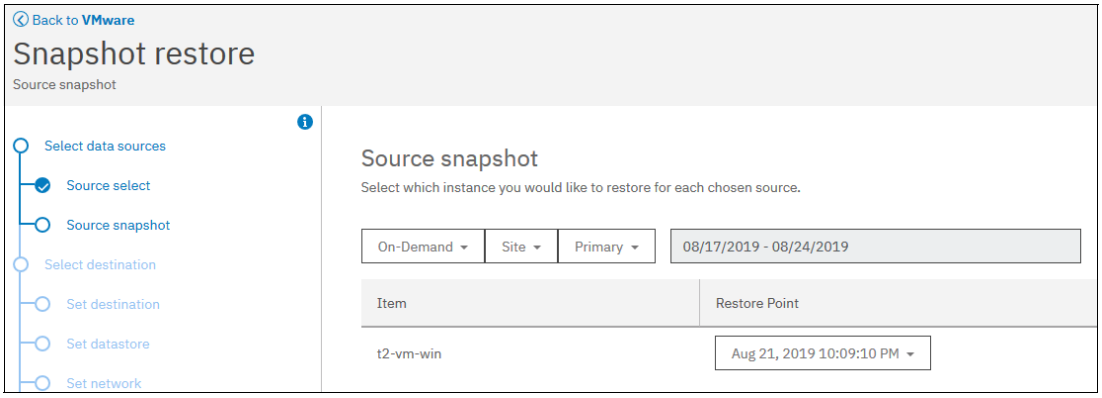

Figure 6-24 Restore wizard for VMware - select the Source Snapshot

Note: Only the schedule of the recurring restore job can be changed later on, as explained in “Scheduler (Jobs and Operations)” on page 9. The content of the job is fixed. If the content must be changed, a restore job must be created.

6.3.3 Restore Source (Location) and Restore Points

Typically, the primary restore source is the one that was used to store the backup initially (for example, the vSnap server from the primary site). In IBM Spectrum Protect Plus, this site is specified in the SLA. As described in Chapter 15, “Replication and additional copies” on page 405, IBM Spectrum Protect Plus also includes different ways to store or replicate backup data in more destinations for protection.

Whatever other destination is chosen to protect the backup data, this destination also is known to IBM Spectrum Protect Plus and can be used as a restore source. The different restore sources can be chosen within the restore wizard, as shown in Figure 6-24 where for example, the Restore Location type Site and the Location Primary are selected by using the suitable pull-down menu.

If you are restoring data from a cloud resource or repository server, select **Use alternate vSnap server for the restore job** to specify an alternative vSnap server. Then, select a server from the Select alternate vSnap server menu. When you restore data from a restore point that was off-loaded or archived to a cloud resource or repository server, a vSnap server is used as a gateway to complete the operation.

By default, the vSnap server that is used to complete the restore operation is the same vSnap server that is used to complete the copy to the cloud resource operations. To reduce the load on the vSnap server, you can select an alternative vSnap server to serve as the gateway.

Based on restore source and restore item that is chosen, IBM Spectrum Protect Plus lists the available Restore Points (with time stamp) from which the best option for the restore operation can be selected. In the example that is shown in Figure 6-24 on page 193, the VM Item `t2-vm-win` is selected with a Restore Point from Aug 21, 2019 10:09:10 PM.

Because each backup destination can have its own retention settings (dissimilar backup policies between primary backup target and other copies), different restore scenarios are supported. For example, recovery of the latest version from a VM is taken from the primary vSnap server, whereas the restore of a VM from a version before the previous operating system patch rollout is taken from the Cloud Offload.

If you define a Recurring restore job, as explained in 6.3.2, “Restore schedule”, IBM Spectrum Protect always uses the latest backup version from the source location that you configured for this job.

6.3.4 Restore destination

In addition to the default option to restore to the Original Host or Cluster, an Alternate Host or Cluster can also be chosen. Selecting the original ESX host implies that the original datastore and network are used for the restore. If you want to choose a different datastore or network on the original ESX host or cluster, select the **Alternate Host or Cluster** option.

Note: With the Alternate Host or Cluster option, you can change the ESX server, datastore, and network that is used by the VM.

For VMware, a third option called ESX host if vCenter is available. In other restore scenarios, actions are completed through the VMware vCenter. If vCenter is unavailable, this option restores the vCenter VMs that the vCenter is dependent upon.

For more information about backing up of a VMware vCenter VM, see IBM Knowledge Center.

6.3.5 Restore methods

IBM Spectrum Protect Plus offers the following restore methods:

- Test mode
- Clone mode
- Production mode
- Instant Access for VDisk restore

Each of the methods are explained next.

Test mode

Test mode creates temporary VMs for development or testing, snapshot verification, and Disaster Recovery (DR) verification on a scheduled, repeatable basis without affecting production environments. The temporary VMs accesses storage from a vSnap server, which is mounted as temporary datastore by way of NFS to the Hypervisor.

Test machines are kept running as long as needed to complete testing and verification and are then cleaned up, moved to production, or cloned. VMs that are created in test mode are also given unique names and identifiers (UUIDs) and MAC addresses to avoid conflicts within your production environment. A test mode restore is also called *Instant Access Restore*.

Clone mode

Clone mode creates copies of VMs for use cases that require permanent or long-running copies for data mining or duplication of a test environment in a fenced network. The VMs access storage from the vSnap server and start copying data to the datastore.

VMs created in clone mode are also given UUIDs and MAC address to avoid conflicts within your production environment. With clone mode, you must be sensitive to resource consumption because clone mode creates permanent or long-term VMs. A clone mode restore is also called an *Instant Recover*.

Production mode

Production mode enables DR of VMs at the local site from primary storage or a remote DR site, replacing original machine images with recovery images. They access storage from the vSnap server and start copying data to the datastore. All configurations are carried over as part of the recovery, including names and UUIDs, and all copy data jobs that are associated with the VM continue to run. A production mode restore is also called Instant Recover.

Instant access for instant disk restore

If only a VDisk (VMDK or VHDX) is selected for a restore operation, IBM Spectrum Protect Plus automatically presents the option for an Instant Disk restore job, which provides instant writable access to data and application restore points. An IBM Spectrum Protect Plus snapshot is mapped from the vSnap server to a target server where it can be accessed or copied as required.

During the restore job creation process, you can move the VDisk to permanent storage and clean up temporary resources from the vSnap server. This process can be done by selecting the **Make IA clone resource permanent** restore option in the restore wizard. Alternatively, you can choose later on from the Jobs and Operations view in the Active Resources tab to make the disk permanent or to End Instant Disk Restore (Cleanup). For more information about restoring a VDisk, see 6.3.7, “Restoring a VDisk” on page 196.

6.3.6 Restore Options

By using the restore wizard, the following options are available to edit the names of the restored VMs:

- Rename the VMs (you can specify an individual name for each VM)
- Append a suffix to all VMs
- Prepend a prefix to all VMs

If you decide to rename the VM, it must be done for each VM that is part of the restore job. However, appending a suffix to the name is done automatically for each VM in the job.

Note: With the rename option, you can change the display name of the VM, but not the host name. If you must change the host name, you can use the Test Mode Restore (for example); then, adjust the host name and decide if the VM should be cloned or moved to production.

In addition, the following options are available that can be selected in the restore wizard:

- Power on after recovery
- Overwrite VM
- Continue with restore even if it fails (default)
- Allow to overwrite and force clean up of pending old session (default)

- Restore VM tags (default)
- Pre-Script
- Post-Script
- Continue job/task on script error

6.3.7 Restoring a VDisk

In this section, we describe how to restore a VDisk on the original host. In our scenario, we restore Hard disk 2 from `t3-vm-win` (Windows VM) and Hard disk 1 from `t3-vm-lx` (Linux VM), as shown in Figure 6-25.

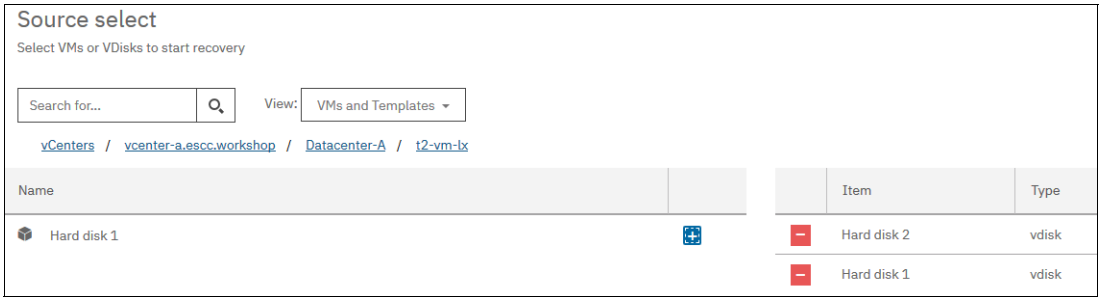

Figure 6-25 Select VDisks for a restore within the restore wizard

We select an **On-Demand** restore from the Primary Site (vSnap server) by using different Restore Points, as shown in Figure 6-26.

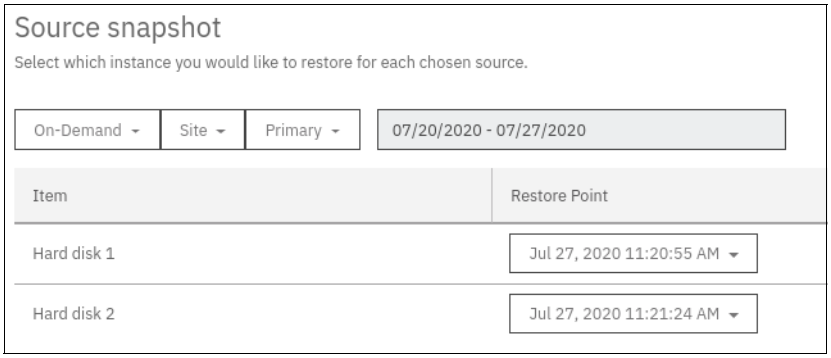

Figure 6-26 Select restore source snapshot for the VDisk restore

The restore should be mapped to the original VMs; therefore, we select **Original Host or Cluster**. We do not make the Instant Access (IA) clone resource permanent yet, so we use temporary resources from the vSnap server.

If you must restore vDisks to an alternative VM, select **Alternate Host or Cluster**, which after the host selection gives you the **Destination Virtual Disk** option to specify the VM to be used and other disk options, as shown in Figure 6-27.

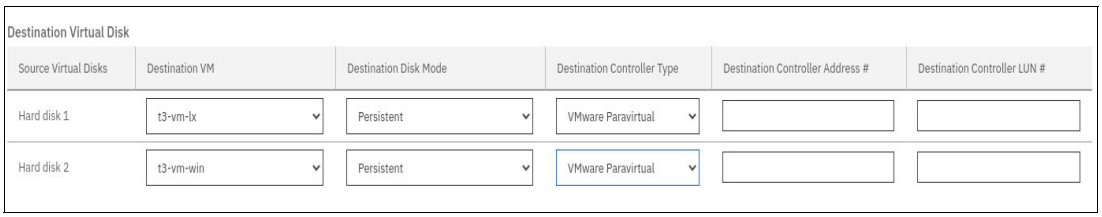

Figure 6-27 Choose an alternate VM for the VDisk restore in the Destination VDisk section

Before the restore job for the vDisks is submitted, you can review your settings, as shown in Figure 6-28.

Figure 6-28 Review the restore settings for the VDisk restore

After the restore job is submitted, you can follow the progress in the Jobs and Operations view in the Running Jobs tab. If mounting the vDisks to the VMs was successful, both disks are listed in the Active Resources tab under Hypervisor in the Jobs and Operations view, as shown in Figure 6-29 on page 198.

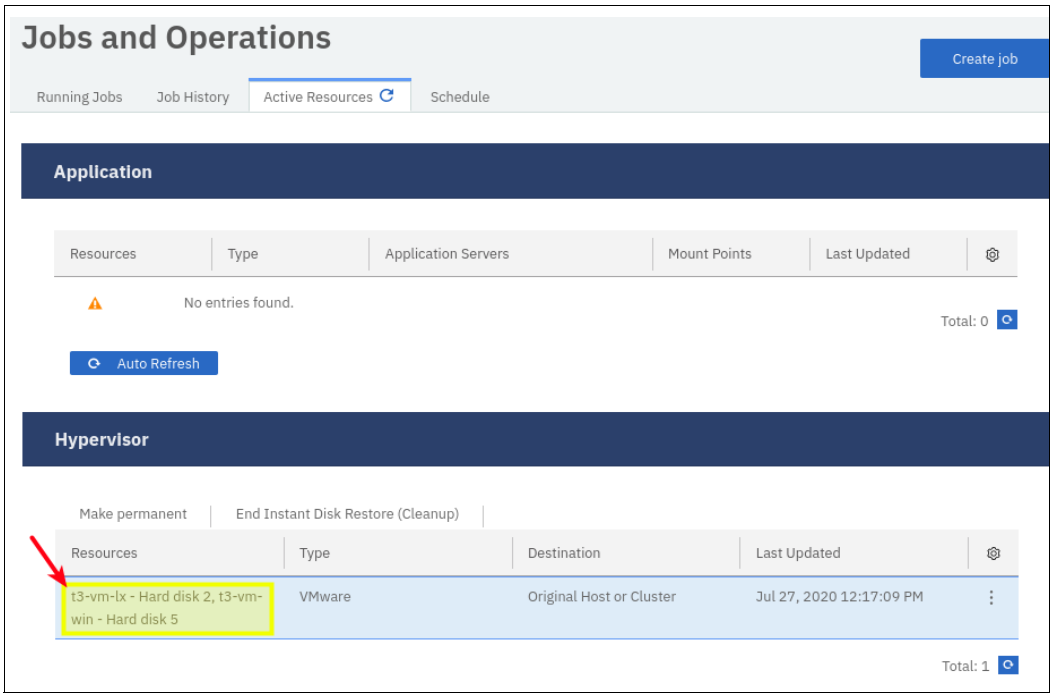

Figure 6-29 Active VDisk resources in the Jobs and Operations view

To use the vDisk in Windows, it must be set to **Online**, as shown in Figure 6-30, by using the Computer Management interface in Windows. This interface is accessible by right-clicking the Windows **Start** button and choosing **Computer Management**. After setting the disk to online, the disk is assigned a drive letter and can be accessed by using Windows Explorer.

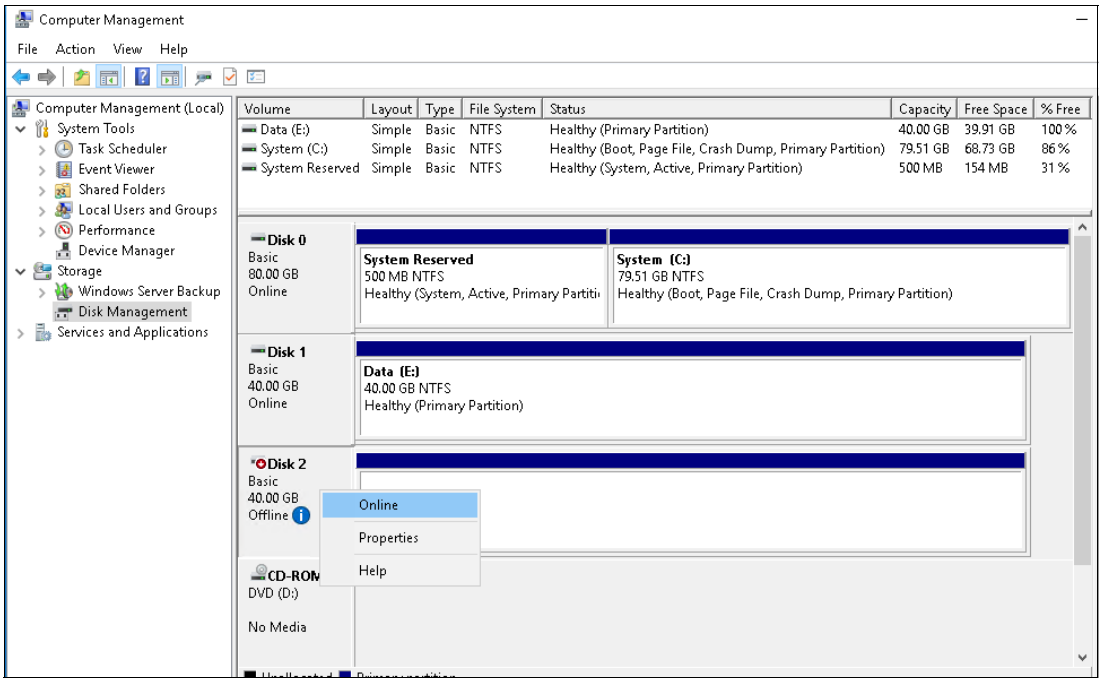

Figure 6-30 Set VDisk Online in the Windows Computer Management interface

After the tasks on the vDisk are completed, we clean up the restore mount, as shown in Figure 6-31.

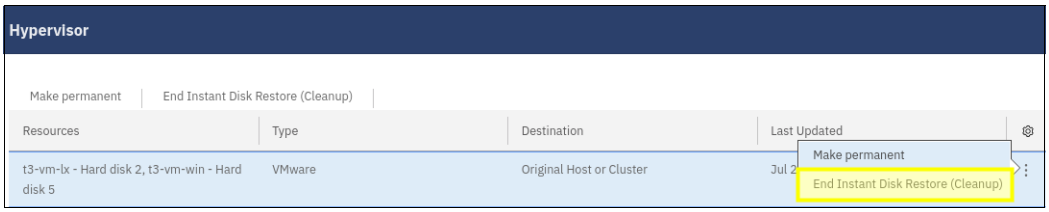

Figure 6-31 Cleanup VDisk restore mount

6.3.8 Restoring a VM from primary site

In this example, we restore a complete VM to a previous state from primary site. We use the following options:

- Change the datastore during restore
- Restore the VM into a different network segment
- Use a Test Mode restore and create a clone later

In our example, we restore the VM `t2-vm-win` to the restore point that was created on 17th of August, as shown in Figure 6-32.

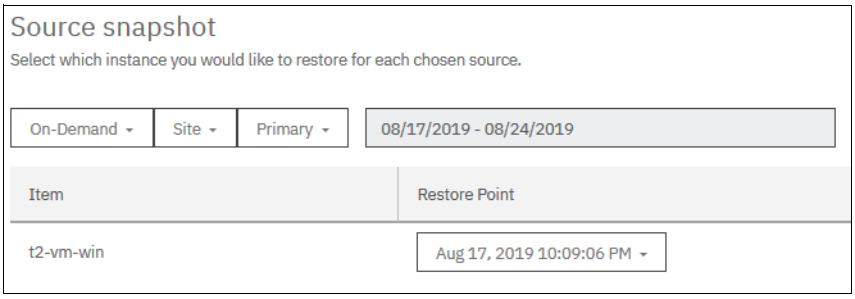

Figure 6-32 Select the source snapshot to restore the VM

We use the **Alternate Host or Cluster** selection to select the same host as in our example that was described in 6.3.7, “Restoring a VDisk”. However, we use a different datastore as the restore destination, as shown in Figure 6-33.

Tip: You must use the alternative host or cluster, even if you intend to restore to the same host or cluster because this option is the only option to get the advanced network settings, as shown in Figure 6-34 on page 200.

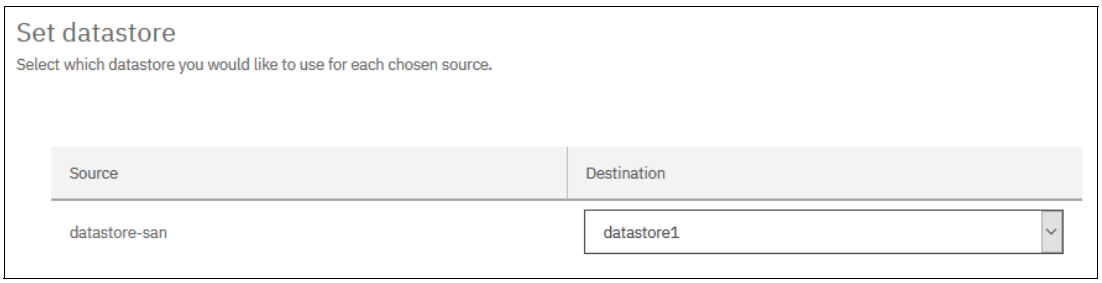

Figure 6-33 Select a different datastore as the restore destination

For the network, we select a different, isolated network but keep the IP configuration to not interfere with the original VM, as shown in Figure 6-34.

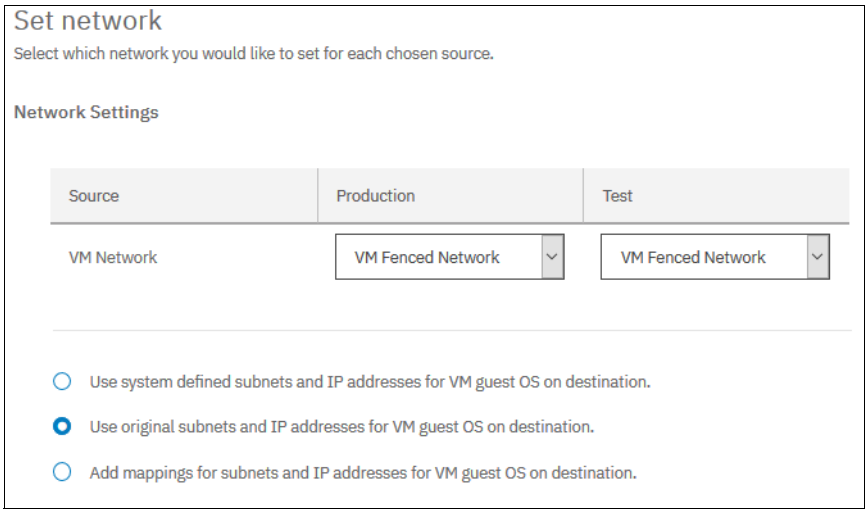

Figure 6-34 Network settings for VM restore

Our restore plan is to check the VM after the test restore and, if everything worked well, create a clone of the VM. In the review pane that is shown in Figure 6-35, you can verify all settings again before submitting the restore job.

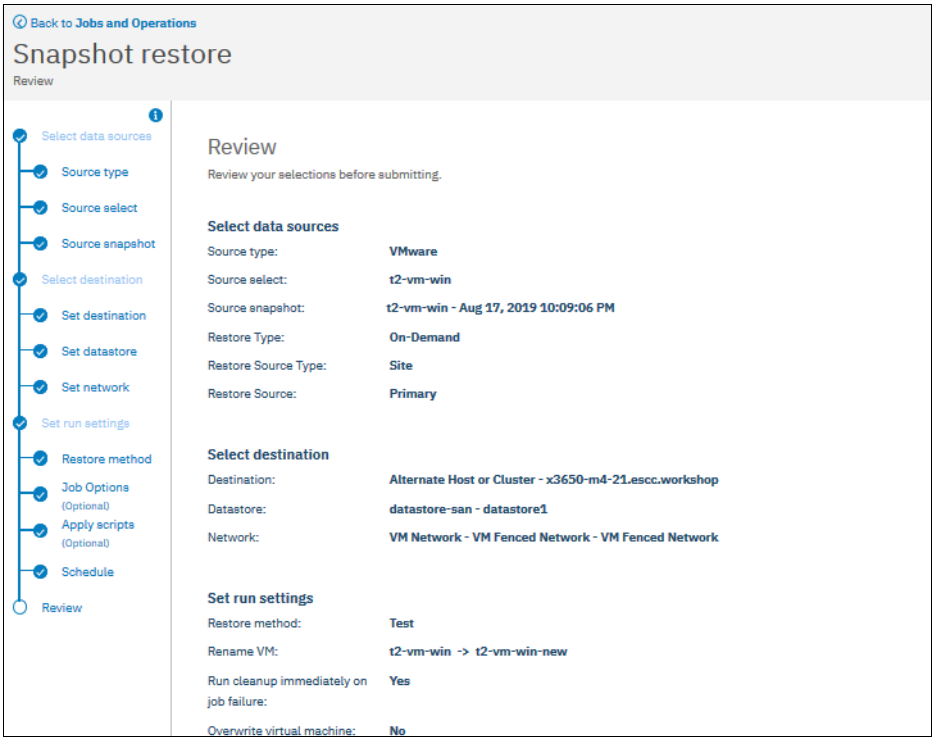

Figure 6-35 Review pane to verify the settings for the restore job

After the restore job is submitted, you can follow the progress in the Jobs and Operations view in the Running Jobs tab. When the test restore completes mounting the restore point from the vSnap server and creating the VM, the VM is listed in the Active Resources tab under Hypervisor in the Jobs and Operations view.

In the vCenter User Interface (UI), we see the VM that was created by using the isolated network and temporary storage resources from the vSnap server, as shown in Figure 6-36.

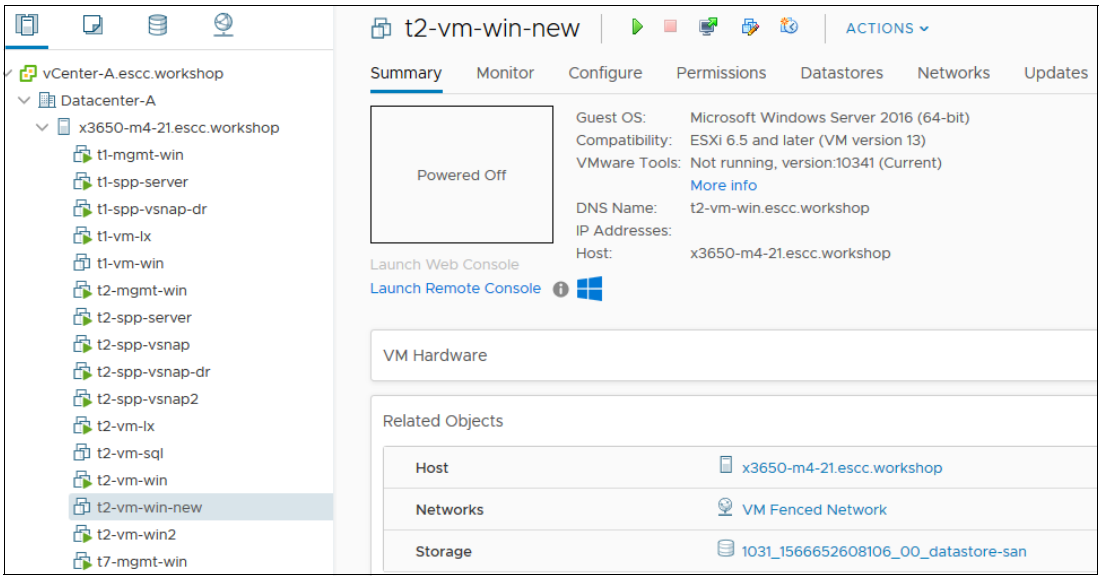

Figure 6-36 Restored VM is created using the fenced network and storage from the vSnap server

The VM can now be started. After starting the Windows operating system, we log in and verify whether the VM was configured as expected. If everything is OK, we go to the Active Resources tab in the Jobs and Operations view and start the Clone operation for this VM, as shown in Figure 6-37. The clone operation starts a VMware vMotion process, which moves the data from the temporary storage of the vSnap server to the final datastore specified in the restore job. While the clone operation is running, you can continue to work with the VM.

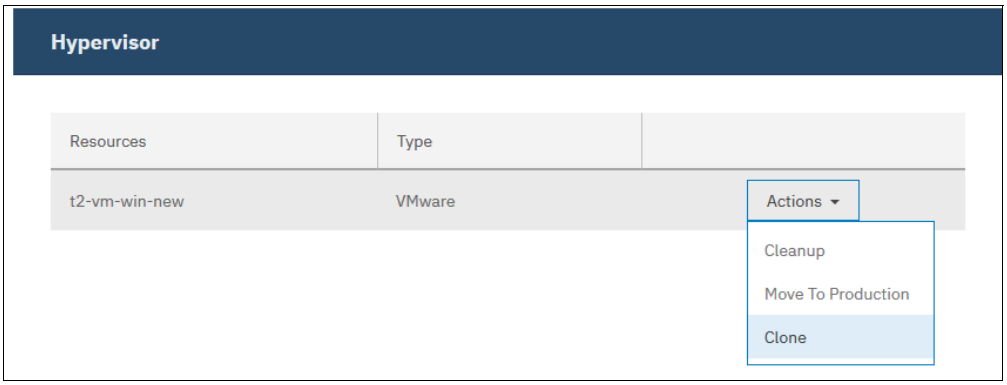

Figure 6-37 Start a clone operation if the VM should be kept

If the restored VM is not the correct version or configuration, we can choose the Cleanup option for the VM. This action removes all configurations from the VMware vCenter, and we can restore another version of the VM, if required. The advantage of the Test Mode restore is that data is not copied to access the VM. Therefore, you can easily check first if the restored VM is configured as expected before deciding to move it to production, or to clone it.

Note: For more information about restoring data when vCenter Server or other management VMs are not accessible, see IBM Knowledge Center.

6.3.9 Restoring a VM from secondary (replication) site

In this section, we describe how to restore a complete VM to a previous state from secondary (replication) site. In addition, we use the following options:

- Restoring to a different datastore location
- Renaming the VM
- Performing a Test Mode restore

In our example, we restore the `t5-vm-lx` (Linux VM) to the restore point that was created on June 30, as shown in Figure 6-38.

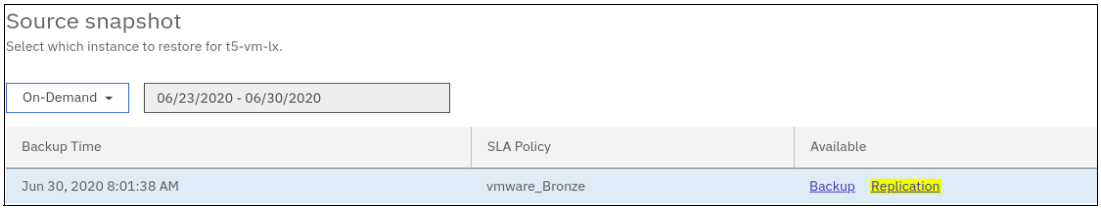

Figure 6-38 Select the source snapshot to restore the VM from replication side

We select the **Alternate Host or Cluster** option, as shown in Figure 6-39.

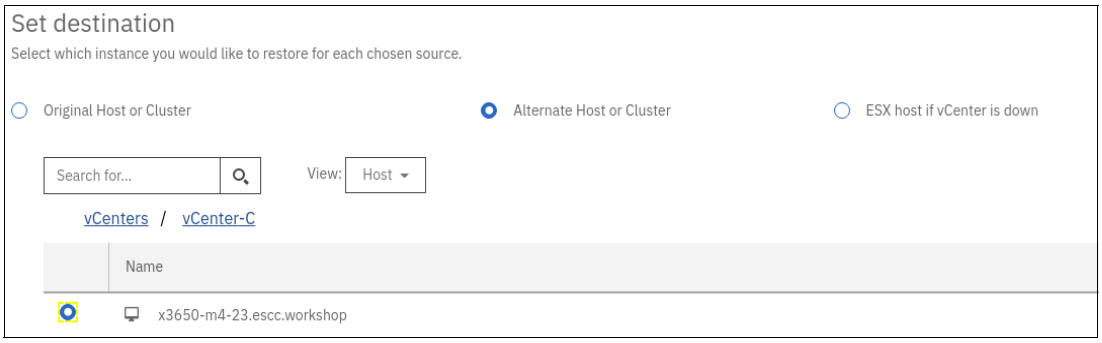

Figure 6-39 Select the Alternative Host or Cluster to restore the VM from replication side

We use a different datastore as the restore destination, as shown in Figure 6-40.

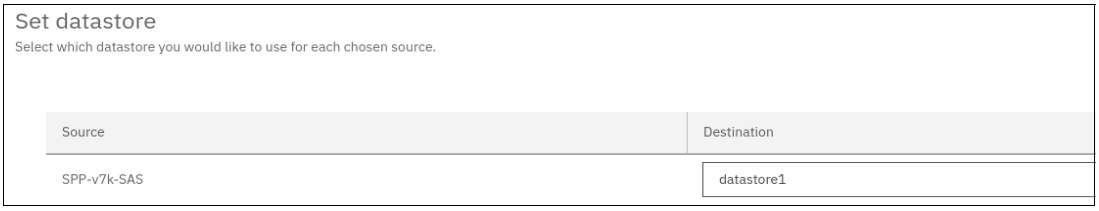

Figure 6-40 Select a different datastore as the restore destination

For the network, we select a VM network, as shown in Figure 6-41.

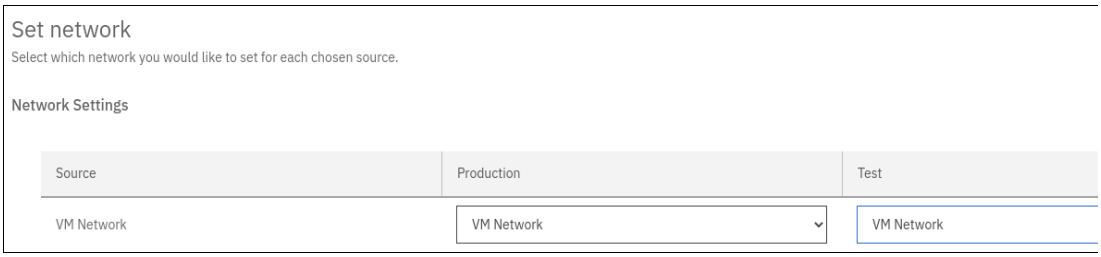

Figure 6-41 Network settings for VM restore

In Restore method window, we select **Test** and select the **Rename VM** option to rename the restored VM with the name `t5-vm-lx-repl` (see Figure 6-42).

Figure 6-42 Select a Restore method

You can verify all of settings again before submitting the restore job in the review pane that is shown in Figure 6-43. After the restore job is submitted, you can follow the progress in the Jobs and Operations view by selecting the **Running Jobs** tab. When the test restore completes mounting the restore point from the vSnap server and creating the VM, the VM is listed in the Active Resources tab under Hypervisor in the Jobs and Operations view.

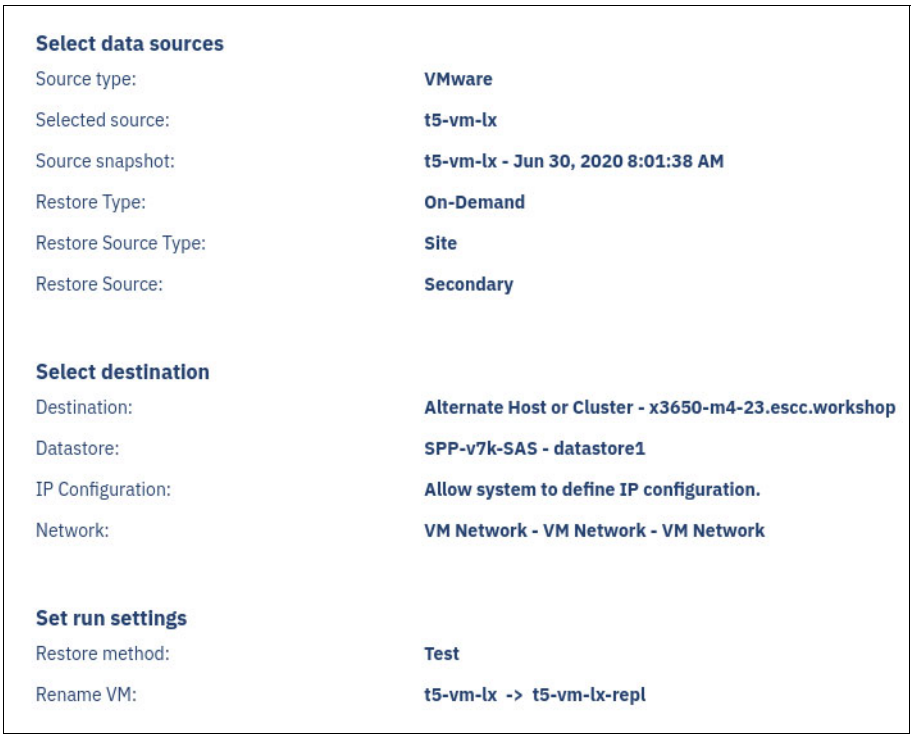

Figure 6-43 Review pane to verify the settings for the restore job

When the VM is restored, you can choose the Cleanup, Move to production, or Clone option, as shown in Figure 6-37 on page 201.

6.3.10 Restoring a VM and changing static IP address on one NIC

In this section, we describe how to restore a VM and change the static IP address on one Network Interface Controller (NIC) for a VM with a Linux operating system. Also, we rename the VM and use a Test Mode restore. In our example, we restore the `t5-vm-lx` (Linux VM).

First, the `root` user must be configured. Select **Manage Protection** → **Virtualized Systems** → **VMware** to see a list of VMs. Select a VM and then, click **Select Options**, as shown in Figure 6-44.

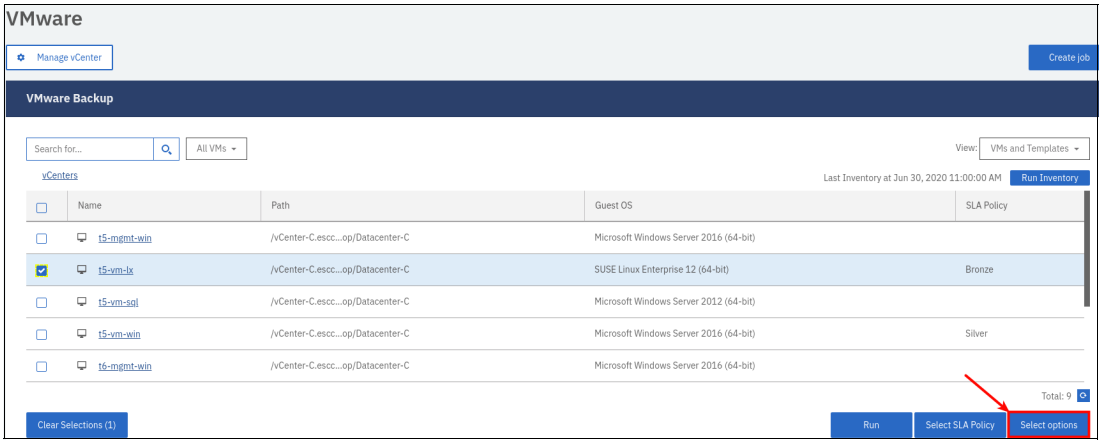

Figure 6-44 Select Options menu

Configure the Guest OS Username `root` and Guest OS password.

Note: The IP configuration for a Linux VM is changed during a restore. The user account `root` must be used as IBM Spectrum Protect Plus agent user (Guest OS Username).

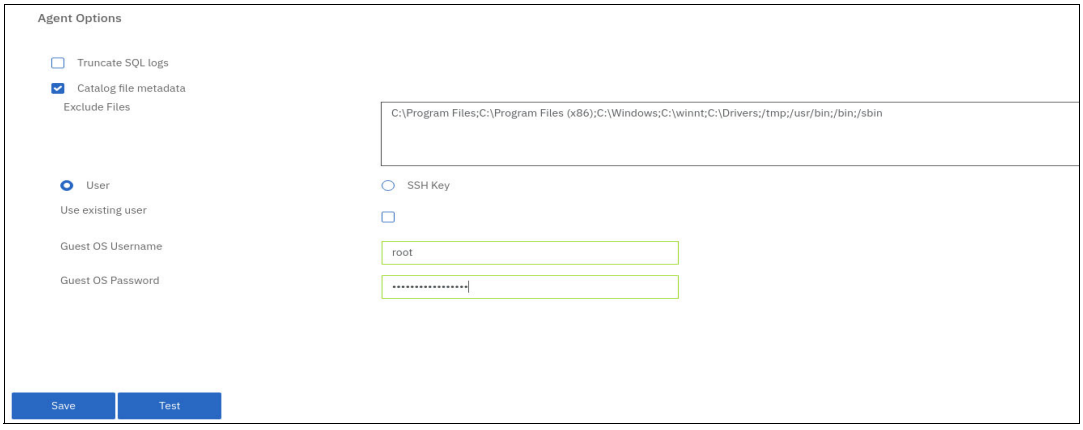

Figure 6-45 Configuration of root user for a VM

To validate the configuration, run a verification test by clicking **Test**. After testing, save configuration for the `root` user (see Figure 6-46).

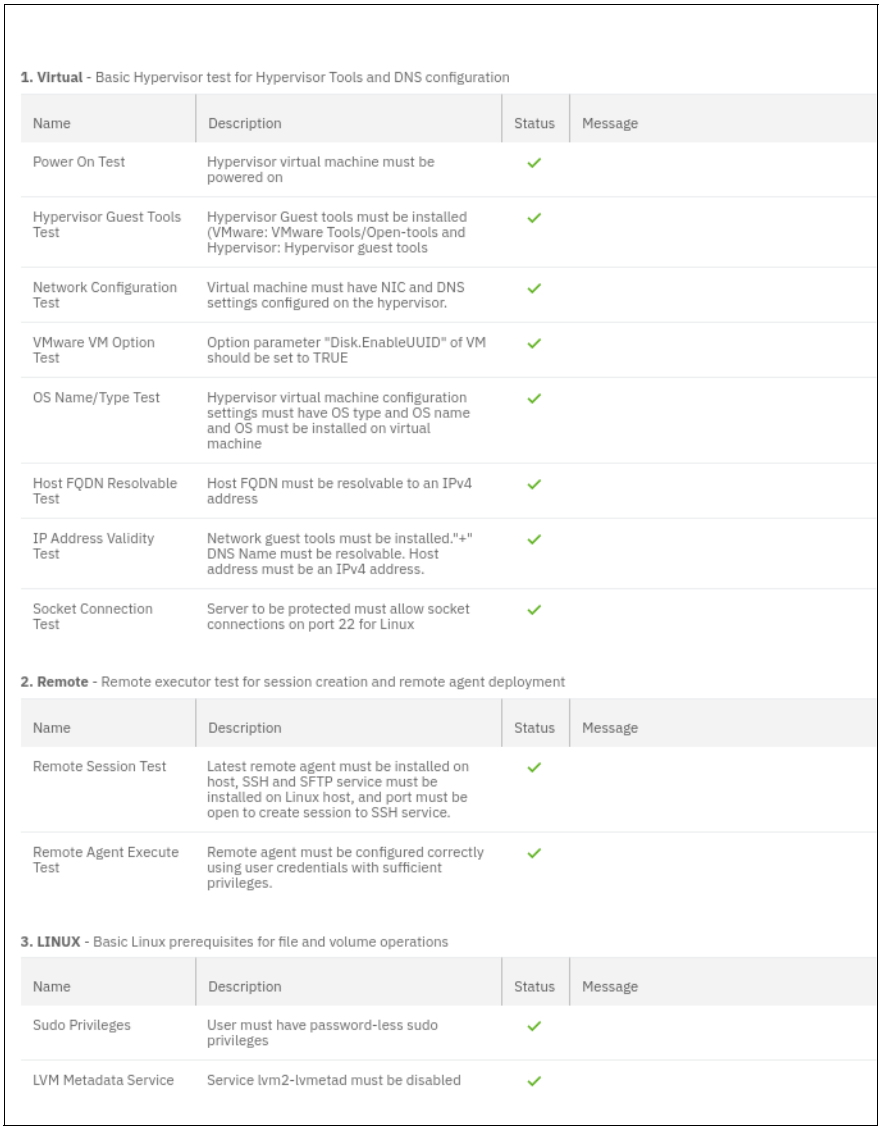

1. Virtual - Basic Hypervisor test for Hypervisor Tools and DNS configuration

Name	Description	Status	Message
Power On Test	Hypervisor virtual machine must be powered on	✓	
Hypervisor Guest Tools Test	Hypervisor Guest tools must be installed (VMware: VMware Tools/Open-tools and Hypervisor: Hypervisor guest tools	✓	
Network Configuration Test	Virtual machine must have NIC and DNS settings configured on the hypervisor.	✓	
VMware VM Option Test	Option parameter "Disk.EnableUUID" of VM should be set to TRUE	✓	
OS Name/Type Test	Hypervisor virtual machine configuration settings must have OS type and OS name and OS must be installed on virtual machine	✓	
Host FQDN Resolvable Test	Host FQDN must be resolvable to an IPv4 address	✓	
IP Address Validity Test	Network guest tools must be installed."+" DNS Name must be resolvable. Host address must be an IPv4 address.	✓	
Socket Connection Test	Server to be protected must allow socket connections 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	✓	
LVM Metadata Service	Service lvm2-lvmetad must be disabled	✓	

Figure 6-46 Test result of VM

In our example, restore the `t5-vm-lx` (Linux VM) restore point that was created on June 30, as shown in Figure 6-47.

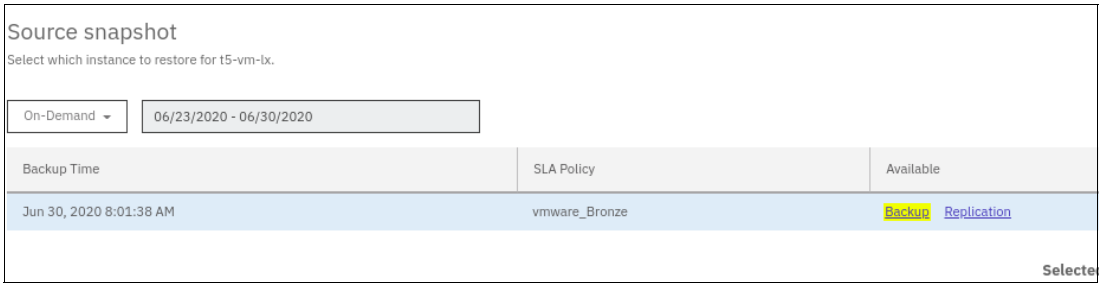

Figure 6-47 Select the source snapshot to restore the VM

Select **Alternate Host or Cluster**, as shown in Figure 6-39 on page 202. Use the same datastore, as shown Figure 6-48.

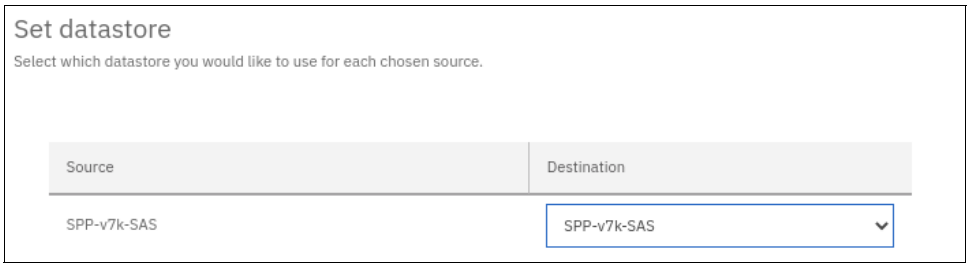

Figure 6-48 Select a datastore as the restore destination

In the Network section, set the network setting as the current IP address (for VM `t5-vm-lx`) and new IP address (for VM `t5-vm-lx_newIP`), as shown in Figure 6-49. The following settings were used in our example:

- Current IP address: 10.0.250.55
- New IP address: 10.0.250.59
- Subnet mask: 255.255.255.0
- Gateway: 10.0.250.9
- DNS: 10.0.250.210

Note: The current IP address must be specified to identify the network interface that is changed. In addition, the new IP address, Subnet mask, and Gateway are mandatory fields that must be completed (the DNS field is optional). Instead of specifying a new IP address configuration, you also can select **DHCP** for the interface.

After the network settings are completed, click the plus (**+**) symbol.

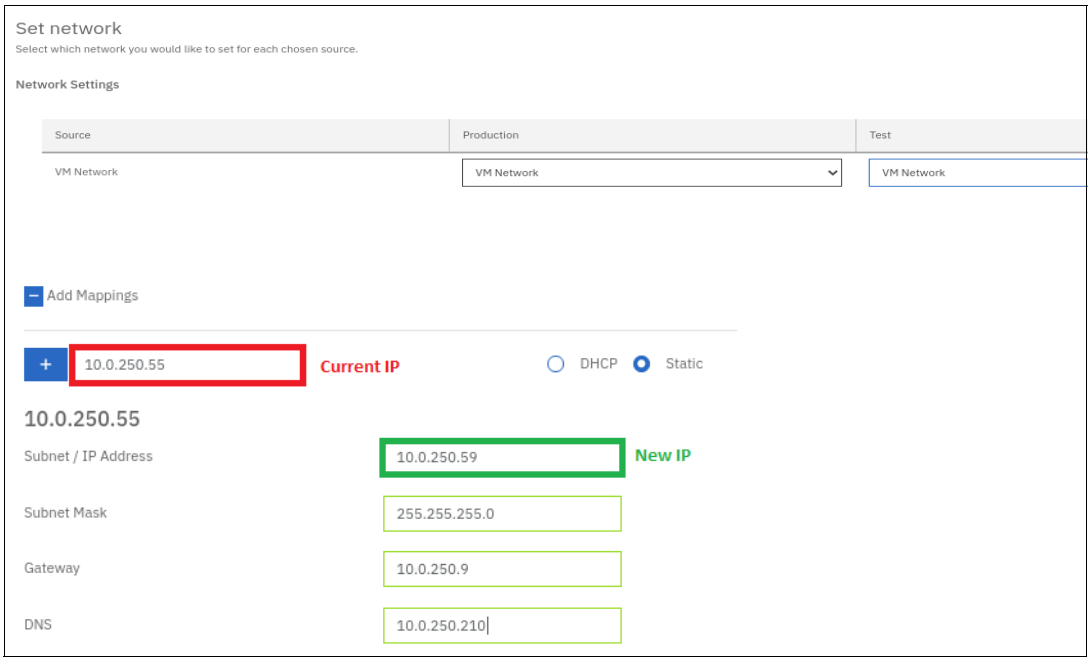

Figure 6-49 Network Setting Configuration

Note: During restore, the VM must be started to configure these network settings.

Choose a Restore method as **Test** and `t5-vm-lx_newIP`, as shown Figure 6-50.

Figure 6-50 Restore Method Test

Verify all settings and submit the restore job (see Figure 6-51).

Figure 6-51 Review pane to verify the settings for the restore job

You can follow the progress of the restore in the Jobs and Operations view under the Running Jobs tab, as shown in Figure 6-52.

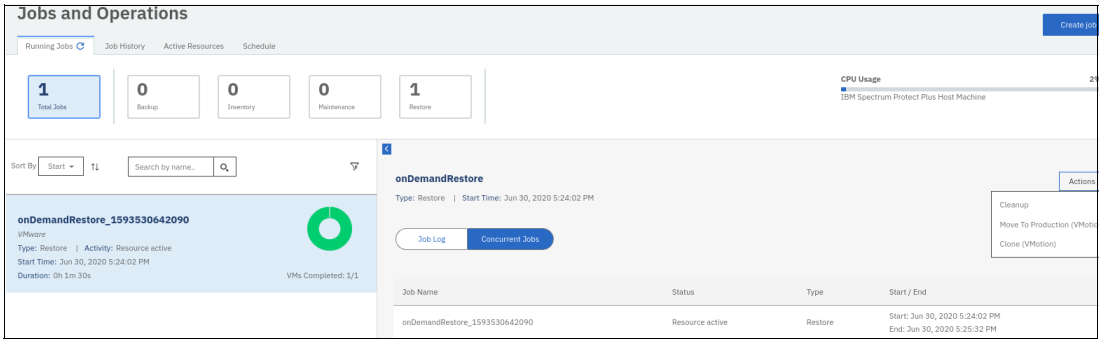

Figure 6-52 Job and Operation view

The VM restored VM is listed in the Active Resources tab under Hypervisor in the Jobs and Operations view (see Figure 6-53).

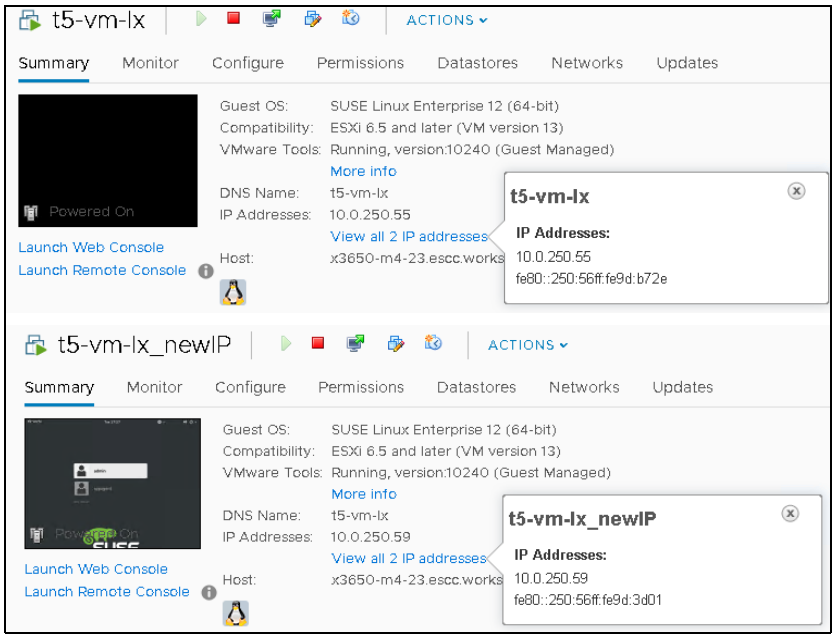

Figure 6-53 Original VM t5-vm-lx and Restored VM t5-vm-lx_newIP with changed IP address

After the restore operation completes, the restored VM can now be started. We can see in Figure 6-53 that the restored VM uses the new IP address. You can decide what you will do with the VM; for example, you can select the **Move to Production** option if the restored VM is configured as expected.

6.3.11 Restoring a VM and changing static IP addresses on two NICs

We now restore a Windows VM and change static IP addresses on two NICs. We also rename the VM and use a Test Mode restore. In our example, we restore the `t5-vm-win` (Windows VM) to the restore point that was created on July 1, as shown in Figure 6-54.

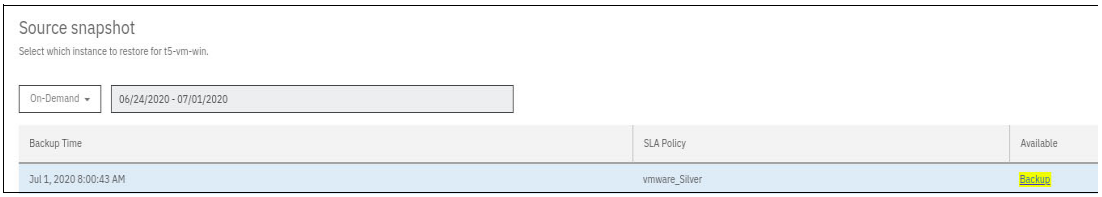

Figure 6-54 Select the source snapshot to restore the VM

Note: To change the IP configuration of a Windows Guest Operating System, you must ensure that you installed Microsoft Visual C++ 2008 Service Pack 1 Redistributable Package MFC Security Update before backing up the VMs so that the package is available in case of a restore with a required IP re-configuration. You can download VC 2008 SP1 Redistributable (9.0.30729.6161) at this web page.

We select the **Alternate Host or Cluster** option in the Set datastore panel (see Figure 6-48 on page 206). In the Set network panel (see Figure 6-49 on page 206), we specify the following network parameters:

- First NIC:
  - Current IP address 10.0.250.54
  - New IP address 10.0.250.59
  - Subnet mask 255.255.255.0
  - Gateway 10.0.250.9

  After the network setting are complete, click plus (**+**) symbol.
- Second NIC:
  - Current IP address 172.0.0.54
  - New IP address 172.0.0.59
  - Subnet mask 255.255.255.0
  - Gateway 172.0.0.1

After network settings are complete, click the plus (**+**) symbol. You can then review the Network parameters, as shown in Figure 6-55.

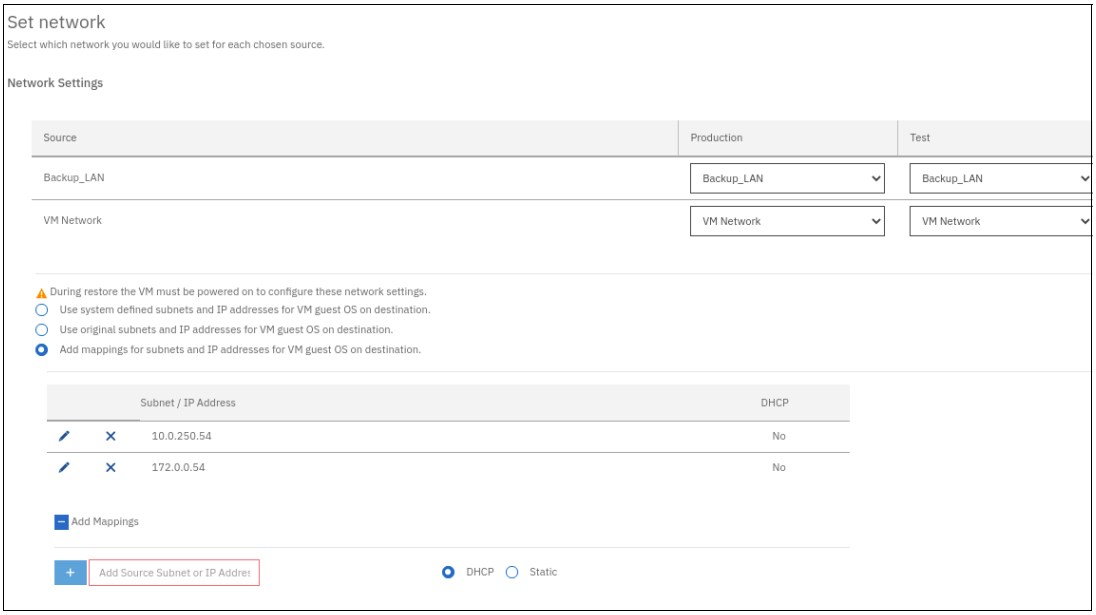

Figure 6-55 Network Setting Configuration

In the Restore method panel, choose **Test** and name `t5-vm-win_newIP`. In the review pane that is shown in Figure 6-56 on page 210, verify all the settings before submitting the restore job.

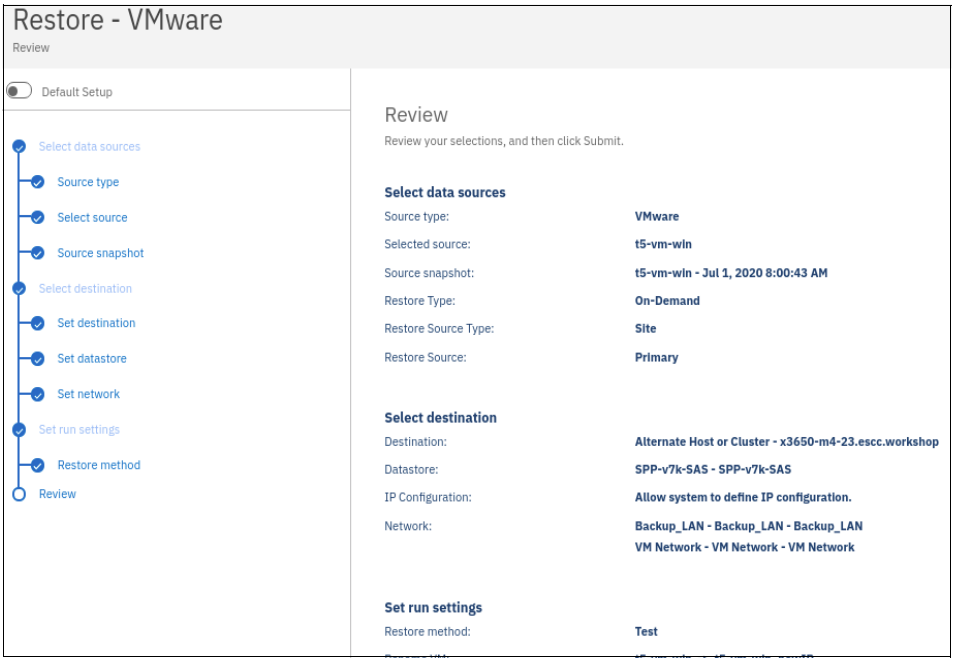

Figure 6-56 Review pane to verify the settings for the restore job

Follow the progress in the Jobs and Operations view in the Running Jobs tab.

In the vCenter User Interface (UI), we see the VM that was created by using the temporary storage resources from the vSnap server. After the restored VM is started, we can see that it uses the changed IP addresses on both NICs, as shown in Figure 6-57.

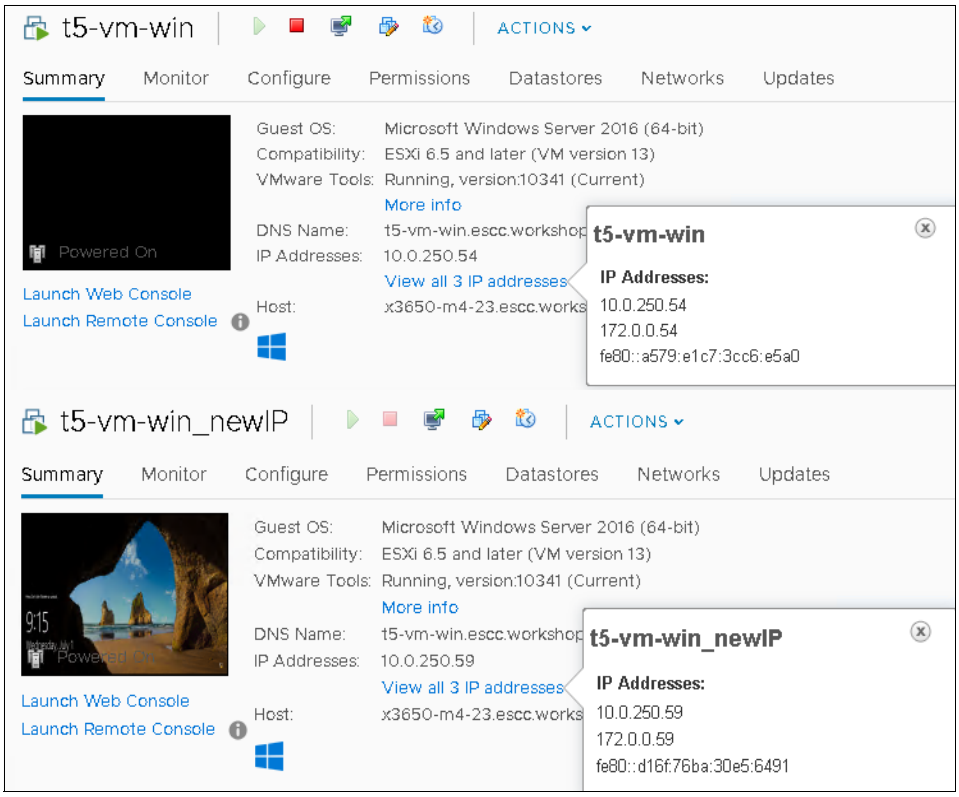

Figure 6-57 Original VM t5-vm-win and Restored VM t5-vm-win_newIP with changed IP addresses

6.4 Protecting and recovering Amazon EC2 data

In this section, we describe a new function of the backup and recovery Amazon EC2 data that is stored in Amazon Elastic Compute Cloud (EC2) managed by Amazon Web Services (AWS). VMs in AWS use Amazon's Elastic Block Storage (EBS). They also feature allocated space for the VDisks.

By using IBM Spectrum Protect Plus, you can create a backup by creating a snapshot of the EC2 instances by protecting VDisk data on the EBS. This new feature was introduced with version 10.1.6. Figure 6-58 shows an example of the hybrid environment implementation of the IBM Spectrum Protect Plus and Amazon EC2.

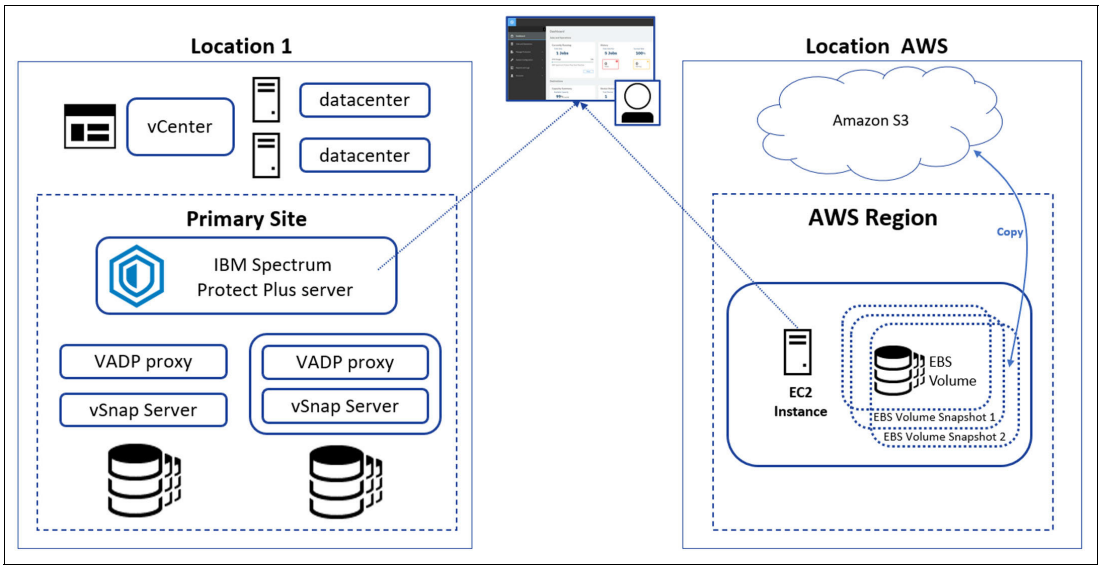

Figure 6-58 IBM Spectrum Protect Plus Amazon EC2 integration

The new Hypervisor provider type Amazon EC2 was designed and integrated the same way as the other providers (Vmware and Hyper-V). In IBM Spectrum protect Plus version 10.1.6, this new function can be found in the Manage Protection Panel. Select **Virtualized Systems** → **Amazon EC2**.

6.4.1 Amazon EC2 requirements and account management

Ensure that the basic requirements to enable the functionality of data protection for Amazon EC2 data by using IBM Spectrum Protect Plus are met. For more information about these requirements, see IBM Knowledge Center.

Creating and registering IAM user in AWS

In AWS, open the Identity and Access MAnagement console (IAM) to create a user for EC2. An IAM user can be created with Administrative permission only.

For more information about this process, see IBM Knowledge Center.

During the initial setup in the AWS console, you must set up users with an Access key.

By clicking the **JSON** tabs, you must enter the action that is shown in Figure 6-59 that defines specific policy and permissions for EC2 users.

```
{
   "Version": "2012-10-17",
   "Statement": [
       {
           "Sid": "VisualEditor0",
           "Effect": "Allow",
           "Action": [
               "ec2:DetachVolume",
               "ec2:AttachVolume",
               "ec2:DeregisterImage",
               "ec2:DeleteSnapshot",
               "ec2:DescribeInstances",
               "ec2:CreateVolume",
               "ec2:DescribeTags",
               "ec2:CreateTags",
               "ec2:RegisterImage",
               "ec2:DescribeRegions",
               "ec2:RunInstances",
               "ec2:DescribeSnapshots",
               "ec2:DescribeAvailabilityZones",
               "ec2:CreateSnapshots",
               "ec2:DescribeVolumes",
               "ec2:CreateSnapshot",
               "ec2:DescribeSubnets",
               "iam:PassRole"
           ],
           "Resource": "*"
       }
   ]
}
```

Figure 6-59 Defining specific policy and permissions for EC2 users

Registration of the Amazon EC2 account in IBM Spectrum Protect Plus

After the successful creation and registration of the IAM user in the AWS portal, the next step is to register the created AWS account by using IAM user's access key and secret key in IBM Spectrum Protect Plus.

To register the EC2 account in IBM Spectrum Protect Plus, select **Manage Protection** → **Virtualized Systems** → **Amazon EC2** → **Manage Accounts** and click **Add Account**. You are prompted to name the account and provide the Access and Secret Key.

If the user is in Spectrum Protect Plus, it can be selected from the drop-down menu by selecting the **Use existing access key** option, as shown in Figure 6-60 on page 213. For the new registration process, you must complete the mandatory fields. The registration process completes when you click **Save**.

In the Account Properties, enter the following information:

- Account Name
- Access Key
- Secrete Key

You can register multiple accounts for different users and manage them separately. The registration process is finished by clicking **Save**.

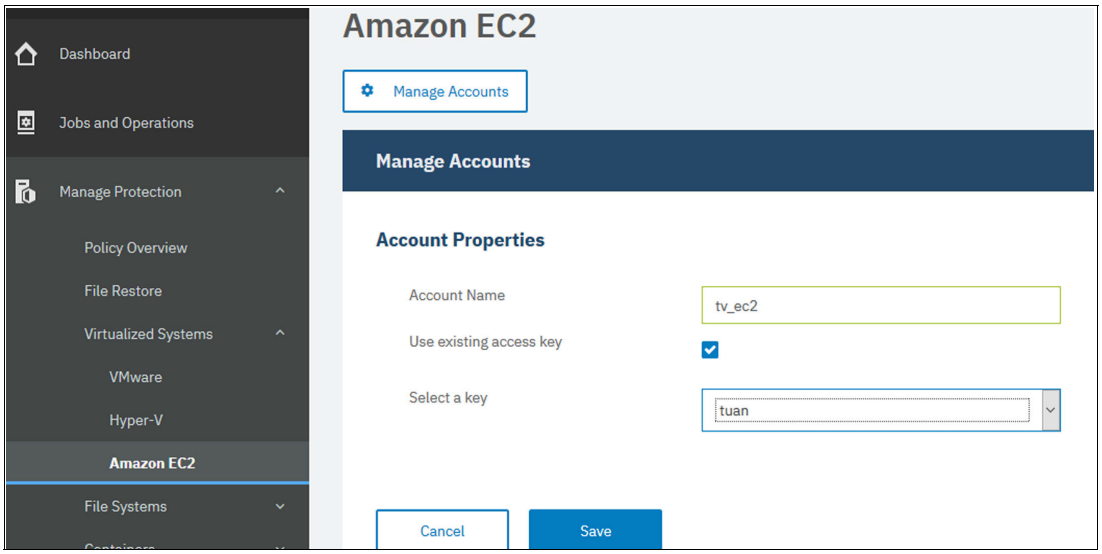

Figure 6-60 Amazon EC2 Manage Accounts

In Spectrum Protect Plus, select **Manage Protection** → **Virtualized Systems**. In the **Amazon EC2 backup** tab, click **Run inventory**, as shown in Figure 6-61.

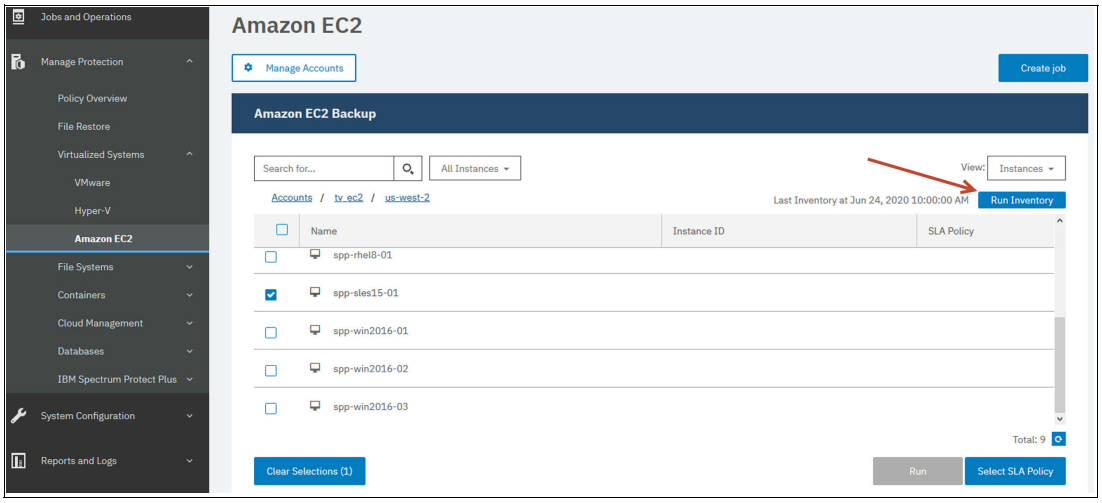

Figure 6-61 Run Amazon EC2 Inventory

The Inventory process starts against AWS that discovers regions and instances. This process is used to update the inventory for AWS EC2.

Discovered regions in AWS are treated as hypervisors. During the inventory process, the information that is related to the detected EBS volumes and EC2 instances on AWS is saved in the IBM Spectrum protect Plus catalog.

After the inventory completes successfully, instances become visible, and you can select the EC2 instances that you want to associate with an SLA policy.

6.4.2 Amazon EC2 data protection and recovery

In this section, we describe how to configure Amazon EC2 data protection and recovery.

Amazon EC2 SLA policy

Create an SLA policy or reuse an existing SLA policy that is used for EC2 snapshots, with a specification of the frequency and retention parameter. Create or adjust SLA Policy by selecting **Manage Protection** → **Policy Overview** → **Add SLA Policy**.

Select the policy type **Amazon EC2**, as shown in Figure 6-62, and enter the policy name, retention, and the frequency settings. You can also specify the start time and a Snapshot Prefix that is added to the beginning of EBS snapshots names.

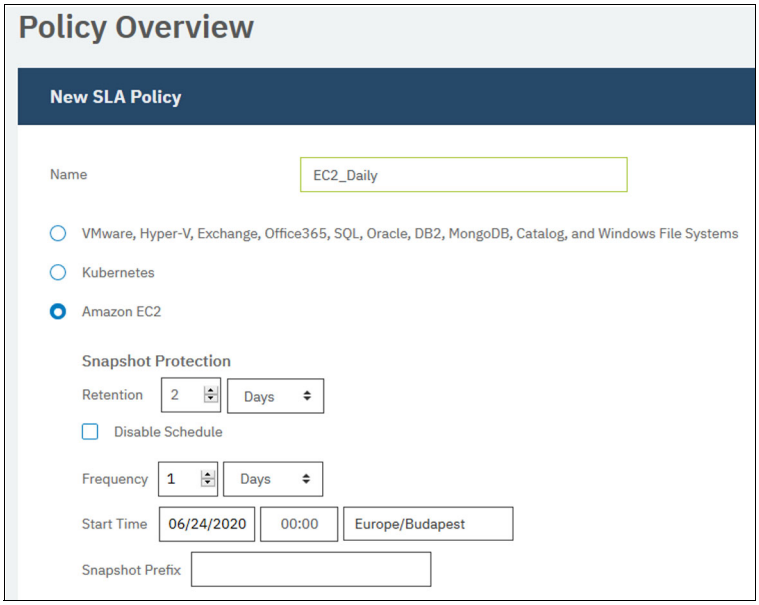

Figure 6-62 EC2 SLA Policy

To assign EC2 Instances to an SLA policy, select **Manage Protection** → **Virtualized Systems** → **Amazon EC2**. Select one or multiple EC2 instances that you want to assign to an SLA Policy and click **Save**, as shown in Figure 6-63.

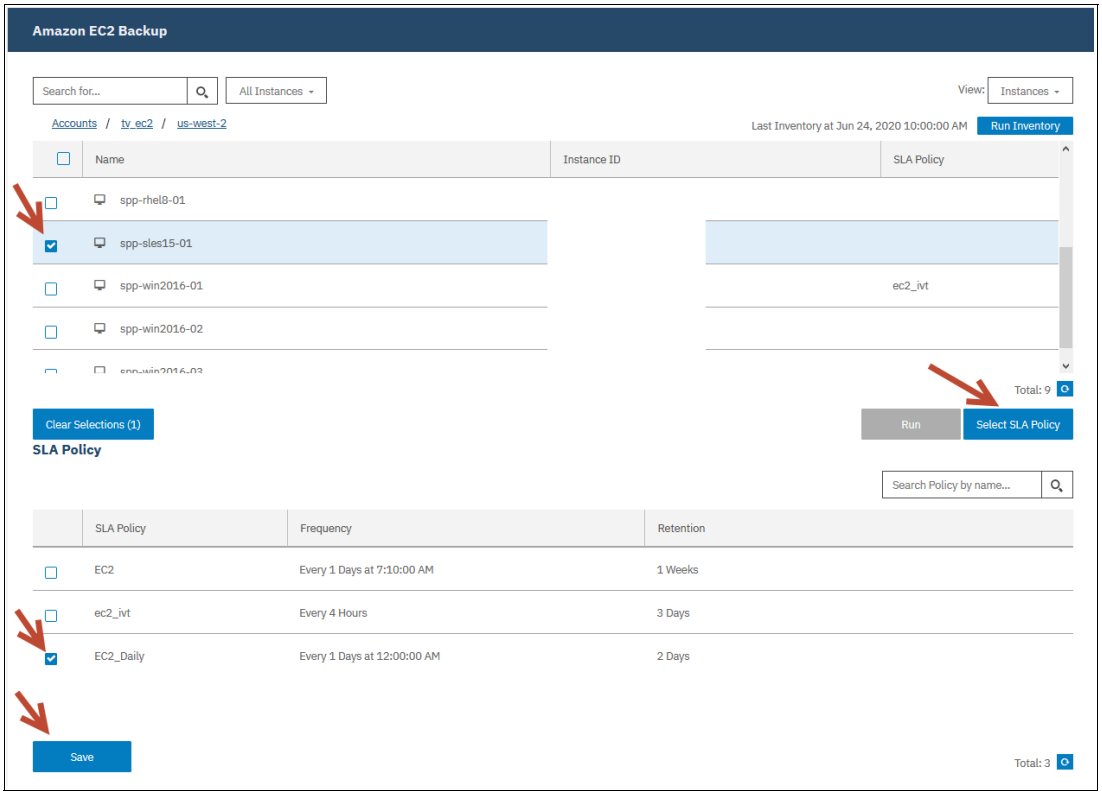

Figure 6-63 Assign EC2 SLA policy

Amazon EC2 instance backup operation

Following the EC2 SLA policy assignment, you can manually run the policy by selecting the **SLA Policy Status** tab → **Actions** → **Start**. The snapshot starts for all EC2 instances with the assigned SLA Policy, as shown in Figure 6-64.

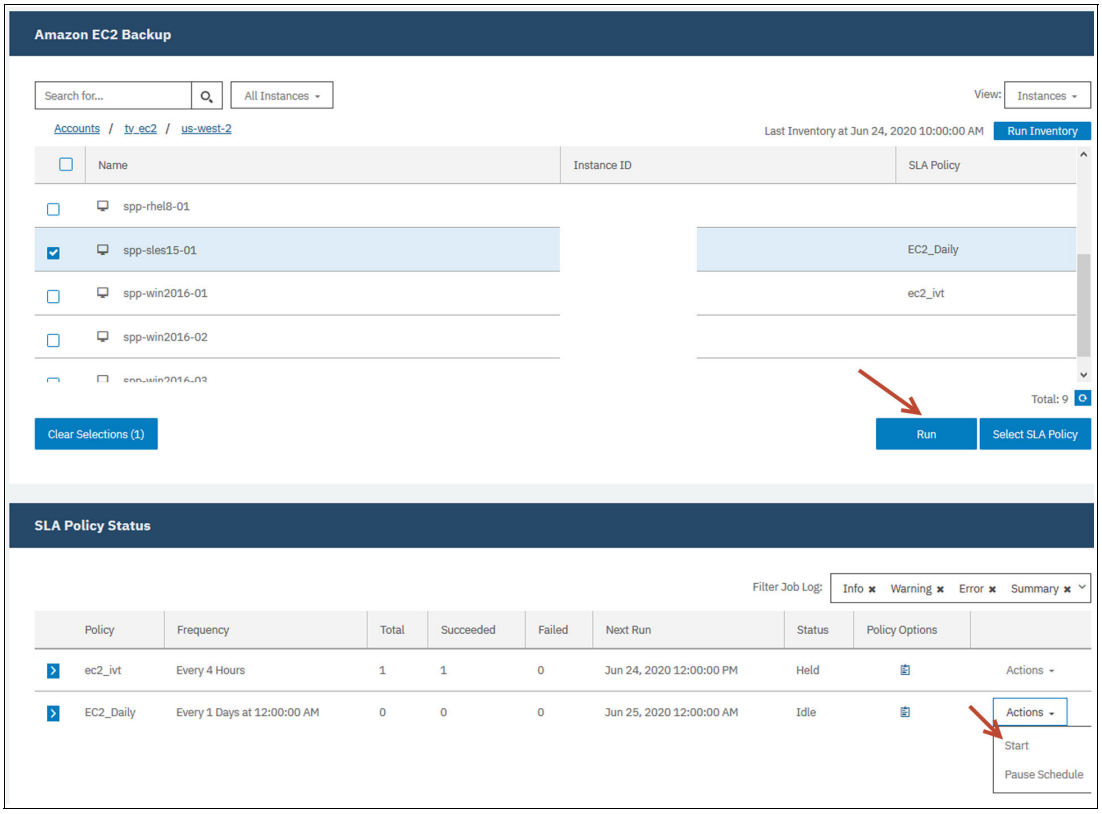

Figure 6-64 Amazon EC2 Instance Snapshot

For more information about how to back up Amazon EC2 data, see IBM Knowledge Center.

IBM Spectrum Protect Plus interacts directly with Amazon services without a proxy or extra agent. EC2 instances snapshot data is copied incrementally to Amazon S3 object storage from the Amazon Elastic Block Store (EBS) volume. The copies of the EBS snapshots are always stored in the user's Amazon S3 object storage.

For EC2 instances that are assigned to an SLA policy, an ad-hoc snapshot is also available. It can be started by selecting the EC2 instance and clicking **Run**, as shown in Figure 6-64.

Another way to start the ad-Hoc backup is by selecting **Manage Protection** → **Virtualized Systems** → **Amazon EC2**. Click **Create Job**, as shown in Figure 6-65.

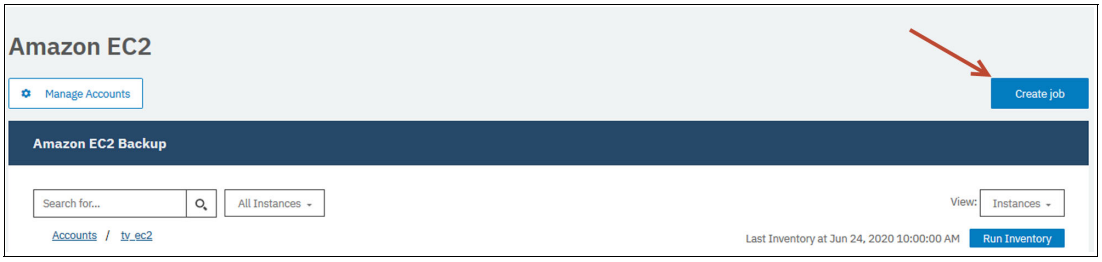

Figure 6-65 Amazon EC2 Create job

Then, select **Ad hoc Backup** and click **Select**, as shown in Figure 6-66.

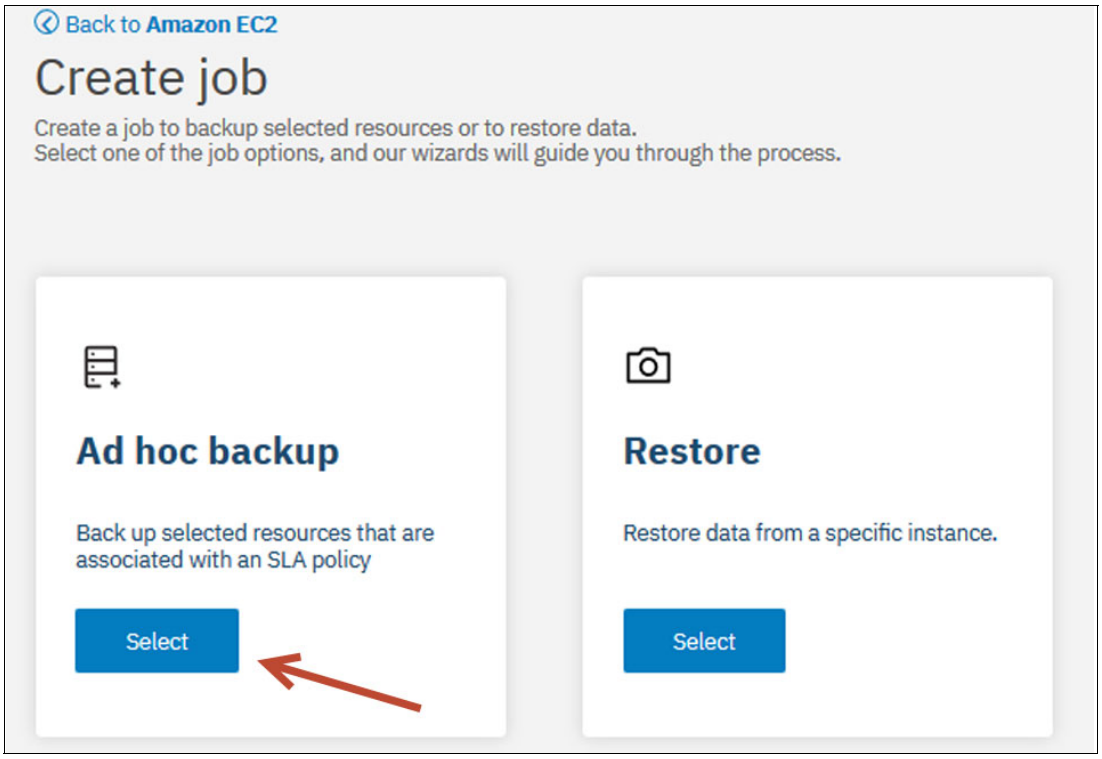

Figure 6-66 Amazon EC2 Ad hoc backup

The next step is to select an SLA Policy from the navigation panel, as shown in Figure 6-67.

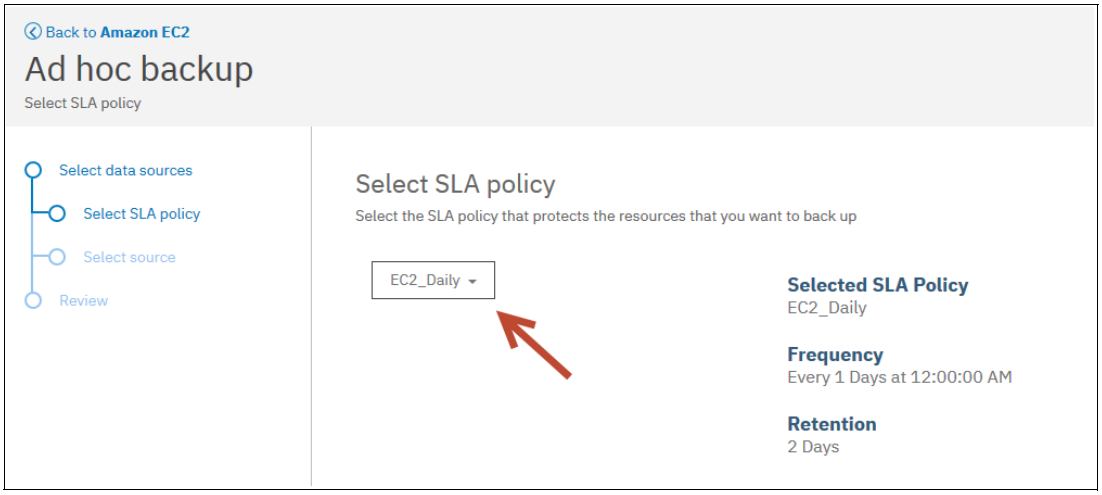

Figure 6-67 Amazon EC2 Ad hoc backup

By clicking the "**+**" button, select the Amazon EC2 instance that you want to add to the Ad hoc backup, as shown in Figure 6-68 on page 218. Continue by clicking **Next**. In the next window, review your choice and then, click **Submit.**

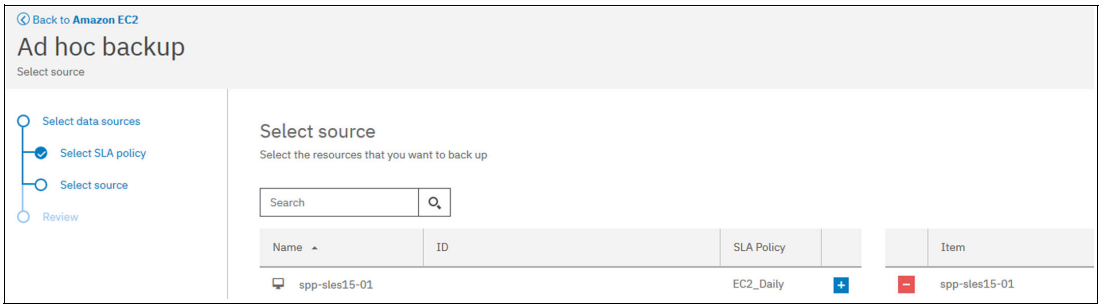

Figure 6-68 Amazon EC2 Ad hoc backup

Amazon EC2 restore operation

The restore operation of the Amazon EC2 instance can be started only after a successful EC2 backup job operation is completed. From the previously created snapshots, you can create the Restore job of the EC2 instance.

The EC2 instance restore operation follows steps similar to the operation for restoring a VM in VMware or Hyper-V hypervisors. How to Restore Amazon EC2 data is also found in IBM Knowledge Center for IBM Spectrum protect Plus under the Restoring Amazon EC2 data option.

Select **Manage Protection** → **Virtualized Systems** → **Amazon EC2**. In the top right corner, click **Create Job**, as shown in Figure 6-65 on page 216. Then, choose to **Restore** to restore data from specific instance. Then, click **Select**, as shown in Figure 6-69.

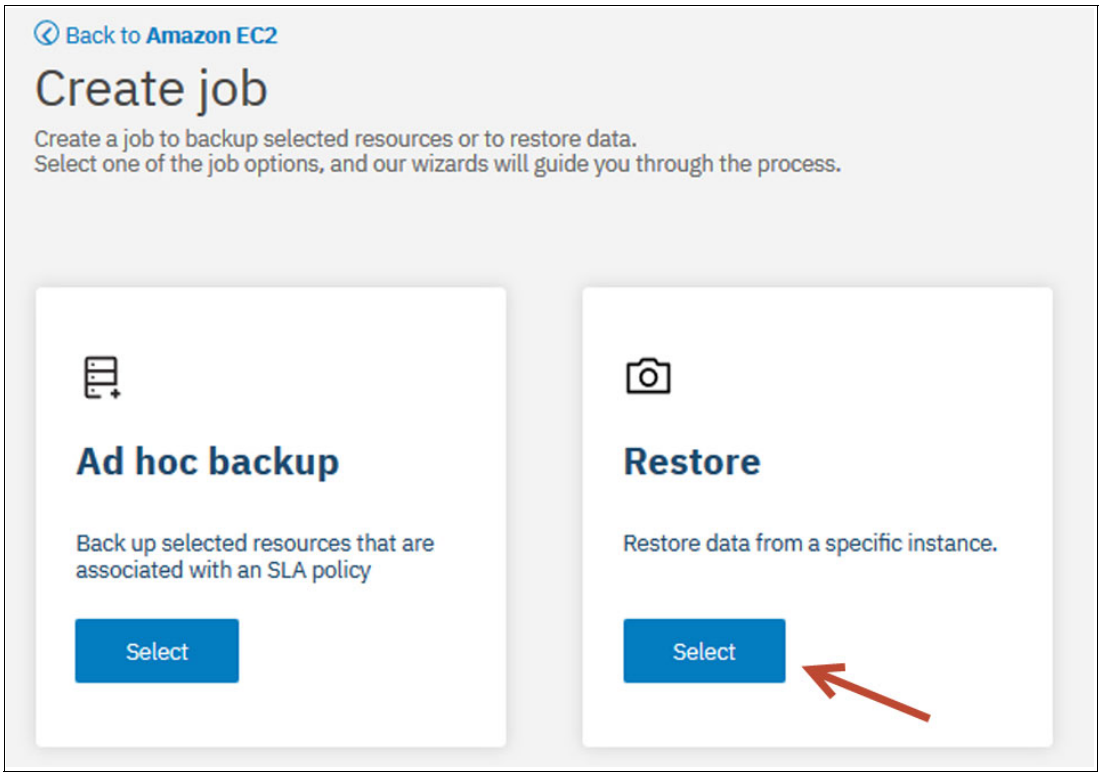

Figure 6-69 Amazon EC2 Restore

The restore job can also be created by selecting **Jobs and Operation** → **Create Job** → **Restore** and select the source type **Amazon EC2**, as shown on Figure 6-70. Continue by clicking **Next**.

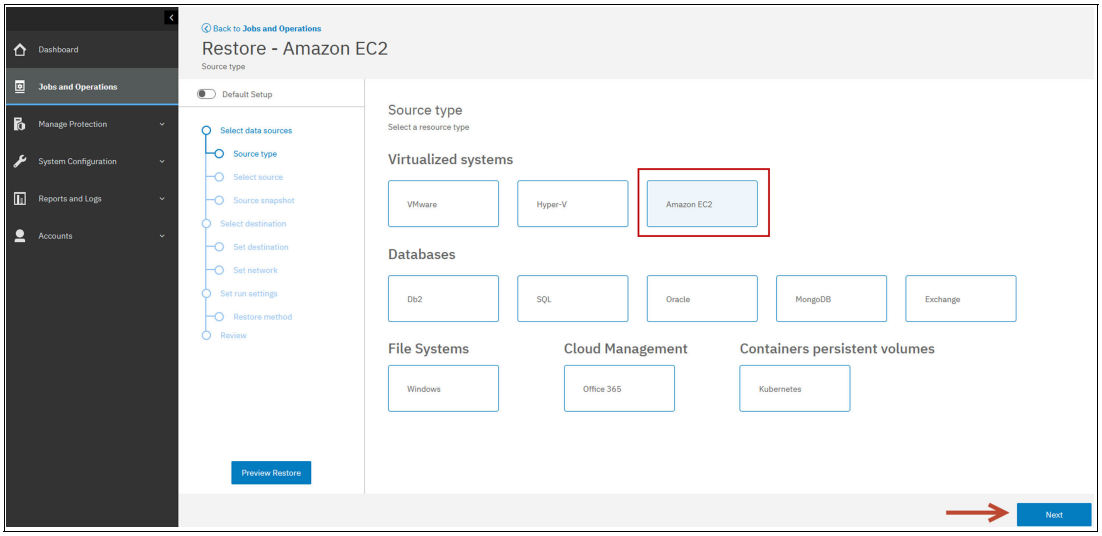

Figure 6-70 Amazon EC2 Restore window opened from Jobs and Operation Navigation Panel

The source Amazon EC2 is among the available resources that are on the panel. To select the instance to recover, click the plus **"+"** symbol, as shown in Figure 6-71. Continue by clicking **Next**. You can select one or multiple Instances from the same region only. If you want to remove Instances or volumes from the list, click the minus **"-"** symbol.

Note: During the restore operation, any volumes that are attached to the instance must be selected separately. It is not possible to select instances and their attached volumes for the same restore.

By default, the restore wizard is opened in Default Setup mode. To continue the restore operation, run the wizard in Advanced Setup mode by clicking the icon for **Advanced Setup** mode, as shown on Figure 6-71.

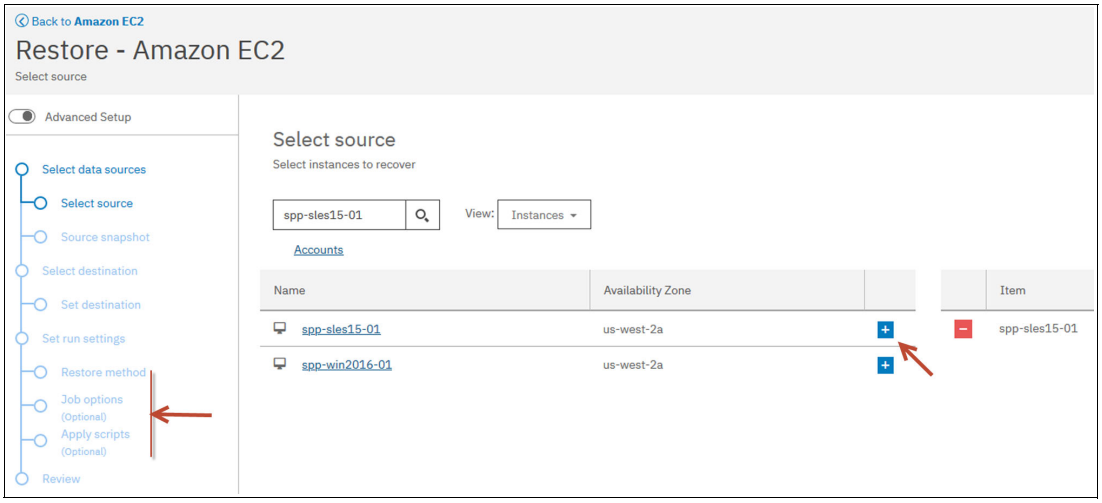

Figure 6-71 Amazon EC2 Restore Wizard opened in Advanced Setup Mode

The Restore wizard includes the following set destination options:

- Original Availability zone

  The EC2 instance snapshot is restored in the same zone as the original (see Figure 6-72).

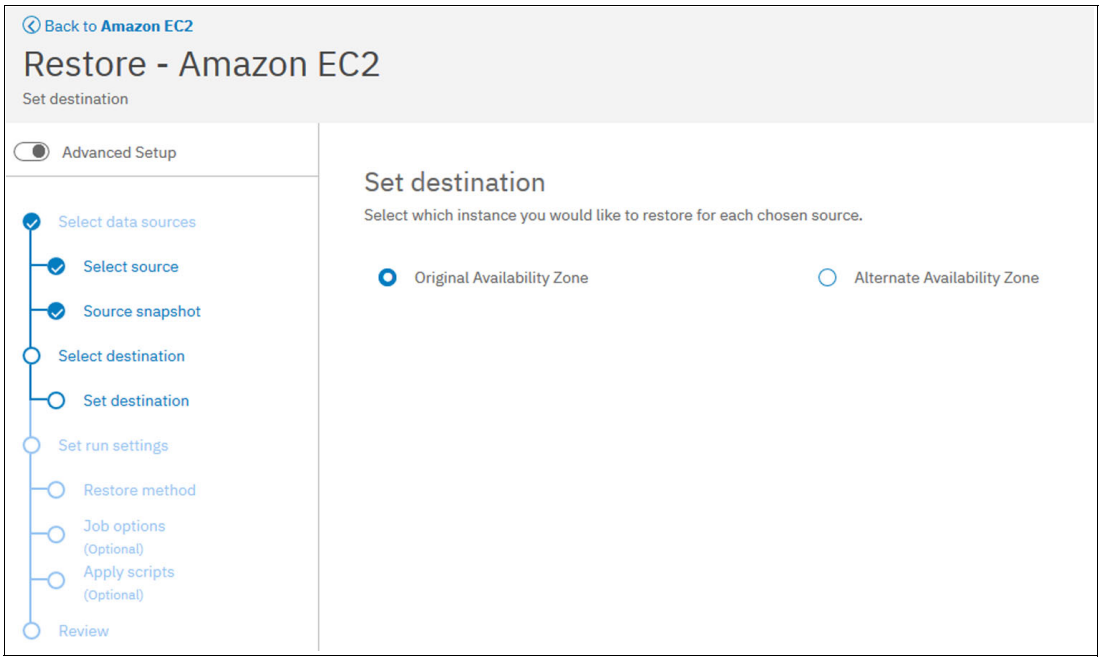

Figure 6-72 Amazon EC2 Set destination option

- Alternate Availability Zone

  The EC2 Instance can be restored in an Alternate Availability Zone (see Figure 6-73). You can select an alternative location from the accessible resources. You must specify the subnet for the Instance to be restored. The Alternate Availability Zone subnet must be in the same region as the instances selected for the restore.

  If you want restore an attached volume in the Alternate Availability Zone, select a destination attachment field.

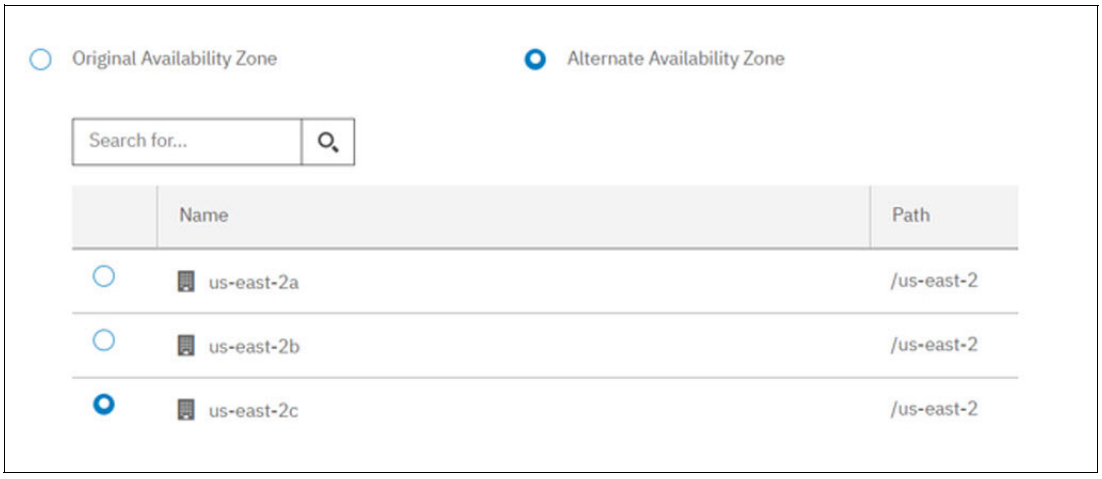

Figure 6-73 Amazon EC2 Restore to Alternate Availability Zone

All EC2 restores run in “clone” mode, which creates a copy of the instance from the EBS snapshots.

The restore job process includes the following steps:

1. Registering a temporary custom Amazon Machine Image (AMI) from the snapshot and instance metadata in inventory.
2. Running an instance from that AMI and then cleaning up the AMI after the instance is created.

In the Restore method window, you can rename the restored Amazon EC2 Instance, as shown in Figure 6-74. Selecting Rename Instance is optional.

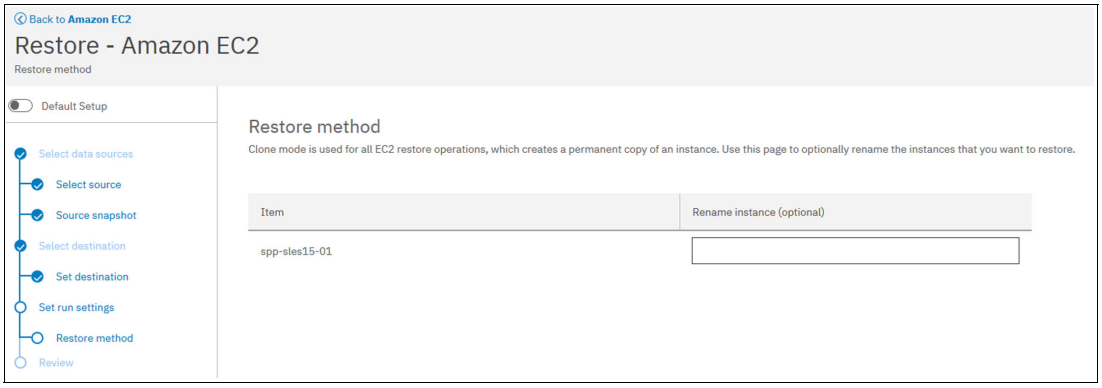

Figure 6-74 Amazon EC2 Restore method

By choosing **Default Setup mode**, as shown in Figure 6-75, you can restore the selections by clicking **Submit** to run the restore job.

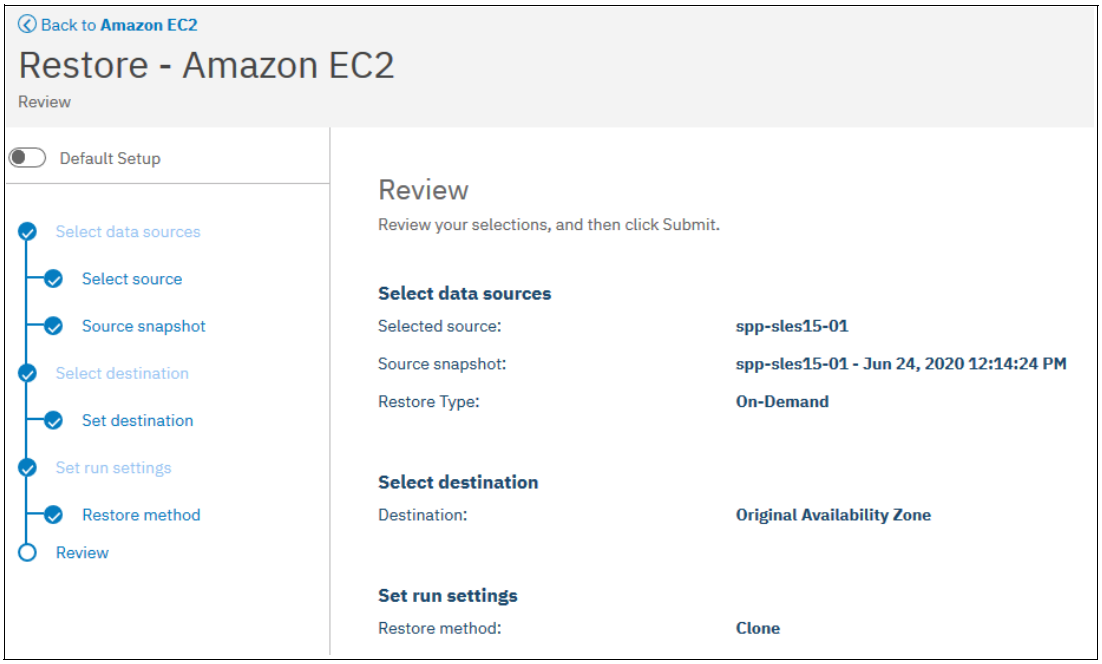

Figure 6-75 Amazon EC2 review Restore selection Default Setup mode

If the Advance Setup mode is selected, as shown in Figure 6-76, the following extra options are available:

- Power on after recovery

  Starts the instance after the recovery action is completed.

- Continue with restore even if it fails

  If this option is enabled, the recovery process of the EC2 instances continues in series, even if the previous instance recovery fails. If you choose to disable this option, the restore operation is stopped, and restoration of the instance fails.

- Run cleanup immediately on job failure

  If the restore job fails, this option automatically enables the clean up of all allocated resources as part of the restore job operation.

- Restore instance tags

  By selecting this option, you also enable restore tags that are applied to the instances through vSphere.

- Prepend prefix to instance name

  The prefix you enter is added to the names for the restored instances.

- Append suffix to instance name

  The suffix that is entered is added to the names for the restored instances.

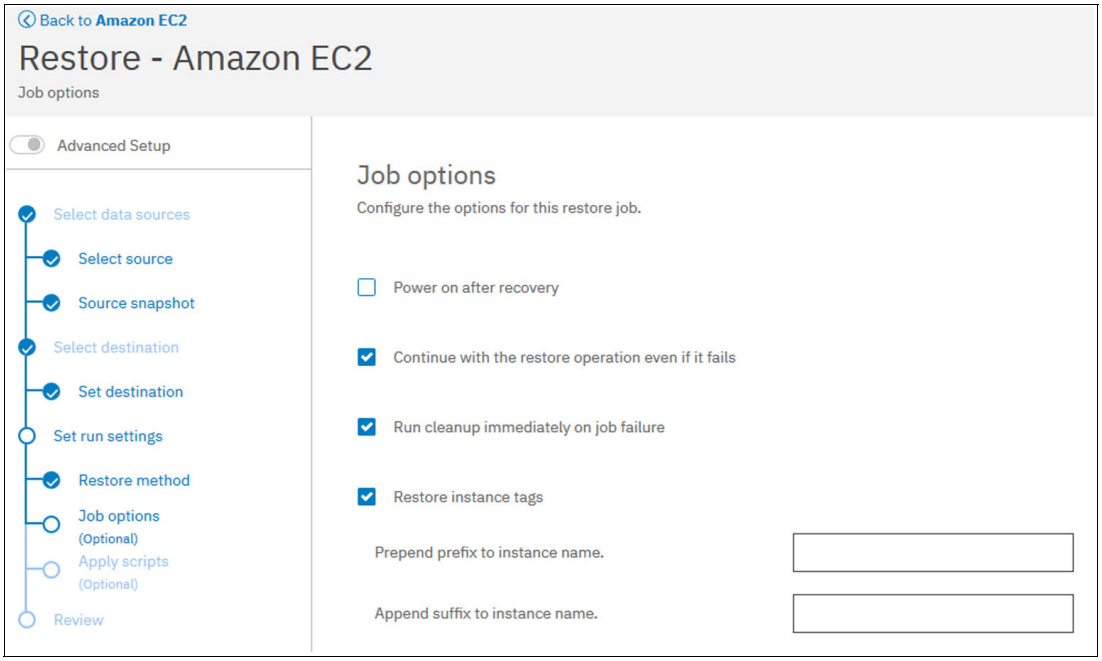

Figure 6-76 Recovery Amazon EC2 Advanced Setup mode

Optionally, in Advanced mode setup, you can also select pre- or post-scripts to run with the Amazon EC2 restore job, as shown in Figure 6-77 on page 224:

- Pre-script

  This option allows you to choose an uploaded script and an application or script server where the prescript runs. If you want to use a pre-script from the server, in the Selection panel, select the text that was uploaded or configured in **System Configuration** → **Script** page.

  If you decide not to use the script server, deselect the **Use Script Server Script** option**.** Then, select the script you uploaded in **System Configuration** → **Script** page and specify the **Application Server** field.

- Post-script

  By selecting this option, you can choose an uploaded script an application or script server where the prescript runs. If you want to use a pre-script from the server, in the Selection panel, select the text that was uploaded to or configured in **System Configuration** → **Script** page.

  If you decide not to use the script server, deselect the **Use Server Script** option. Then, select the script you uploaded in **System Configuration** → **Script** page and specify the **Application Server** field.

- Continue job/task on script error

  By selecting this option, the job continues to run, even if the associated pre- or post-script task status completes with a non-zero return code. At the end of the job, the pre- or post-script task status returns COMPLETED. If you do not select this option and the restore job does not run, the pre- or post-script task responds with a FAILED state.

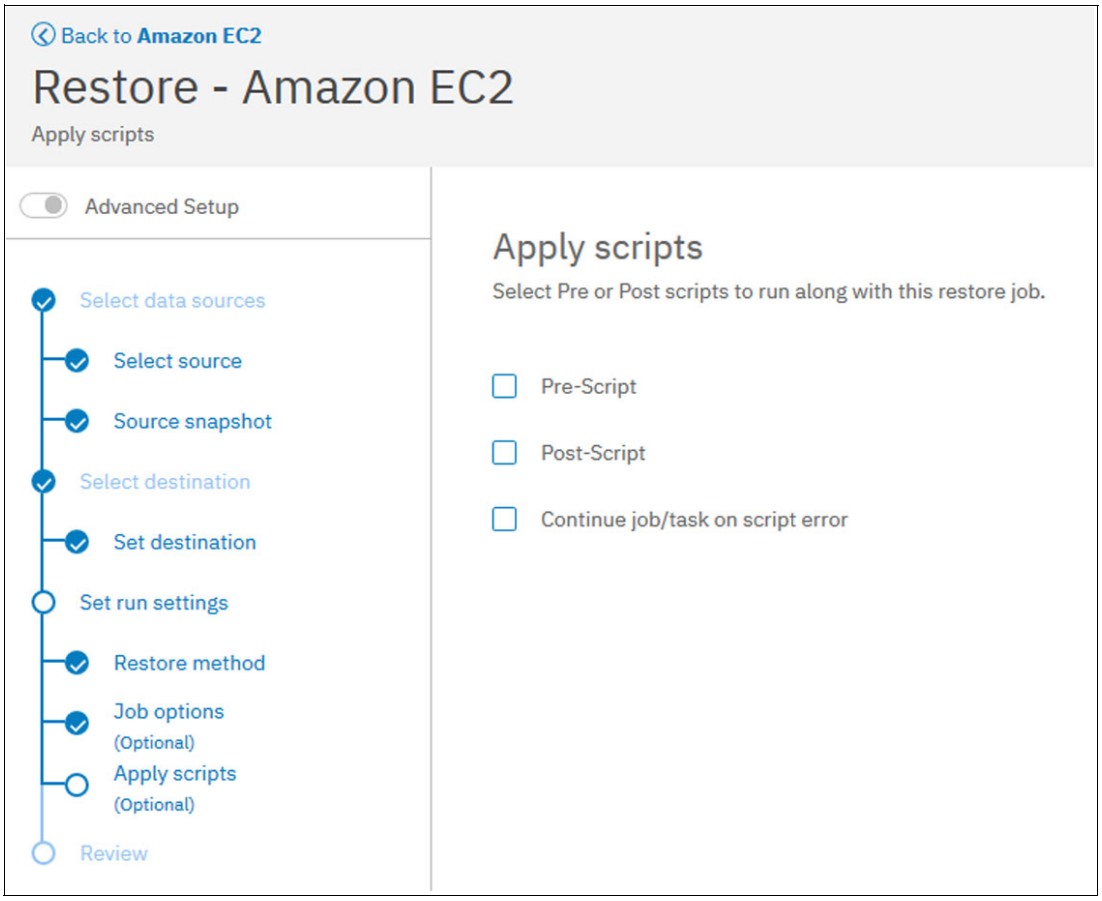

Figure 6-77 Amazon EC2 Restore Advanced Setup mode Scripts options

Calcite **Next** and you are redirected to the Review page. Carefully review the selected Restore job selections, and then, click **Submit** to create the restore job, as shown in Figure 6-78 on page 225.

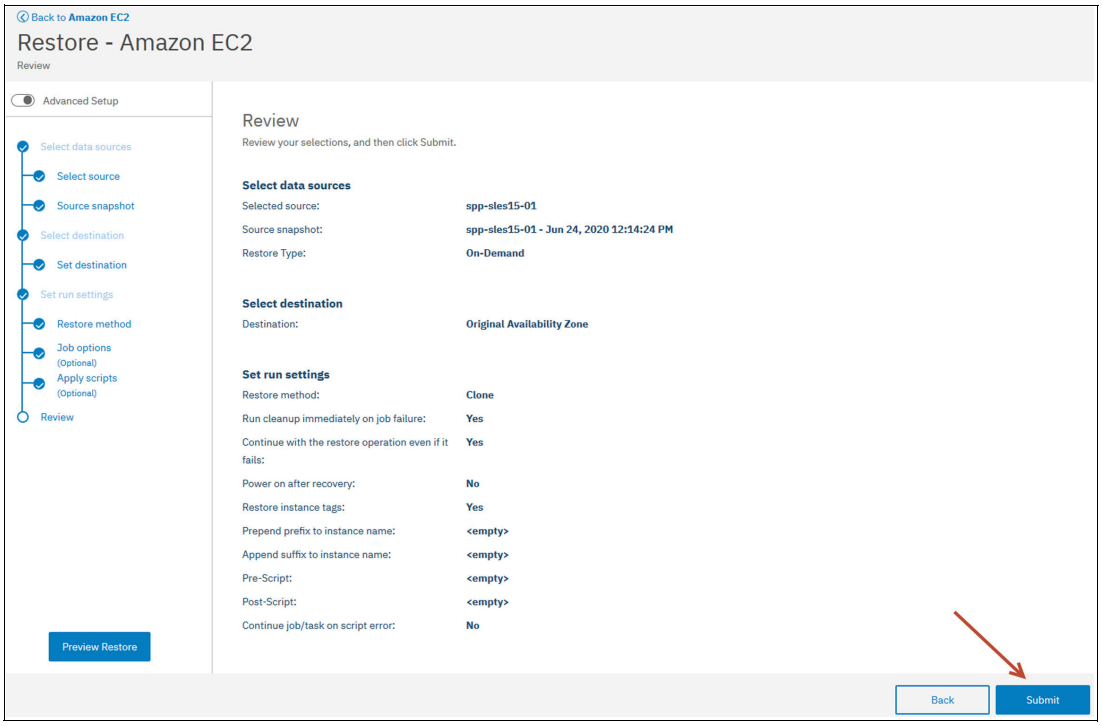

Figure 6-78 Amazon EC2 Restore Review screen

7

Backing up and restoring Windows file system data

Starting with version 10.1.6, IBM Spectrum Protect Plus can protect Microsoft Windows file system data that is stored on physical or virtual servers.

This chapter describes how to back up and restore Microsoft Windows file system data and includes the following topics:

7.1 Supported platforms and browser requirements

For more information about IBM Spectrum Protect Plus Windows file systems backup and restore requirements, see this web page.

The following versions of the Microsoft Windows (64-bit kernel) operating systems are supported:

- Windows 2012 R2 (Standard and Datacenter editions)
- Windows 2016 (Standard and Datacenter editions)
- Windows 2019 (Standard and Datacenter editions)

The following Microsoft Windows file systems are supported as local file systems:

- Microsoft Windows ReFS (IBM Resilient® File System)
- Microsoft NTFS (New Technology File System)

The following browsers are supported to use the File Systems Restore GUI:

- Firefox 55.0.3 and later
- Google Chrome 60.0.3112 and later
- Microsoft Edge 40.15063 and later
- Microsoft Edge HTML 15.15063 and later

7.2 Prerequisites and configuration

Back ups of Windows file systems in Spectrum Protect Plus can be scheduled through assigned SLA policy or manually triggered as ad hoc backups.

The SLA Policy can be created and assigned to the single file systems or to the entire file server.

7.2.1 SLA policy configuration and requirements

Before you start configuring IBM Spectrum Plus for file system backup and restores, all prerequisites must be met. The first prerequisite is to create a new or reuse the existing Service Level Agreement (SLA).

Complete the following steps:

1. From the Spectrum Protect Plus GUI, select **Manage protection** → **Policy Overview**.
2. Scroll down to see defined SLA policies. For File Systems backups, you can create a new or reuse one of the existing SLA policies in agreement with your business needs for retention and backup frequency.
3. Figure 7-1 on page 229 and Figure 7-2 on page 230 show the creation of a new SLA policy. Enter the following information:
   - Policy name and type

     The same SLA backup policy type that was created for file system protection can also be used for VMware, Hyper-V, Microsoft 365, databases, and application (Catalog) protection.

– Replication Policy

Enable the Replication Policy to protect the primary Storage backup data to the secondary vSnap server.

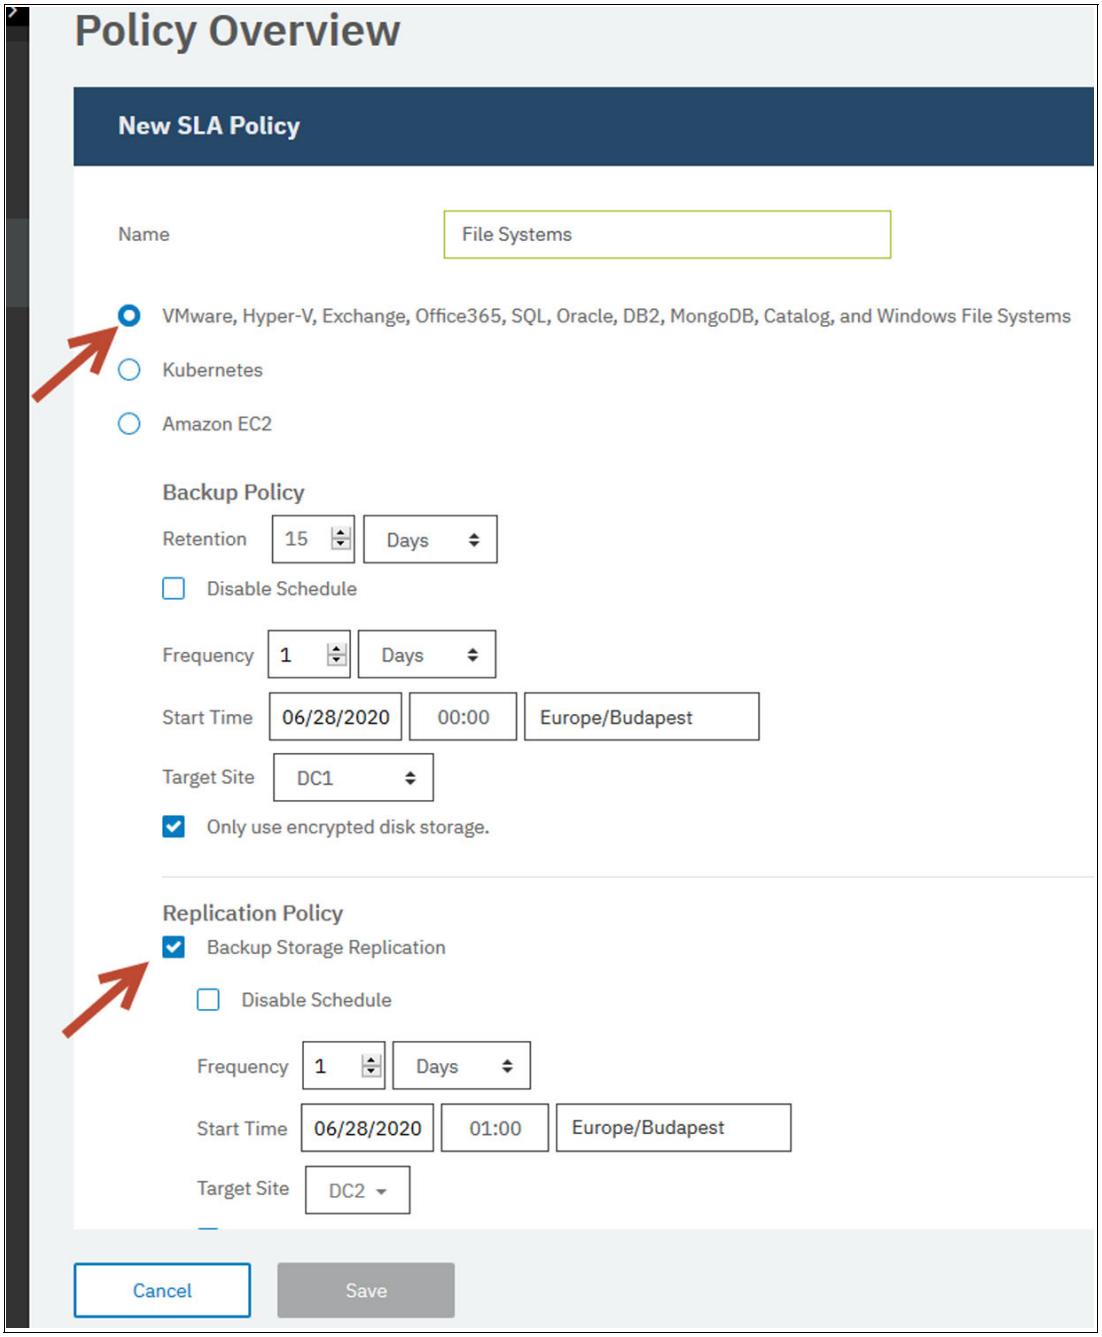

Figure 7-1 Service Level Agreement (SLA) policy for Files Systems

– Additional Copies

By selecting the additional copies option, you can add copies to the repository server, IBM Spectrum Protect server, or the Cloud services (S3).

– Archive Object Storage

By using this option, you can archive a full copy of the data to the cloud service or, as in the example that is shown in Figure 7-2, to the Repository servers to tapes with attached tape Library.

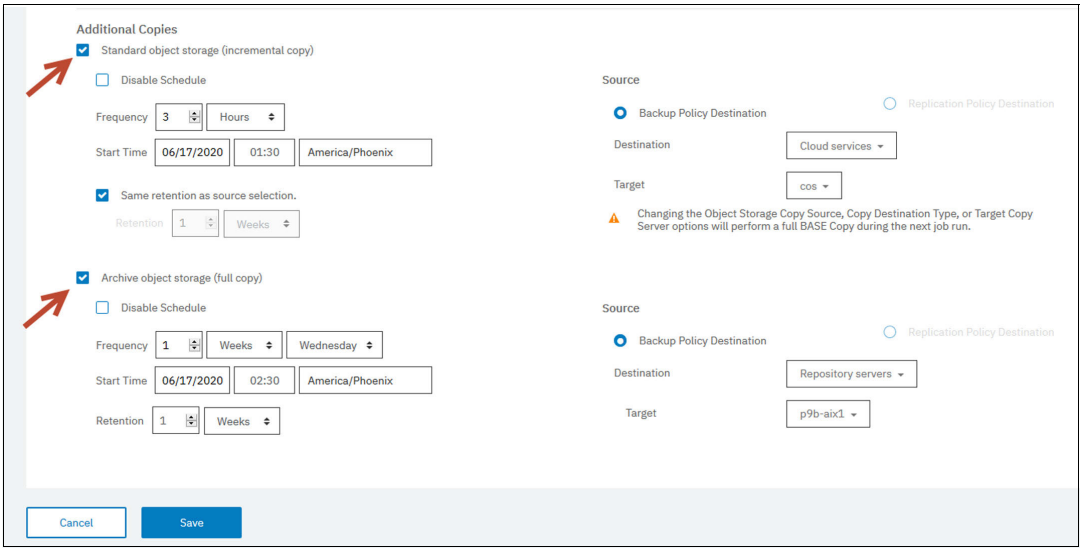

Figure 7-2 Enable Additional copies (incremental copy) and Archive object storage (full copy)

Connectivity requirements

Verify or set up the following connectivity requirements:

- Open firewall ports 5985 and 5986 between IBM Spectrum Protect Plus and Microsoft Windows File Systems server to allow IBM Spectrum Protect Plus server to connect to the file systems server by using the Microsoft Windows Remote Management (WinRM).
- Open firewall port 9085 to enable the IBM Spectrum Protect Plus File Systems Restore UI to connect to the restore service.
- On the file server, the WinRM service must be set up and running.

For more information about connectivity and other requirements, see the following resources:

- Windows file systems backup and restore requirements: IBM Spectrum Protect Plus V10.1.6
- System requirements: IBM Spectrum Protect Plus V10.1.6
- IBM Knowledge Center

The architecture overview that is shown in Figure 7-3 on page 231 shows the communication flow and interactions between components that are needed for the file systems restore and firewall ports that must be opened for communication. The following main components also are shown:

- IBM Spectrum Protect plus server
- IBM Spectrum Protect Plus Browser Restore User Interface
- IBM Spectrum Protect Plus vSnap server with metadata database
- Microsoft File server with file systems

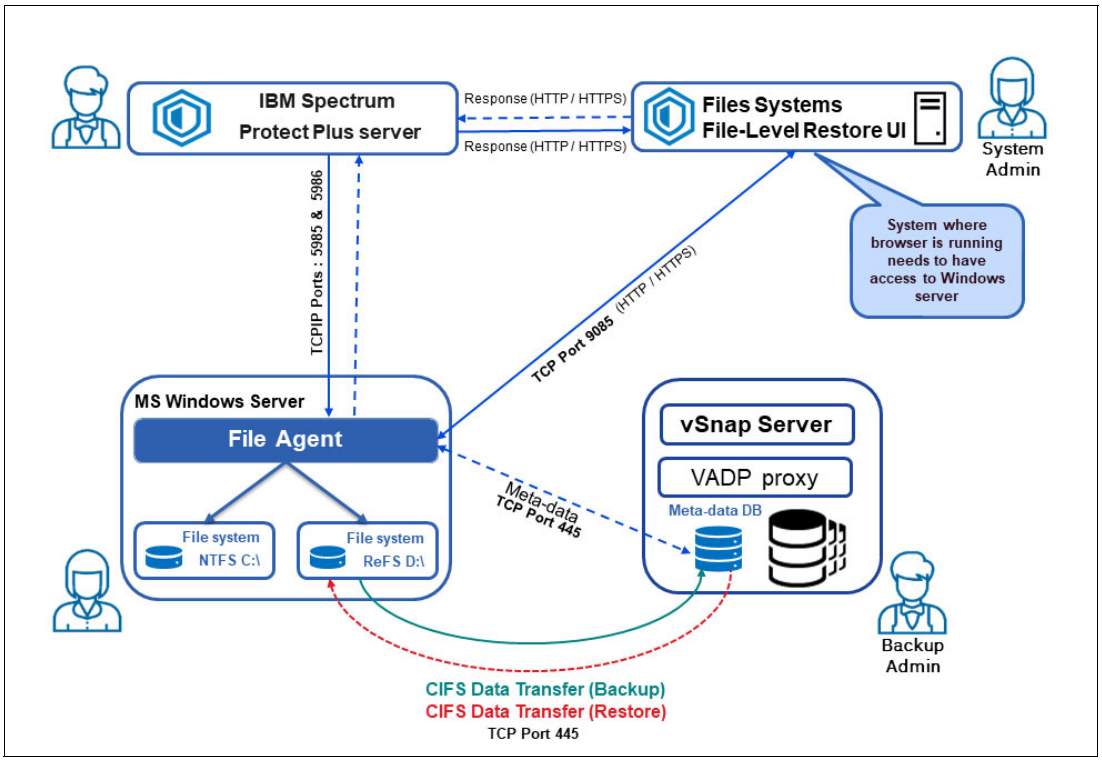

Figure 7-3 File system architecture overview

7.2.2 Microsoft Windows File Systems backup configuration

This section describes how to configure Microsoft Windows File Systems backup.

From the IBM Spectrum Protect Plus GUI, select, **Manage Protection** → **File Systems** → **Microsoft Windows**.

To add a system, click **Managed file servers**, as shown in Figure 7-4.

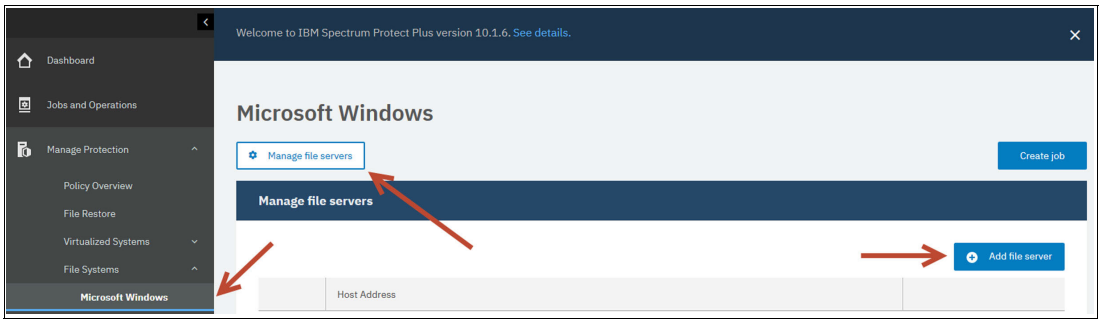

Figure 7-4 Add New Microsoft Windows File server

Next, click **Add new file server,** and the window that is shown in Figure 7-5 on page 233, is displayed. Enter more configuration details under the Edit file server properties section.

Specify the following information:

- Host Address

  Enter the Domain Name System (DNS), Hostname or the IP address of the file server.

Important: If DNS Names is used, ensure that the following conditions are met:

- DNS names must be resolvable by IBM Spectrum Protect Plus and vSnap servers.
- DNS names must also be resolvable by the host from which the File System Restore UI is used.
- All IBM Spectrum Plus configuration components must be resolvable by their DNS names.

- Use the existing user

  This option allows you to select the identity from the drop-down menu that you registered before. If a user does not exist, leave the selection empty to specify a User ID and Password.

- User ID

  Select the Local or Active Directory Domain administrator ID

  When entering the User ID, it is not necessary to specify the domain.

  The following options are available for the user ID for registering the file servers:

  – The Local System Administrator user account with the User Account Control (UAC) security component disabled.

    With this user, you must access the User Account Control Settings dialog in your Windows system control panel, and move the slider to **Never**.

  – A user who is a member of the Local Administrator Group with the Admin Approval Mode security policy setting disabled.

    With this user you must access the Local Security Settings dialog on your Windows system and disable the **User Account Control: Run all administrators in Admin Approval Mode policy** setting.

    Ensure that your Local Administrator Group includes the Log on as Service policy option

- Password

  Enter Password of the File Server

- Maximum parallel file systems

  Enter the maximum number of the file systems. The default value is set to use a maximum of 10 parallel mounts for each file system. This value is adjustable. If the file server configuration has more than 10 file systems resources, it is best to increase the default value to accommodate a new parallel process to speed up the backup and restore operations.

After completing the Server Properties configuration fields, click **Save** to complete the registration of the file server configuration, as shown in Figure 7-5.

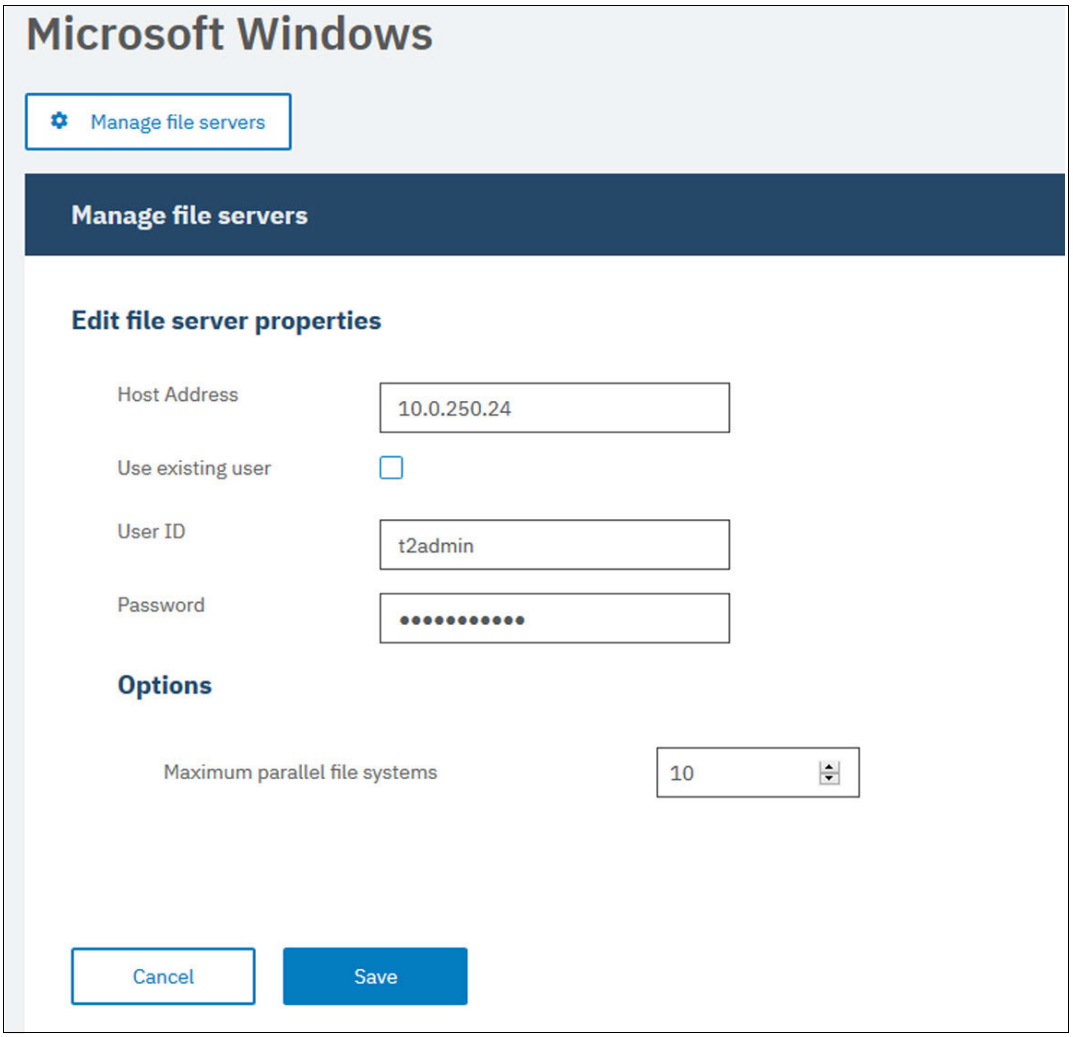

Figure 7-5 Add File Server configuration detail properties

As part of the file server registration process, the application agent is pushed from the IBM Spectrum Protect Plus server to the file server and other I Microsoft Windows services that are needed for the IBM Spectrum Protect Plus file system protection agent are set up.

After completion of the file server registration, the initial inventory process is automatically started and discovers all of the file systems (all drives and the directory structure, including files). You can validate and monitor Inventory process from the Jobs and Operation tab, as shown in Figure 7-6.

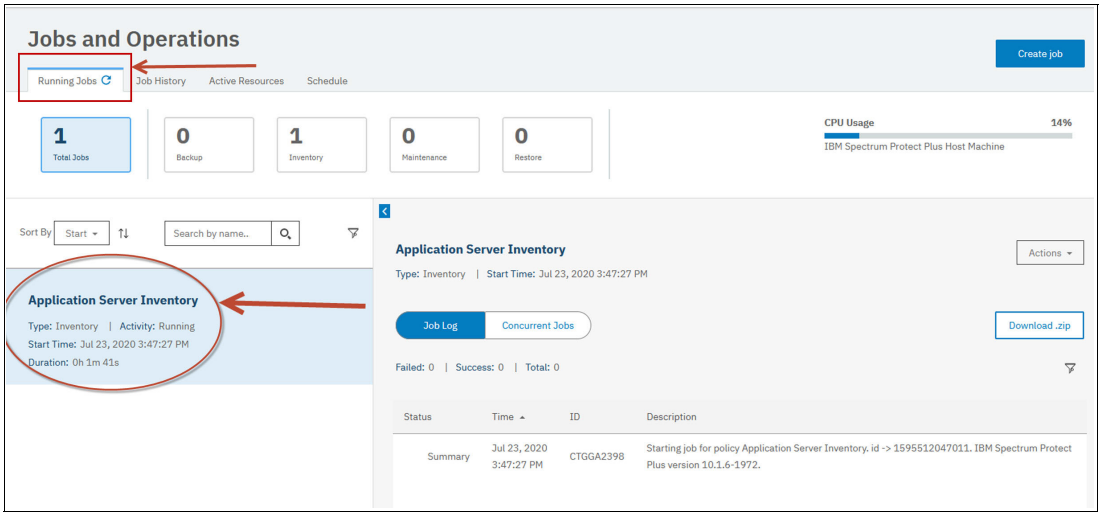

Figure 7-6 Progress of the Application Server Inventory

In the **Manage file servers** tab, select file server. From **Actions** menu, you can start the inventory manually or select the **Test** action, as shown on Figure 7-7.

.

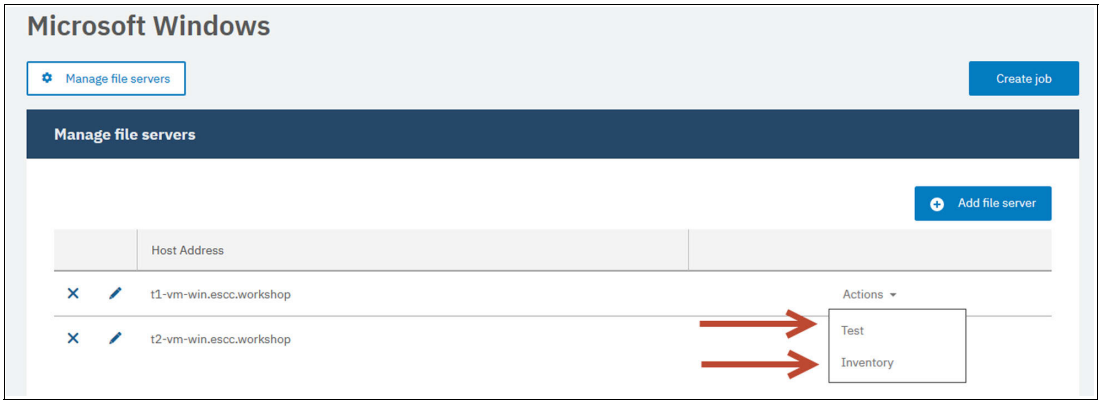

Figure 7-7 File system action menu: Inventory

The Test option consists of the following test scenarios:

- Network Connection test

  This test also includes a Socket Connection test to verify whether an open connection to access port 5980 exists on the Windows Server. An example of this test result is shown in Figure 7-8 on page 235.

- Remote Session test

  This test includes the following components:

  – Windows Remote session test to validate agent deployment, configuration, WinRM, and Required Administrator Privileges.

  – Remote Agent Execute test to validate the credentials and required Administrator's rights.

- Pre-Requisites test

  This test validates the local administrator's privileges and HTTPS connection to the IBM Spectrum Protect Plus appliance.

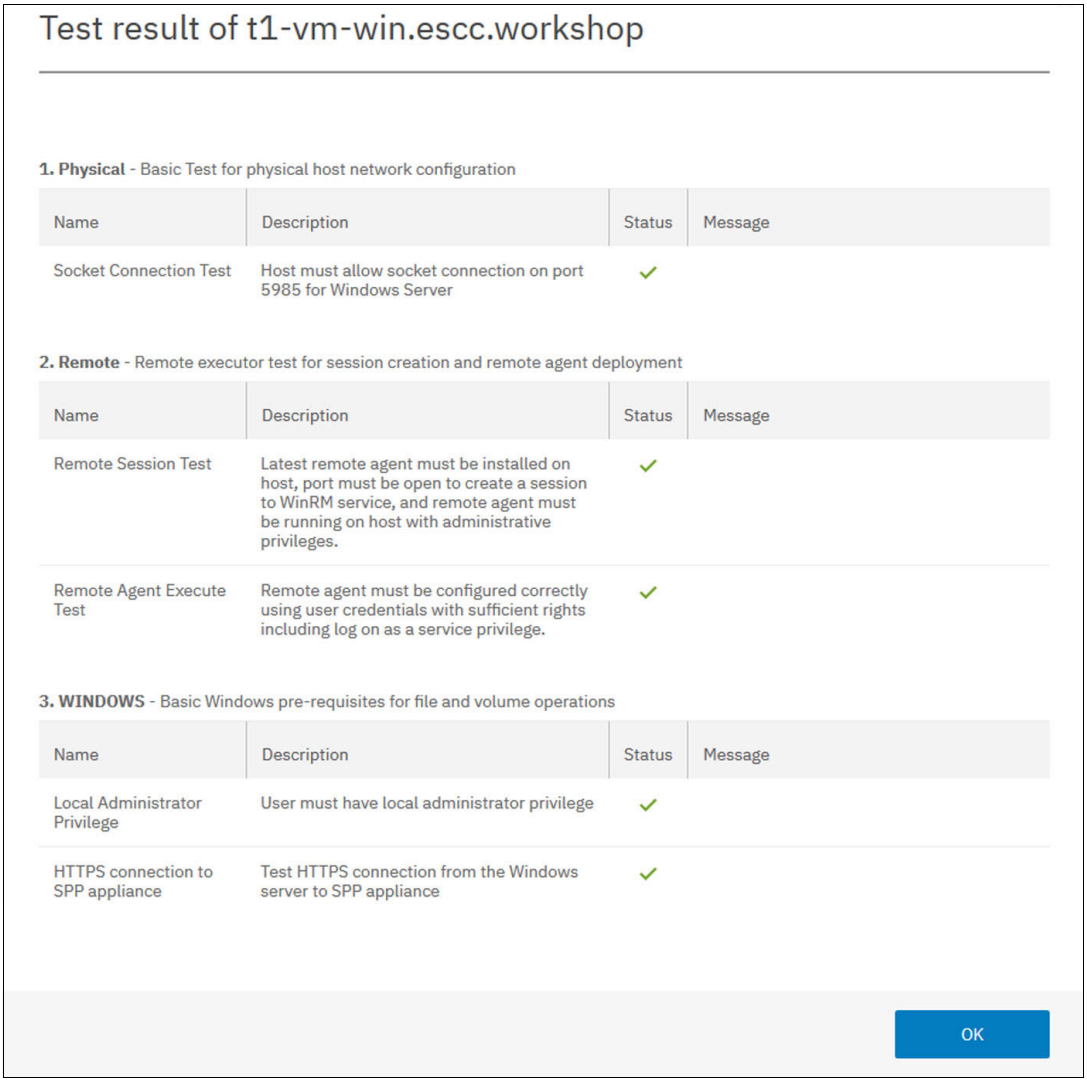

Figure 7-8 Example of the test results check of the file server

Note: On the same system, a configured backup agent cannot be used for file systems back ups and IBM Spectrum Protect Plus application backup protection at the same time.

The Microsoft Windows tab shows an example of the successfully registered file servers, as shown on Figure 7-9. Select **Run Inventory** to automatically start an inventory for all registered Microsoft Windows File Servers.

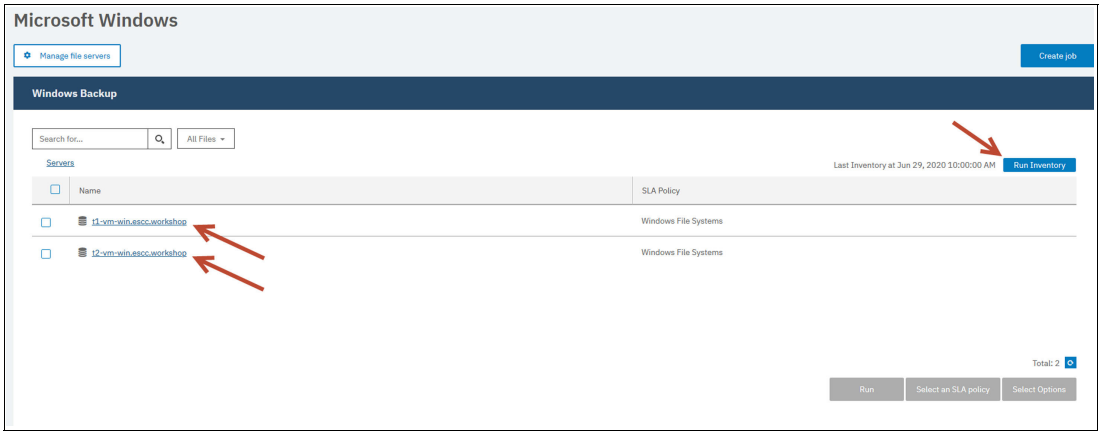

Figure 7-9 Successfully registered File Servers

To view the output of the inventory, select **Jobs and Operations** → **Job History**, as shown in Figure 7-10.

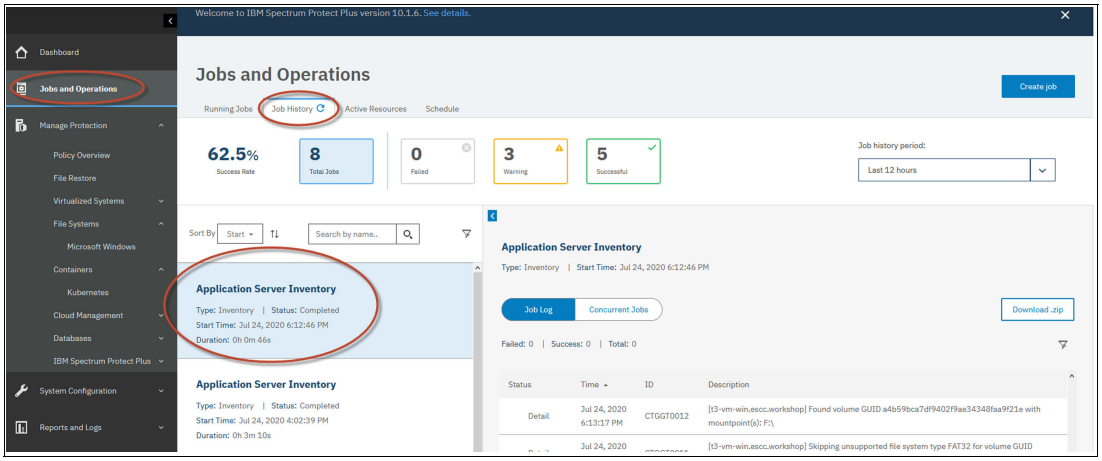

Figure 7-10 Example History of the job log inventory

Limitation: After successful registration, the file server's drives might not appear under the registered file server. If such a server includes an assigned SLA policy during the backup process, the Assigned SLA policy is removed and warning messages are shown in the job log. Such behavior is described and documented when the VM file server instance was created from a template or cloned from other VMs.

If you are experiencing similar issues on VMmware, you can validate it by running following commands:

- To identify the VMs with the same UUID, run the following command at the CLI by way of SSH to verify the UUID:

  ```
  vmkfstools -J getuuid <vmname>.vmdk
  ```

- If you identified the VMs with the same UUID, run the following command to change the ddb UUID:

  ```
  vmkfstools -J setuuid <vmname>.vmdk
  ```

- On the operating system, you can open PowerShell with admin rights and run the following command. You do not need to modify it. The command returns a list of all volumes and their GUIDs:

  ```
  GWMI -namespace root\cimv2 -class win32_volume | FL -property DriveLetter,
  DeviceID
  ```

For more information about IBM Spectrum Protect Plus known issues and limitations in the file system application server, see this web page.

7.2.3 Exclude syntax rules for the file systems

It is possible to exclude specific drives, directories. or files from file systems backups. The exclude list rules are applied during backup and restore operation. The restore process does not restore the excluded drives, directories, or files that are listed as part of the exclude rules.

To create new or manage existing exclude lists, from the IBM Spectrum Protect GUI select: **Manage protection** → **File systems** → **Microsoft Windows.** The Windows Backup window is displayed.

Select one or more registered file servers, and click **Select Options**, which opens another view in which you can select **Modify Exclude List** for the selected Microsoft file server. The exclusion list is processed from the top-down based on the definition statement. Update the Exclude list with excluded syntax and click **Save** to submit the changes.

The exclude rules syntax supports various wildcards, characters, and patterns.

The exclude rules can be set to apply to all file servers and their file systems. You also can specifically customize the exclude list for each of the file servers and each file system separately, as shown in the example in Figure 7-11 on page 238.

Complete the following steps:

1. Select one or multiple file servers.
2. Select **Options**.
3. Select **Modify Excluded Files**.
4. Define Exclude rule syntax.

For more information about the exclude rules syntax, see IBM Knowledge Center.

Note: By the default, the following files are excluded from all drives:

- `DIR ?:\$RECYCLE.BIN`
- `DIR ?:\System Volume Information`
- `DIR ?:\Users\*\AppData\Local\FSPA`
- `DIR ?:\ProgramData\SPP`
- `DIR ?:\Program Files\IBM\IBM Spectrum Protect Plus`

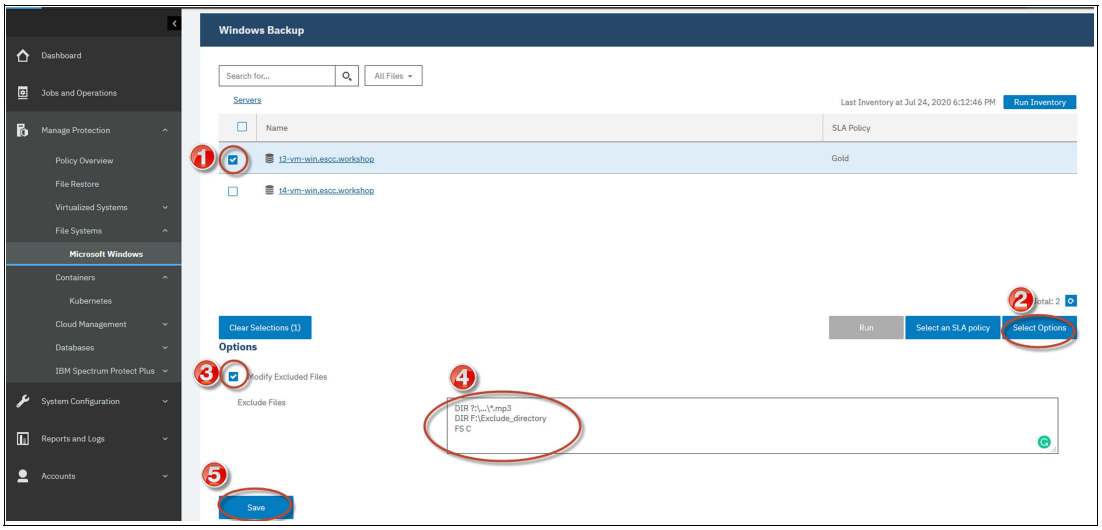

Figure 7-11 Exclude example statement for file server

Important: As shown on Figure 7-12 on page 239, file system types that are not supported, such as FAT32 or GPFS, are grayed out. These types cannot be assigned to the backup SLA policy.

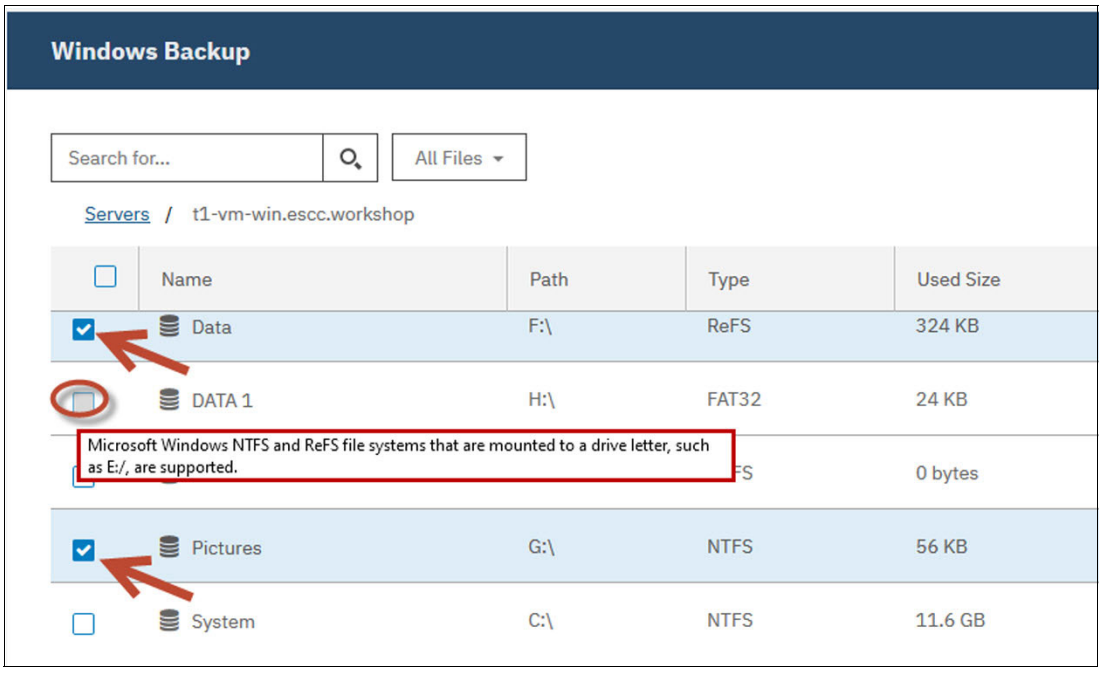

Figure 7-12 Exclude syntax rule for selected file systems drives

The Spectrum Protect Plus File System Data Protection is checking the Registry Key of the file server and automatically excluding from backup and recovery actions, all the files and directories present in the folder **Files Not to Backup.**

Note: As shown in Figure 7-13 on page 239, the Microsoft Windows registry key information for files not to be backup can be found in the following directory:

`SYSTEM\CurrentControlSet\Control\BackupRestore\FilesNotToBackup`

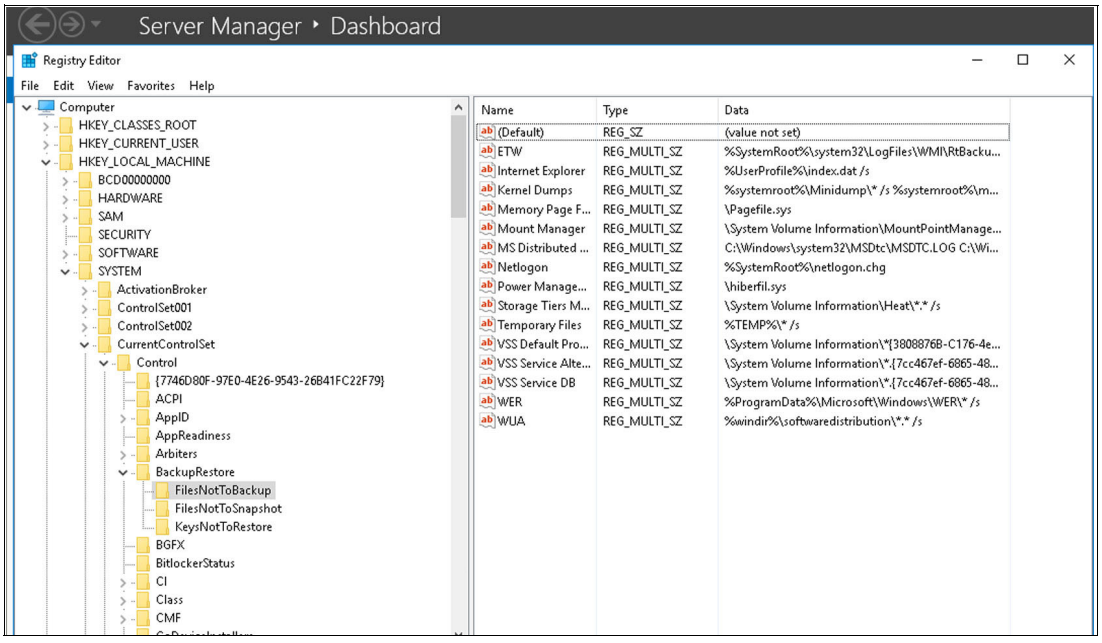

Figure 7-13 Example of the Microsoft Windows registry key

7.3 File systems backup with IBM Spectrum Protect Plus

This section describes how to enable backups for file systems with or without exclude rules.

In the **Windows Backup** window, proceed with the backup configuration by selecting a Microsoft Windows file server. Then, click **Select an SLA policy**, as shown in Figure 7-14. Select the SLA Policy and then, click **Save**.

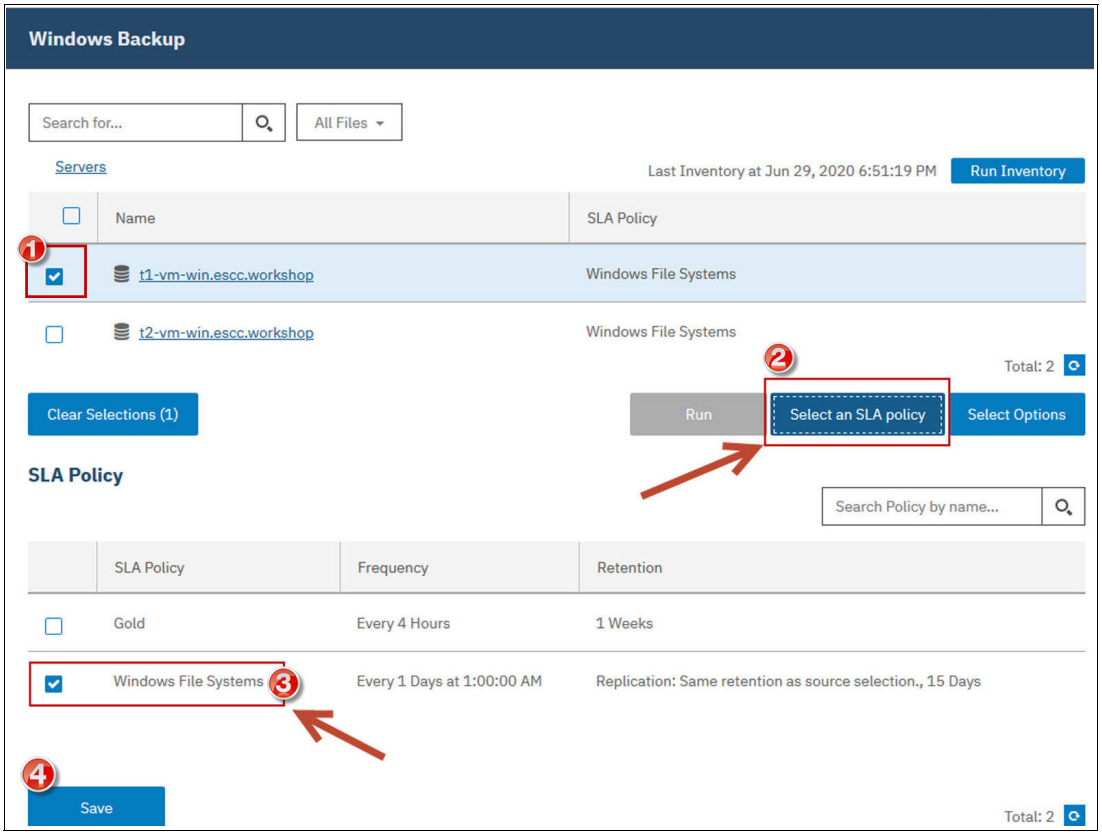

Figure 7-14 Select and assign SLA policy to the file server

In the **SLA Policy Status** tab, select **Action** → **Manage Protection** → **File Systems** → **Microsoft Windows**. The backup process starts for the all systems with the selected assigned SLA policy (see Figure 7-15).

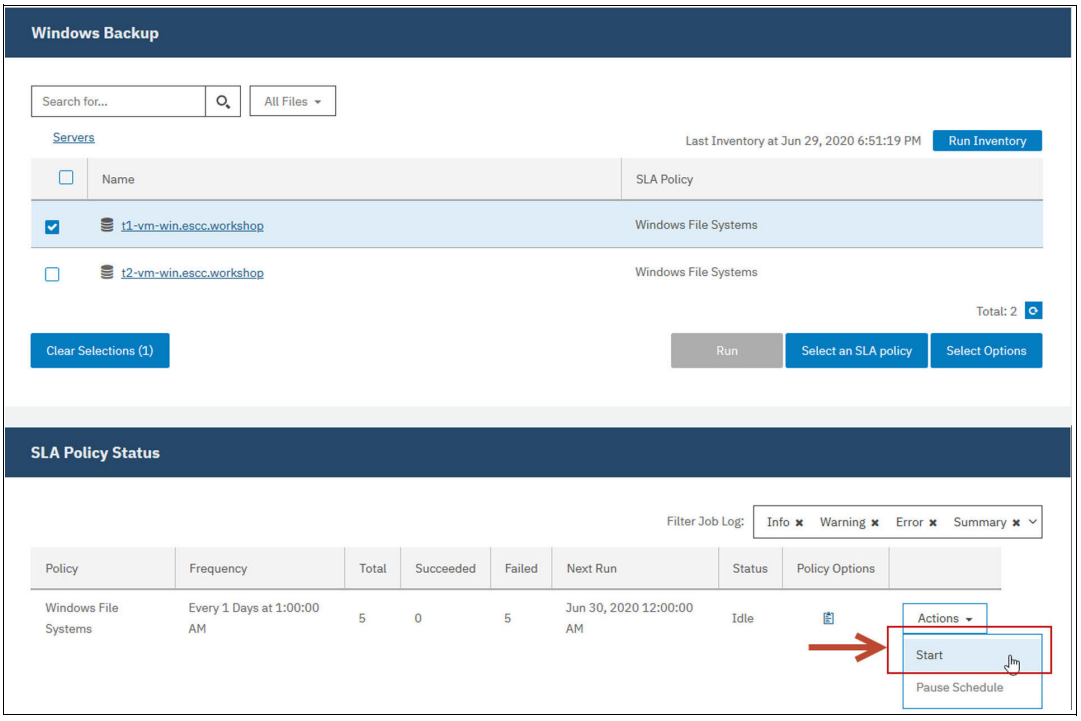

Figure 7-15 Manual Start of the file systems backup

On-demand backups can be started individually for selected file system separately under file server by selecting the file system drive and clicking **Run**, as shown Figure 7-16.

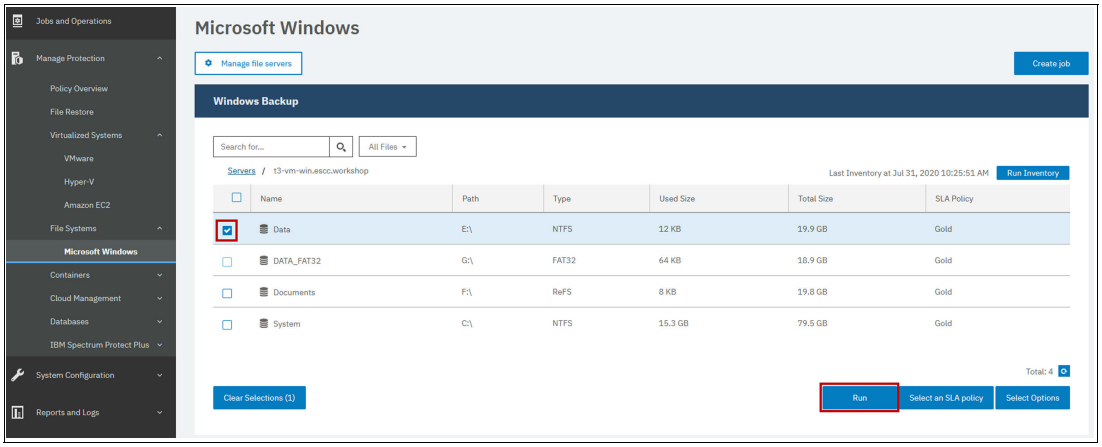

Figure 7-16 Run backup for single file system

If the SLA policy is configured with the replication option, clicking **Start** (as shown in Figure 7-15 on page 241) opens the **Start Options** window in which you can select the **Backup to vSnap** option (as shown on Figure 7-17 on page 242) or **Replicate** to a secondary vSnap server.

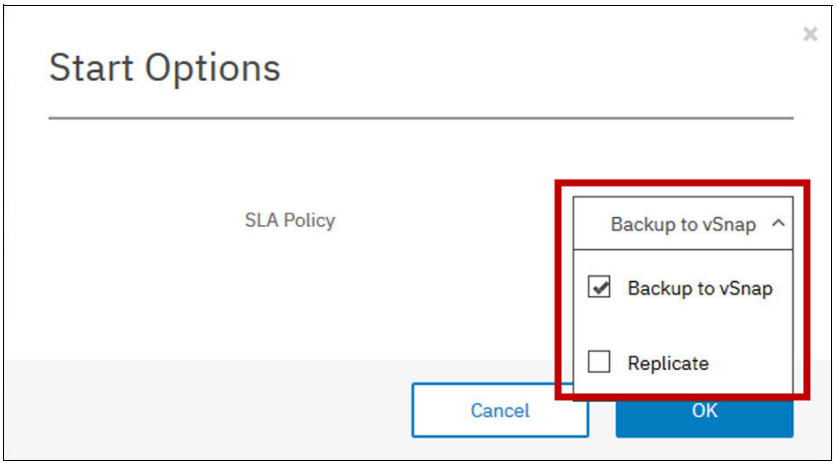

Figure 7-17 Select backup to vSnap

Select the **Backup to vSnap** option and click **OK** to confirm. The backup starts immediately in accordance with the selected SLA policy.

The first file system backup is always a full backup. All subsequent backups are incremental forever backups. During the backup operation, CIFS shares from the vSnap server are mounted to the file system. Mounted CIFS shared activity that is started in background is not visible to the Administrator user.

During the first full backup operation, a scan of the file systems is performed and the result of the scan is stored in the newly created database that contains metadata information for each of the files. For any subsequent backup, files are compared against metadata information; that is, if files were deleted, they are marked as expired. If files were created, they are included in the scope of the backup task.

In the following action, SPP takes a consistent snapshot of all data and metadata, which is saved as point in time for the backup recovery.

For subsequent backups, the same CIFS file shares are mounted and the database is updated with new metadata entries, including the files that were created or changed. After the operation is completed, a new snapshot is taken.

The metadata database holds information about each individual file's attributes, such as size, security, and reparse information. The vSnap server contains only a copy of the files content. Files without content (such as symbolic links), do not have a copy on the vSnap server; those files are included as empty directories.

Progress of the running backups operation can be monitored from the following places:

- SLA policy Status view, as shown on Figure 7-18.

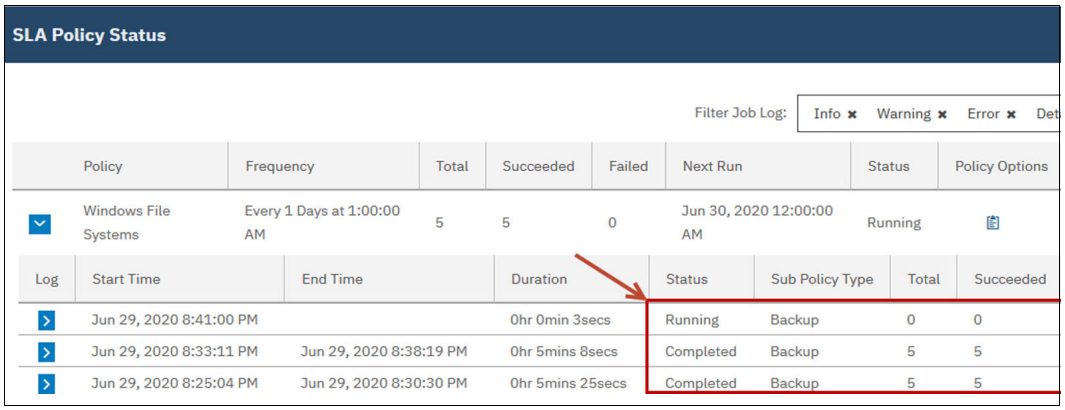

Figure 7-18 SLA Policy Status overview about backup progress

Note: In the status from the job log view, a message informs the administrators about non-supported file systems that were skipped, as shown on Figure 7-19.

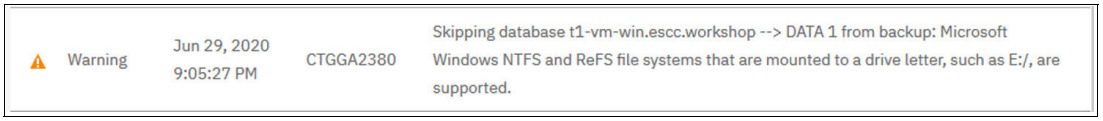

Figure 7-19 Warning message about not supported file systems types

- In the Jobs and Operation tab by selecting the **Running Jobs** status, as shown in Figure 7-20. The Overall Progress column shows all active file systems backups.

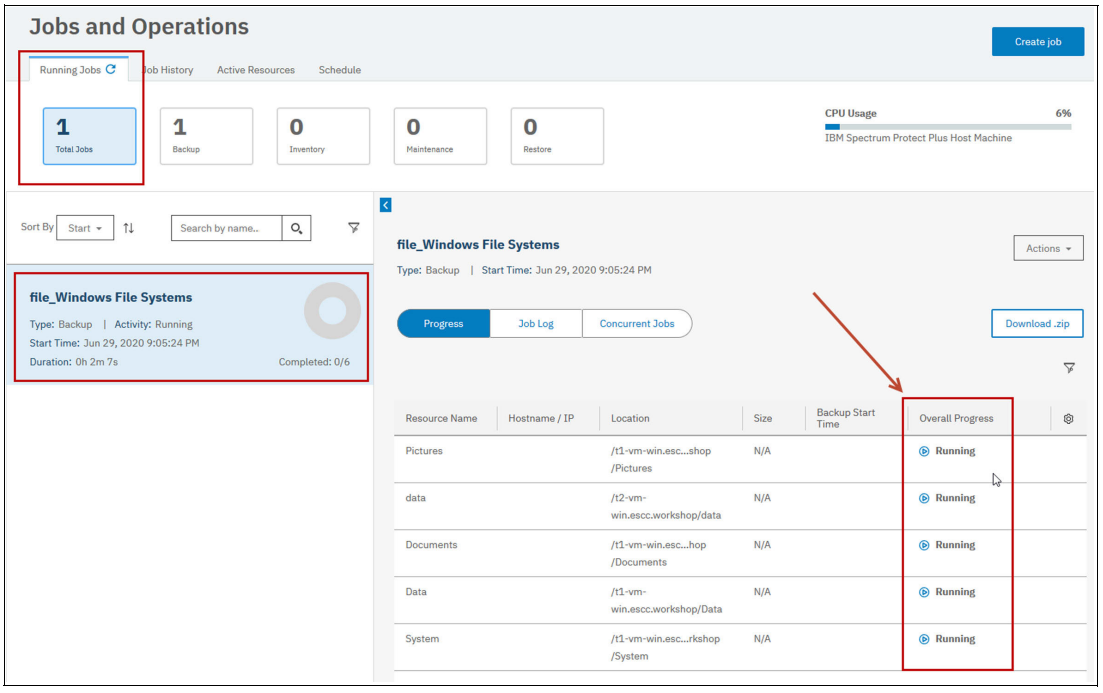

Figure 7-20 Detailed Jobs and operation overview

Drive letters that are changed after full backup: Avoid making any change in the file system path name or drive letter on the Windows server between the full backup and the following incremental backups. Such changes can result in failed incremental backups.

If you want continue sequences of the Incremental forever backups, you must revert the file system path name or the drive letter on the Windows platform back to the original.

Alternatively, you can run a new full backup by removing the SLA policy from the file system and reassign it again by rerunning the backup job operation.

7.3.1 Microsoft Windows files systems ad hoc backup

The following options are available to start an ad hoc backup:

- Select **Jobs and operations** → **Create job** → **Ad hoc Backup**. Then, select **File Systems (Windows),** as shown on Figure 7-21. Proceed with the ad hoc backup configuration steps, as shown in Figure 7-23.

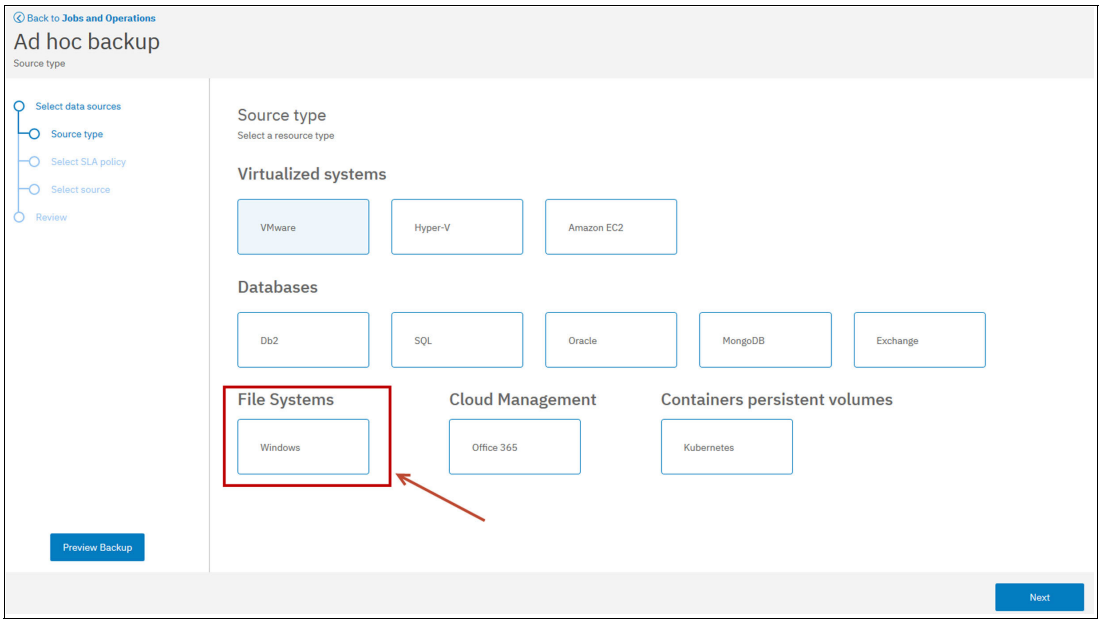

Figure 7-21 How to create Ad hoc backup from Job and operation window

- Select **Manage Protection** → **File Systems** → **Microsoft Windows**. Then, select **Ad hoc backup** as shown in Figure 7-22.

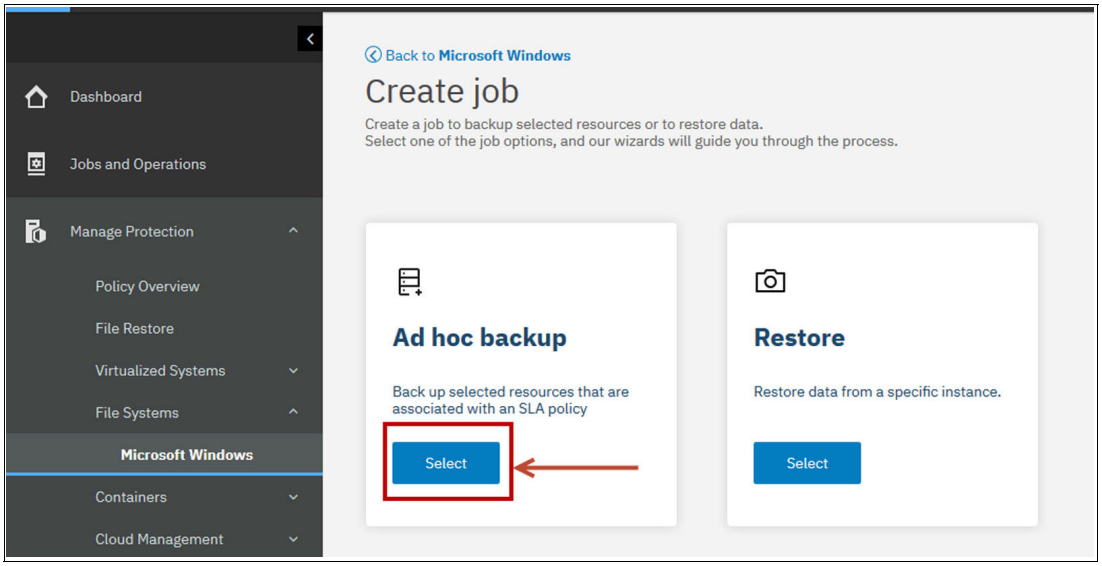

Figure 7-22 Create a job with Ad hoc backup

Using either option opens the same window that is shown in Figure 7-23 in which you can configure the ad hoc backup.

To configure the ad hoc backup, scroll down and click **Select SLA policy**. In our example, we choose the SLA Policy called **Windows File Systems**, as shown in Figure 7-23. Click **Next** to continue.

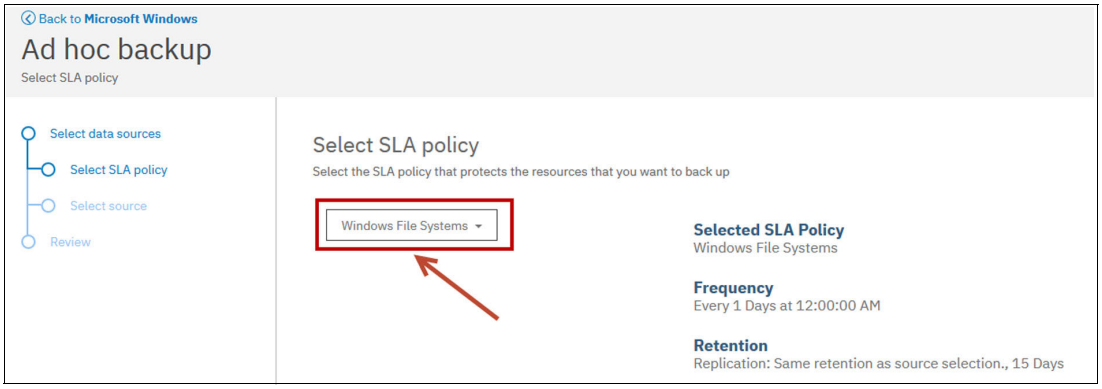

Figure 7-23 Ad hoc backup select available policy

The next window (see Figure 7-24) presents you with the add option [+] to individually select each file system to be added to the ad hoc backup job. After you click the add sign, your selected file system is added to the Item list that is on the right side of the window.

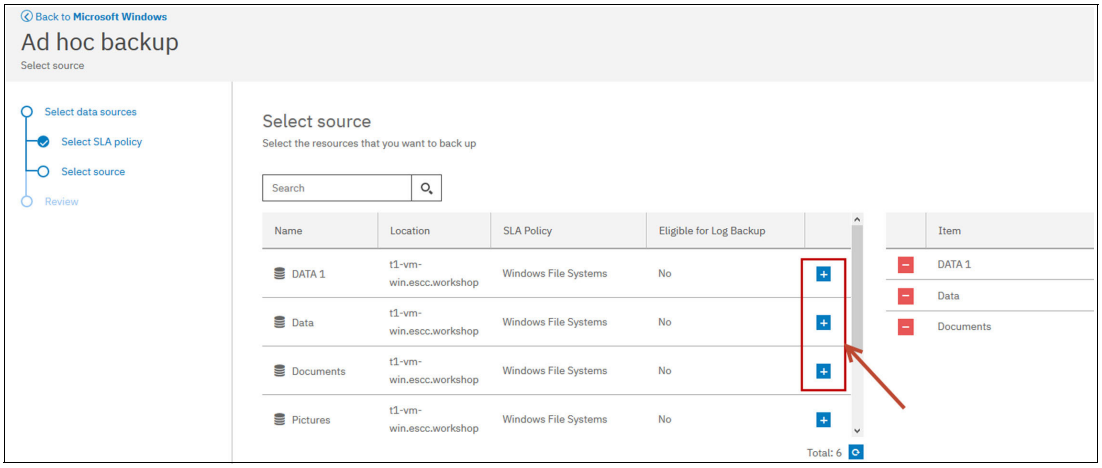

Figure 7-24 Select file systems for Ad hoc backup

When your selection is completed, click **Next** to be presented with a summary view of the ad hoc backup operation, as shown in Figure 7-25.

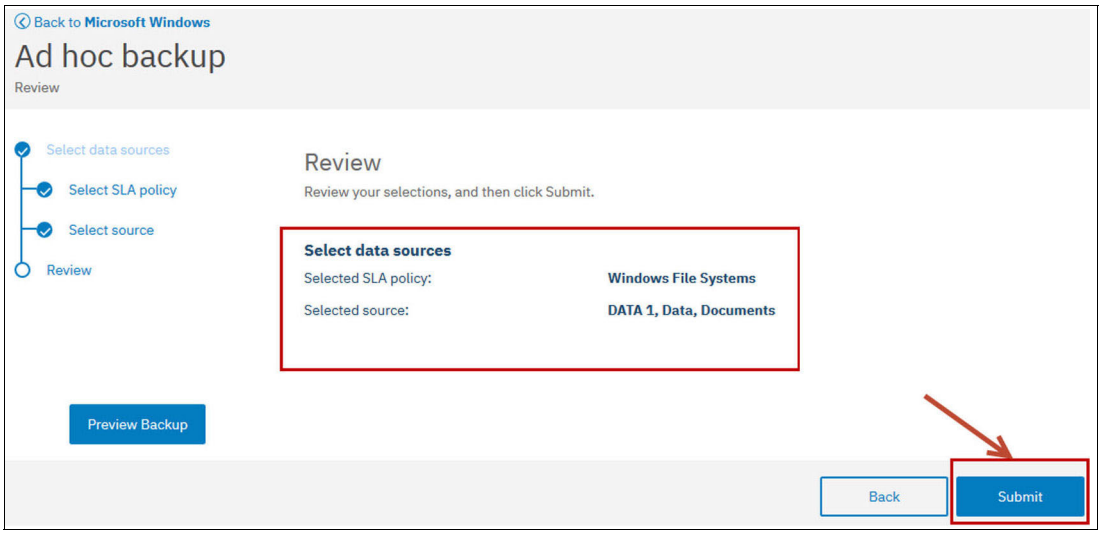

Figure 7-25 Review and Submit Ad hoc backup process

Review the summary and then, click **Submit** to start the backup. You are redirected to the Jobs and Operation window view, as shown in Figure 7-20. In this view, you can monitor the progress of the backup process.

7.4 File systems restore with IBM Spectrum Protect Plus

This section describes how to restore the Microsoft Windows server file systems. To start the restore operation, navigate in the GUI as described in7.3, “File systems backup with IBM Spectrum Protect Plus” on page 240 until you reach the window that is shown in Figure 7-26. Select the **Restore** option to open the restore wizard.

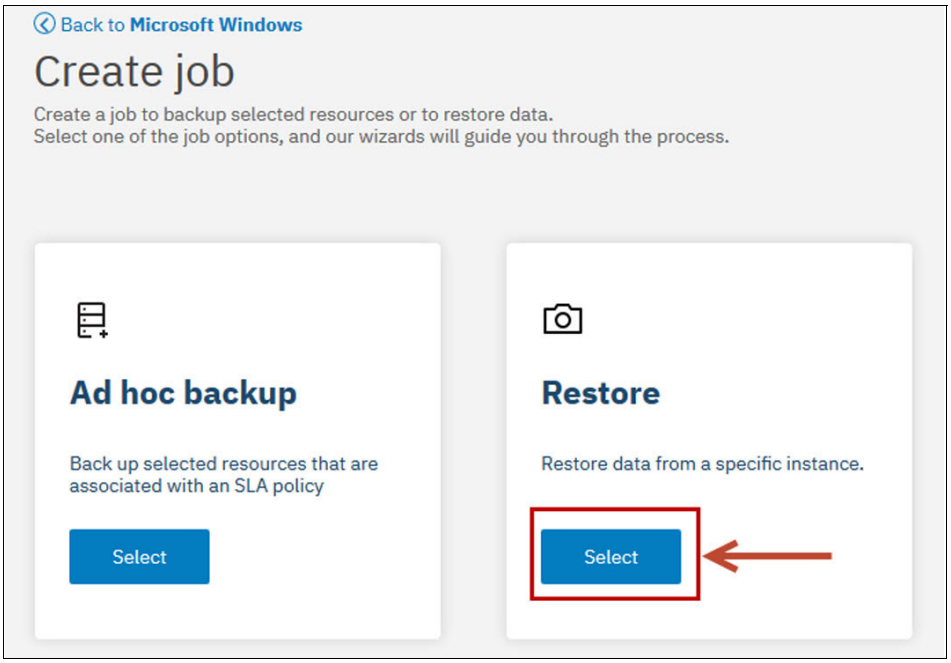

Figure 7-26 Create job and Select Restore operation

The restoring file systems is done in two steps:

1. Restore sequence

   Create a Job and start the restore operation for the selected file server and file system. This action prepares a file system mount that becomes available for the next recovery processes.

2. Recovery process of the files and directories

   In this step, you access the restore UI to specify the location where to restore the files and directories. That location can be same location as the original files or an alternative location.

These steps are described next.

7.4.1 Step 1: Restore sequence of the file systems recovery process

The Restore wizard can be opened in default mode or in advanced mode. Using the default process in the Select Source option, select the source file server by clicking its host name. Click the plus icon to select a specific file system to restore. Your selection is added to the Item list on right side of the window, as shown on Figure 7-27 on page 248. Continue to **Source snapshot** window and click **Next**.

Note: For the restore operation, you can select only one file system at a time.

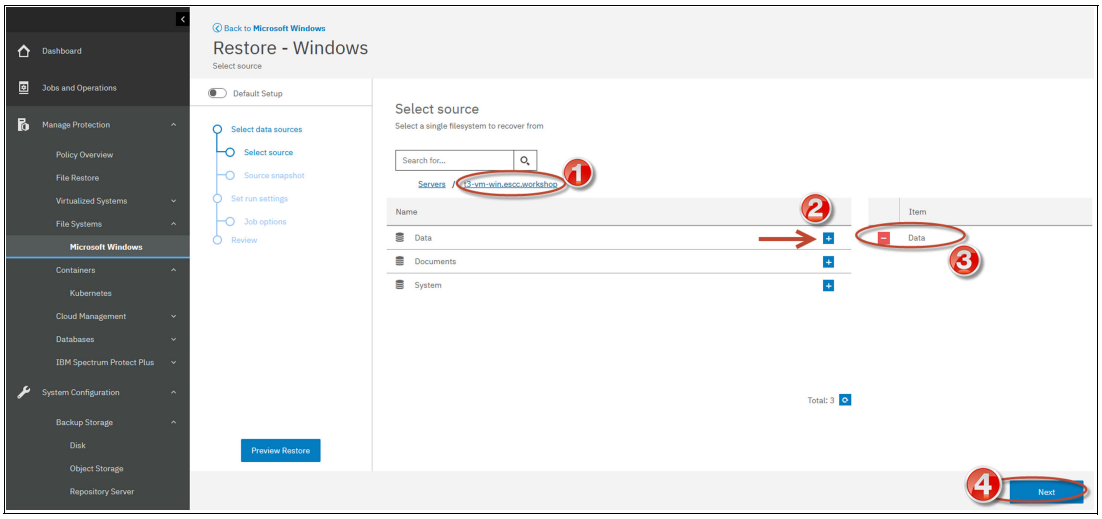

Figure 7-27 Select sources server and file system for recovery

Using the search field allows you to locate a specific file system to restore by entering the name of the file system, such as “data”, as shown on Figure 7-28. or partial string information like “dat”. The search results display all file systems that match the specified search string, including host names of the file servers.

Figure 7-28 Use Search option to display file systems for restore operation

By clicking the date, as shown on Figure 7-29, a new window opens in which you specify the date range of the recovery point you want to use. These windows provide all of the recovery points dates that are available for the selected file system.

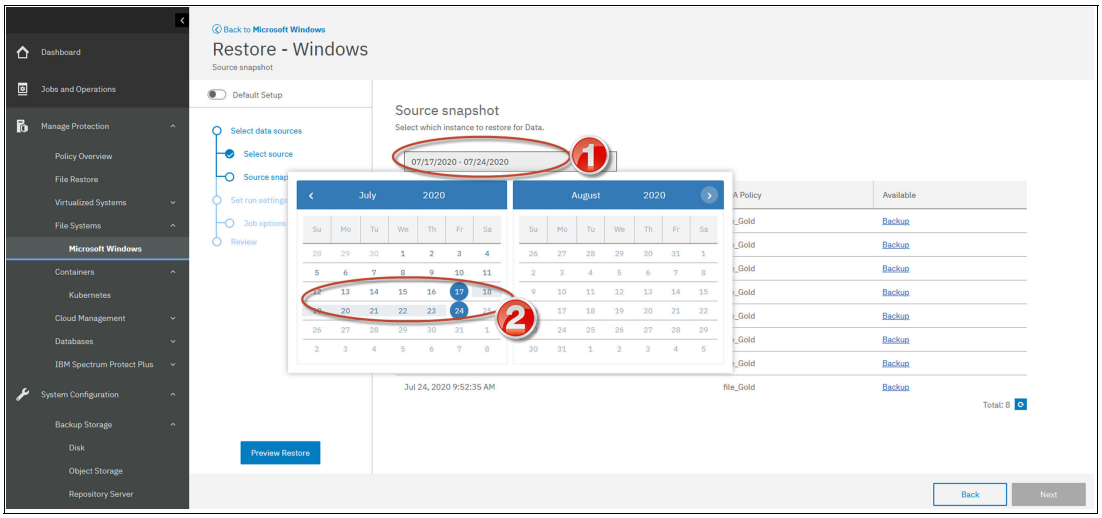

Figure 7-29 File systems Recovery point overview

After selecting the date range, the window closes and returns you to the selection window in which you select a backup recovery point. The backup is highlighted for the restoration operation that was taken by a snapshot, as shown on Figure 7-29. Click **Next.**

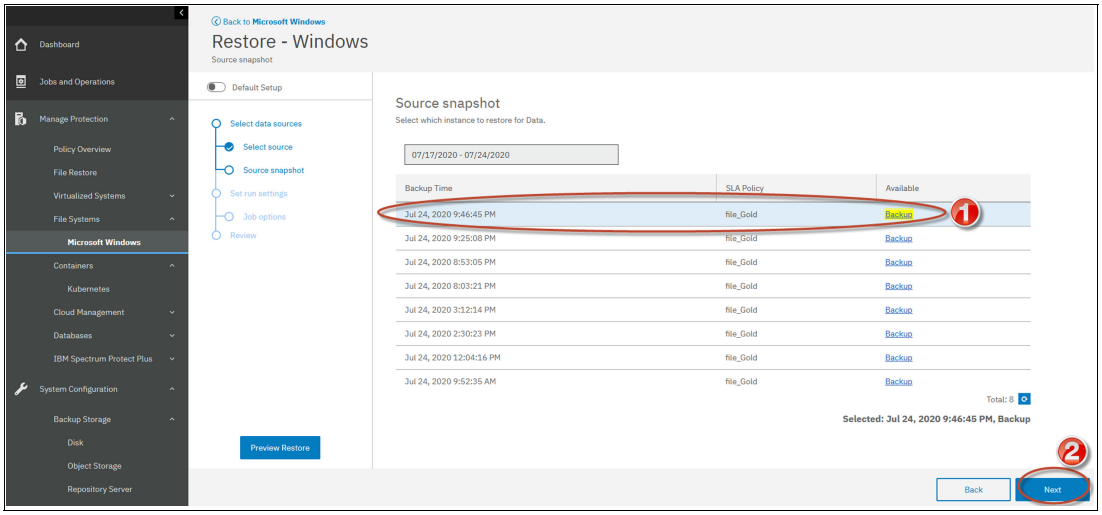

Figure 7-30 File systems Recovery point selection

Tip: This option is the default and it is best to keep it selected if issues occur when mounting a file share for the restore operations. This option allows the system to automatically clean up any left over mounted file shares.

In the next window, as shown on Figure 7-31 on page 250, the Run cleanup immediately on job failure option is available.

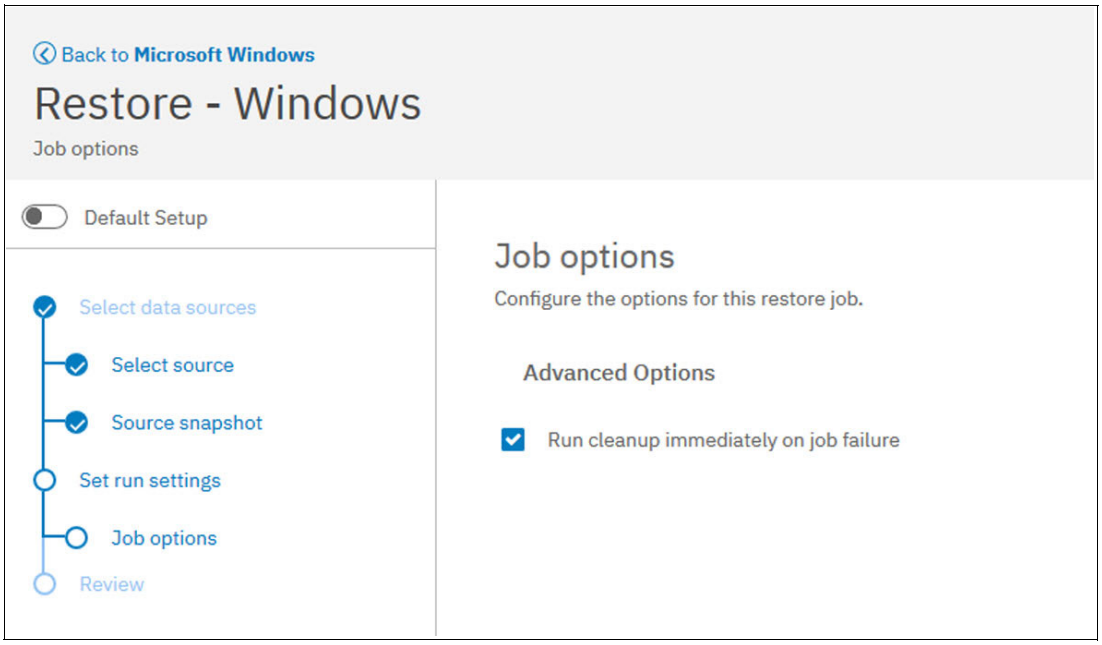

Figure 7-31 Select option to run automatic cleanup of the mount share

Before submitting the restore process operation, review the details of the selected sources and snapshot date, as shown on Figure 7-32. Start the restore by clicking **Submit**.

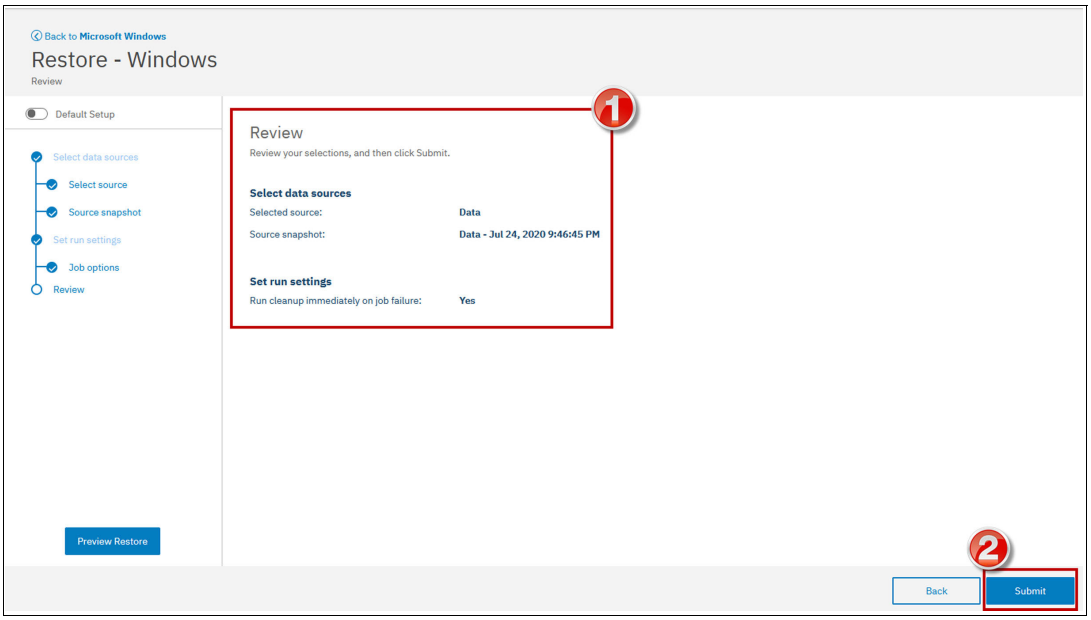

Figure 7-32 Review and submit restore operation

When you click **Submit**, the pop-up window that is shown in Figure 7-33 opens. It shows information for the preparation steps to follow to start the file systems restore operation.

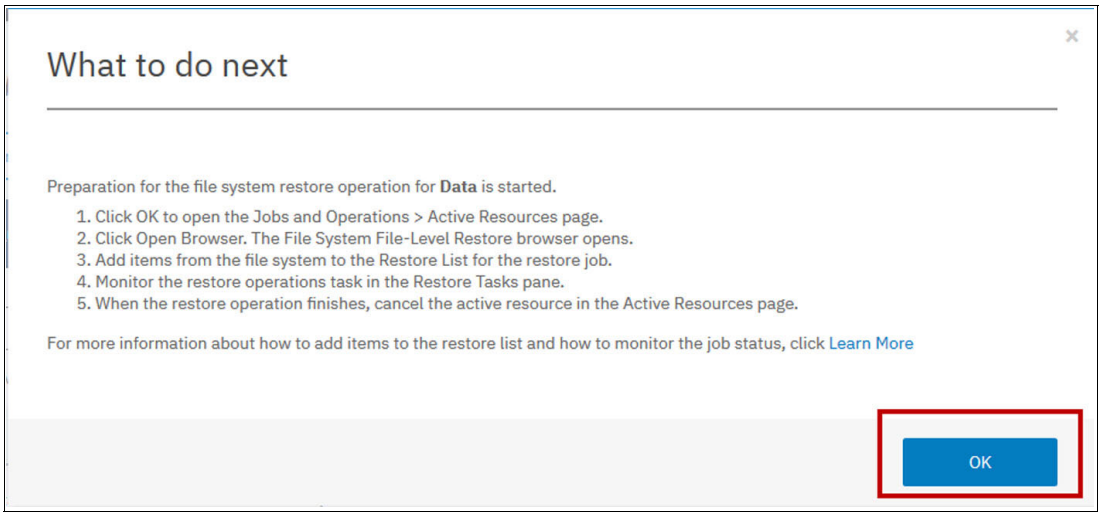

Figure 7-33 Information message about restore preparation

The guidance that is provided corresponds to the ongoing activity that is carried out. After reviewing the details, click **OK**. You are redirected to the Jobs and Operations page.

7.4.2 Step 2: Restoring sequence of the file systems recovery process

This section presents the final steps of the file systems recovery process. By selecting **Jobs and Operations** → **Running Jobs** tab, you can see the progress of the activity and the current active jobs, as shown in Figure 7-34. You can then switch to the Active resource task window.

This task stays open until you cancel the operation after the successful item restore is completed.

Be patient:. The preparation for the restore activity can take several minutes to complete until the mounted share is ready for the item recovery process.

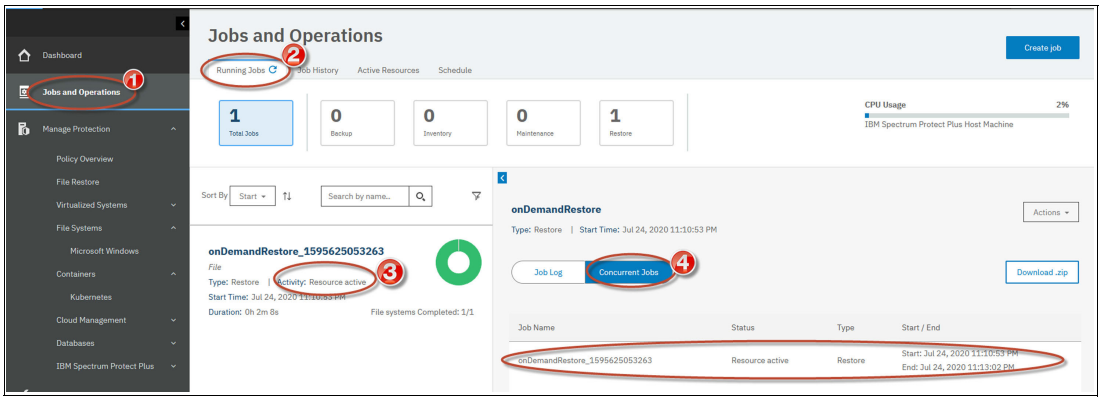

Figure 7-34 Status of the running Jobs for Item recovery

Select **Active Resources** tab → **File Systems** → **Open Browser** tab. Next, click the **File (i)** information icon, as shown in Figure 7-35.

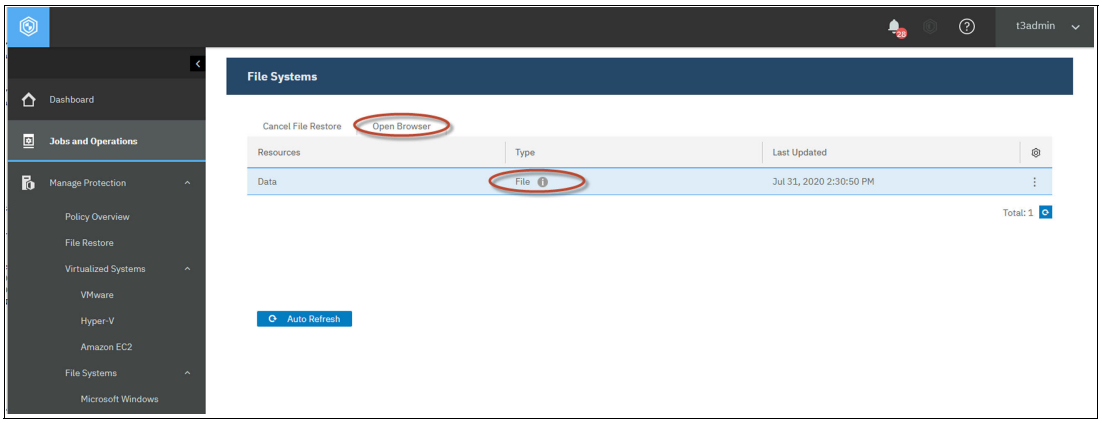

Figure 7-35 Active resources tab information

The Info window that is shown in Figure 7-36 is displayed. The link that is provided can be copied and used to open the browser restore interface. This option is advantageous when you cannot directly access the environment and you must use an example jump host for the restore operation. If you have an issue with DNS, the Host entry can be replaced with an IP address.

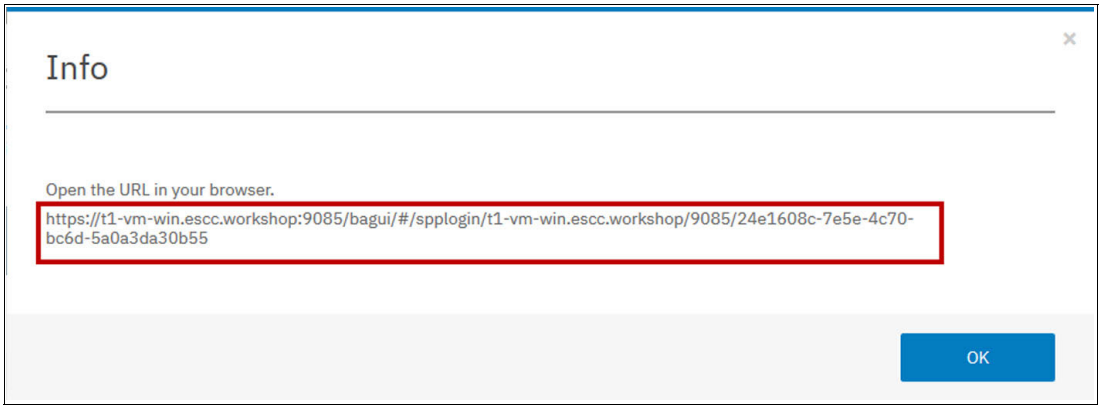

Figure 7-36 Information message with link point to the browser restore interface

In the File System Restore interface that is displayed in the browser, you can also search through the directory structure. By clicking the **arrow** icon, you can select individual directories with full content or choose single files that were backed-up.

The files or directories that you selected move to the right side in the Restore List pane, as shown in Figure 7-37 on page 253. You can restore the data to their original location or you can specify an alternative location with the fully qualified path where files should be restored by clicking **Options**. Start the restore process by clicking **Restore**. After the restoration task completes, the Restore Tasks pane is updated with the restore operation details.

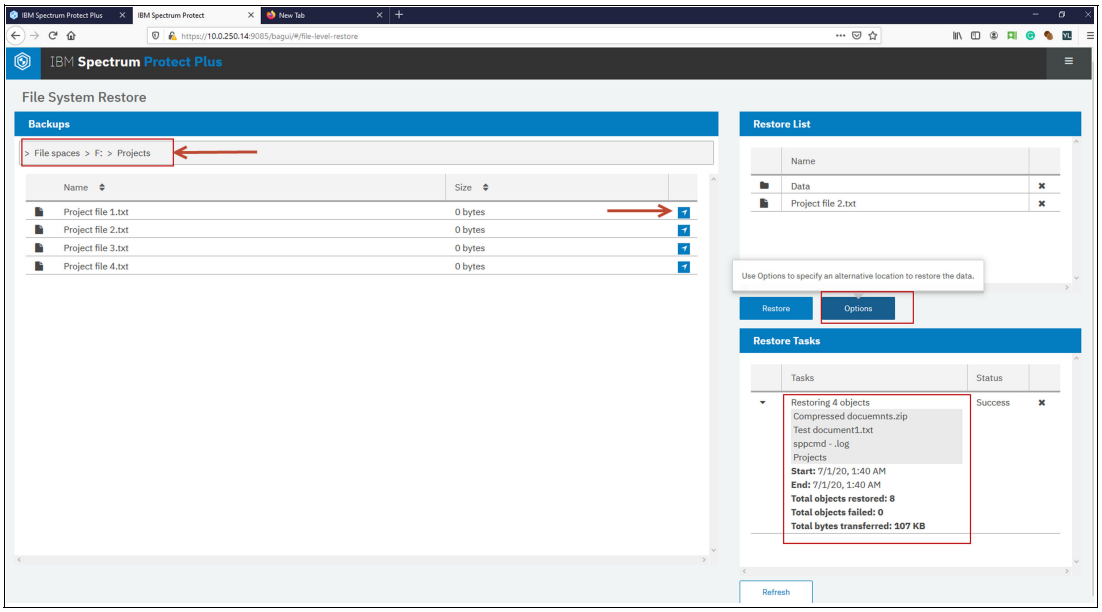

Figure 7-37 File System Restore UI

After the restoration process is complete successfully, from the top right corner, click the menu icon and select **Sign Out**, as shown in Figure 7-38.

Note: A default time-out function also is available that automatically signs you out of the restore interface if you do not manually sign out.

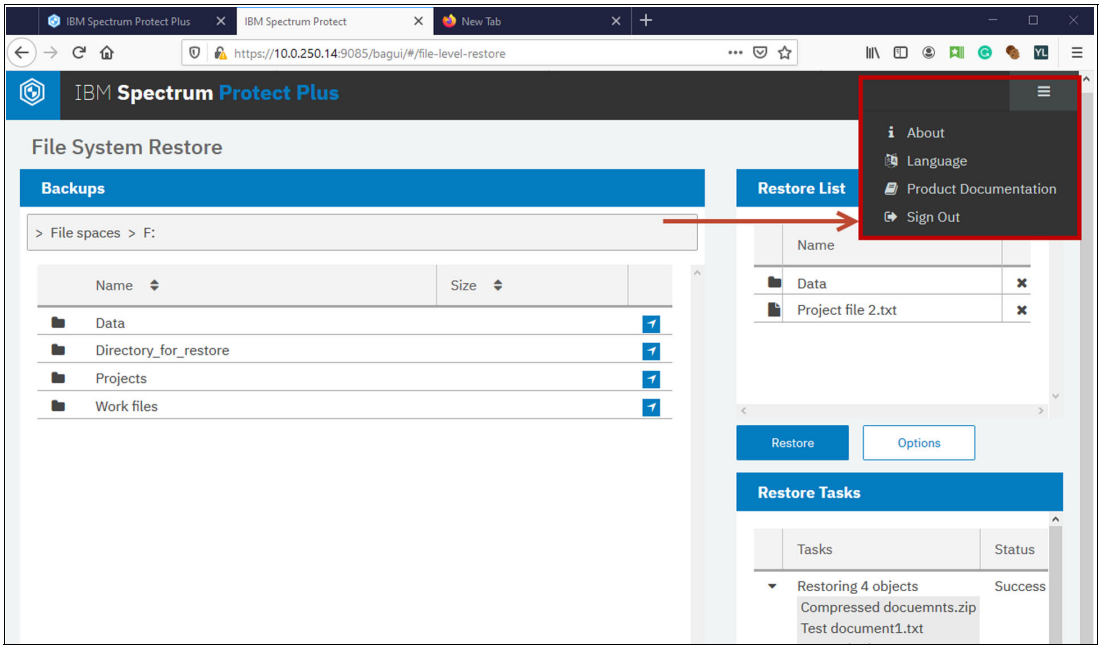

Figure 7-38 Sign out of the Browser restore interface

After the restore operation is completed, select **Jobs and Operations Panel** → **Active Resources** tab and click **Cancel File Restore**, as shown in Figure 7-39.

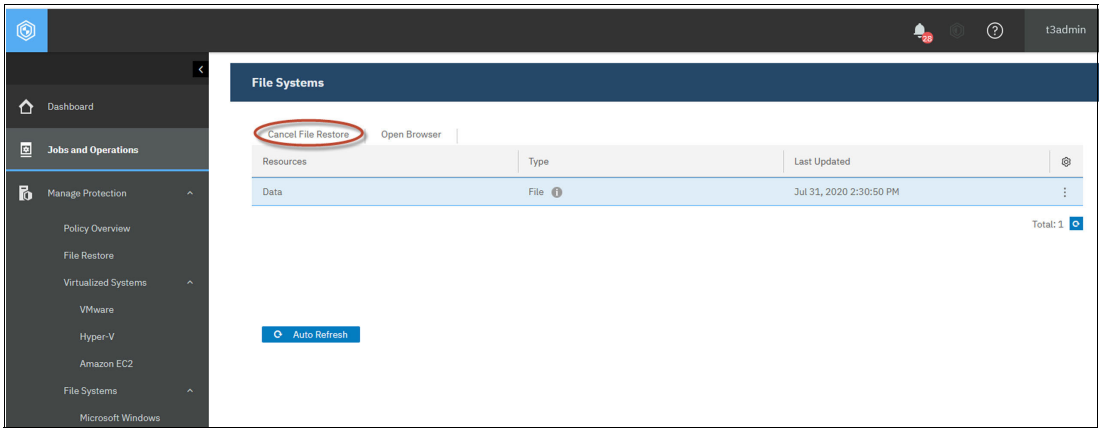

Figure 7-39 Cancel File Restore

Note: During any active restoration operation, the Cancel File restore function is inactive to avoid conflict with or interruption of the ongoing restoration activity.

The Job History tab is updated with a status of “Completed” for the recovery process, as shown in Figure 7-40.

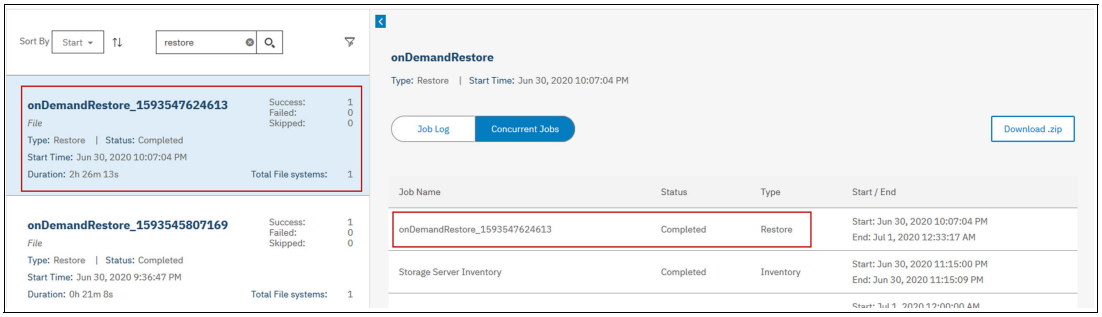

Figure 7-40 Job History

In the operating system directory structure view, as shown in Figure 7-41, the restored file names include timestamps that indicate the date when the restore was completed.

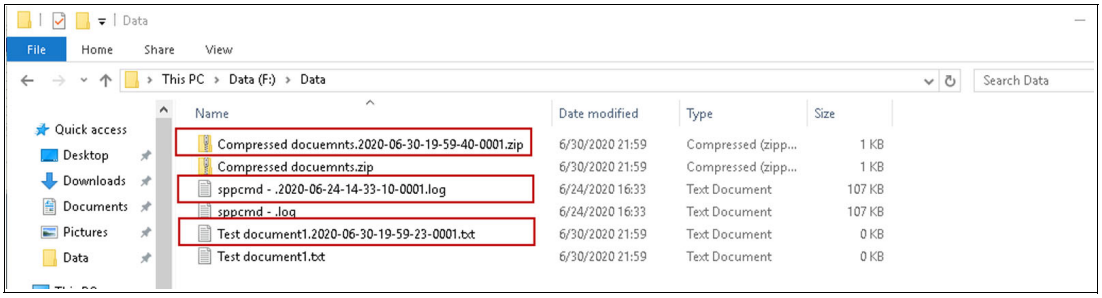

Figure 7-41 Operating system view with restored files overview

8

Backing up and restoring databases

This chapter discusses databases management by using IBM Spectrum Protect Plus. 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.

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

This chapter describes backup and restore of databases and data reuse by using IBM Spectrum Protect Plus. It includes the following information:

Oracle database examples are used in this chapter. For more information about specific databases, see the following chapters:

- Chapter 9, "Backing up and restoring MongoDB databases" on page 279
- Chapter 10, "Backing up and restoring Db2 databases" on page 291
- Chapter 11, "Backing up and restoring SQL Server" on page 309

This chapter includes the following topics:

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.

8.1 Database backup configuration basics

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

Backup, restore, and data reuse handling functions for the supported (relational) databases are all similar in IBM Spectrum Protect Plus. Therefore, the examples in this section are valid for all supported databases to a large extent. However, the examples and figures apply to a specific database in some cases. For this chapter, we chose Oracle Database.

The details in Chapter 9, “Backing up and restoring MongoDB databases” on page 279, Chapter 10, “Backing up and restoring Db2 databases” on page 291, and Chapter 11, “Backing up and restoring SQL Server” on page 309 in this Redbooks publication provide additional information about data protection of these particular databases.

For more information about supported databases and environments, see *IBM Spectrum Protect Plus Installation and User’s Guide*.

8.1.1 Creating an Identity

An operating system user is required to register a database application server, and discover the databases that exist on this 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 8-1.

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.

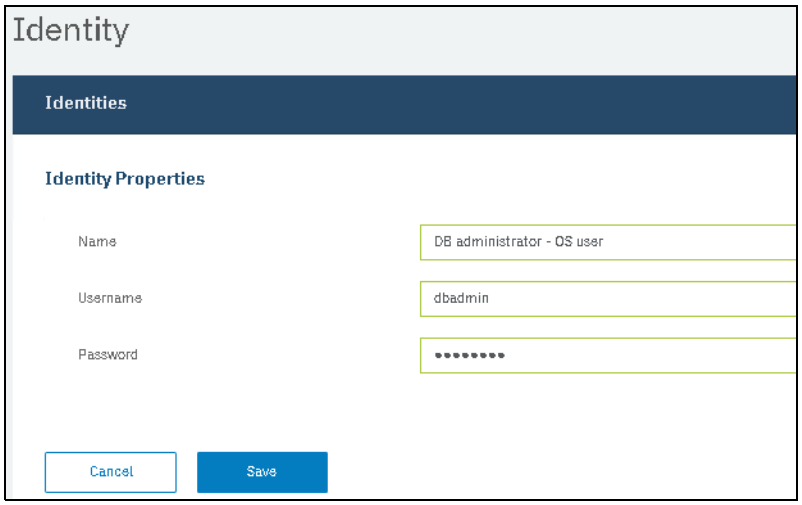

Figure 8-1 Creating an identity for an operating system user

With a list of identity entries, as shown in Figure 8-2, you can see that there are default system identities (LocalvSnapadmin, vsnapadmin or serveradmin) and identities explicitly created for database backup and restore (DB administrator, operating system user, Mongo DB user, and Oracle DBA).

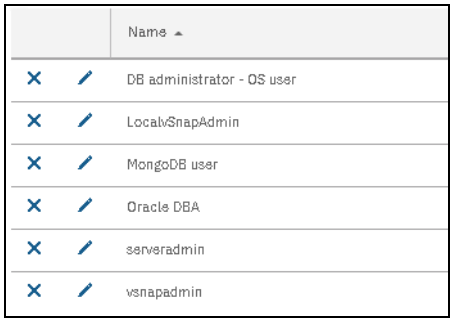

Figure 8-2 Identity entries

8.1.2 Adding an application server

In the IBM Spectrum Protect Plus GUI, select **Manage Protection** → **Databases** → ***<your database type>***. Then, click the **Manage application servers** and **Add application server** buttons. Enter the database server's IP name or address, enter the database administration user or select an identity that you defined earlier. Click the **Get databases** button to start a database discovery job on the server.

Figure 8-3 shows an Oracle database example. The name of the discovered database is SPP.

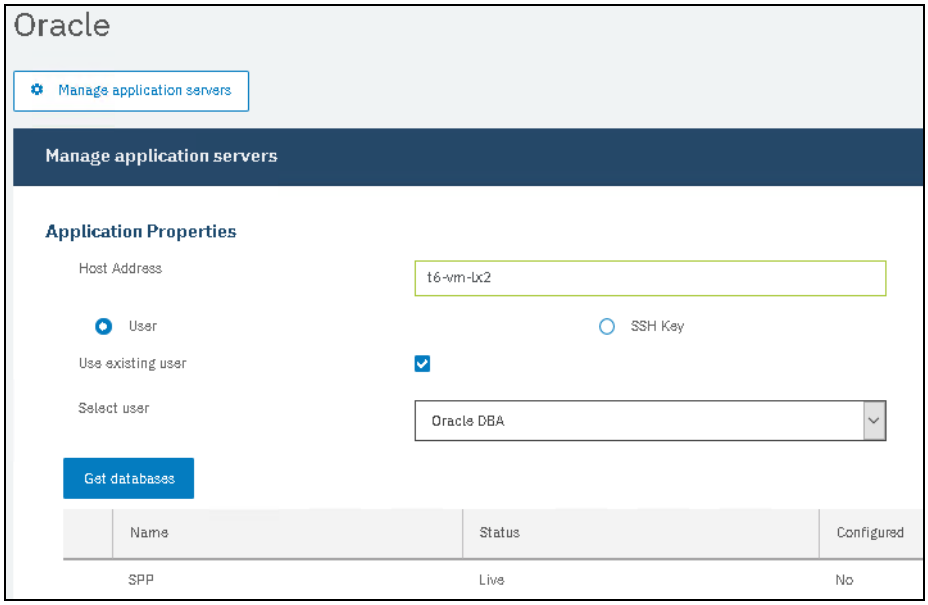

Figure 8-3 Add an application server

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. Switch to the **Jobs and Operations** menu to check the job results.

8.1.3 Assigning an SLA policy

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

Section 1.3, “SLA backup policies” on page 19 discusses this topic in more detail.

8.1.4 Running a backup

A complete database backup includes the data files, metadata (such as Oracle 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 rollforward to a 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.

Click the **Select options** button to configure a log backup schedule, as shown in Figure 8-4.

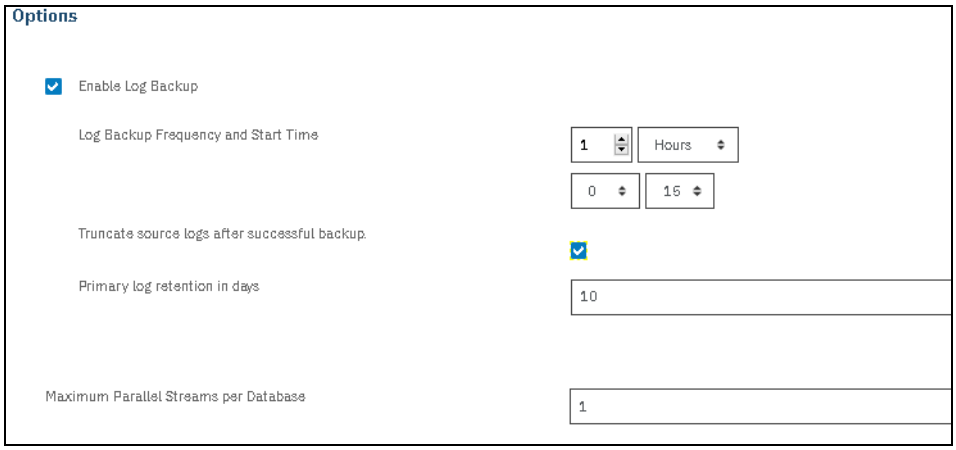

Figure 8-4 Enable log backup for a database

If you configure log backups for the database, IBM Spectrum Protect Plus performs the following actions:

- Mounts another target volume for log backups to the database server. This volume remains mounted to the database server at all times.
- Schedules a cron job that regularly starts the log backup. Example 8-1 shows a sample crontab entry.

Example 8-1 Crontab entry

```
# Added by SPP
15 * * * * /opt/IBM/SPP/logbackup/SPP/3524c51a79a84693c152356ad4aed42c/logbackup.sh
```

Note: This example is an Oracle example. The IBM Spectrum Protect Plus implementation of database log backup varies depending on the database type.

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 8-5).

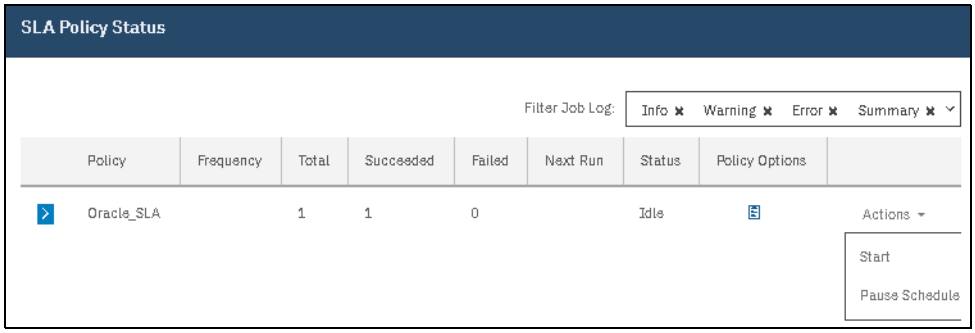

Figure 8-5 Manually start a database backup

Switch to the **Jobs and Operations** menu to check the job results.

8.2 IBM Spectrum Protect Plus database restore and data reuse

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 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 show 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 with-in the list of backup date and time)
  - 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 rollforward: 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 8-6.

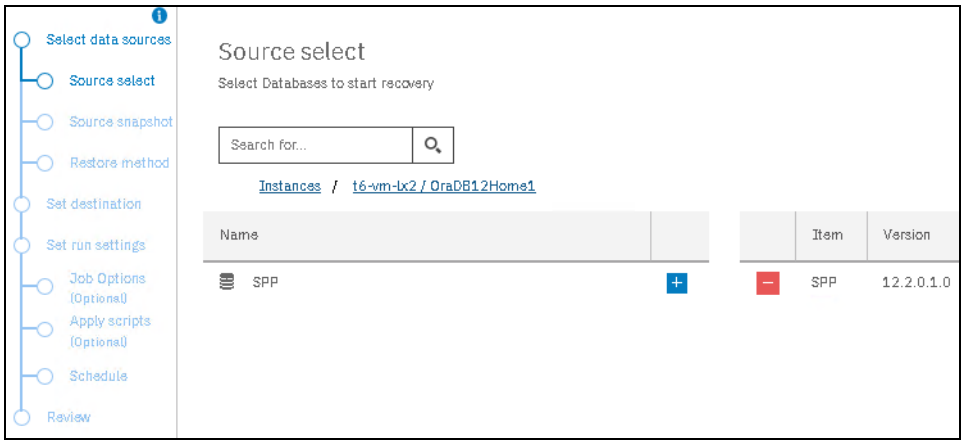

Figure 8-6 Select the source for a database restore

8.2.1 Test restore

This section describes a possible test restore use case: Every Monday morning, a development tester or application key 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 8-6. 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 8-7.

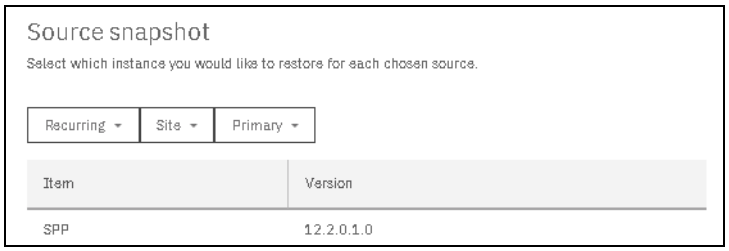

Figure 8-7 Selecting a site to create the new database instance

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 8-8 and Figure 8-9 show the corresponding parameter selections.

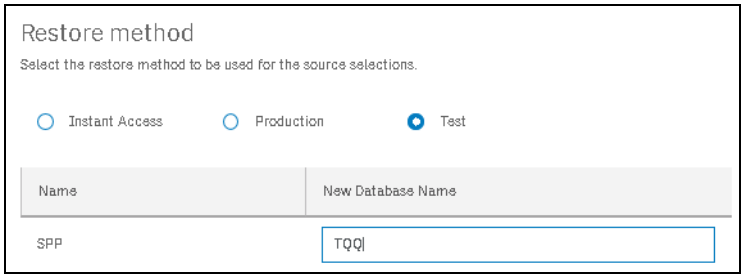

Figure 8-8 Selecting the Test restore method

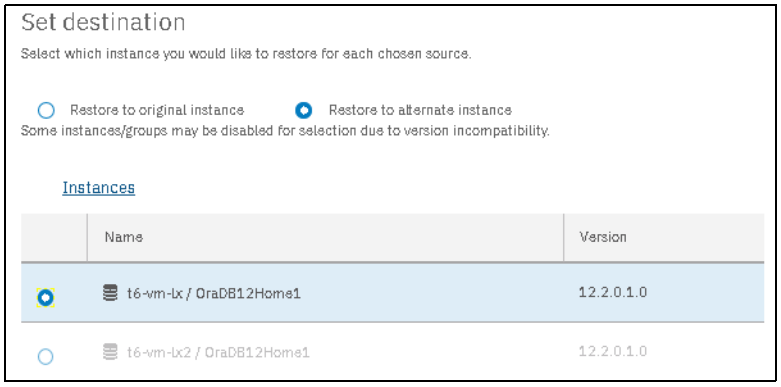

Figure 8-9 Selecting a restore to an alternative instance

Finally, we define the schedule: every Monday at 8:00 o'clock. Figure Example 8-10 on page 262 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.

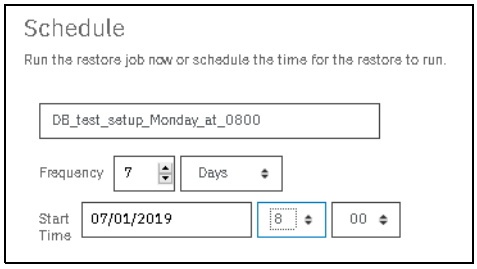

Figure 8-10 Weekly schedule for a database copy

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). See Example 8-2 and Example 8-3.

Example 8-2 Database files on the mounted vSnap server directory

```
t6-vm-lx:/ # df -h|grep vsnap
10.0.250.53:/vsnap/vpool1/fs8   96G  726M   95G   1% /mnt/spp/vsnap/vpool1/fs8

t6-vm-lx:/ # ls /mnt/spp/vsnap/vpool1/fs8/SPP
arch
control01.ctl
controlfile.ctl
controlfile.txt
data_D-SPP_I-2016102274_TS-SYSAUX_FNO-3_38u4doto
data_D-SPP_I-2016102274_TS-SYSTEM_FNO-1_39u4dotv
data_D-SPP_I-2016102274_TS-UNDOTBS1_FNO-4_3au4dou7
data_D-SPP_I-2016102274_TS-USERS_FNO-7_3cu4douc
pfile.txt
redo01.log
redo02.log
tempfile_1.dbf
```

Example 8-3 Database instance started by an IBM Spectrum Protect Plus restore job

```
t6-vm-lx:/ # su - oracle

oracle@t6-vm-lx:~> env|grep SID
ORACLE_SID=SPP
oracle@t6-vm-lx:~> sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production on Thu Jun 27 11:43:20 2019
Copyright (c) 1982, 2016, Oracle.  All rights reserved.

Connected to an idle instance.

SQL> exit
Disconnected

oracle@t6-vm-lx:~> ORACLE_SID=TQQ
```

```
oracle@t6-vm-lx:~> sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production on Thu Jun 27 11:43:35 2019
Copyright (c) 1982, 2016, Oracle.  All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production
```

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**.

8.2.2 Instant access

In Instant access restore mode, IBM Spectrum Protect Plus mounts the volume from the vSnap server repository. The setup uses the restore wizard is similar to the test restore setup that is described in “Test restore” on page 260.

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, as shown in Example 8-4, 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.

Example 8-4 A database backup mounted for instant access

```
oracle@t6-vm-lx:/mnt/spp/vsnap/vpool1/fs5/SPP> hostname
t6-vm-lx

oracle@t6-vm-lx:/mnt/spp/vsnap/vpool1/fs5/SPP> df -h|grep vsnap
10.0.250.53:/vsnap/vpool1/fs5   96G  717M   96G   1% /mnt/spp/vsnap/vpool1/fs5

oracle@t6-vm-lx:/mnt/spp/vsnap/vpool1/fs5/SPP> ls -l
total 733166
drwxr-xr-x 2 oracle oinstall          3 Jun 21 15:57 arch
-rw-r----- 1 oracle oinstall   10698752 Jun 21 15:57 controlfile.ctl
-rw-r--r-- 1 oracle oinstall       5920 Jun 21 15:58 controlfile.txt
-rw-r----- 1 oracle oinstall 1415585792 Jun 21 15:57 data_D-SPP_I-2016102274_TS-SYSAUX_FNO-3_38u4doto
-rw-r----- 1 oracle oinstall  870326272 Jun 21 15:57 data_D-SPP_I-2016102274_TS-SYSTEM_FNO-1_39u4dotv
-rw-r----- 1 oracle oinstall  330309632 Jun 21 15:57data_D-SPP_I-2016102274_TS-UNDOTBS1_FNO-4_3au4dou7
-rw-r----- 1 oracle oinstall    5251072 Jun 21 15:57 data_D-SPP_I-2016102274_TS-USERS_FNO-7_3cu4douc
-rw-r--r-- 1 oracle oinstall        974 Jun 21 15:58 pfile.txt
```

As described in 8.2.1, “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.

8.2.3 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 8-6 on page 260. 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 8-11.

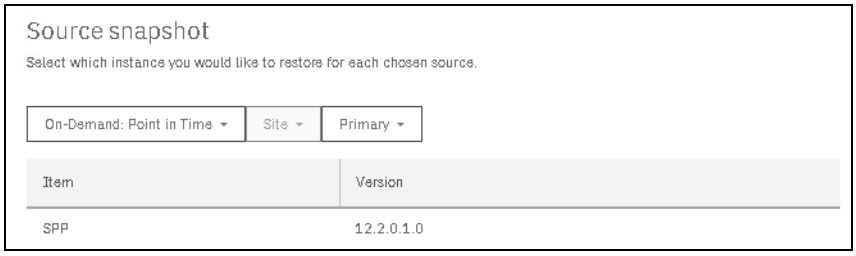

Figure 8-11 Select a site to restore the database

The next two selections indicate what we are trying to achieve; that is, production restore to the original instance, as shown in Figure 8-12 and Figure 8-13.

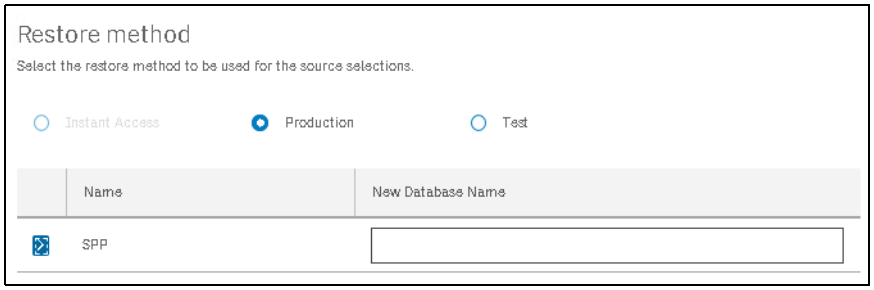

Figure 8-12 Production restore to the original database

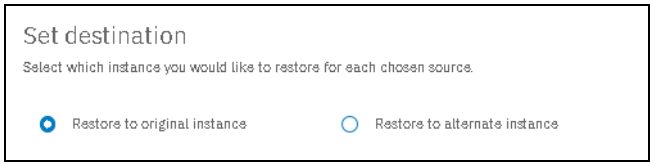

Figure 8-13 Database restore to the original instance

For a restore of a production database, the IBM Spectrum Protect Plus restore wizard assumes a database rollforward 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 8-14 on page 265. IBM Spectrum Protect Snapshot 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.

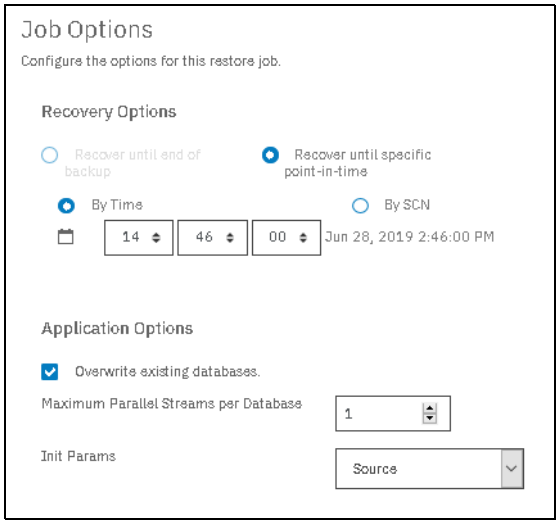

Figure 8-14 Selecting the database rollforward and overwrite options

Carefully review the job summary that IBM Spectrum Protect Plus displays (Figure 8-15) and, if the information describes what you are trying to achieve, run the restore job.

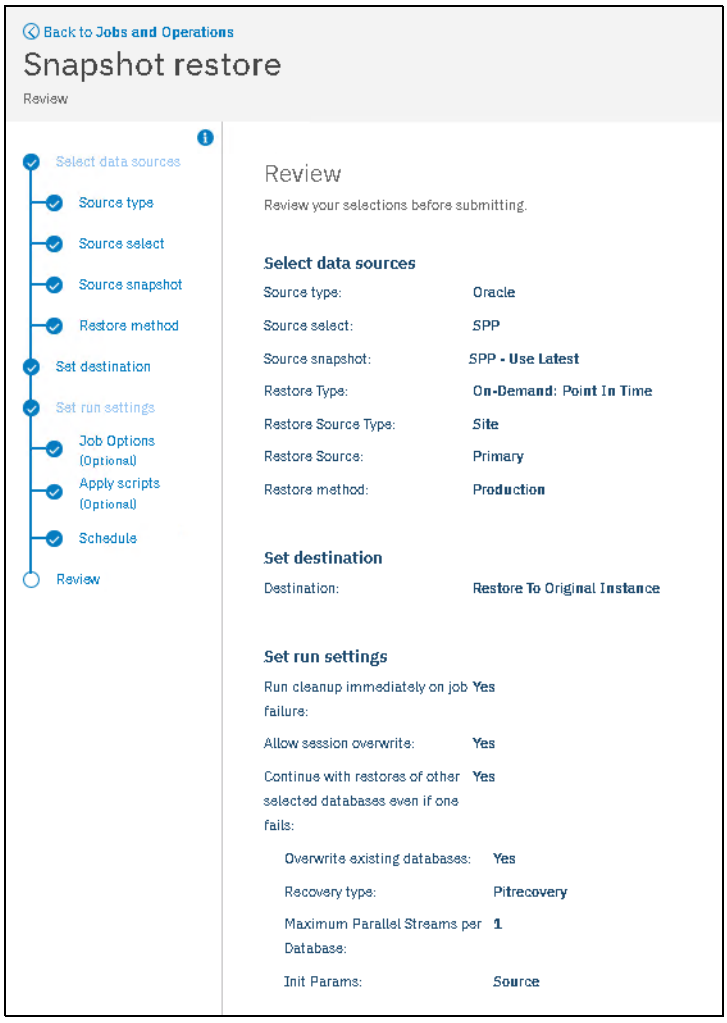

Figure 8-15 Restore job summary

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 8-5.

Example 8-5 Restore job log

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

8.3 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 databases with IBM Spectrum Protect Plus, see the following resources:

- 8.4, “Oracle overview” on page 269
- Chapter 9, “Backing up and restoring MongoDB databases” on page 279
- Chapter 10, “Backing up and restoring Db2 databases” on page 291
- Chapter 11, “Backing up and restoring SQL Server” on page 309
- Chapter 12, “Backing up and restoring Microsoft Exchange data” on page 329

For more information about vSnap server, see 1.2.3, “vSnap Backup Storage server” on page 13.

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

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

8.3.1 Backup operations

For the database that includes log backup enabled, two volumes are created: one volume for the data and one volume for the log management. Each volume has features its own snapshot and recovery points.

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

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 also 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, as shown in Figure 8-16.

Note: Log volumes feature the main characteristics:

- Log volumes stay mounted.
- Application is working directly on that log volume. A backup of these logs is taken when SPP creates read-only snapshots of that vSnap server volume.
- Log volumes are not copied or archived to any S3-mode repositories

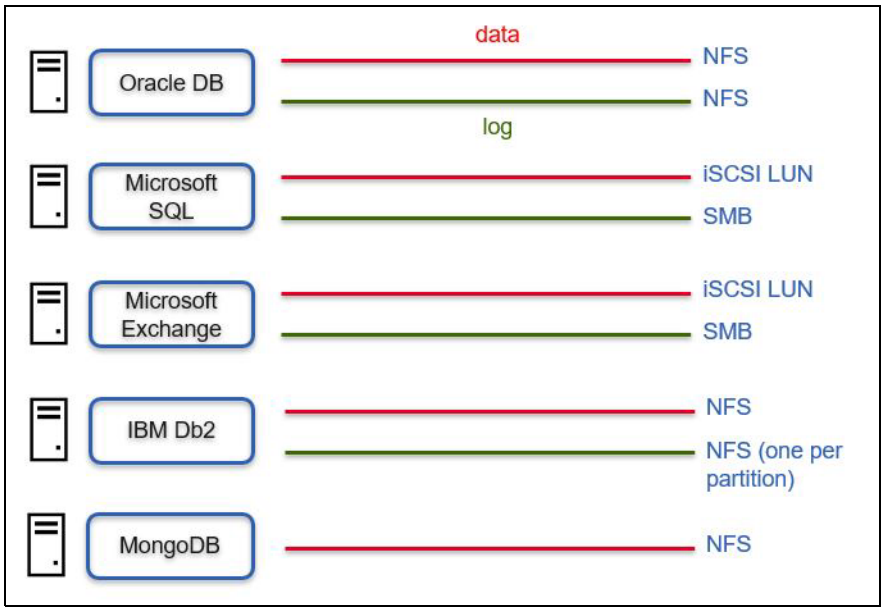

Figure 8-16 How database server access their respective vSnap server volumes

Table 8-1 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 8-1 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

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.

RMAN does not permit to make inconsistent backups when the database is in NOARCHIVELOG mode.

Therefore, IBM Spectrum Protect Plus requires that the Oracle database to be in ARCHIVELOG mode.

8.3.2 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 rollforward 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.

8.4 Oracle overview

This section describes the management of databases with IBM Spectrum Protect Plus in general. However, we chose the Oracle database for our examples. Although Chapter 9, “Backing up and restoring MongoDB databases” on page 279, Chapter 10, “Backing up and restoring Db2 databases” on page 291, and Chapter 11, “Backing up and restoring SQL Server” on page 309 specifically address MongoDB, IBM Db2, and Microsoft SQL Server databases, we mention here some important facts about the configuration and handling for Oracle databases.

8.4.1 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 have database administration permissions and sudo permissions to perform system operations. When registering Oracle Real Application Cluster (RAC) nodes, register each node by using its physical IP name or address.

8.4.2 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 multithreading at the time of backup, the restored database is non-multithreaded. The restored database must be manually reconfigured to use multi-threading.

8.4.3 Oracle backup

IBM Spectrum Protect Plus writes backup information to the Oracle control files. Writing to a remote Oracle Recovery Manager (RMAN) catalog is not supported at the time of this writing.

As listed in the prerequisites, 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 way of IBM Spectrum Protect Plus, select the **Enable Log Backup** option and specify the **Primary log retention in days** parameter in the Backup Options window, as shown in Figure 8-17.

Figure 8-17 Select Options button to access Oracle backup options

Figure 8-18 shows the parameter controlling the local archive log cleaning.

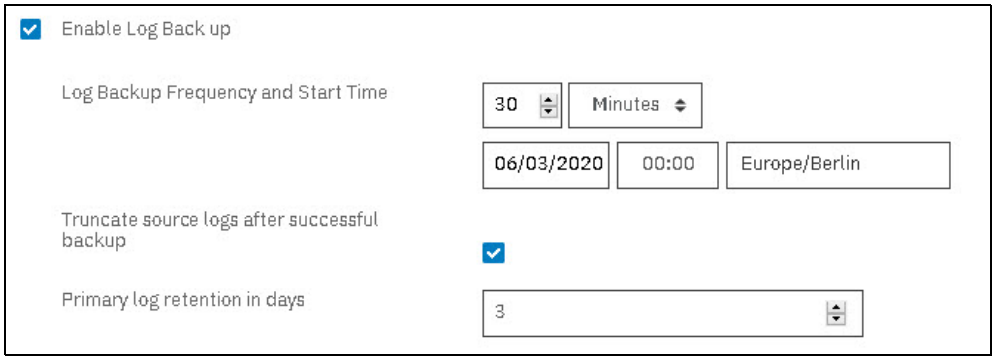

Figure 8-18 Set the local archive log retention to 3 days

For more information about the commands that are used by IBM Spectrum Protect Plus, see 8.4.7, “Oracle commands used by IBM Spectrum Protect Plus” on page 271

Note: For archive log backups, it is important to decide up front which mechanism is preferred to manage the local archive log deletion: controlled by IBM Spectrum Protect Plus through the option that is shown in Figure 8-18 or the database administrator prefers to control it on their side. In the latter choice, the archive log must be kept at least for the time as defined in the Log Backup Frequency and StartTime setting in the IBM Spectrum Protect Plus SLA Policy so that IBM Spectrum Protect Plus can take them for backup.

8.4.4 Oracle Block Change Tracking

IBM Spectrum Protect Plus uses Oracle Block Change Tracking to perform incremental backups. If Block Change Tracking is not enabled, IBM Spectrum Protect Plus enables it automatically during the backup.

8.4.5 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).

8.4.6 Troubleshooting hint

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, *IBM Spectrum Protect Plus Installation and User's Guide*.

8.4.7 Oracle commands used by IBM Spectrum Protect Plus

IBM Spectrum Protect Plus relies on Oracle RMAN to perform backup operations.

Any command that is used for the database backup, archive logs backup, and archive logs cleaning can be found in the Job log when the Detail filter is enabled.

The database backup process includes the following tasks:

- Archive log backup
- Force archive log switch
- Incrementally backup the database
- Archive log backup
- SPfile backup
- Control file backup
- Local Archive log cleaning when enabled

The details that are shown in Example 8-6 - Example 8-9 on page 272 are extracted from job logs where you can see which RMAN commands were used by IBM Spectrum Protect Plus to perform the backup.

Example 8-6 RMAN commands used by IBM Spectrum Protect Plus to do archive log backup

```
RMAN> set echo off;
2> connect target *
3> run {
4> configure controlfile autobackup off;
5> set command id to 'spp_logbackup_1591215736';
6> allocate channel spp1 type disk format
'/mnt/spp/vsnap/vpool1/fs3/10_0_250_32/SPP/%h_%e_973347459.dbf';
7> backup as copy archivelog from scn 4398899;
8> release channel spp1;
9> }
10> exit;
```

Example 8-7 RMAN command used by IBM Spectrum Protect Plus to do database backup

```
connect / as sysdba;
set timing off;
alter system archive log current;
exit;
```

```
connect target /;
run {
configure controlfile autobackup off;
set command id to 'SPP_BACKUP_1006_1591907663';
allocate channel spp1 type disk format
'/mnt/spp/vsnap/vpool1/fs4/10_0_250_32/SPP/%U';
allocate channel spp2 type disk format
'/mnt/spp/vsnap/vpool1/fs4/10_0_250_32/SPP/%U';
backup incremental level 1 for recover of copy with tag 'SPP_BACKUP_1006'
database;
recover copy of database with tag 'SPP_BACKUP_1006';
delete noprompt backup tag 'SPP_BACKUP_1006';
release channel spp1;
release channel spp2;
}
exit;


connect target /;
run {
configure controlfile autobackup off;
set command id to 'SPP_BACKUP_1006_1591907693';
allocate channel spp1 type disk format
'/mnt/spp/vsnap/vpool1/fs4/10_0_250_32/SPP/arch/%h_%e_973347459.dbf';
allocate channel spp2 type disk format
'/mnt/spp/vsnap/vpool1/fs4/10_0_250_32/SPP/arch/%h_%e_973347459.dbf';
backup as copy archivelog from scn 5202016 until scn 5202018;
release channel spp1;
release channel spp2;
}
exit;


connect / as sysdba;
set timing off;
create pfile='/mnt/spp/vsnap/vpool1/fs4/10_0_250_32/SPP/pfile.txt' from spfile;
exit;
```

Example 8-8 RMAN commands used by IBM Spectrum Protect Plus to do controlfile backup

```
RMAN> set echo off;
2> connect target *
3> run {
4> set command id to 'SPP_BACKUP_1006_1591855534';
5> allocate channel spp1 type disk format '/mnt/spp/tmp/SPP_spp_backup_1006.ctl';
6> backup as copy current controlfile reuse tag 'SPP_BACKUP_1006';
7> }
8> exit;
echo set off
```

Example 8-9 RMAN commands used by IBM Spectrum Protect Plus to do local archive log cleaning

```
set echo off;
connect target /;
```

```
run {
delete noprompt force archivelog until time 'SYSDATE-3' like '/home/SPP/arch/%';
}
exit;
```

8.5 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 8-10 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 8-10 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.
3. 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 8-19. 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.

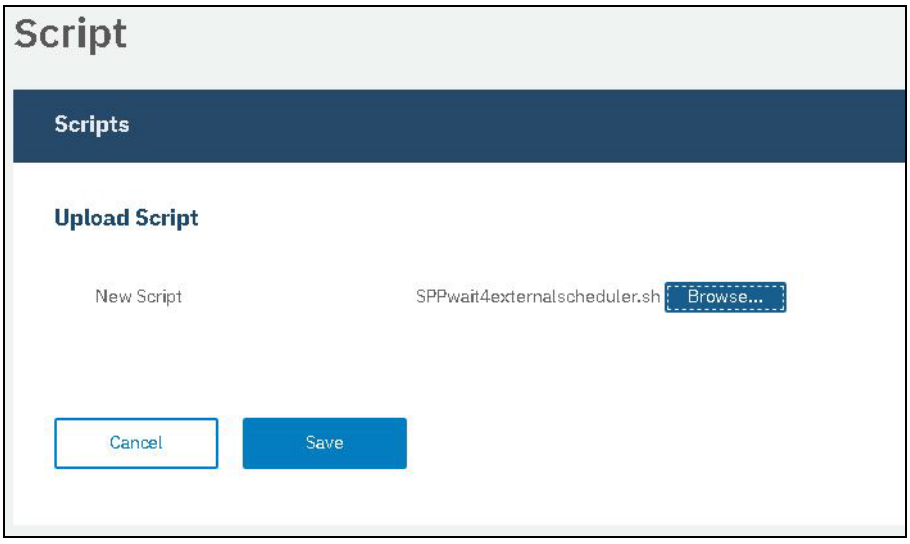

Figure 8-19 Define script by uploading it

The second step is to define the Script Server, as shown in Figure 8-20. Select **System Configuration** → **Script.** Specify the Host address, login credentials, and operating system type for which you plan to use script.

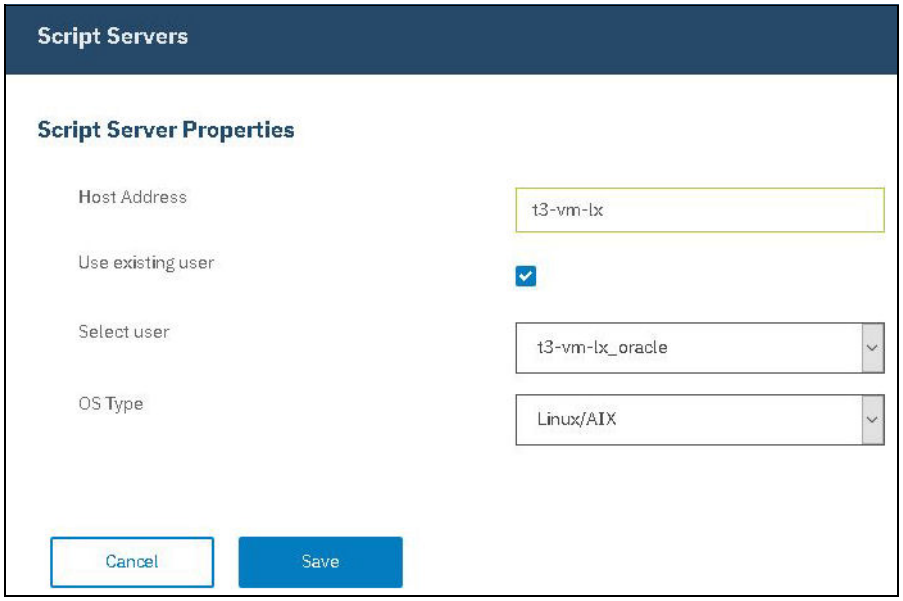

Figure 8-20 Define SCript Server by specifying its address and credentials

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 8-21.

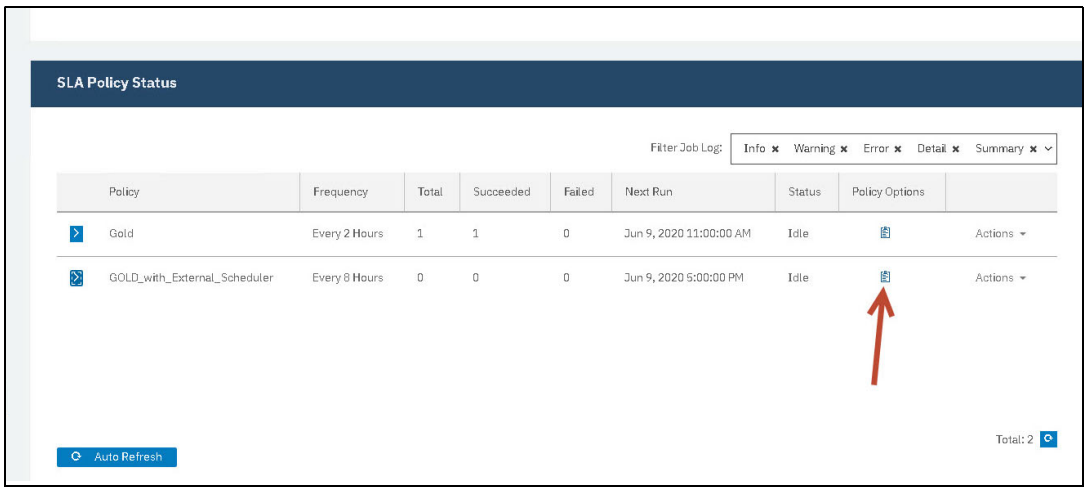

Figure 8-21 Open the Policy Options to enable the use of a script for that SLA Policy run

When you click **Policy Options**, a pop-up window opens (see Figure 8-22) and you must specify whether you want to enable pre-script or post-script, and which script to run.

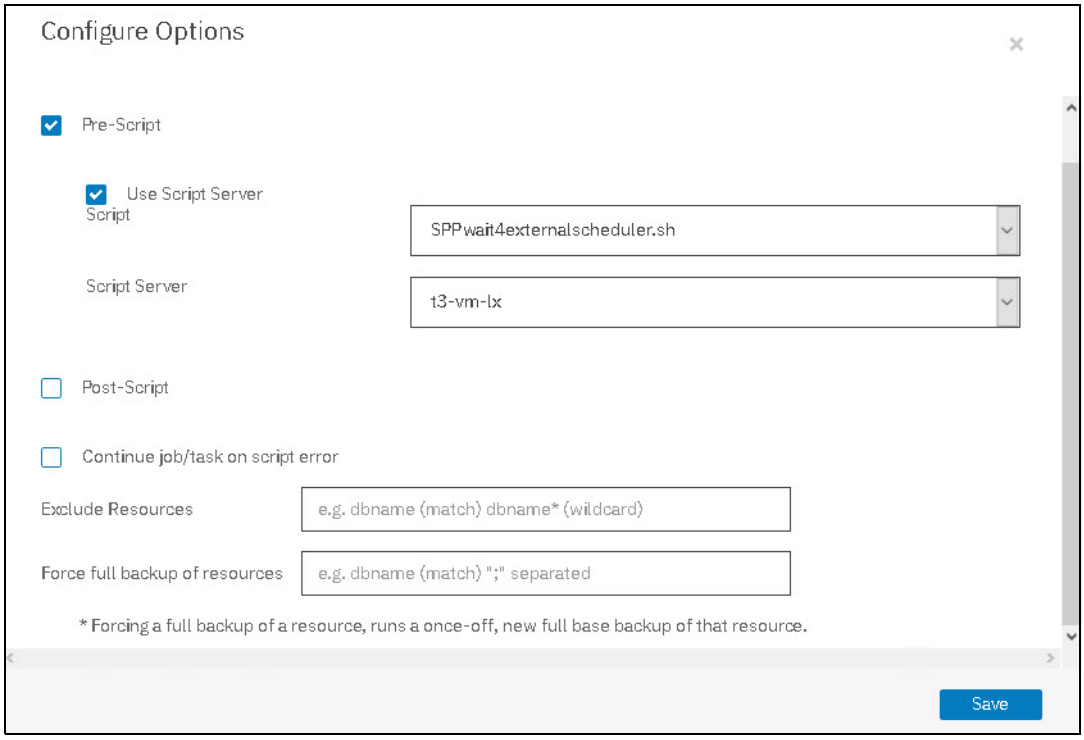

Figure 8-22 Configure Policy Options to use a specific script for pre or post backup tasks

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, 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 8-23 shows the output log that lists the execution of pre-script.

Figure 8-23 Job log showing that pre script action is executed on the target server t3-vm-lx

Figure 8-24 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.

Figure 8-24 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. For more information about how to use REST API, see Chapter 16, “REST API” on page 461.

9

Backing up and restoring MongoDB databases

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

This chapter includes the following topics:

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 10, “Backing up and restoring Db2 databases” on page 291, and Chapter 11, “Backing up and restoring SQL Server” on page 309

For more information about generic test restore or DevOps use cases, see Chapter 8. “Backing up and restoring databases” on page 255. This chapter also describes database backup, restore, and DevOps use cases in general, but refers specifically to an Oracle database whenever necessary.

9.1 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 Knowledge Center.

9.1.1 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 Knowledge Center.

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 9-1 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 9-1 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 sharded cluster instances. A backup always includes all databases in the instance.

9.1.2 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.

9.1.3 “MongoDB databases with authentication enabled” on page 281 describes how to enable MongoDB authentication.

9.1.4 “Register a MongoDB server” on page 283 describes the MongoDB registration in IBM Spectrum Protect Plus.

9.1.3 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 9-2.

Example 9-2 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 9-3. According to the MongoDB manuals, the ClusterAdmin role includes the clusterManager, clusterMonitor, and hostManager roles.

Example 9-3 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 9-4 shows how to start the daemon on a Linux command line.

Example 9-4 Starting mongod on Linux

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

9.1.4 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 9-1 and Figure 9-2 show Identities for an operating system and a MongoDB user. The two user names can be identical.

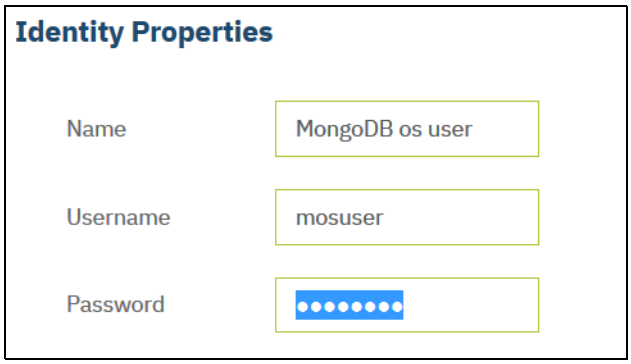

Figure 9-1 Identity definition for an operating system user

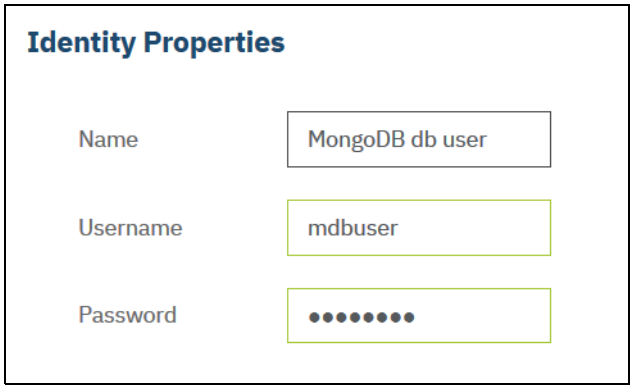

Figure 9-2 Identity definition for a MongoDB user

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 9-3).

For more information about handling of MongoDB databases with authentication, see section 9.1.3.

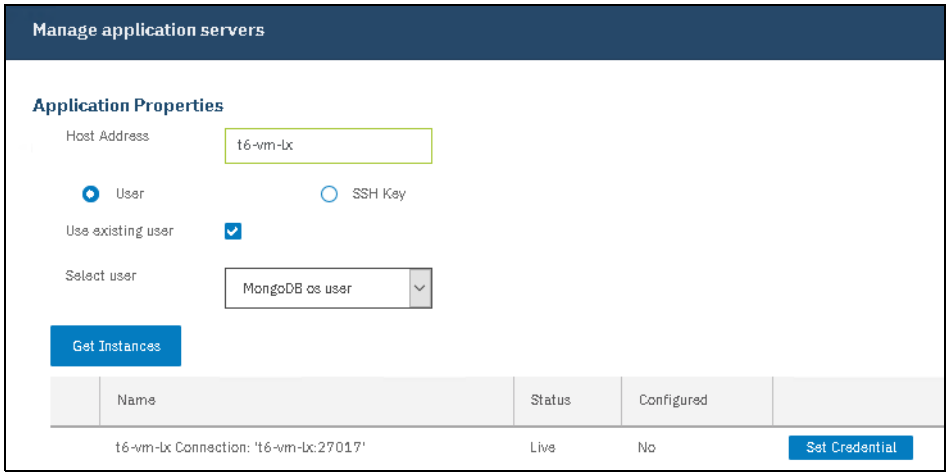

Figure 9-3 Add a MongoDB server

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.

9.2 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 8.2 "IBM Spectrum Protect Plus database restore and data reuse" on page 259.

9.2.1 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.

For more information, see 1.3 "SLA backup policies" on page 19.

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 8.5 “Database backup with pre-script and post-script” on page 274.

Starting the database backup

The SLA policy that you assigned to your database (see Figure 9-4) 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**).

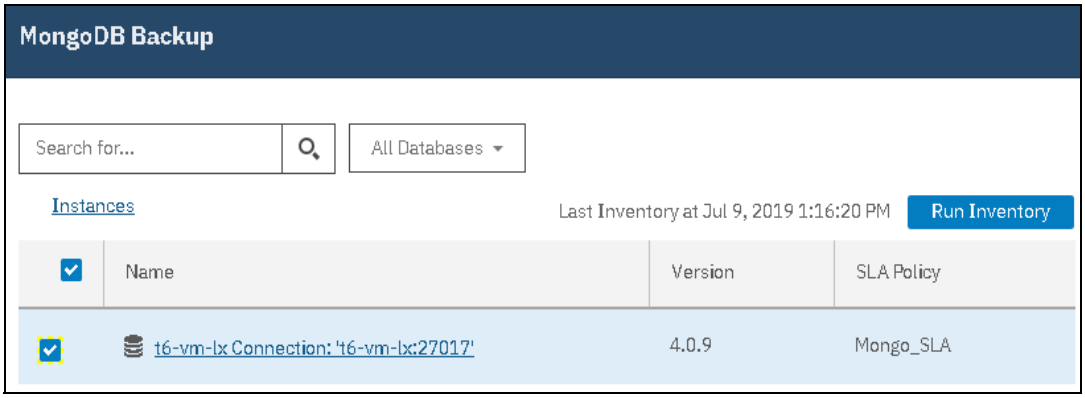

Figure 9-4 MongoDB instance discovered by IBM Spectrum Protect Plus with an SLA policy assigned

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.

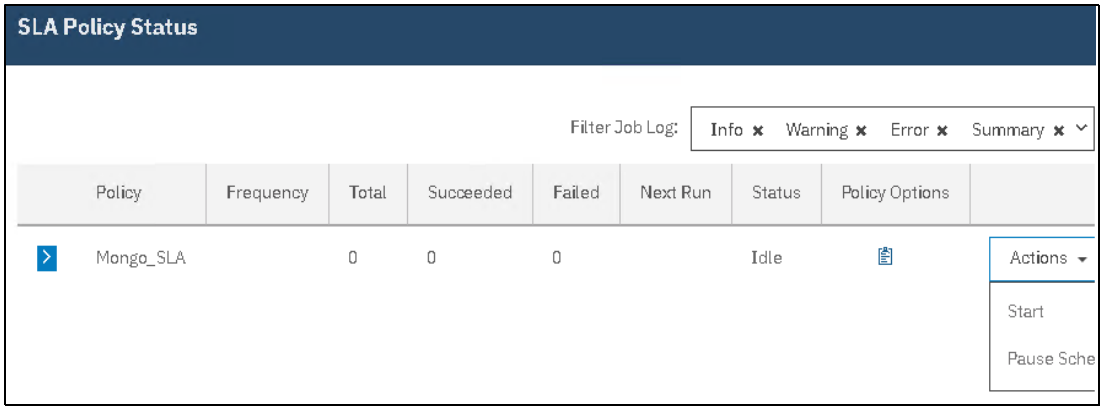

Figure 9-5 Manually start a database backup

IBM Spectrum Protect Plus mounts a vSnap server directory to the database server to copy the backup data (see Example 9-5). During the initial backup operation, IBM Spectrum Protect Plus creates a vSnap server volume and NFS share.

Example 9-5 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-mongolv      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.

9.2.2 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 8.2 “IBM Spectrum Protect Plus database restore and data reuse” on page 259.

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 9-6. Select the available Source Snapshot that needs to be restored, as shown in Figure 9-7.

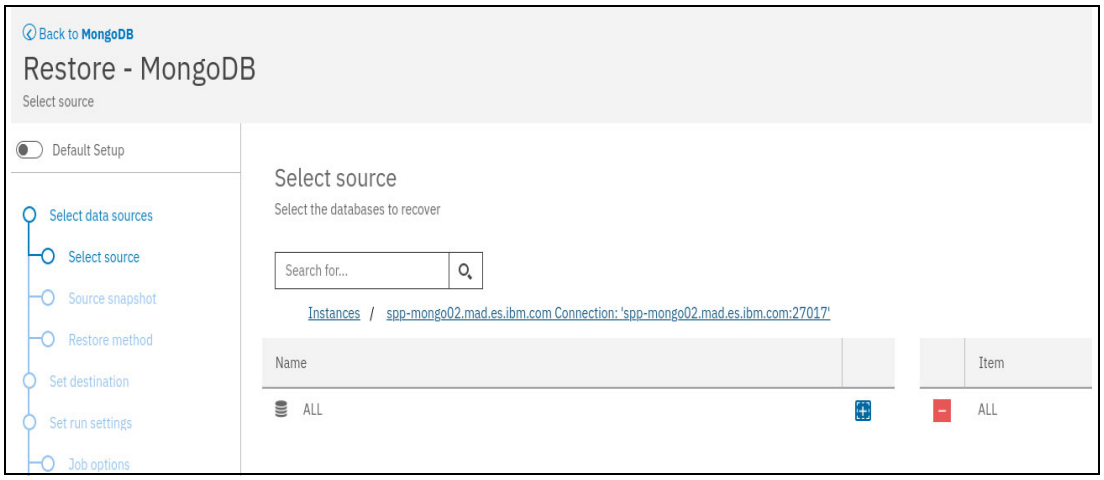

Figure 9-6 Select the source for a database restore

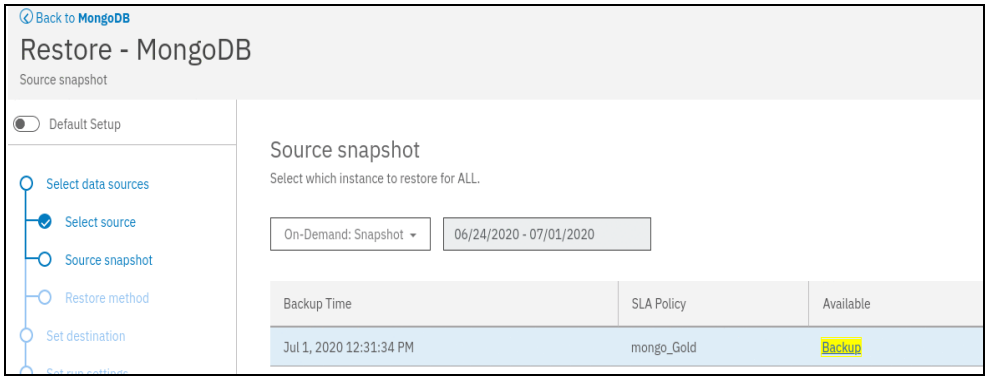

Figure 9-7 Select a site from which to restore the database

The next two selections express what we are trying to achieve: A production restore to the original instance, as shown in Figure 9-8 and Figure 9-9 on page 288.

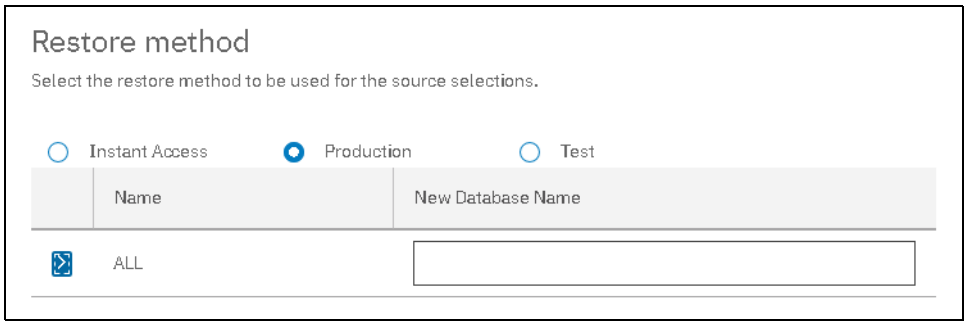

Figure 9-8 Select the restore method

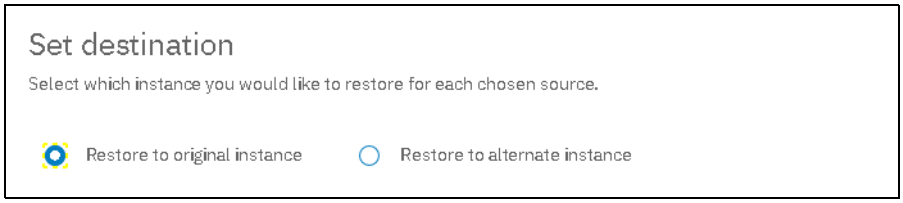

Figure 9-9 Select the restore destination

For an on-demand snapshot restore of a production database, the IBM Spectrum Protect Plus restore wizard assumes a subsequent database rollforward to the end of logs included in the backup (see Figure 9-10).

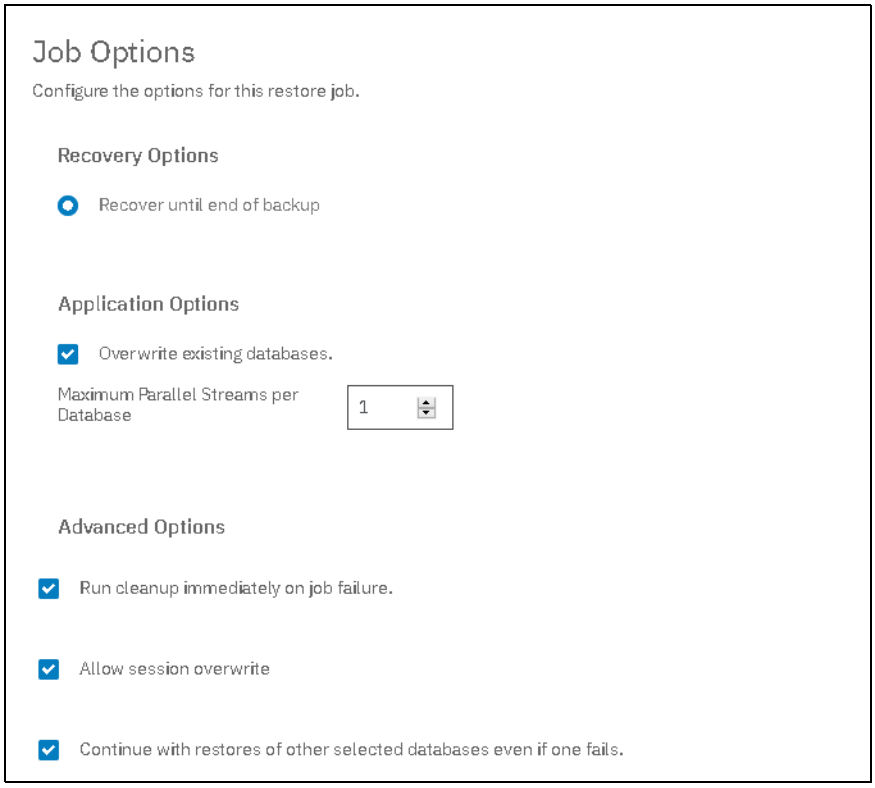

Figure 9-10 Select database overwrite and other restore options

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 9-11.

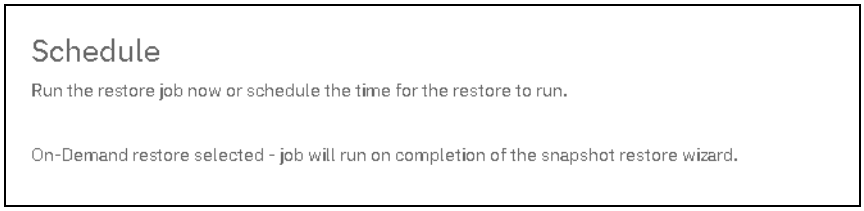

Figure 9-11 Information about a job schedule

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 9-12).

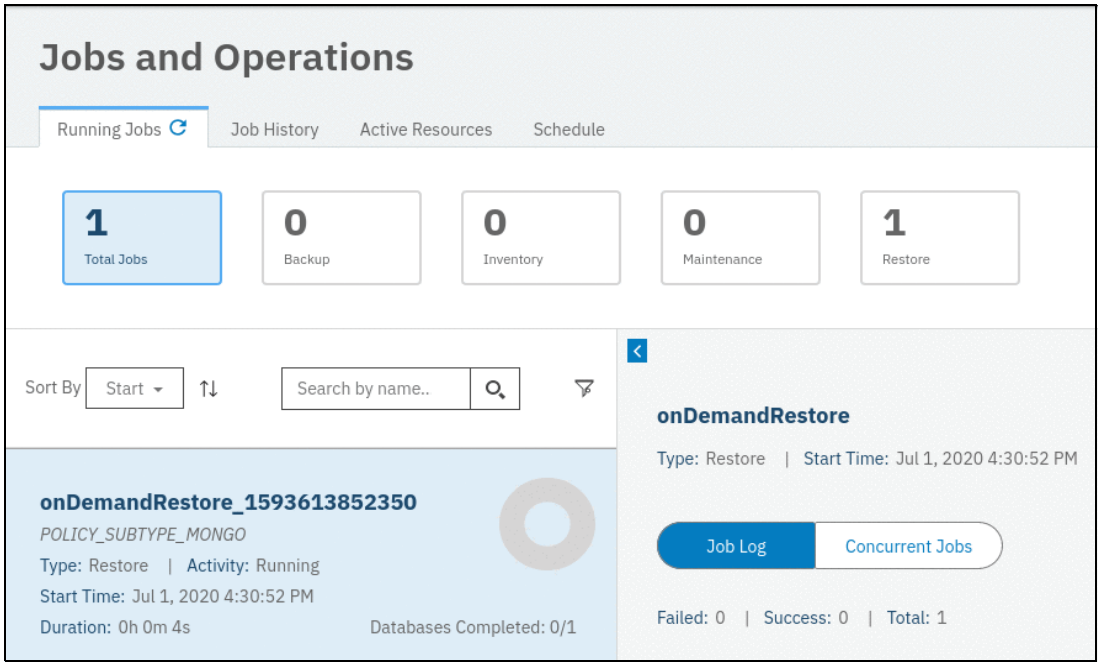

Figure 9-12 Job status view

10

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:

10.1 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 rollforward) modes are available in IBM Spectrum Protect Plus, which includes point-in-time recovery by using the archive logs.

10.2 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 10-1.

Example 10-1 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.

- 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 IBM Redbooks 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 10-1.

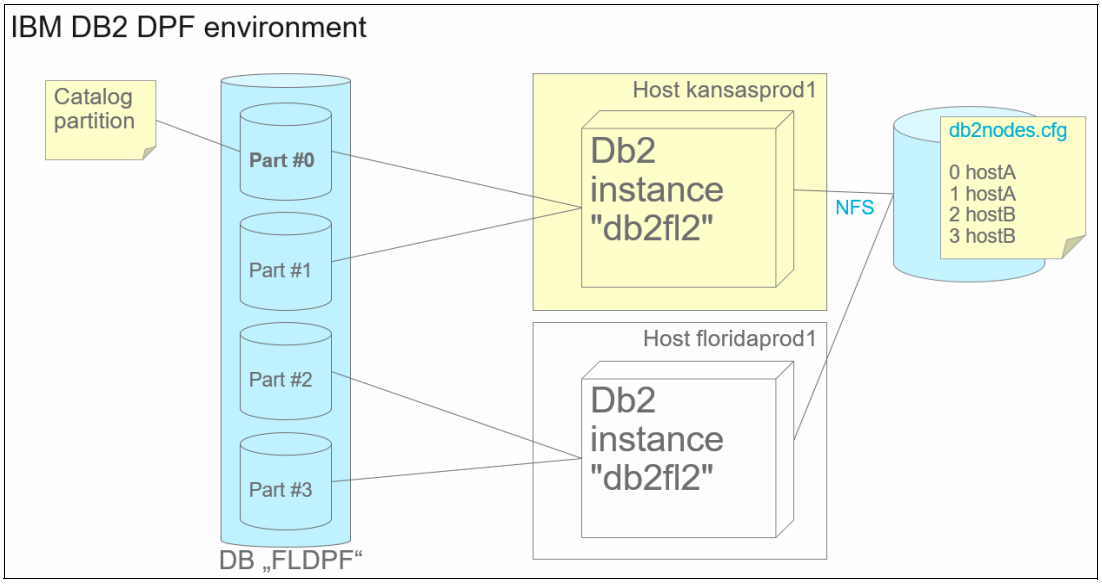

Figure 10-1 Db2 DPF environment

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 10-2.

Example 10-2 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 10-2 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 -all 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.

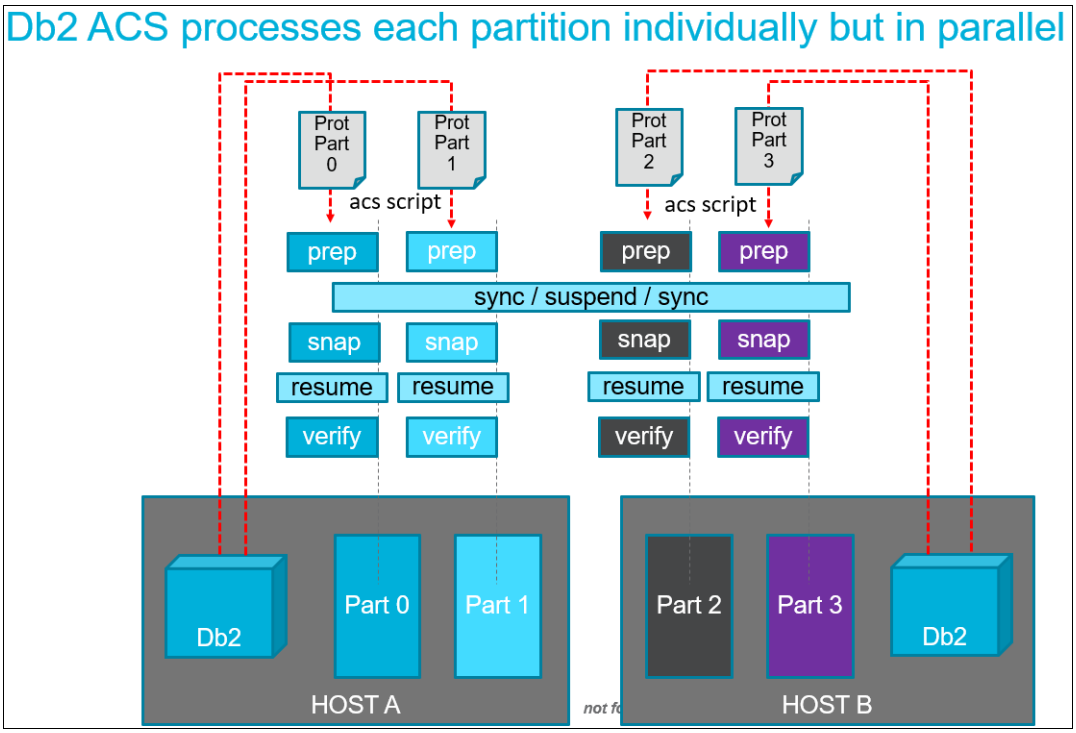

Figure 10-2 Parallel backup mode with Db2 Advanced Copy Services (ACS)

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 10-2. 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.

10.3 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.

10.3.1 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 10-3.

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 sappagent user can be reused for multiple Db2 servers.

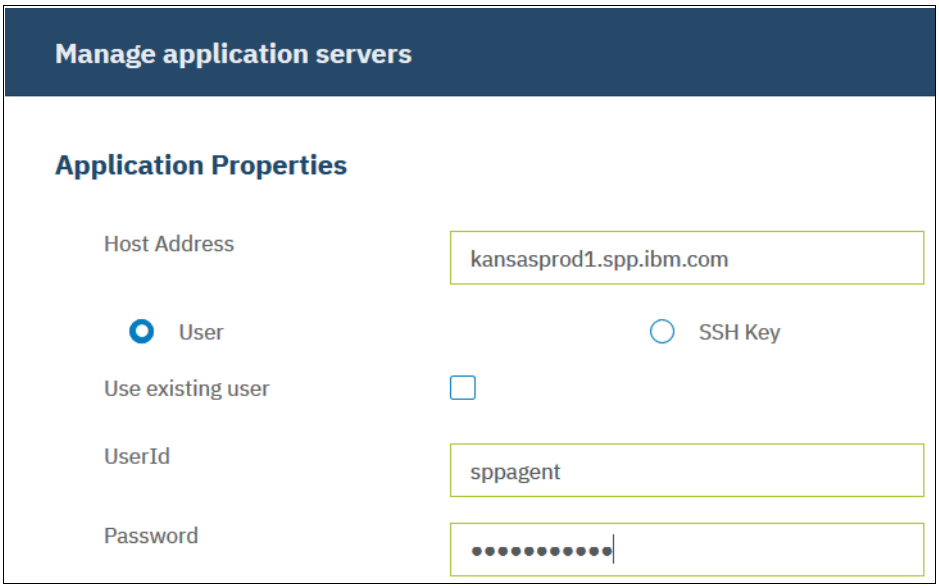

Figure 10-3 Add Db2 application server pane

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 10-4 on page 296.

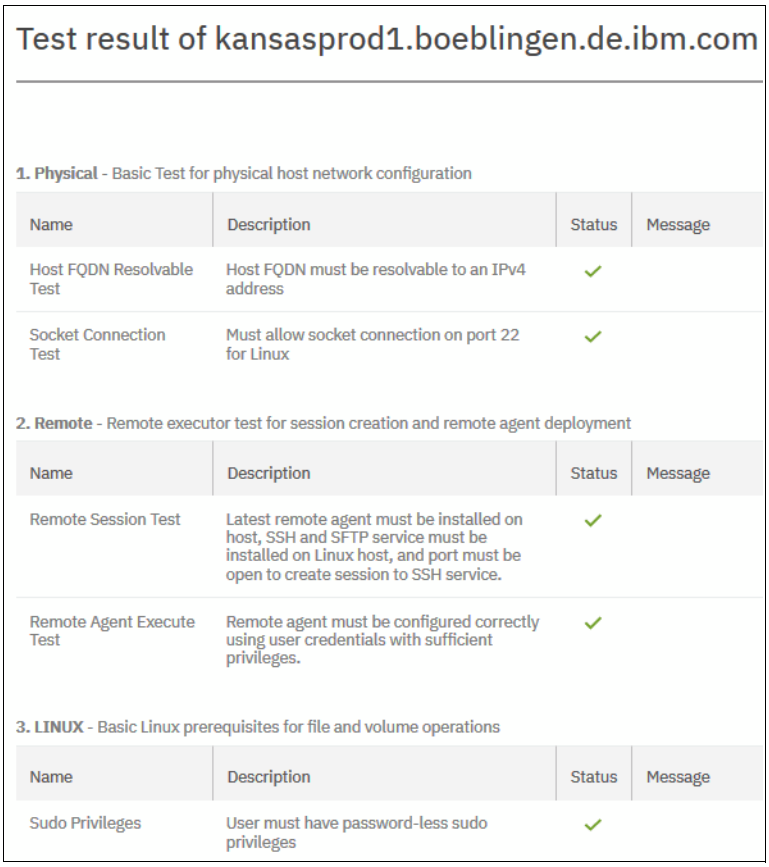

Figure 10-4 Test result pop-up

10.3.2 Backup Db2 data

Before starting a backup of Db2, the Db2 database must be assigned to one or more SLA policies. For more information about defining an SLA policy, see “SLA backup policies” on page 19.

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 10-5, you can optionally click the **Select Options** button, to enable Log Backup, as shown in Figure 10-6, “Select options to enable log backup of Db2” on page 298. With log backup enabled, IBM Spectrum Protect Plus will automatically create a log backup volume and mount it to the application server.

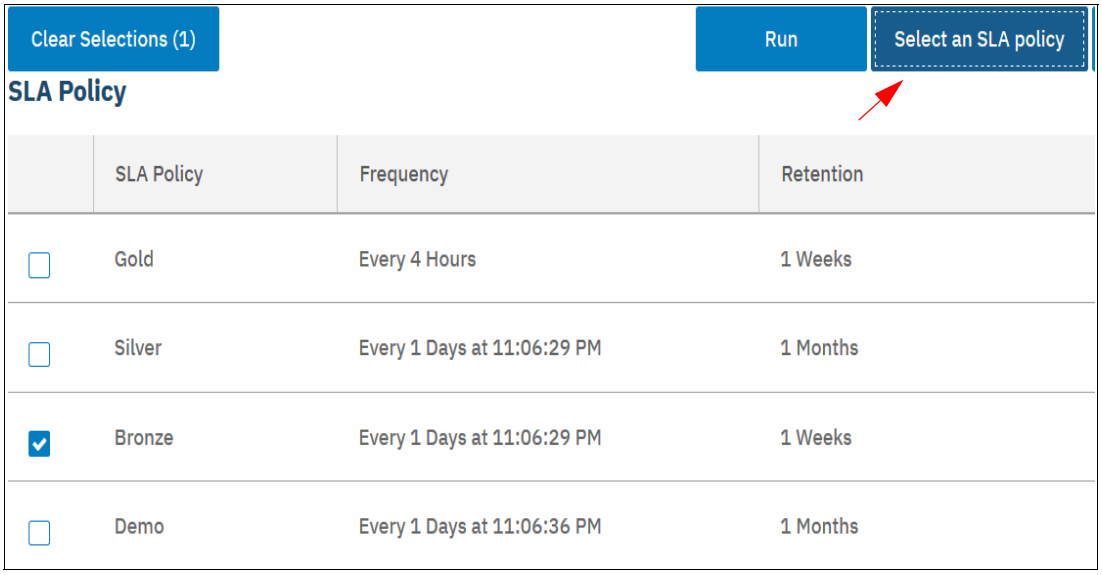

Figure 10-5 Assign a SLA policy to the database

Enable Log backup

Archived logs for databases contain committed transaction data. This transaction data can be used to run a rollforward 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 10-3. 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 10-3 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 10-6, to allow rollforward 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.

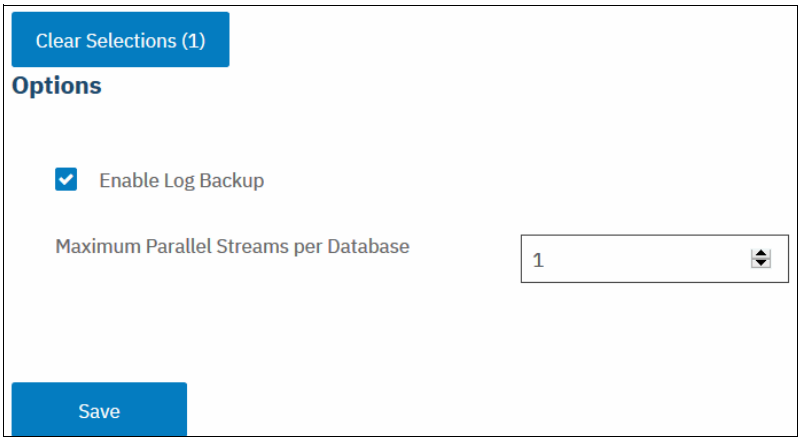

Figure 10-6 Select options to enable log backup of Db2

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 10-4, 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 10-4 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 10-5.

Example 10-5 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 db2iadm1 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 db2iadm1 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 db2iadm1 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 db2iadm1 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 db2iadm1 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 db2iadm1 12288 Jun 30 08:00 S0002034.LOG
-rw-r-----. 1 db2inst1 db2iadm1 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 10-7. 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.

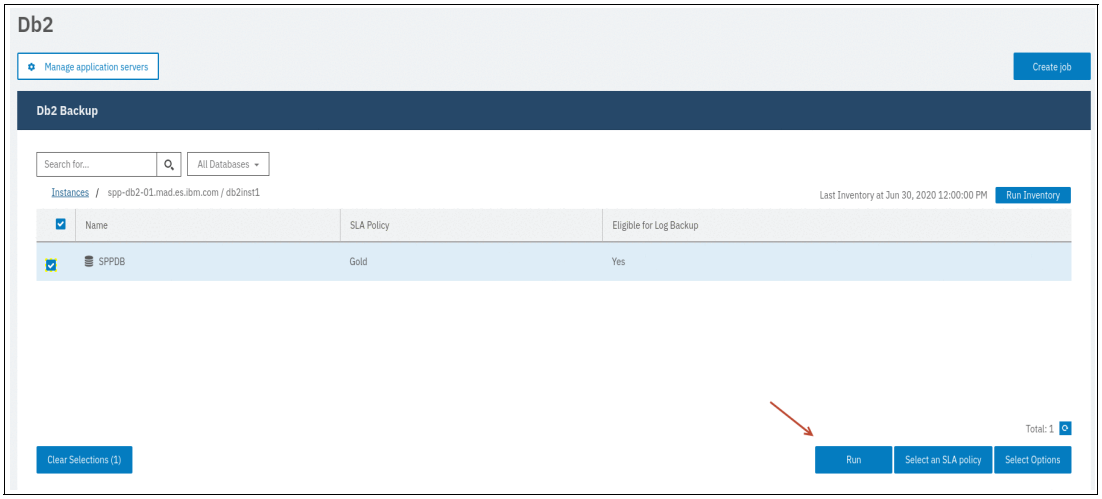

Figure 10-7 Start the Db2 backup

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 10-6, and review the vSnap server volumes.

Example 10-6 vSnap server volumes for data and log backup

```
[root@spp-db2-01 COOOOOOO]# 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 NODEXXXX 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 10-8 on page 301.

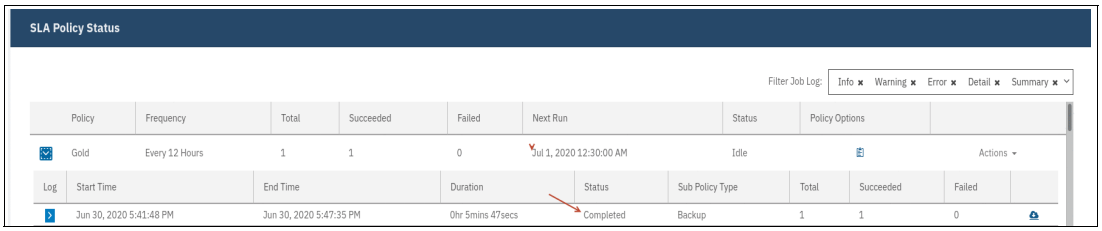

Figure 10-8 Db2 backup job completed

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 8.5, “Database backup with pre-script and post-script” on page 274.

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.

10.3.3 Restoring Db2 databases

IBM Spectrum Protect features a restore wizard (see Figure 10-9) 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 10-9, 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 259.

As shown in Figure 10-9, the user must select the Db2 database that requires a restore.

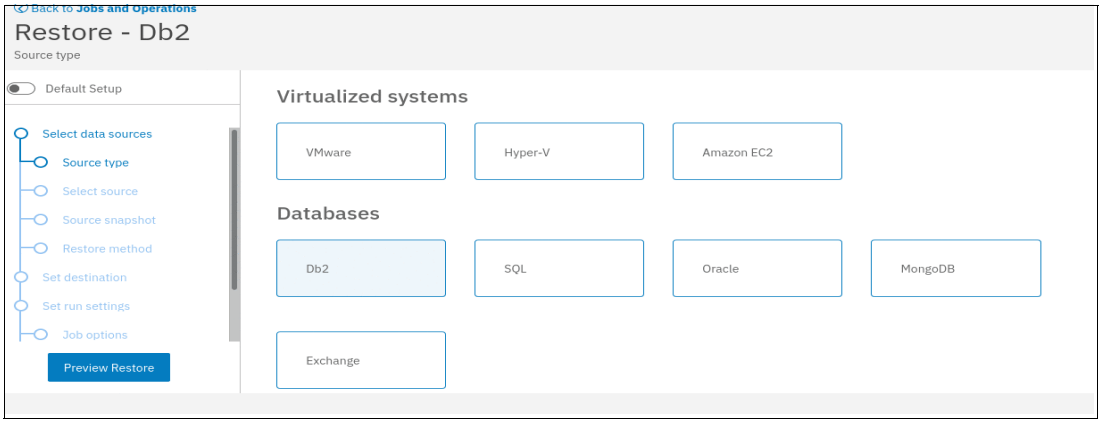

Figure 10-9 Spectrum Protect Plus restore wizard

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

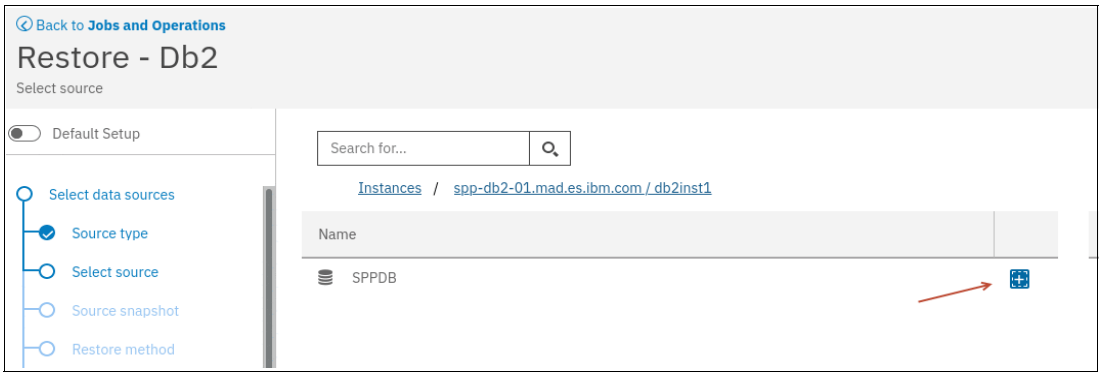

Figure 10-10 Db2 source database

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 10-11. Here, On-demand Point in Time was selected.

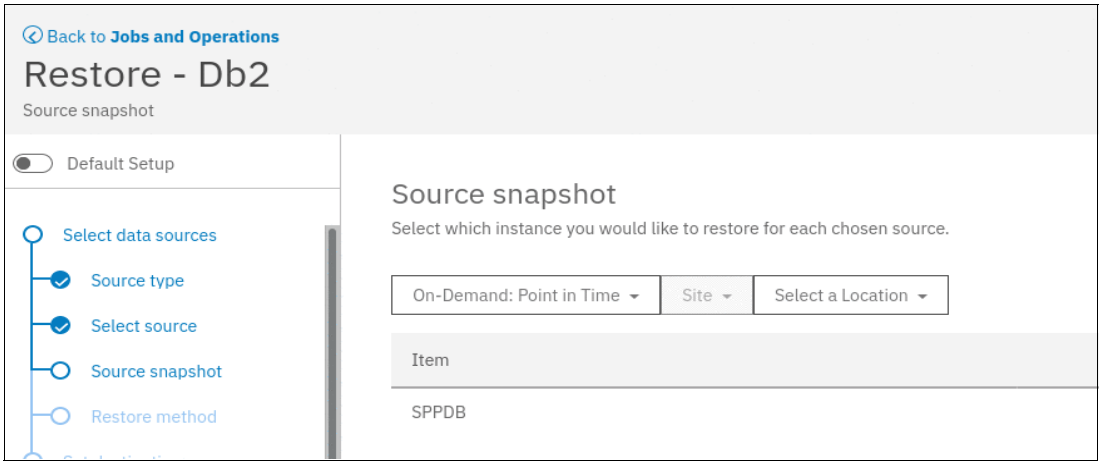

Figure 10-11 Select type of restore: On-Demand: Point in Time

Then, choose a restore location, as shown in Figure 10-12. 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 10-12.

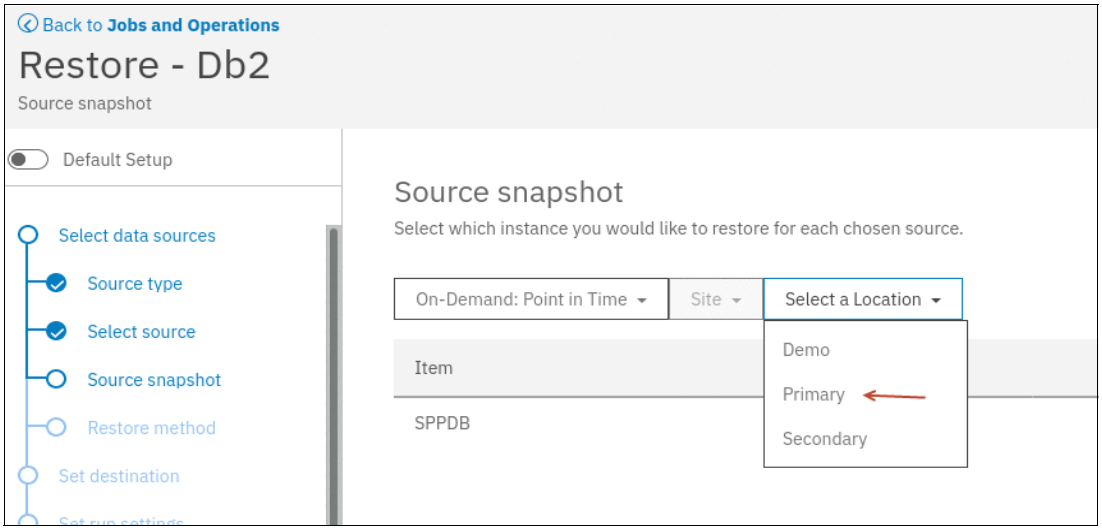

Figure 10-12 Select restore location type

There are three restore methods available, as shown in Figure 10-13. 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.

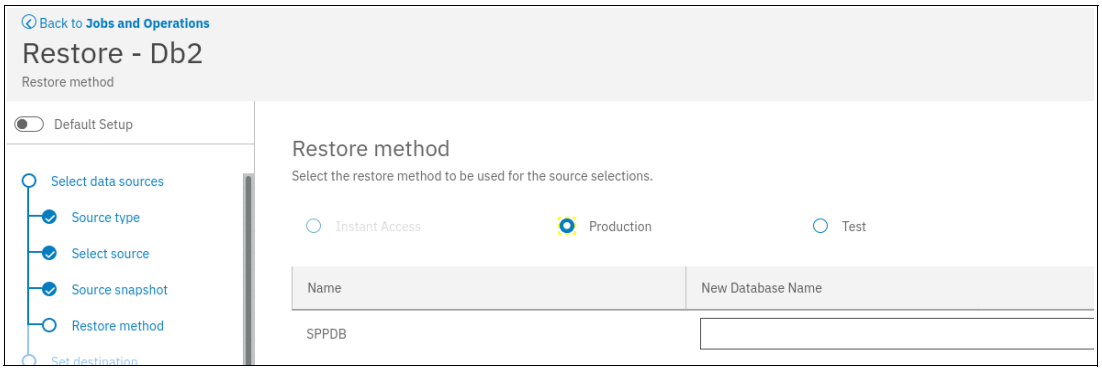

Figure 10-13 Select restore method

As shown in Figure 10-14, click **Restore to original instance** to restore to the Db2 production server.

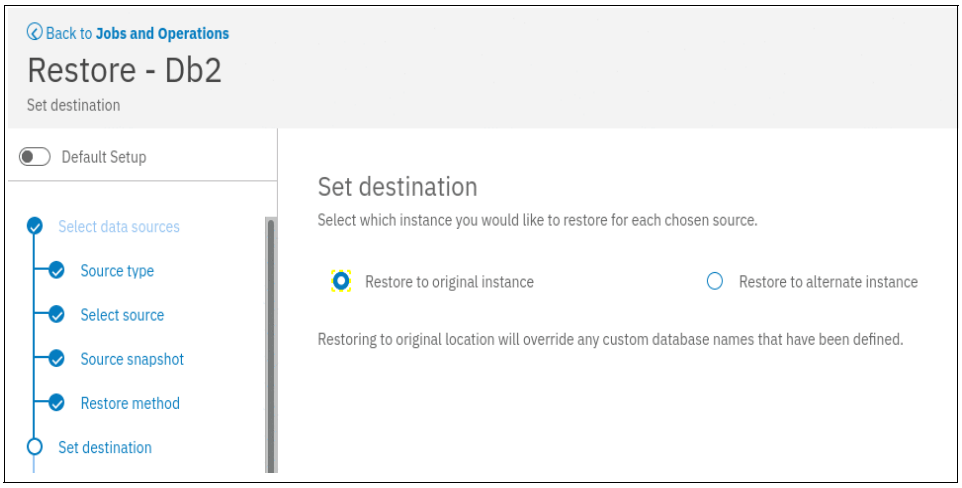

Figure 10-14 Select destination for the restore

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 10-15.

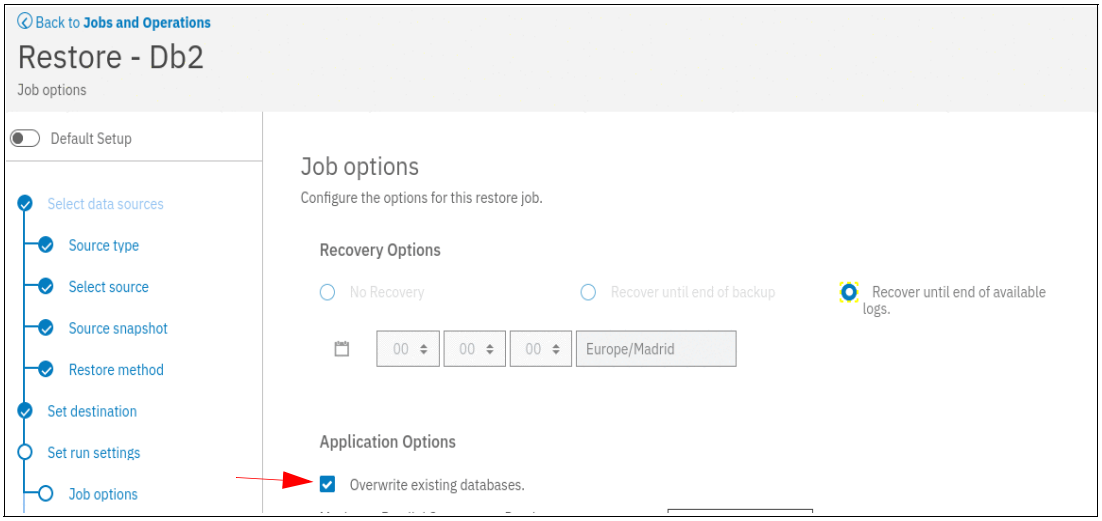

Figure 10-15 Specify restore job options: Overwrite existing databases

Another option, as shown in Figure 10-16, 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.

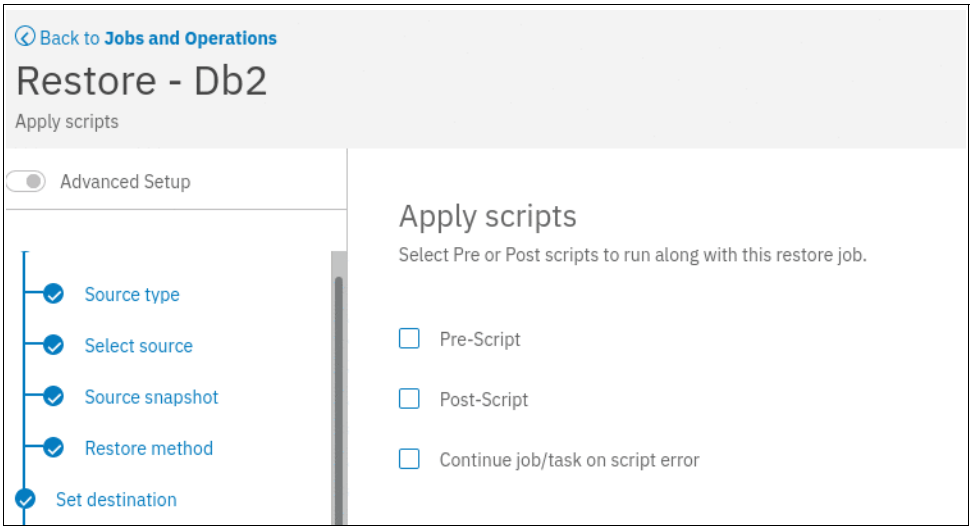

Figure 10-16 Specify scripts for the restore job

Finally, the Review page is displayed and after checking all values, the on-demand restore job can be submitted. See Figure 10-16. To start the on-demand restore job, click **Submit**.

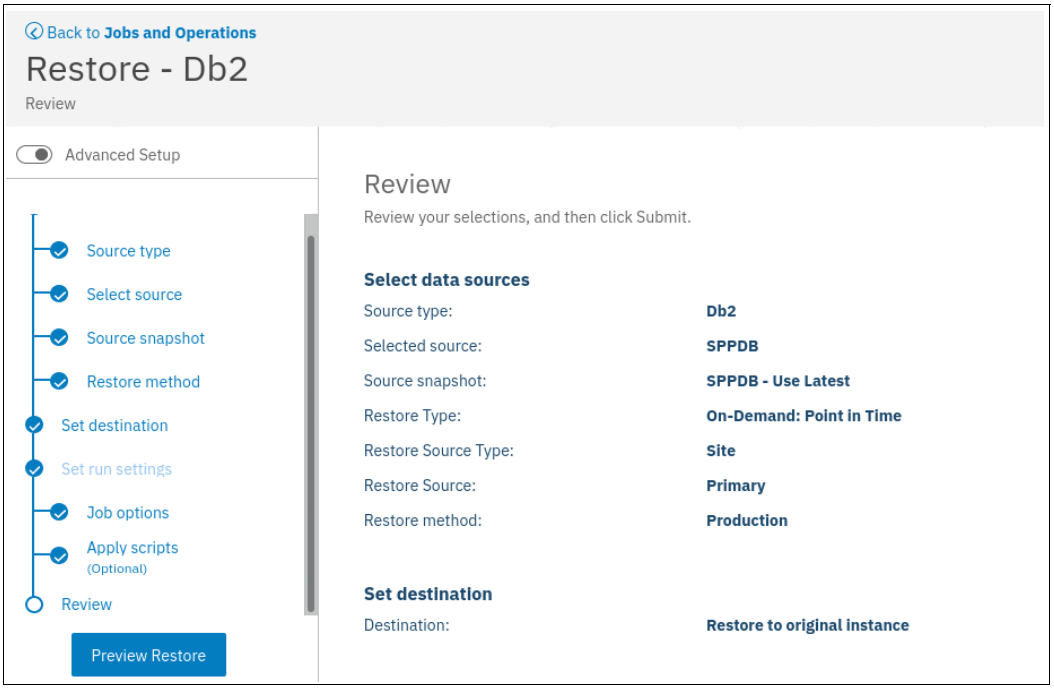

Figure 10-17 Review of restore job parameters.

The restore job can be monitored by selecting **Jobs and Operations** → **Running Jobs**, as shown in Figure 10-18. When the restore job finishes, it is removed from the Running Jobs list.

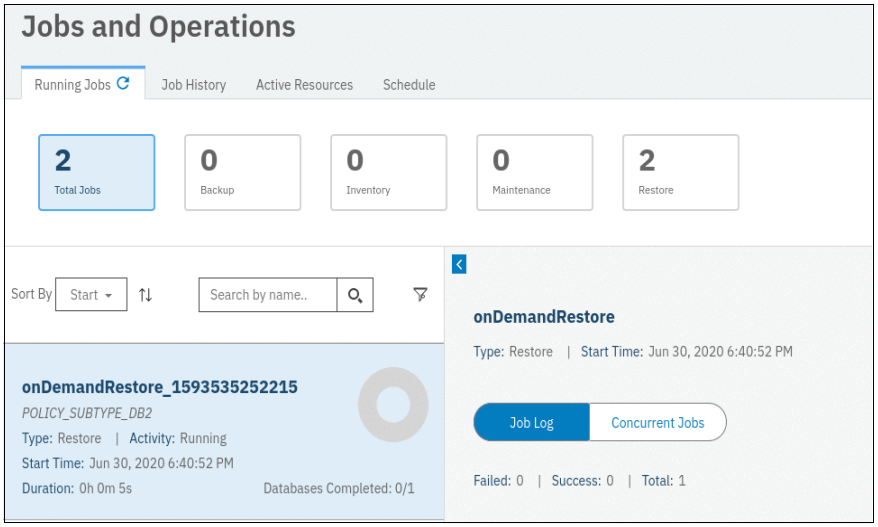

Figure 10-18 Monitor the Db2 restore job

11

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:

11.1 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

11.2 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 11-1.

```
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 11-1 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 11-2.

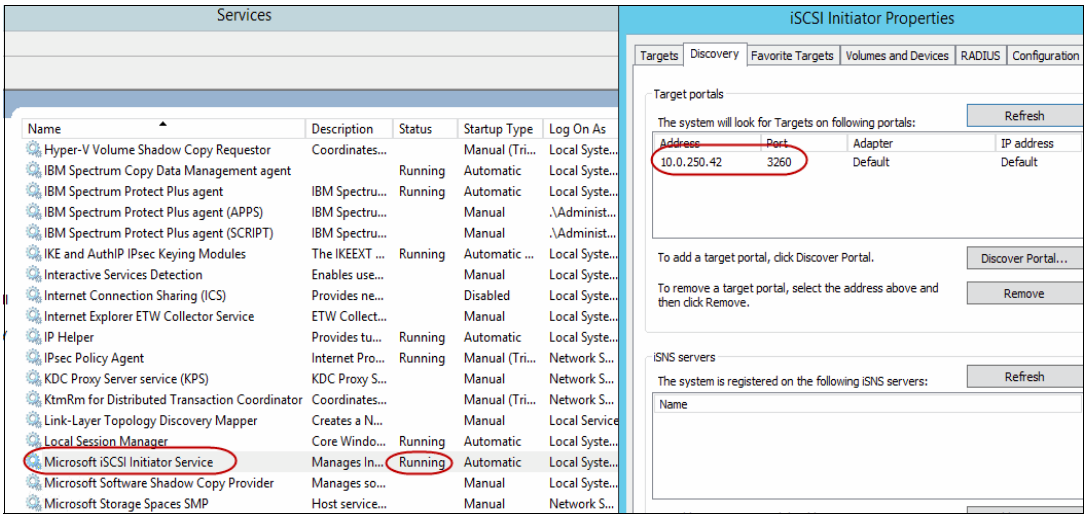

Figure 11-2 Microsoft iSCSI Initiator running on MIcrosoft SQL Server

- 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 11-3.

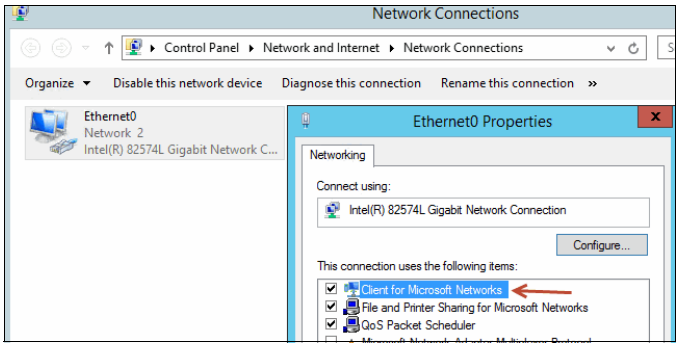

Figure 11-3 Client for Microsoft Networks option enabled into MIcrosoft SQL Server

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

11.3 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.

11.3.1 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 11-4. In this example, the IBM Spectrum Protect Plus admin includes predefined the SQL Server Admin in **Accounts** → **Identity** → **Add Identity**.

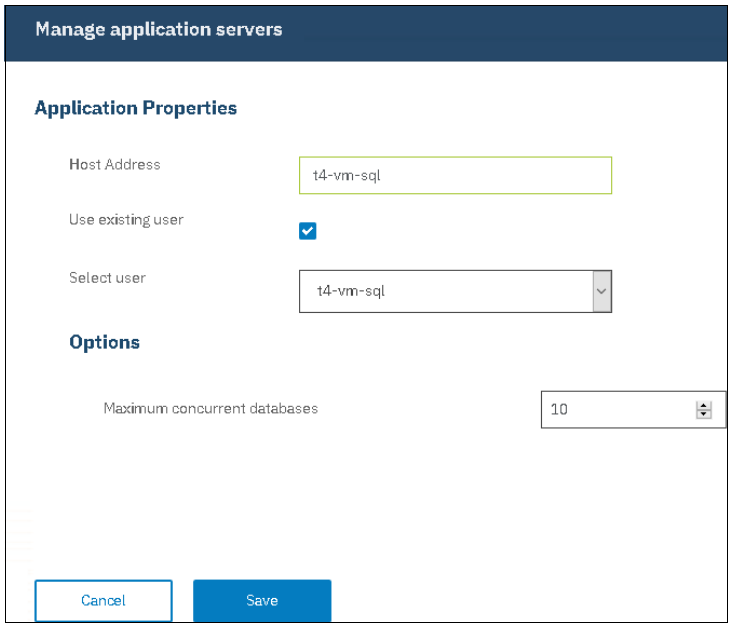

Figure 11-4 Register SQL application server

Perform a configuration test of the newly assigned SQL Server in IBM Spectrum Protect Plus, as shown in Figure 11-5 on page 313.

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.

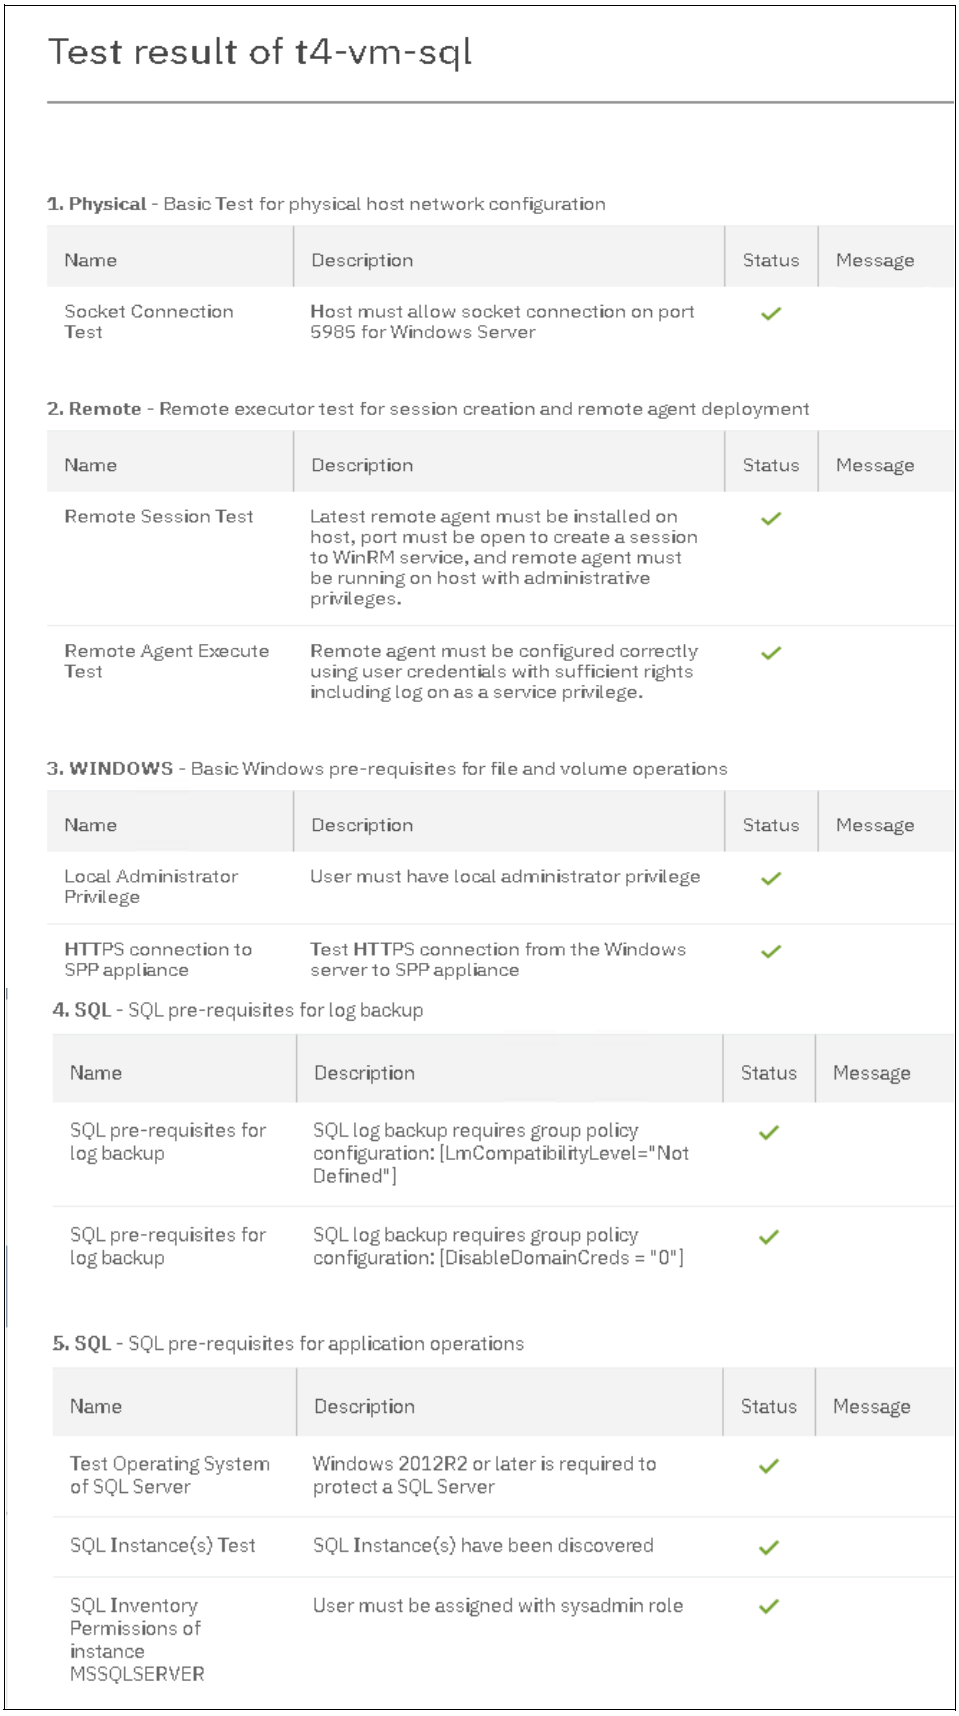

Test result of t4-vm-sql

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 11-5 SQL Server configuration test results

11.3.2 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. How to define an SLA policy is described in “SLA backup policies” on page 19.

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 11-6. 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.

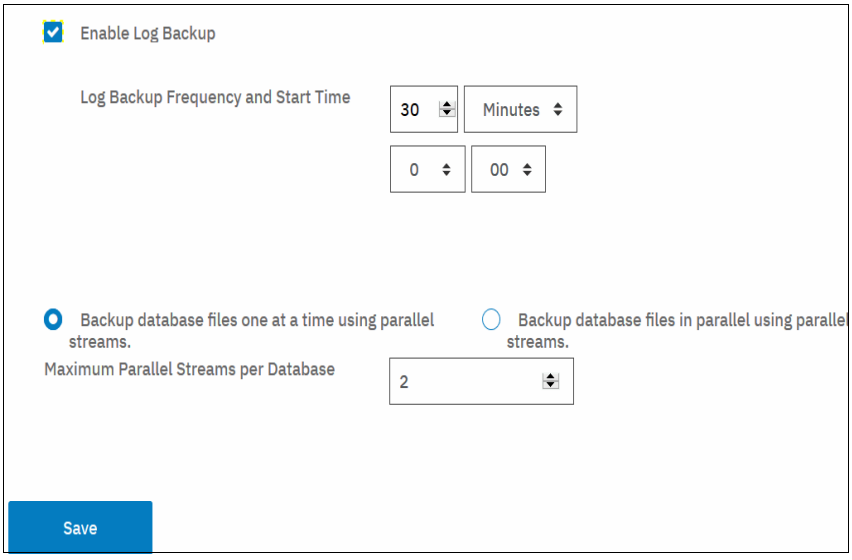

Figure 11-6 Enable SQL Server log backup

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 11-7.

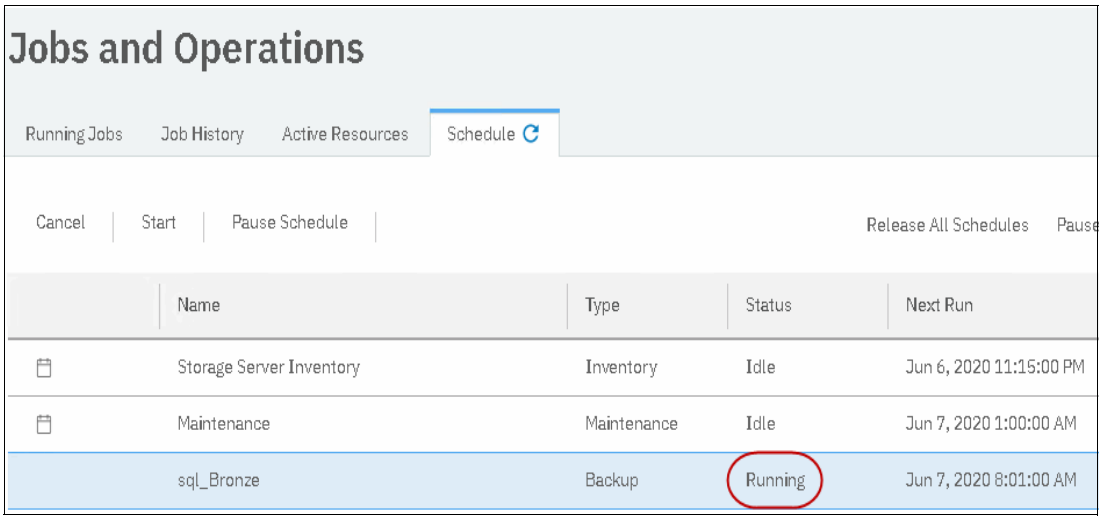

Figure 11-7 SQL Server backup job with Status: Running

It also can be monitored by selecting **Jobs and Operations** → **Running Jobs** → **Progress**, as shown in Figure 11-8.

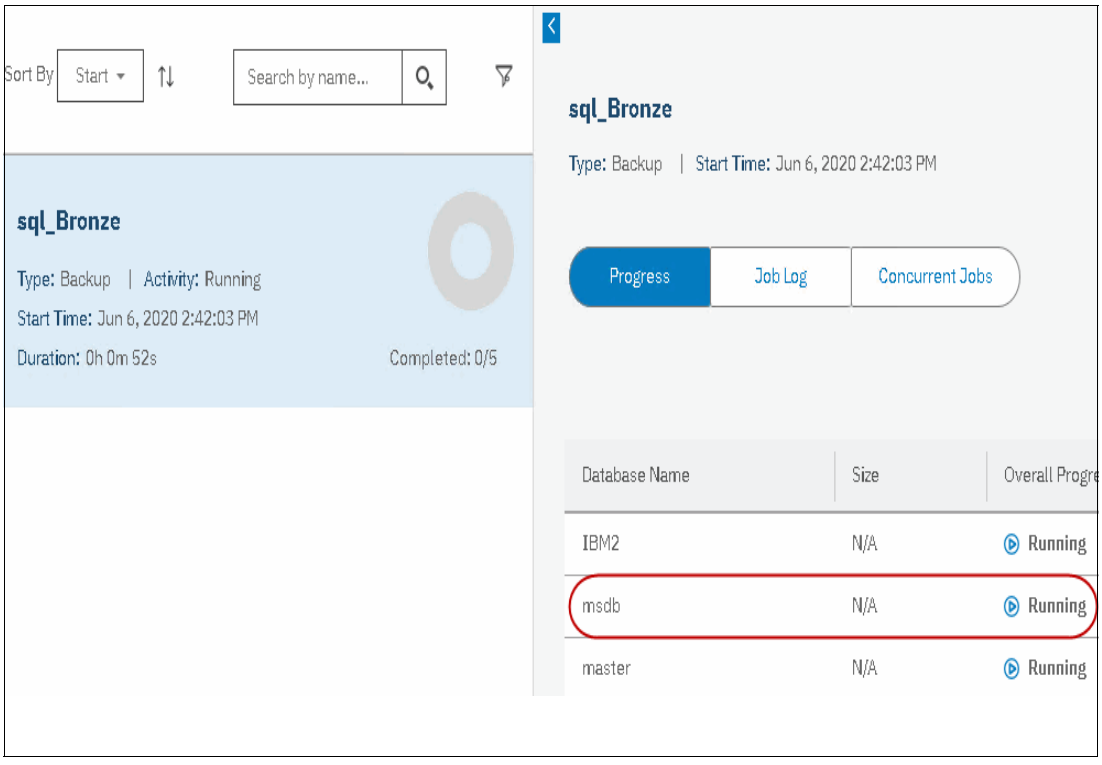

Figure 11-8 SQL Server Databases Backup Status: Running

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.s
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.

11.3.3 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 11-9.

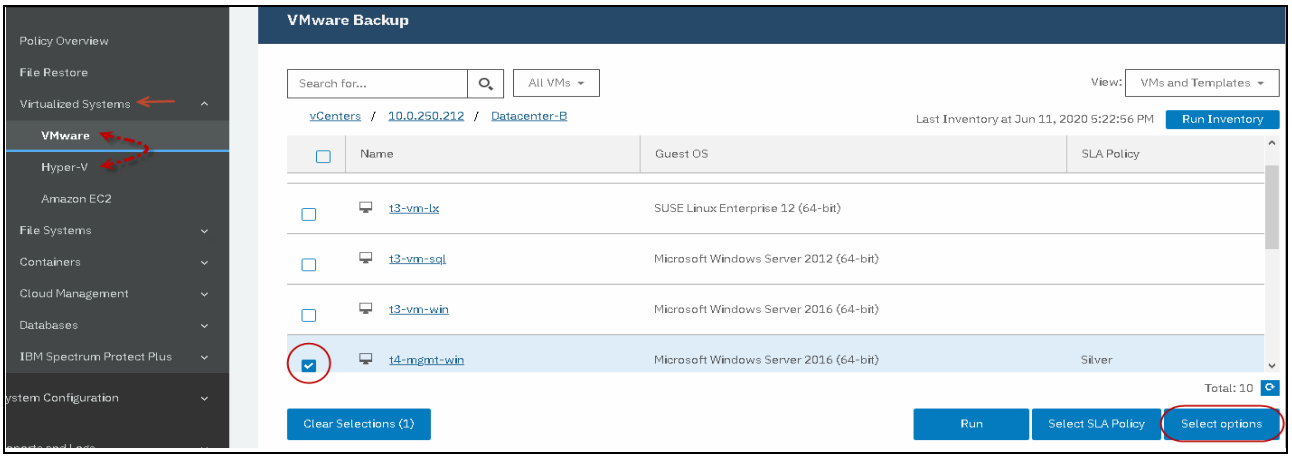

Figure 11-9 Selected options for VM Server

Under **Agent Options,** select the **Truncate SQL Logs** option, as shown in Figure 11-10. Click **Save**.

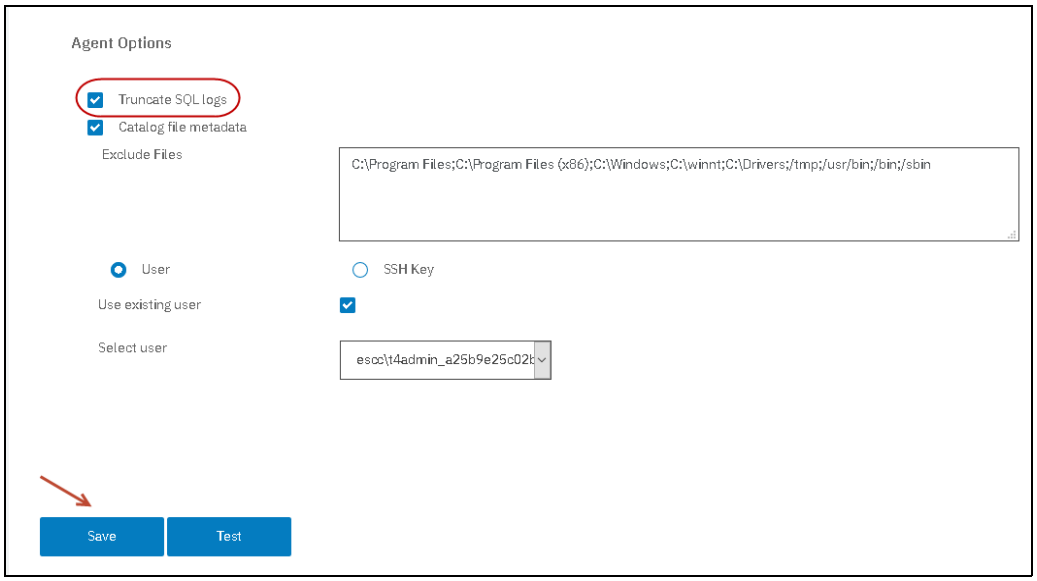

Figure 11-10 Truncate SQL Logs Option Enabled

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 11-11.

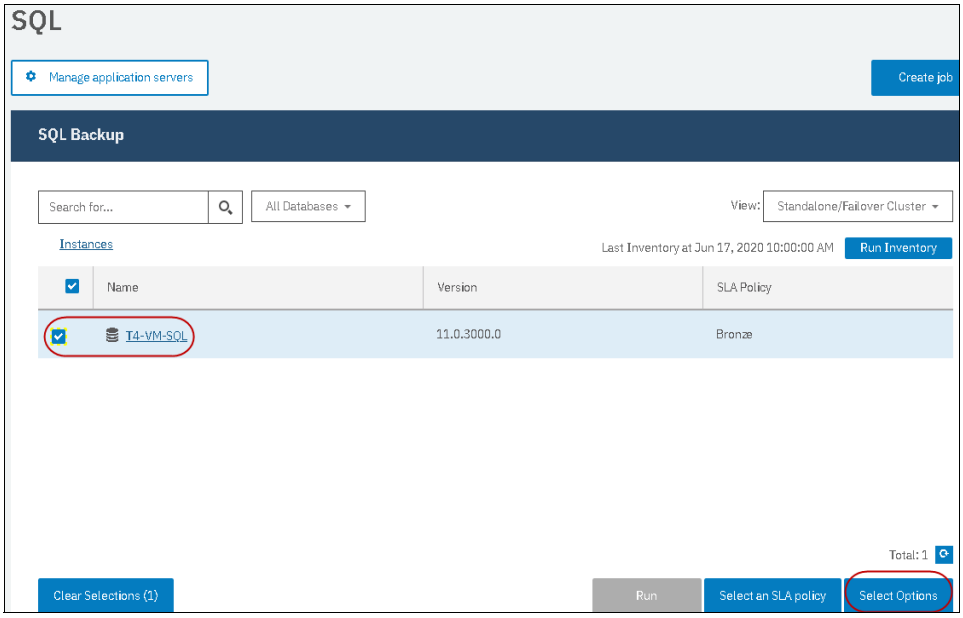

Figure 11-11 Select Options for SQL Instance

Under Options, select **Enable Log Backup** and define a Log Backup Frequency, as shown in Figure 11-12.

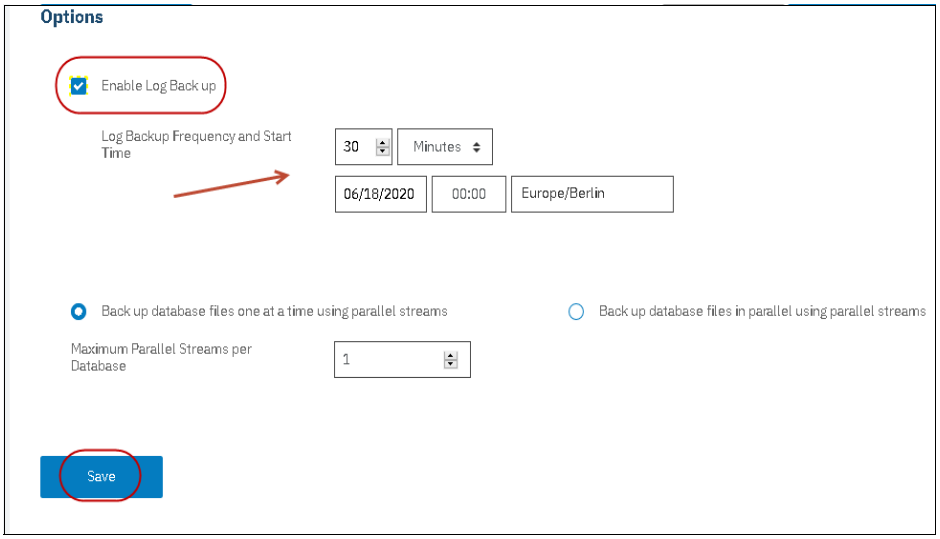

Figure 11-12 Enable Log Backup option enabled

The enabled SQL log backup schedule option can also be reviewed in the Microsoft SQL Server system, as shown in Figure 11-13 on page 320 under **Task Scheduler** → **Task Scheduler Library** → **IBM** → **SPP Windows Agent**.

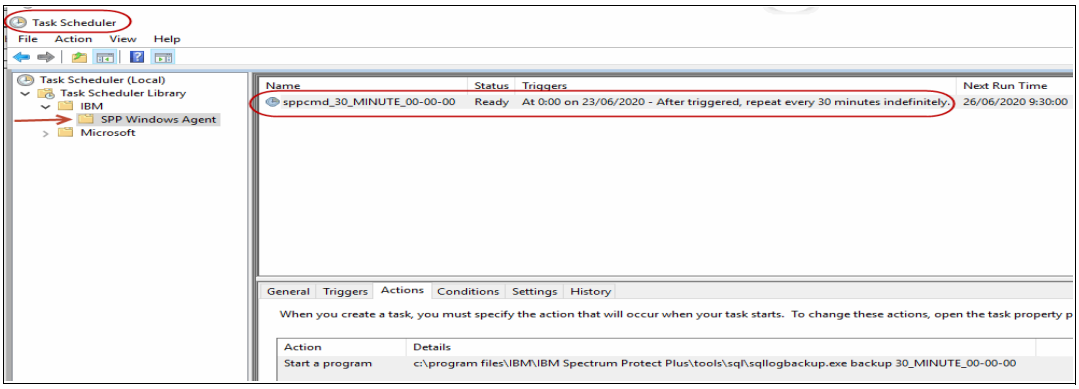

Figure 11-13 SQL Log Backup Scheduler in Microsoft SQL Server

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 available at IBM Knowledge Center.

11.3.4 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 11-1, lists the active share where the Volume ID `1` and the file system name `/vpool1_fs2` can be identified.

Example 11-1 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
```

```
    READ ONLY: No
```

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 11-2.

Example 11-2 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-vsnapdatalv   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]$
```

11.3.5 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 11-14.

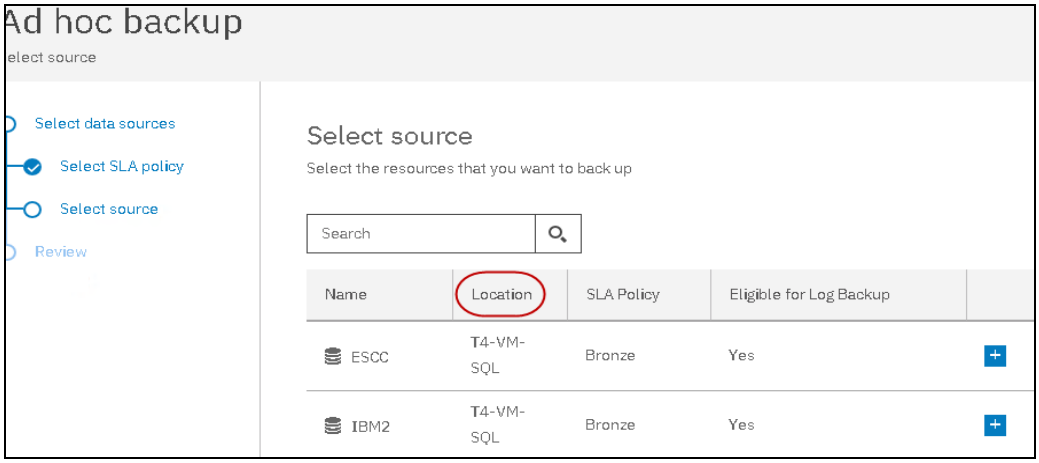

Figure 11-14 Ad Hoc SQL Backup showing Name, Location and SLA Policy

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 11-15.

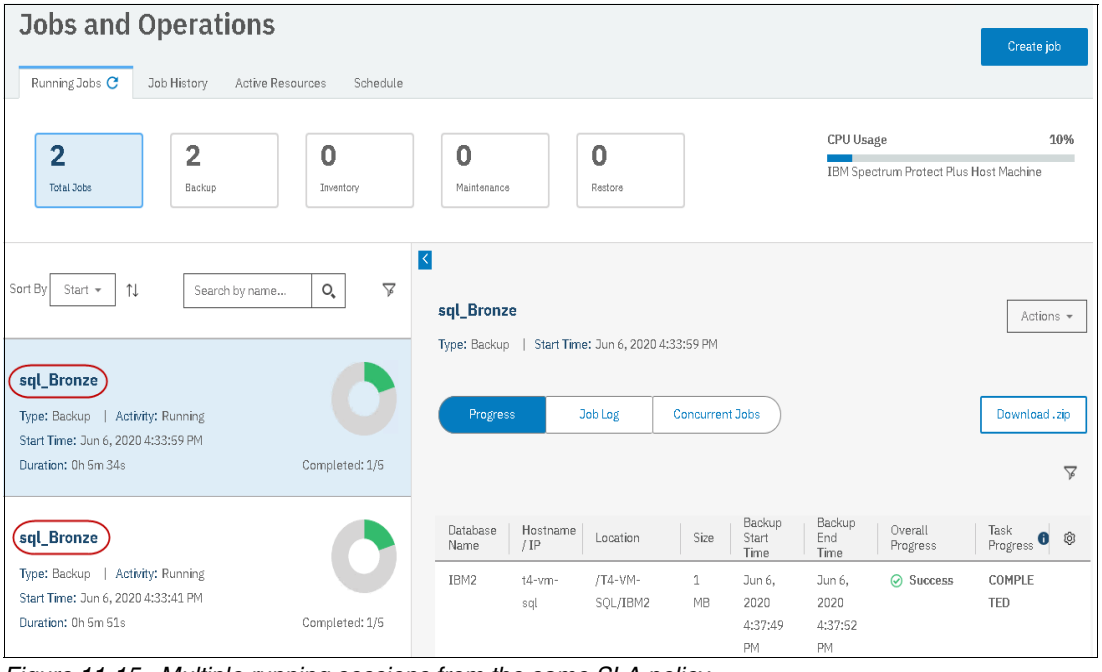

Figure 11-15 Multiple running sessions from the same SLA policy

11.3.6 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 11-16

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.

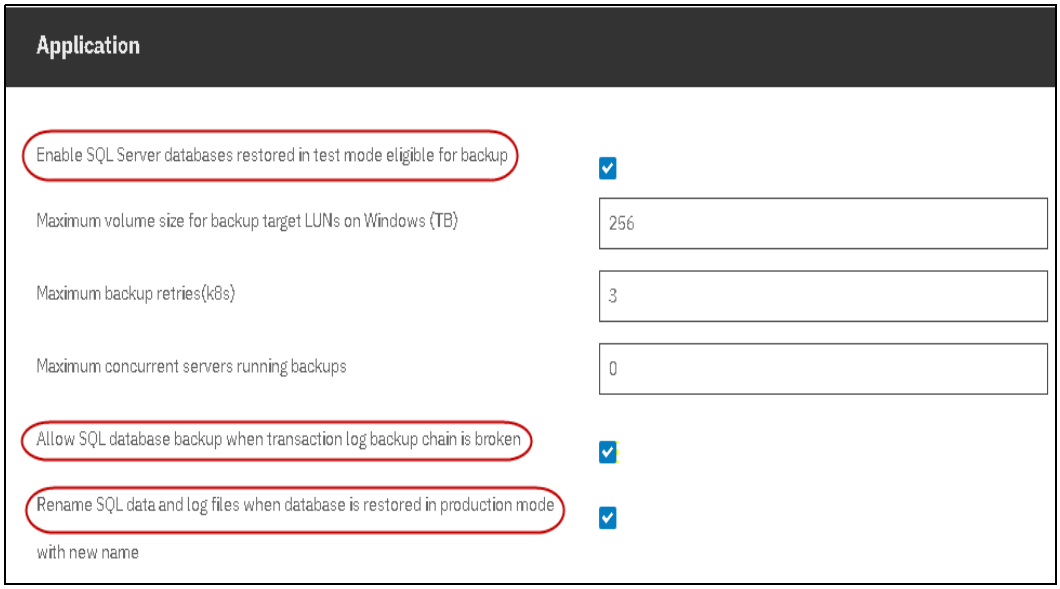

Figure 11-16 Global Preferences SQL Application Options

11.4 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 chose 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 11-17 on page 325, 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 259.

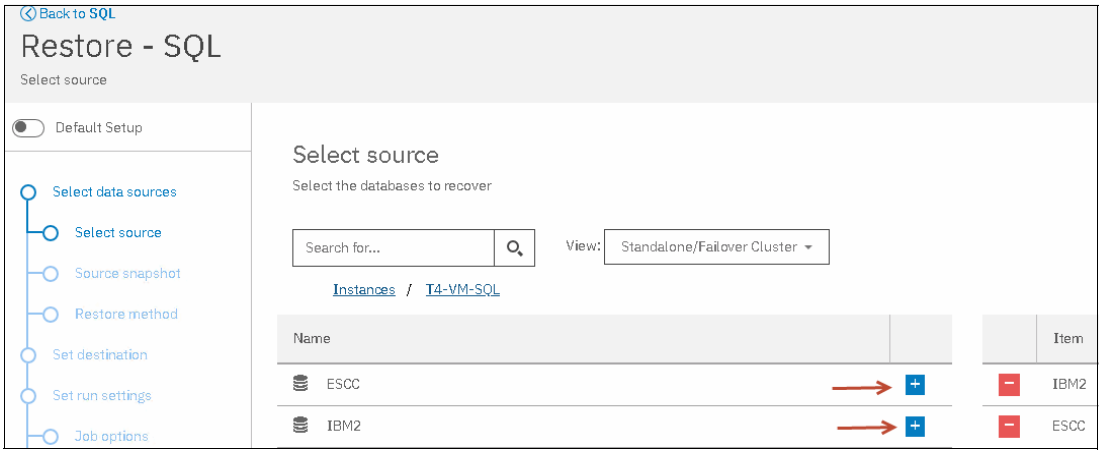

Figure 11-17 Select the SQL Server backup source

Other restore parameters that must be specified are shown in Figure 11-18:

- 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.

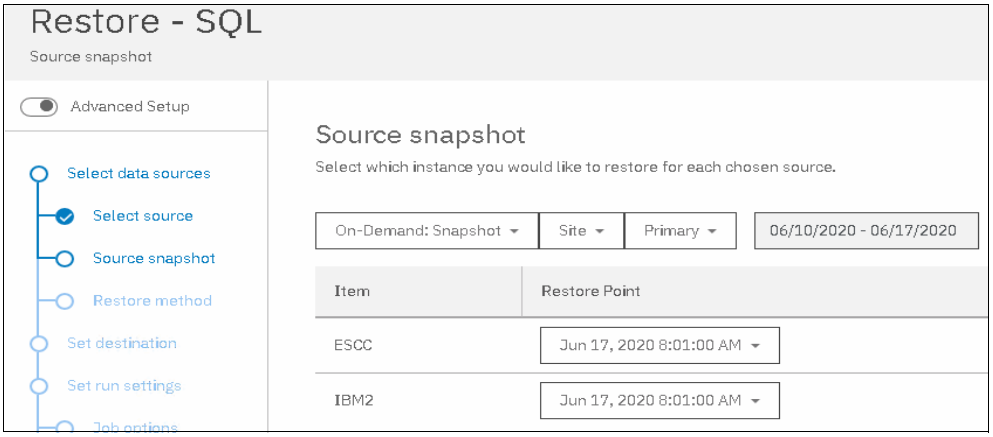

Figure 11-18 Restore parameters

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 11-19 and then, click **Next**.

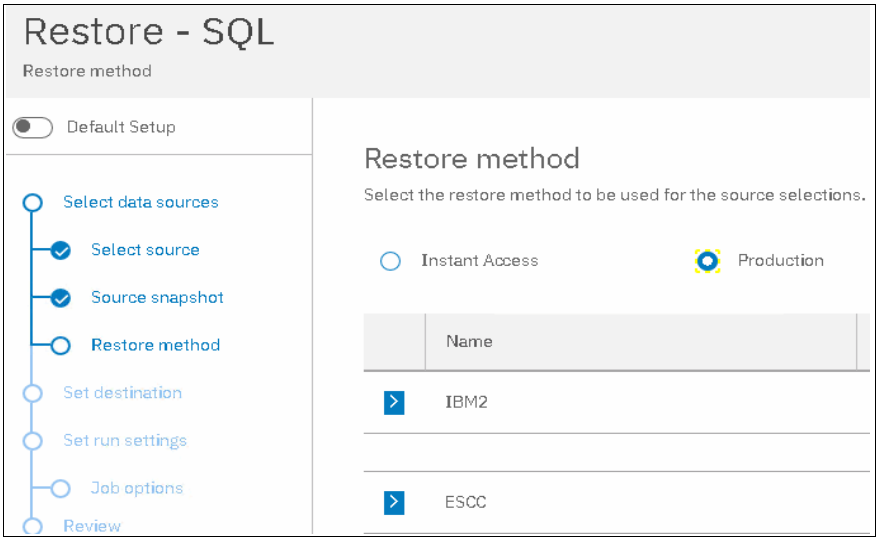

Figure 11-19 Select the restore method

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 11-20.

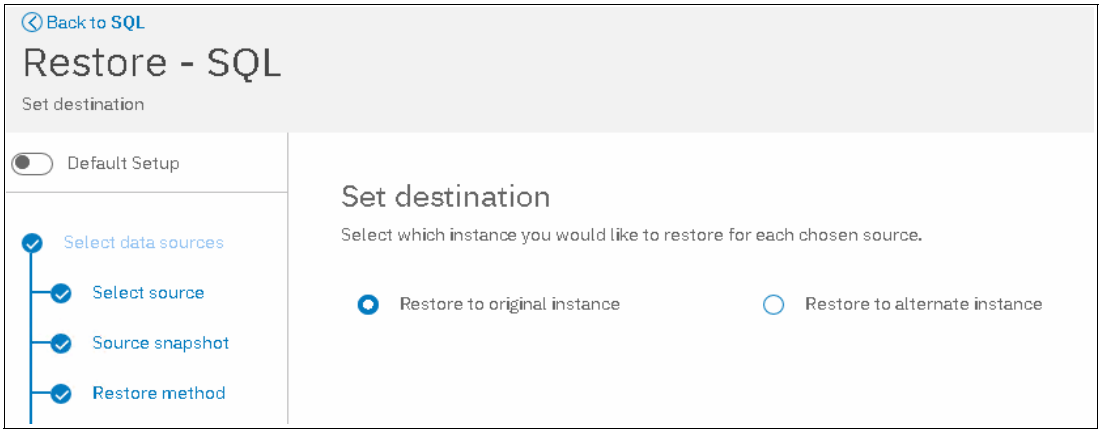

Figure 11-20 Select the restore destination

Enable the restore job to overwrite the selected database. By default, this option is not enabled, as shown in Figure 11-21 on page 327.

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.

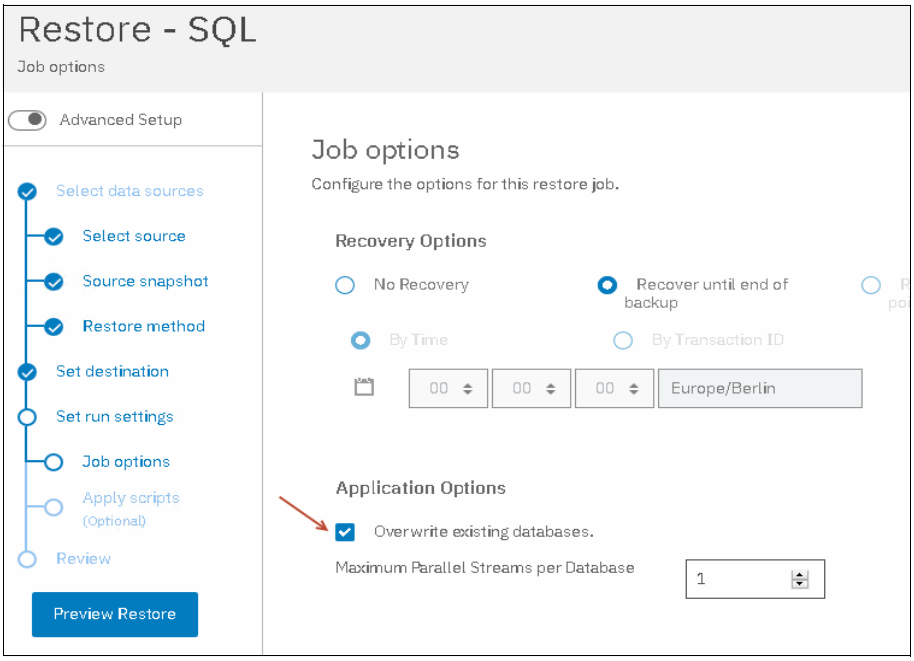

Figure 11-21 Select restore job options

In the Review page, check all entered restore job parameters, as shown in Figure 11-22. Click **Submit** to start the on-demand restore job.

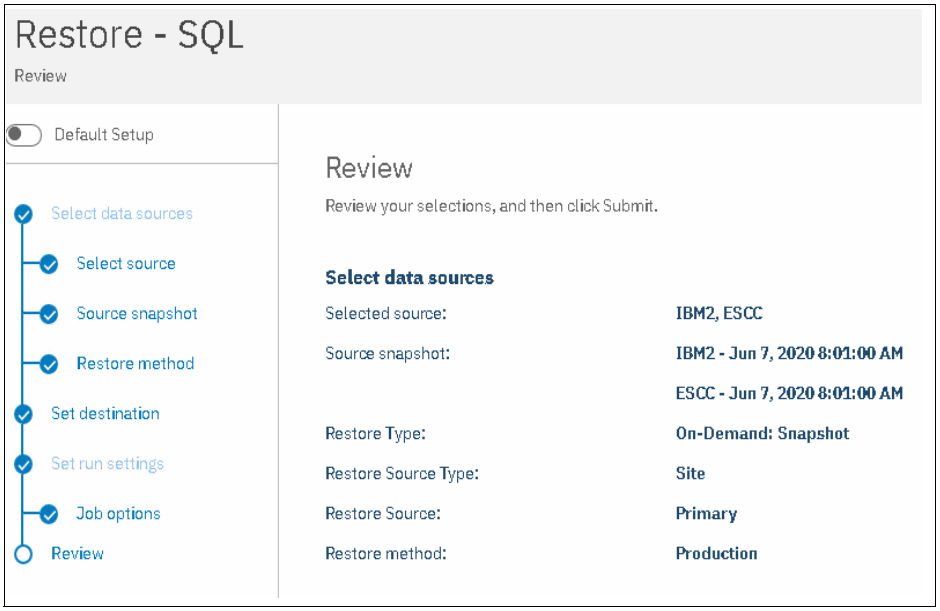

Figure 11-22 SQL Server restore summary

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 11-23, the vSnap server backup volume is mounted as Disk1 during the restore job.

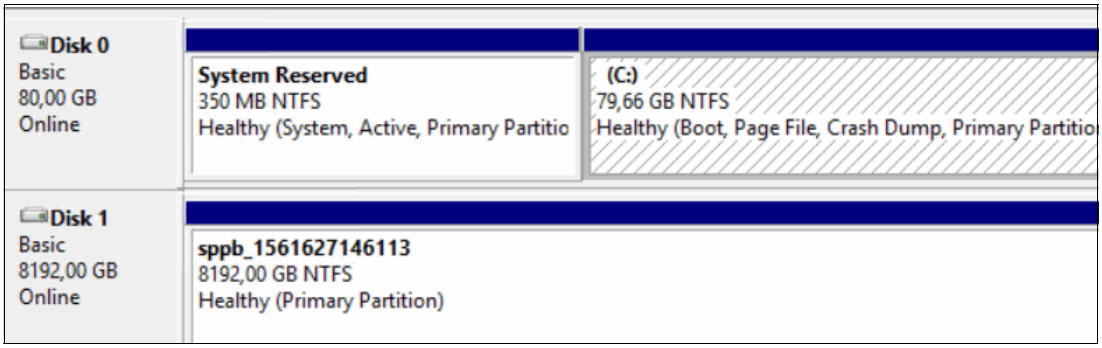

Figure 11-23 Mount of vSnap server volume for the restore on the SQL application server

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 11-24, we select the new Standby mode Job Options and perform the restore similar to the previous example.

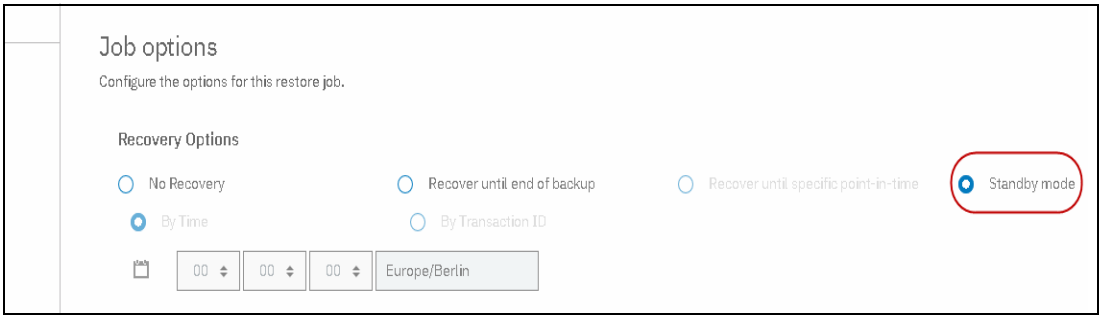

Figure 11-24 Select restore Standby mode job options

IBM Spectrum Protect Plus mounts the vSnap server backup volume at the SQL application server and copies the backup data to the source. Figure 11-25 shows that the new Database `IBM2_TEST` was restored with Standby/Read-Only mode.

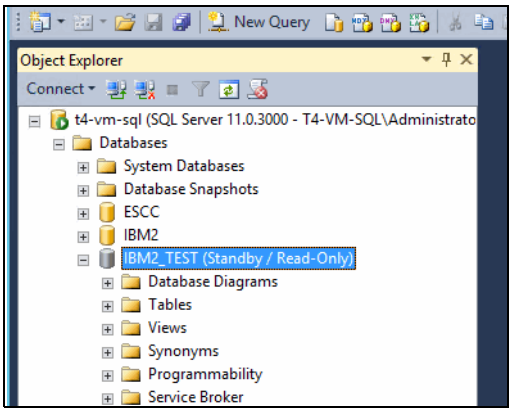

Figure 11-25 Database in standby mode after the restore

12

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:

12.1 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.

12.1.1 Server roles

Depending on the Microsoft Exchange release, different server roles are available that must be protected, as listed in Table 12-1.

Table 12-1 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*.

12.1.2 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.

12.1.3 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.

12.1.4 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.

12.2 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 Knowledge Center.

12.2.1 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.ex`, 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 12-1
- Exchange server (`"epc-exchange.xxxxxxx.lab,192.168.111.167"`), as shown in Example 12-2 on page 335

Example 12-1 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 12-2 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 12-1.

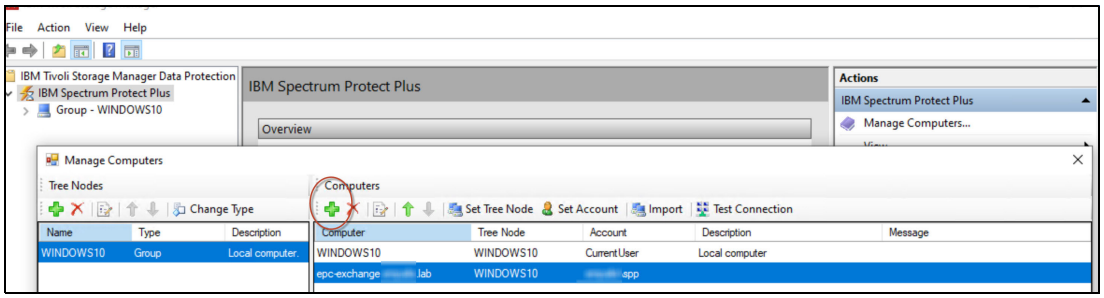

Figure 12-1 Adding the Application Server

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

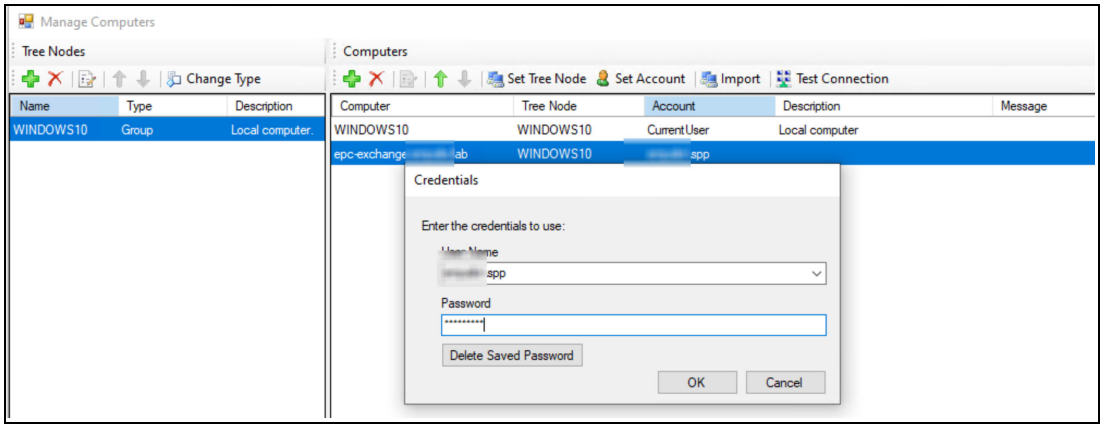

Figure 12-2 Managing credentials

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

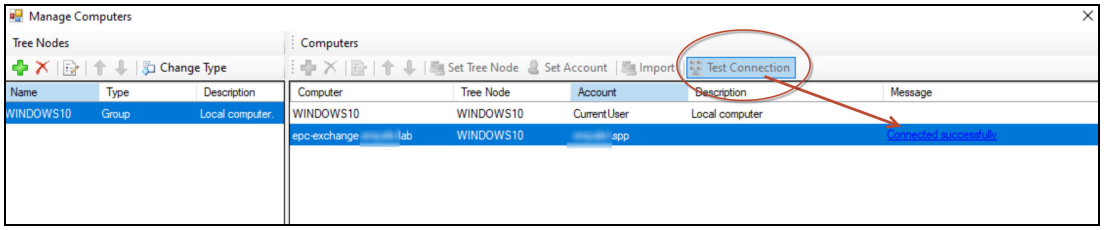

Figure 12-3 Testing the connection

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 12-4.

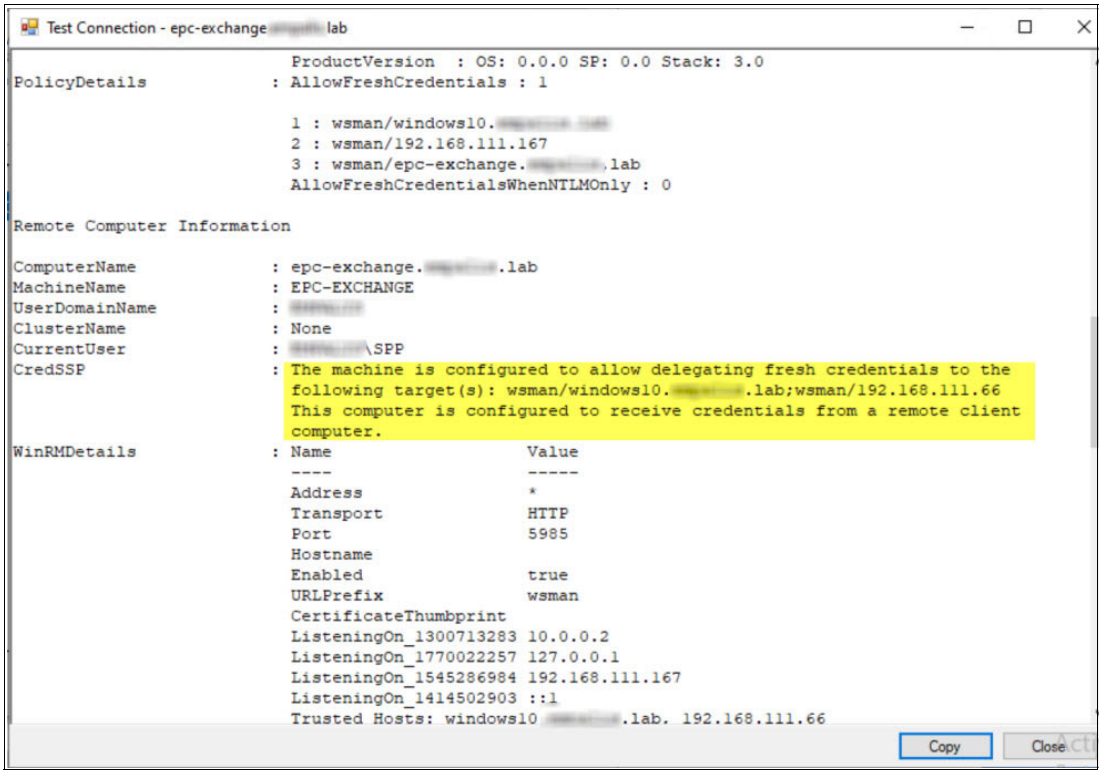

Figure 12-4 Successful connection details

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.

12.3 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 12-5 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.

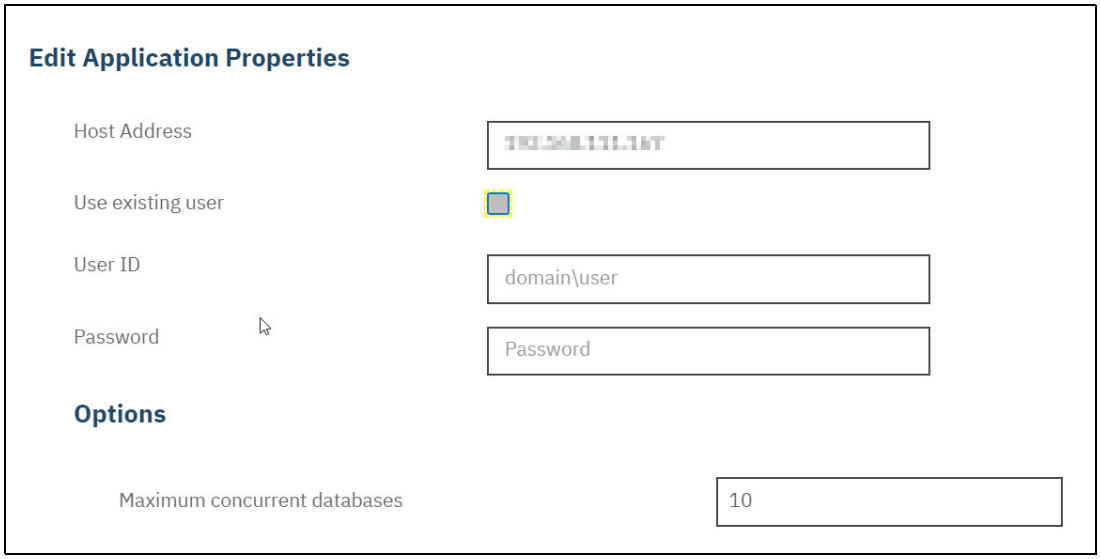

Figure 12-5 Edit Application Properties window

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 12-6 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.

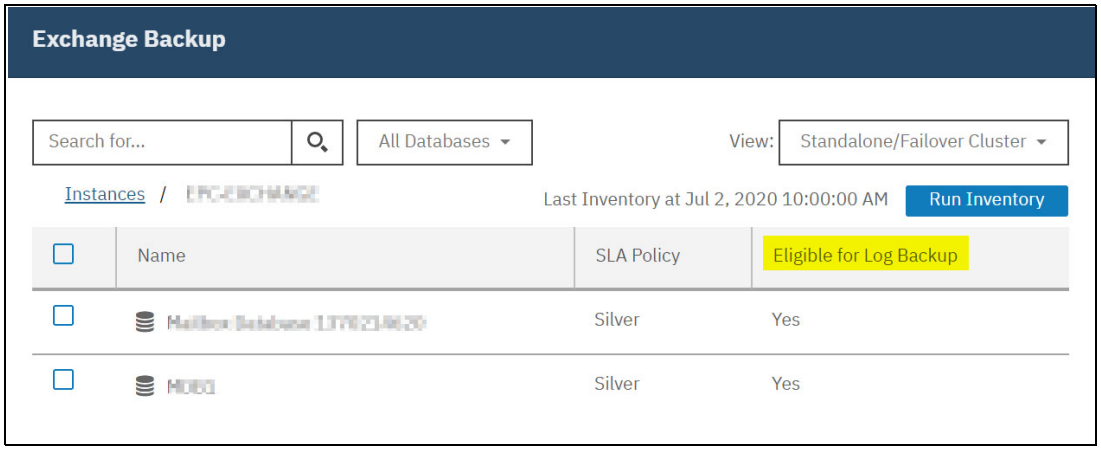

Figure 12-6 Exchange Backup menu

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.

12.3.1 Log backup

By clicking **Select Options**, the Log Backup menu opens, as shown in Figure 12-7. 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.

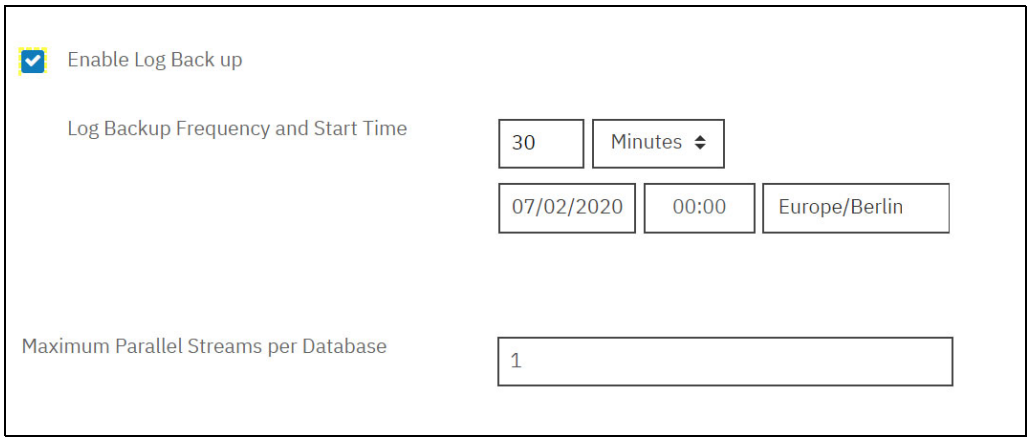

Figure 12-7 Log Back up menu

12.3.2 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.

12.4 Backup jobs overview

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

12.4.1 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.

For more information about SLA policies, see 1.3, “SLA backup policies” on page 19.

12.4.2 Backup types

The following backup jobs are available for Microsoft Exchange Applications:

- Scheduled
- Ad hoc

12.4.3 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 12-8.

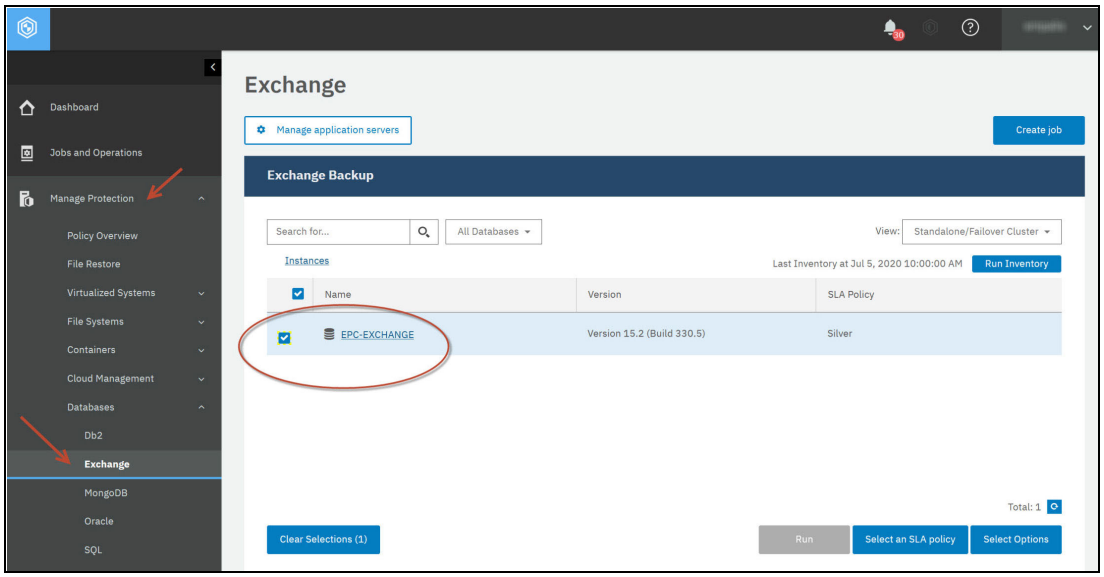

Figure 12-8 Defining a backup job

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 12-9.

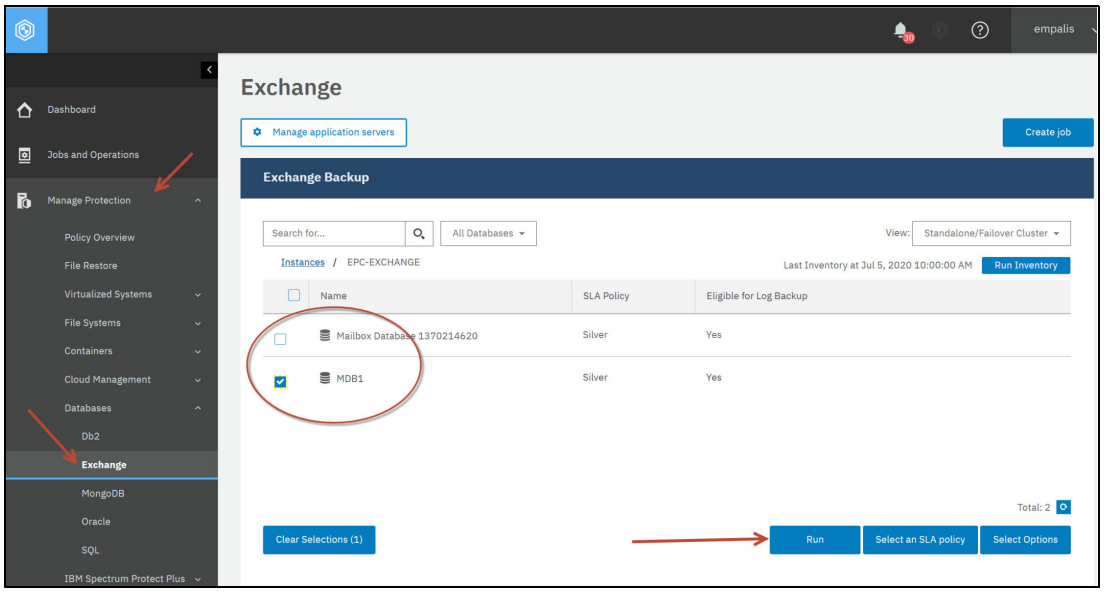

Figure 12-9 Selecting an instance or database

Click **Select an SLA Policy**. Predefined choices are: Gold, Silver, and Bronze. Each choice includes different frequencies and retention rates, as shown in Figure 12-10.

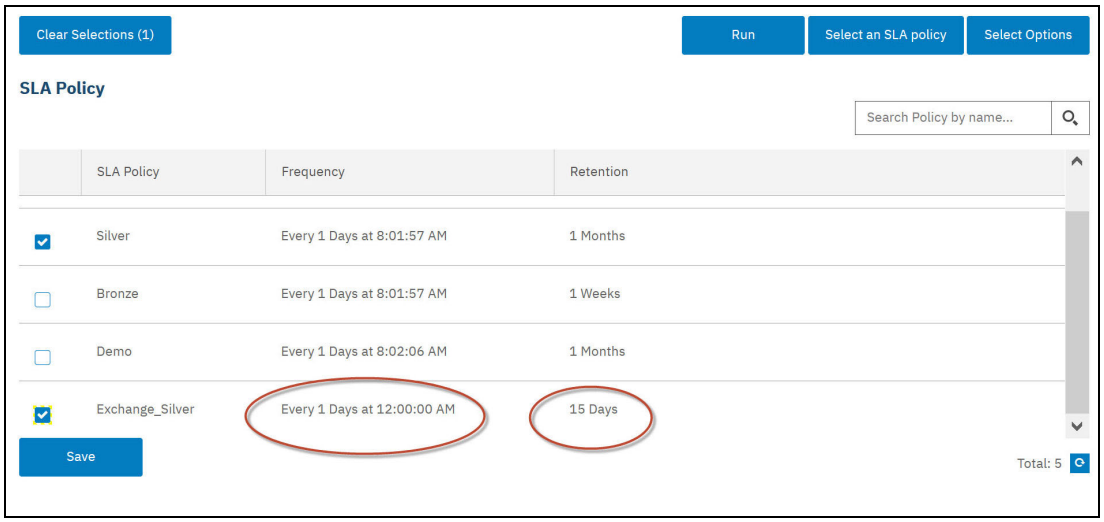

Figure 12-10 Custom SLA policy Exchange_Silver

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 12-10. 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 12-11.

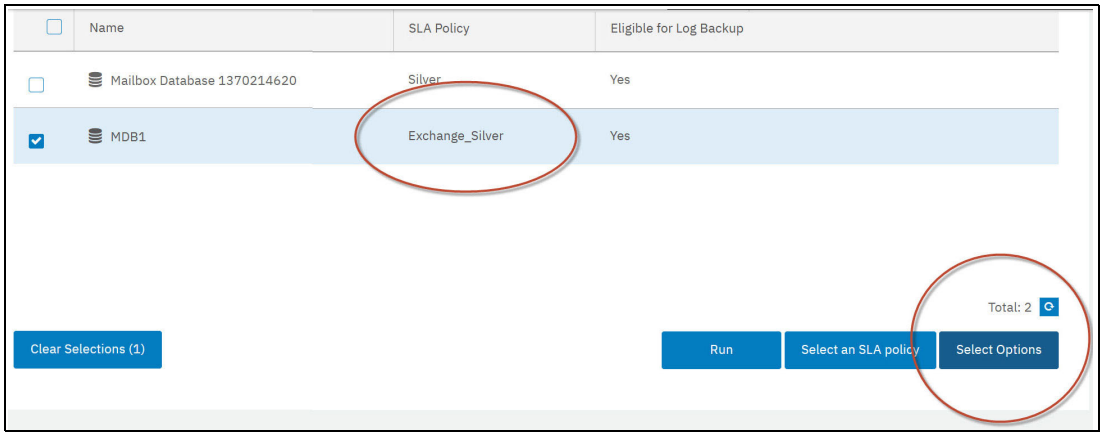

Figure 12-11 Checking SLA selection and selecting options

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 12-12). Click **Save**.

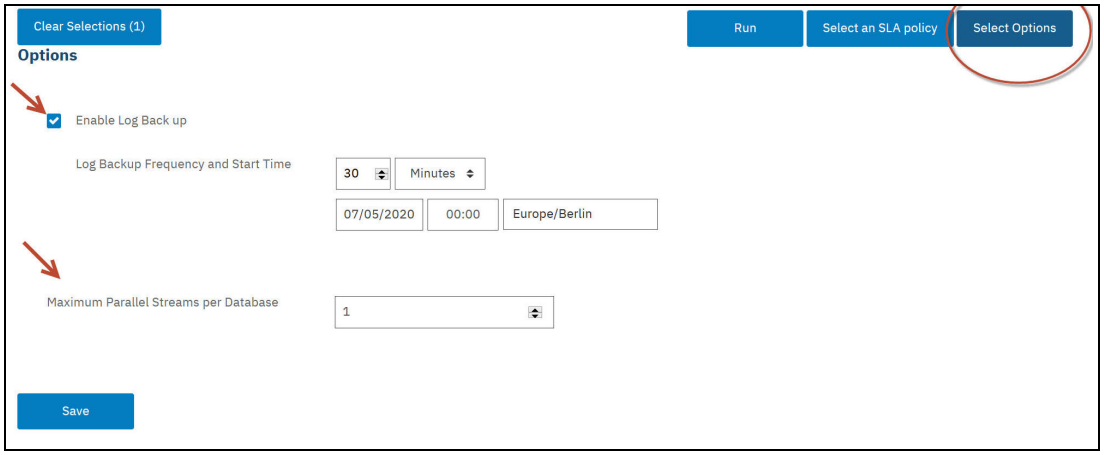

Figure 12-12 Selecting options

Configure the SLA policy by clicking the icon in the **Policy Options** column of the **SLA Policy Status** table, as shown in Figure 12-13.

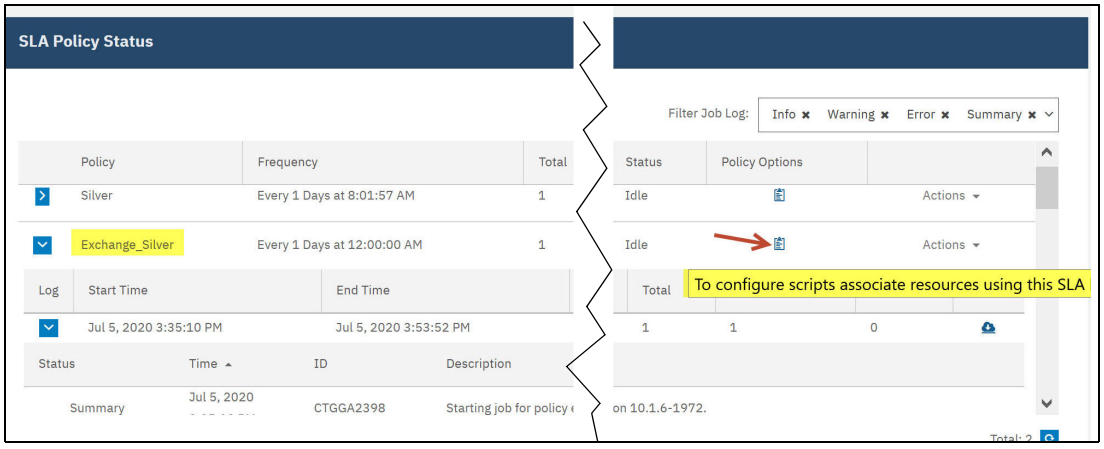

Figure 12-13 Configuring SLA policy

After clicking the icon, a pop-up window appears, as shown in Figure 12-14, in which you can configure more policy options.

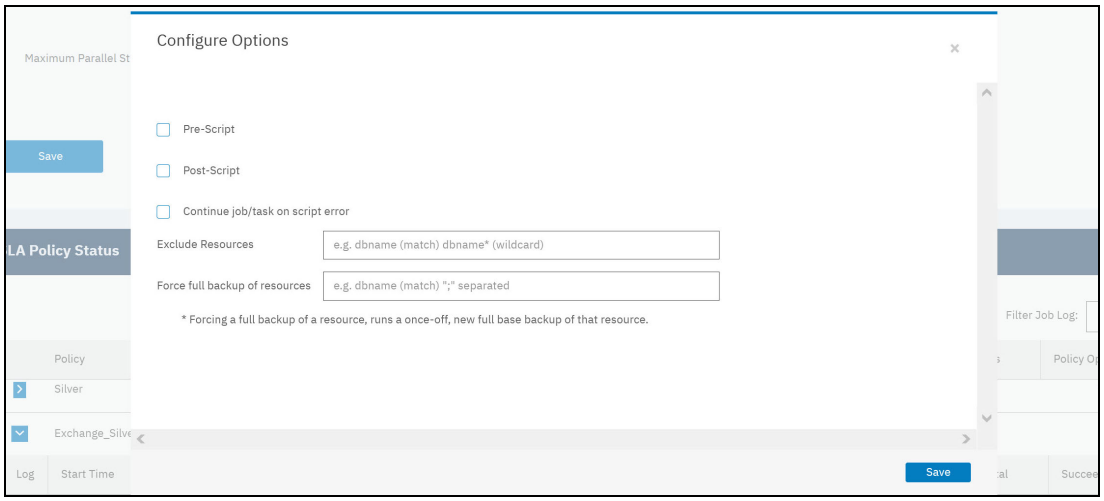

Figure 12-14 Configuring SLA options

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 12-15).

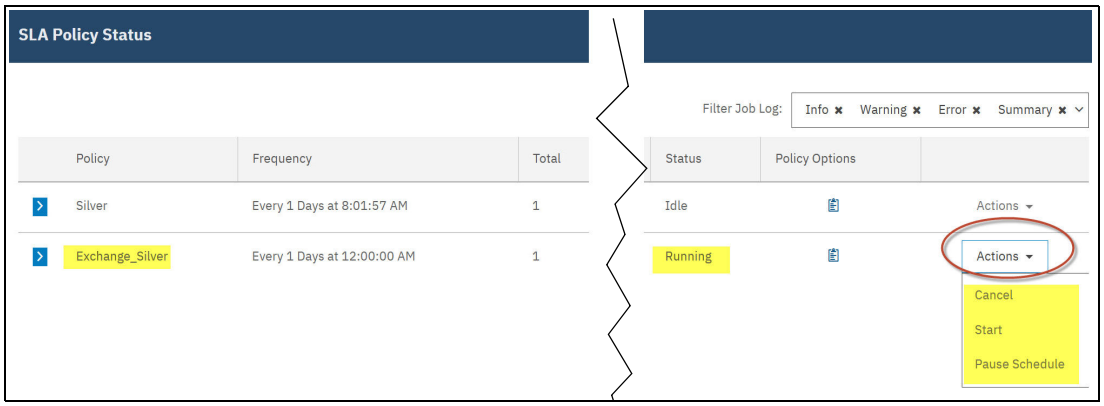

Figure 12-15 Running Jobs and actions

12.4.4 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 12-16.

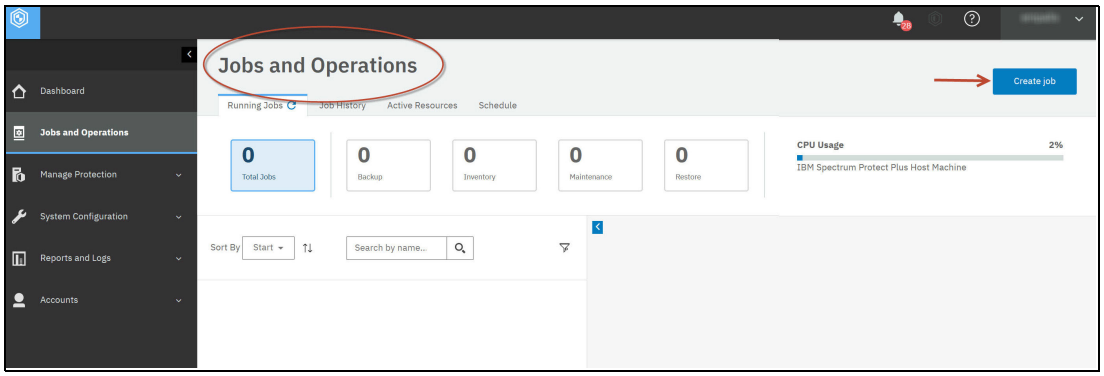

Figure 12-16 Creating a Job

2. You are presented with a choice for Ad hoc backup or Restore. Select **Ad hoc backup** (see Figure 12-17).

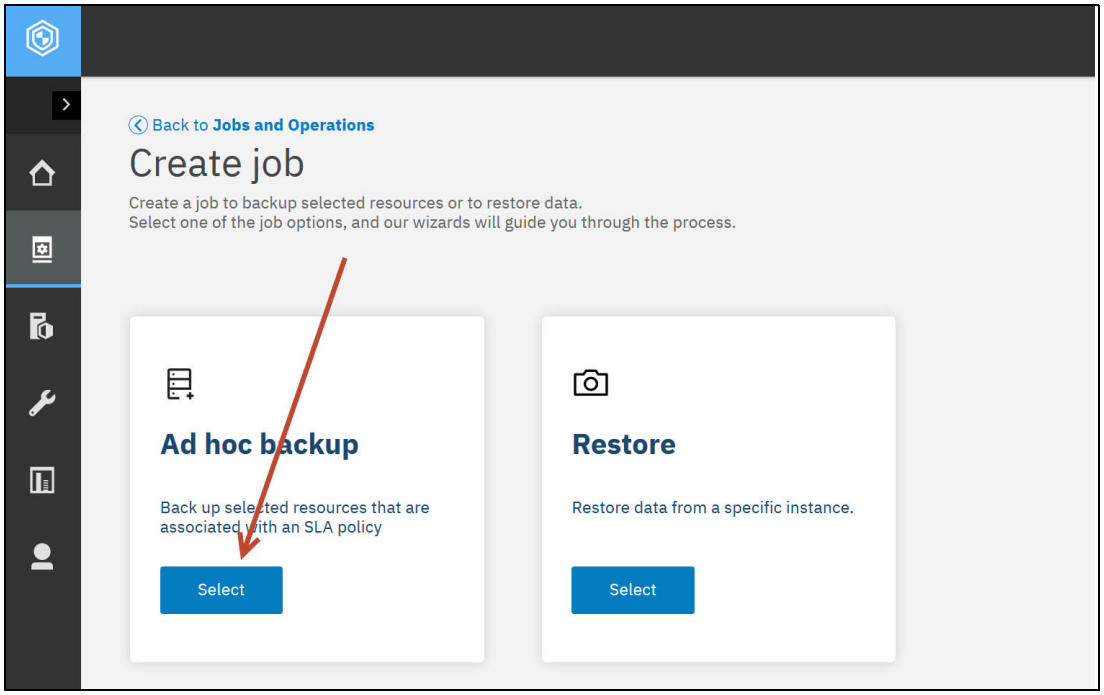

Figure 12-17 Ad hoc backup

3. In the Database selection, select **Exchange**, as shown in Figure 12-18.

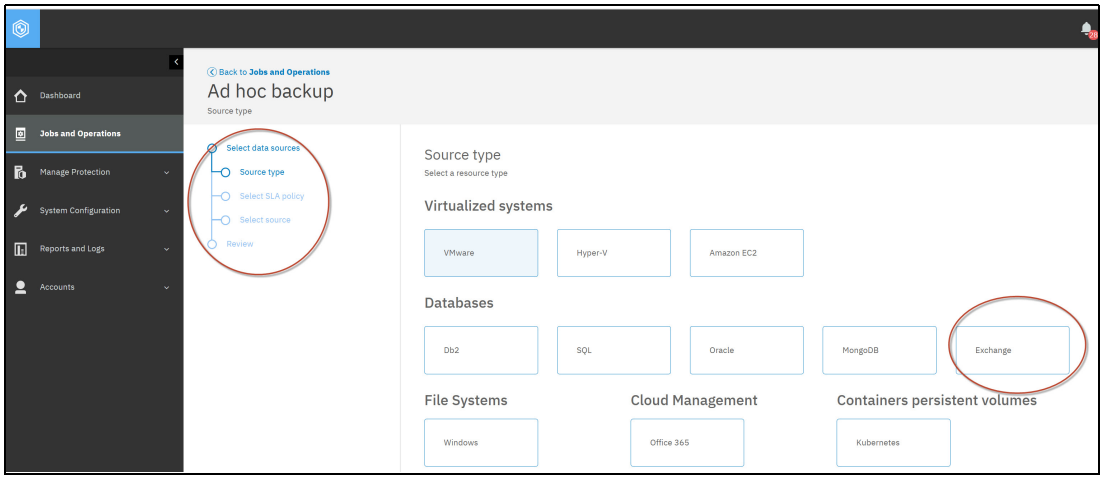

Figure 12-18 Definition diagram and selection

4. Select a predefined SLA policy. After clicking the SLA policy, the defined values for that policy are shown (see Figure 12-19).

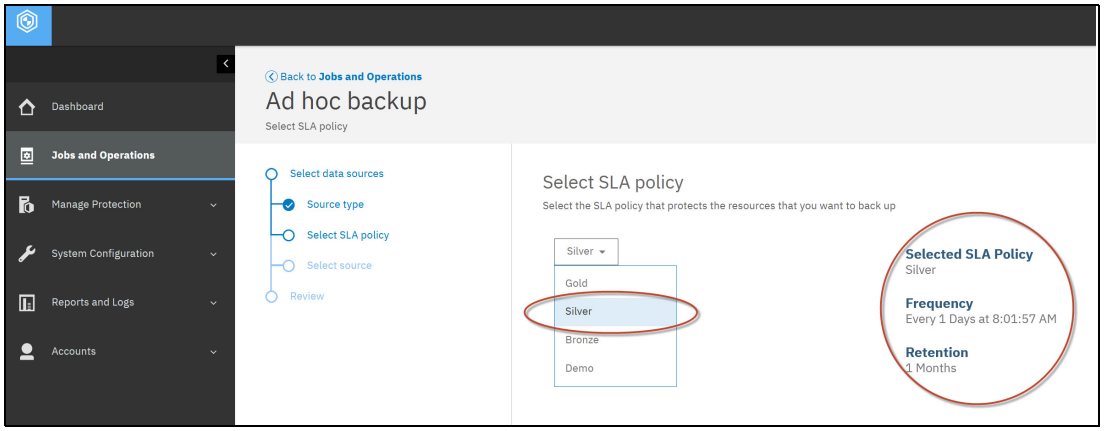

Figure 12-19 Selected SLA policy

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 12-20.

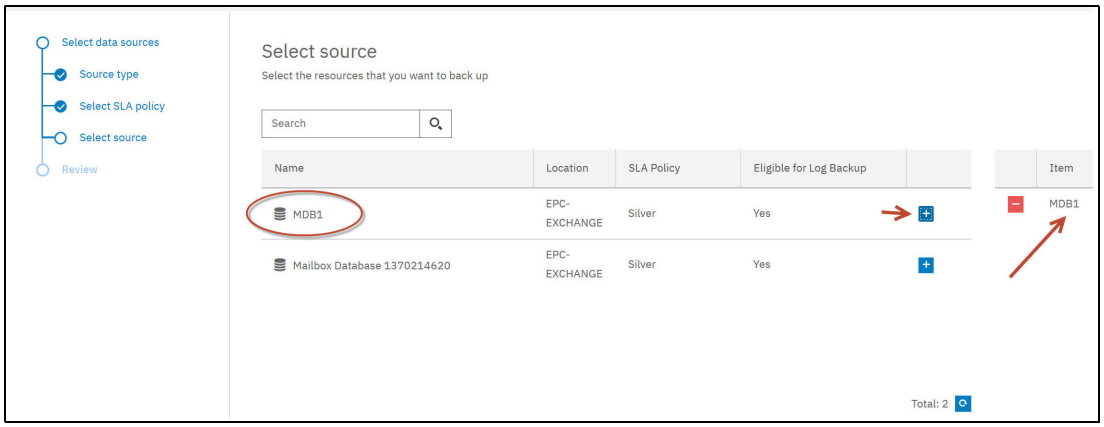

Figure 12-20 Choosing databases for ad hoc backup

6. Review the backup job options. Then, click **Submit** to start the job, as shown in Figure 12-21.

Figure 12-21 Ad hoc backup preview

7. As shown in Figure 12-22, a message is displayed to confirm that the job was submitted. Click **OK** to close the message.

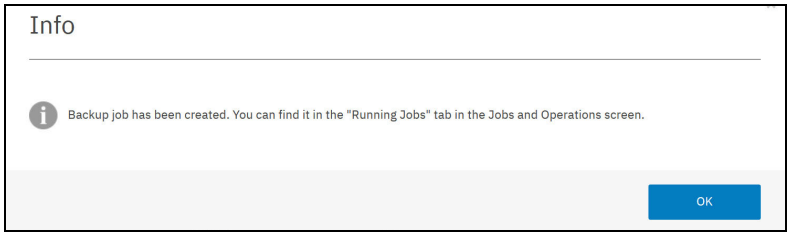

Figure 12-22 Job submission confirmation

The Ad hoc backup job can be monitored under the **Jobs and Operations** pane, as shown in Figure 12-23.

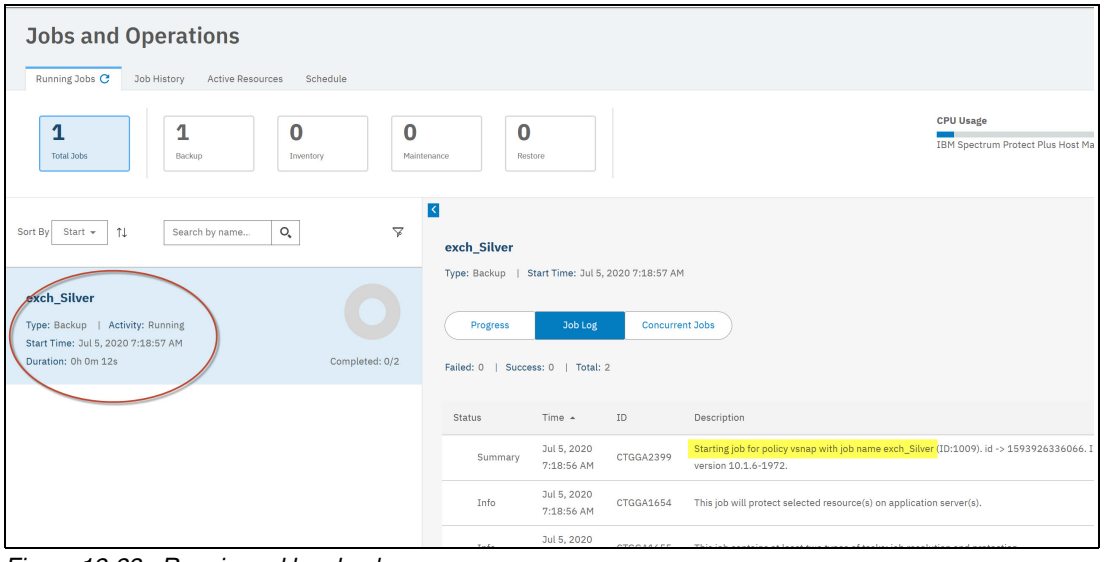

Figure 12-23 Running ad hoc backup

12.5 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 12-24.

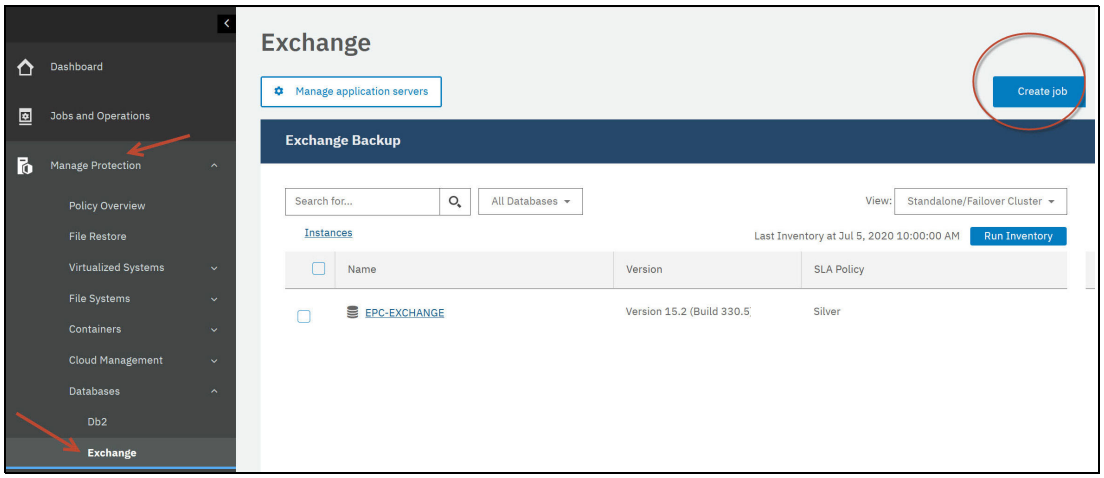

Figure 12-24 Creating Restore job

To create the restore job, select **Create Job** in the **Manage Protection** - **Exchange** menu. Then, select **Restore**, as shown in Figure 12-25.

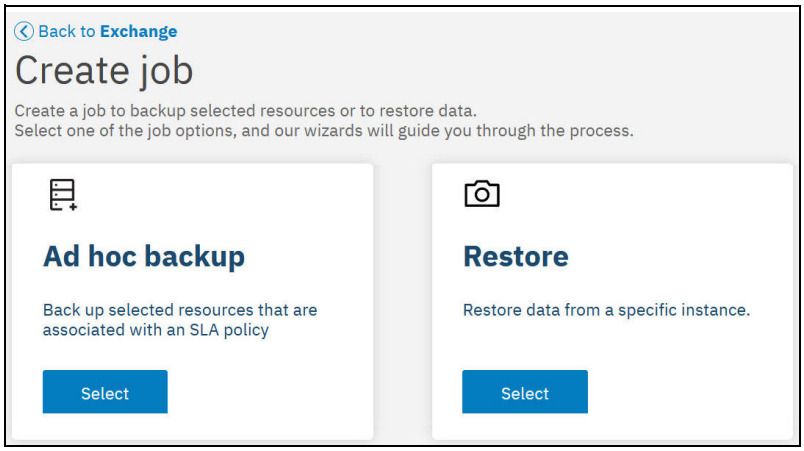

Figure 12-25 Selecting Restore

The restore dialog opens, as shown in Figure 12-26.

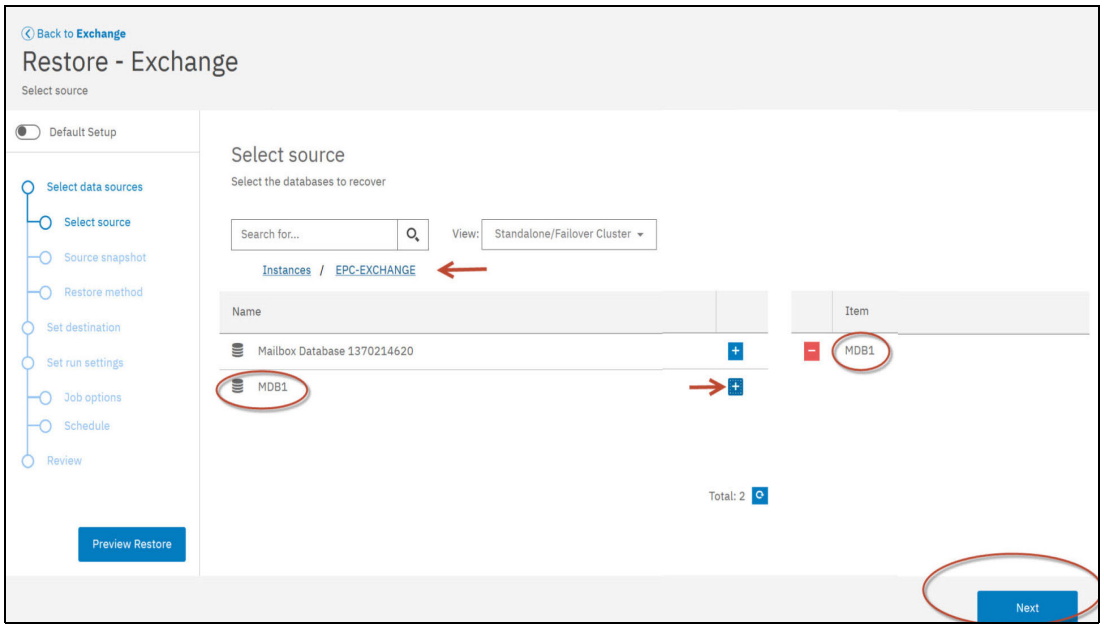

Figure 12-26 Selecting Restore items

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 12-27, the type of restore can be selected.

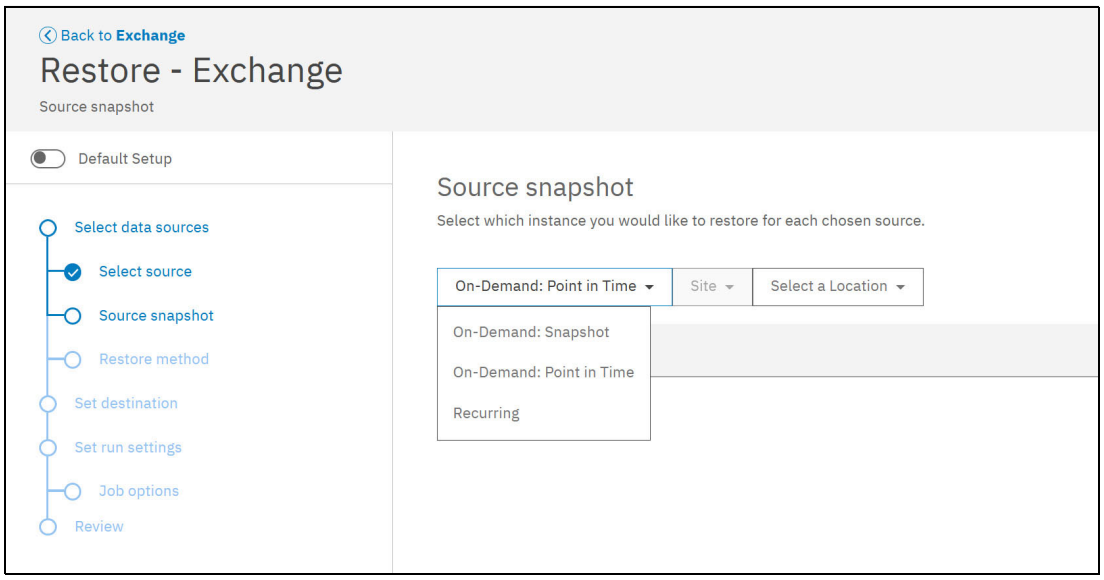

Figure 12-27 Selecting source snapshot

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 12-28.

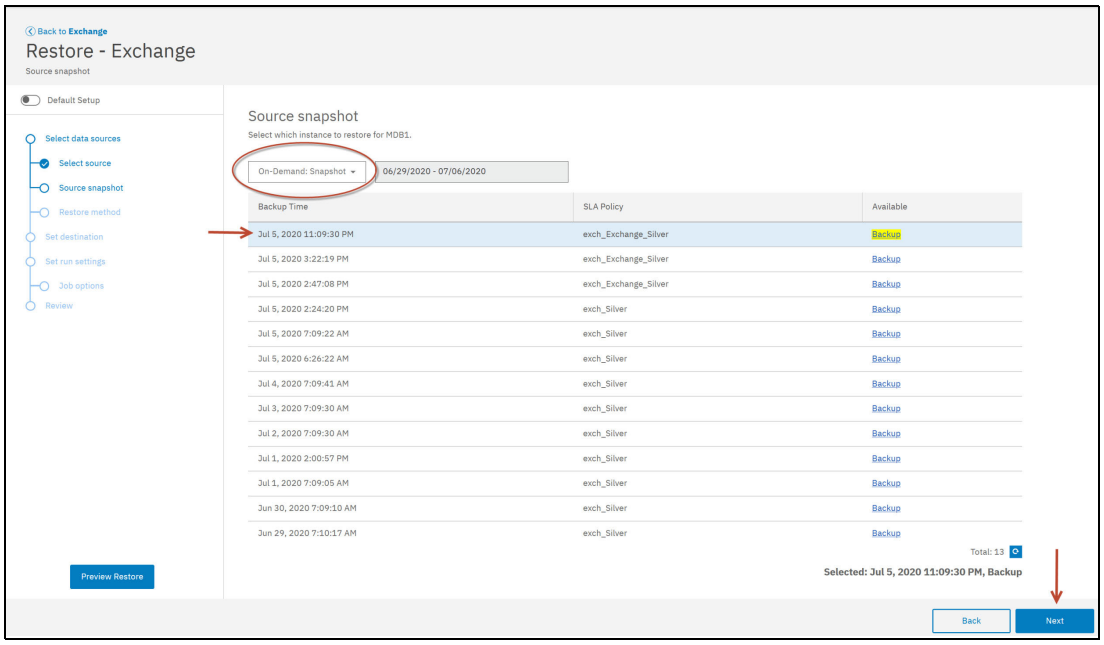

Figure 12-28 Setting date and time

– On-Demand Point-in-Time

  In this case, the available restore location sites are: Demo, Primary, and Secondary (see Figure 12-29).

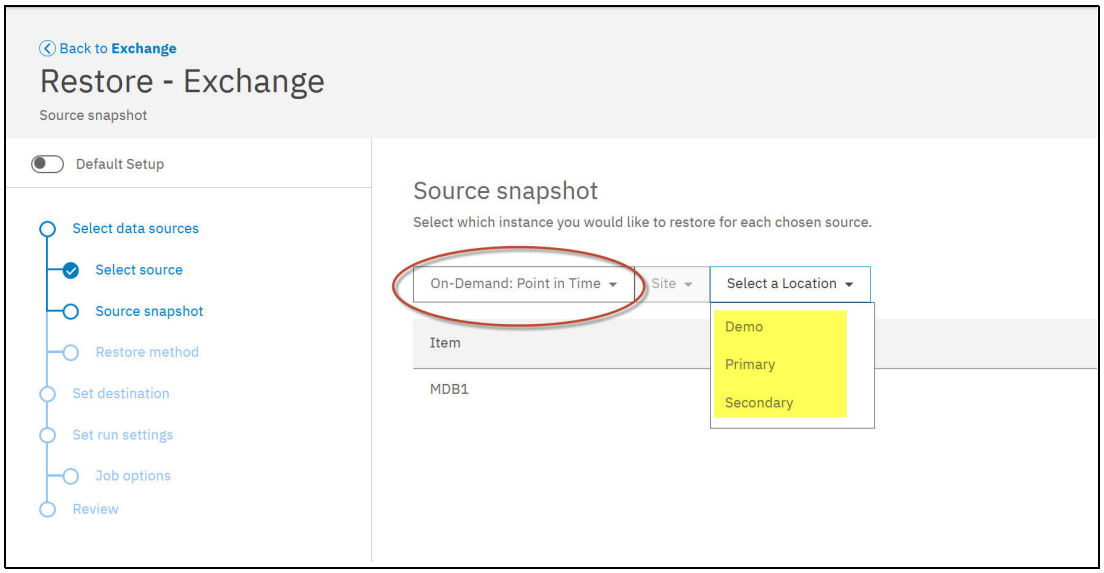

Figure 12-29 Selecting a location for On-Demand Point-in-Time

– Recurring

Choose a restore location type. As shown in Figure 12-30, 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.

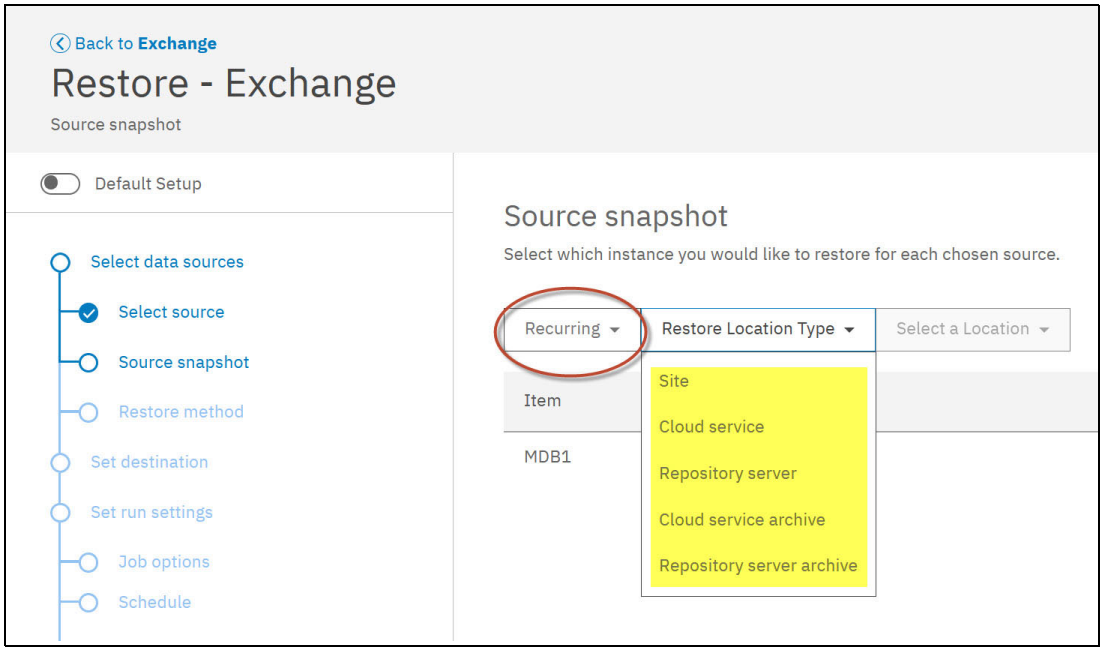

Figure 12-30 Recurring Restore location types

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.

12.5.1 Complete Restore

This section describes how to restore the complete mailbox database and use it in instant access, production and test (see Figure 12-31).

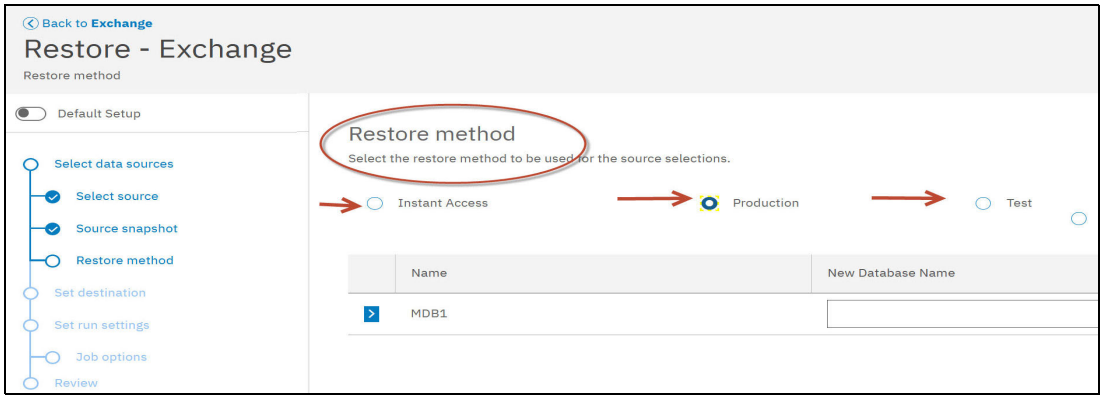

Figure 12-31 Restore method

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 12-32.

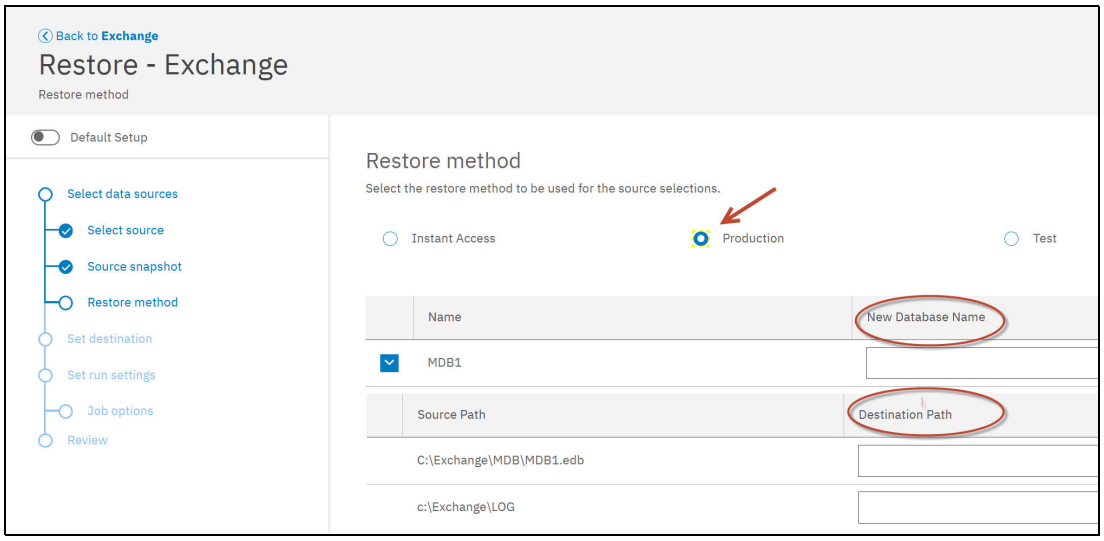

Figure 12-32 Restore method: Production

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 12-33.

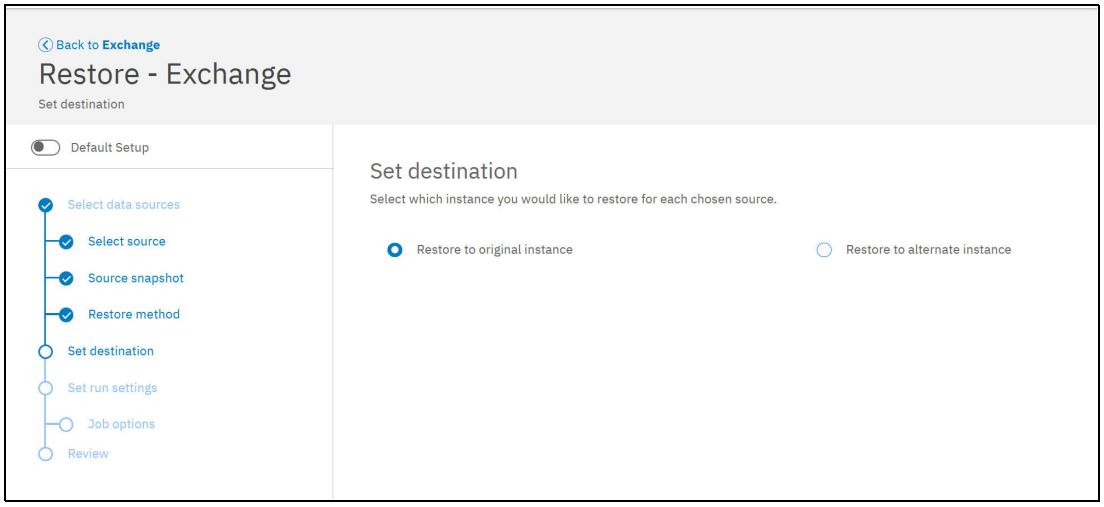

Figure 12-33 Setting destination

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 12-34.

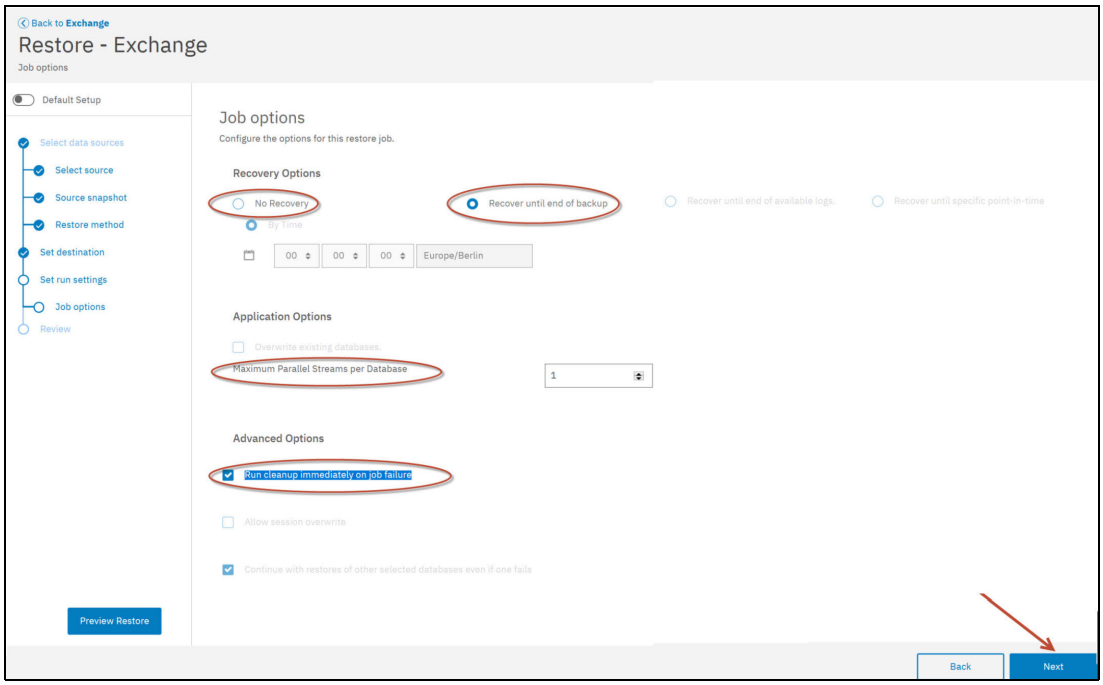

Figure 12-34 Job options

Click **Next** to proceed. The last panel displays a summary for review, as shown in Figure 12-35. Click **Submit** to start the restore job.

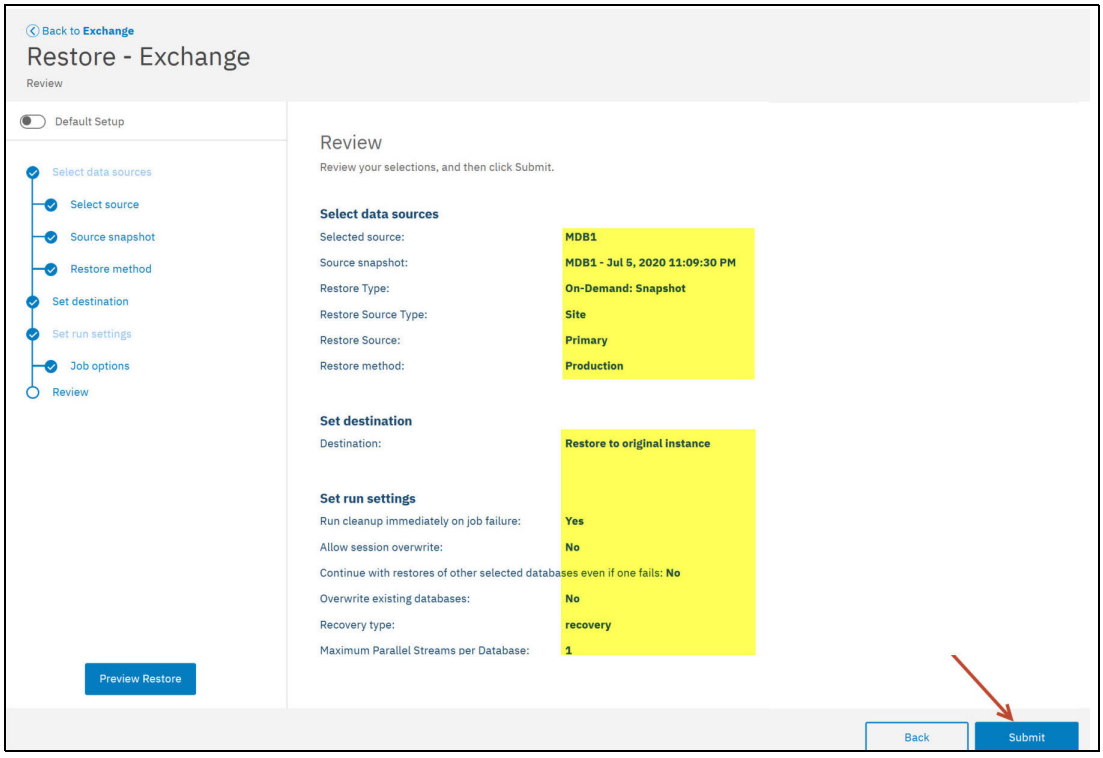

Figure 12-35 Review and submit

Under **Jobs and Operations**, you can follow the progress of the operation, as shown in Figure 12-36.

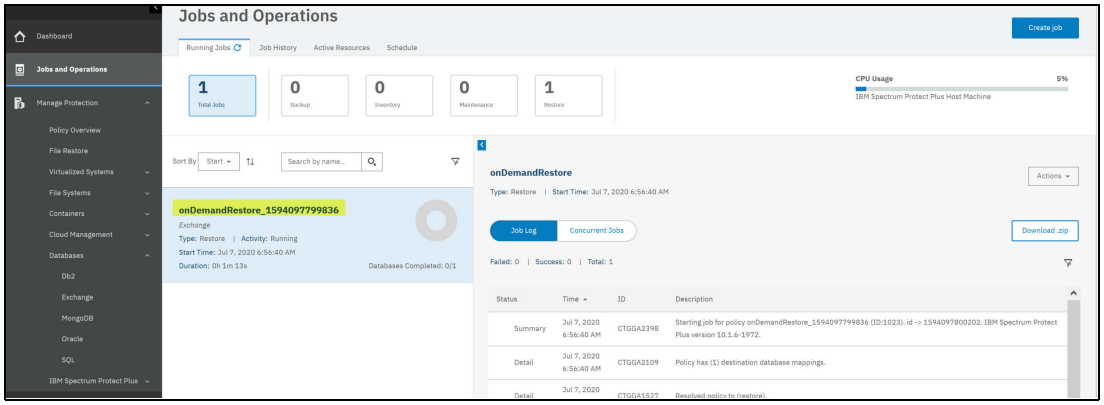

Figure 12-36 Monitoring running job

After completion, the job information is moved into the Job History window, as shown in Figure 12-37.

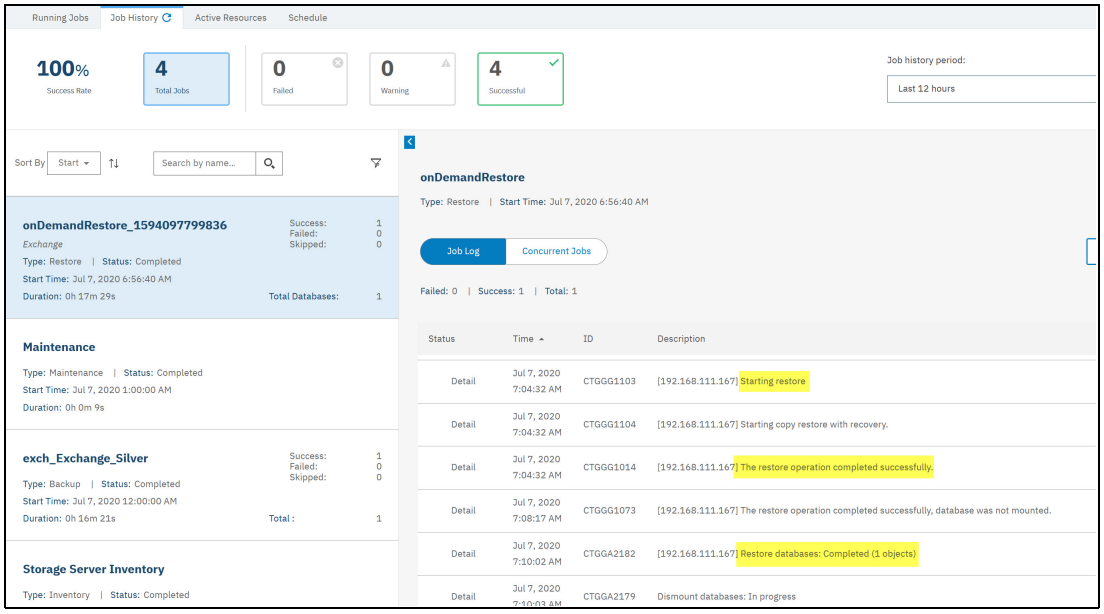

Figure 12-37 Restore Job history

12.5.2 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 12.2.1, “Granular restore remote package installation” on page 332.

In both cases, the restore procedure is started as described in 12.5, “Restore jobs” on page 349.

Select **Granular Restore** as shown in Figure 12-39.

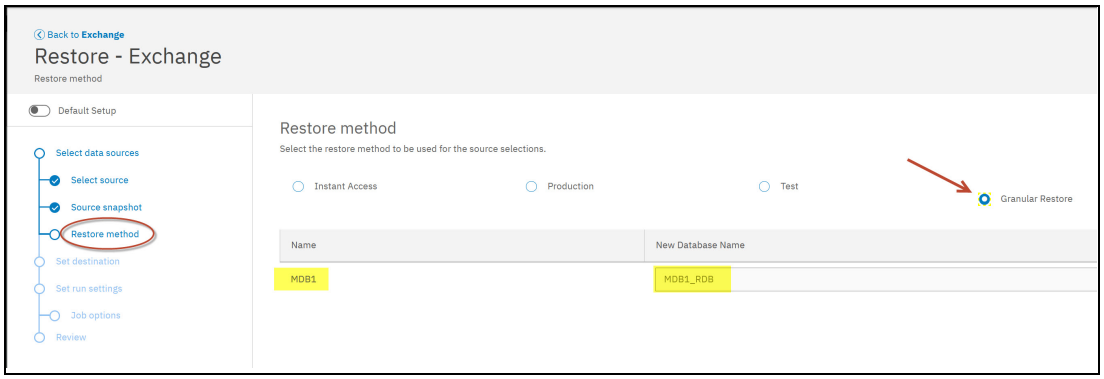

Figure 12-38 Granular restore method

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 12.5.1, “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 login Example 12-3 shows the steps that are performed.

Example 12-3 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 12-39.

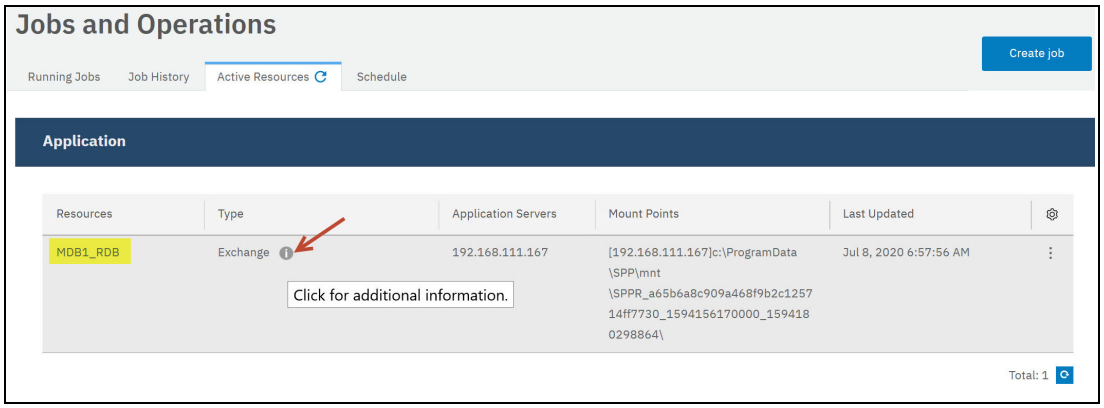

Figure 12-39 Active Resources tab

Clicking the information icon (as indicated by the arrow in Figure 12-39) 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 12-40).

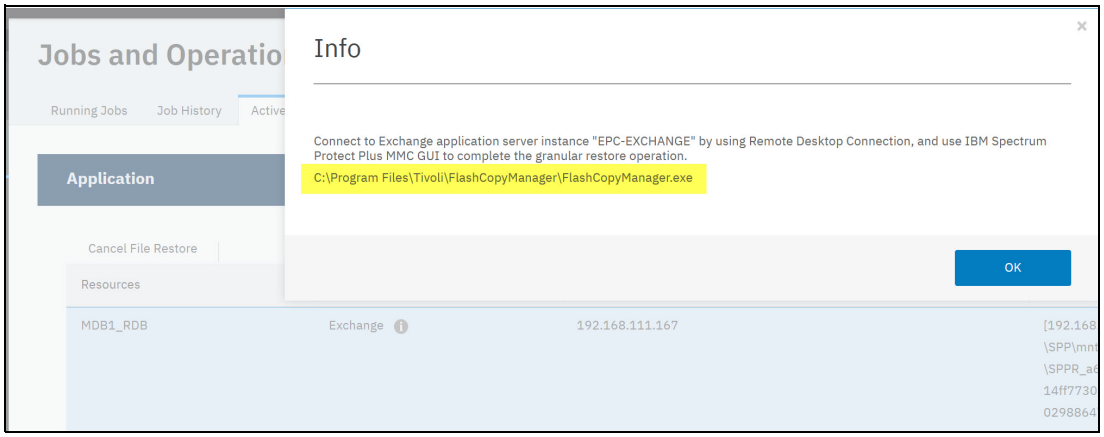

Figure 12-40 Information how to start the IBM Spectrum Protect MMC GUI

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 12-41.

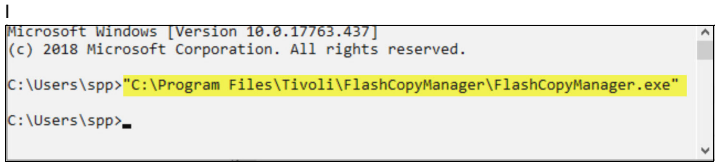

Figure 12-41 Opening IBM Spectrum Protect Plus MMC GUI

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 12-42.

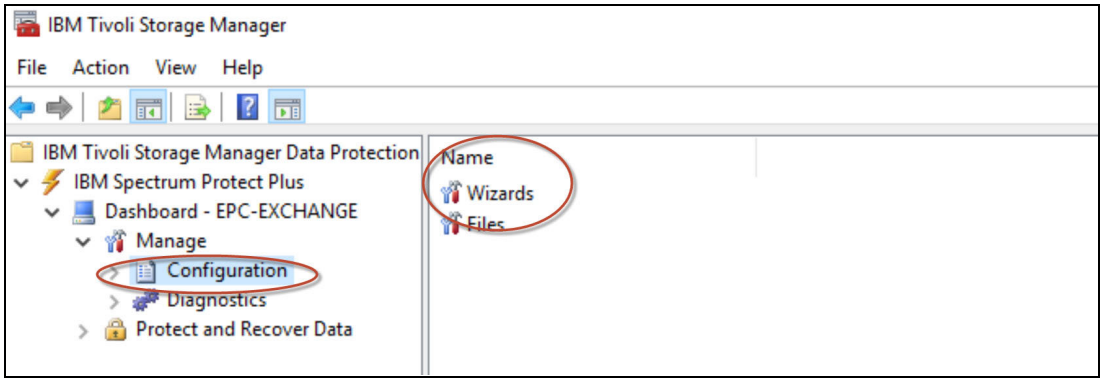

Figure 12-42 Starting the Configuration wizard

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 12-43.

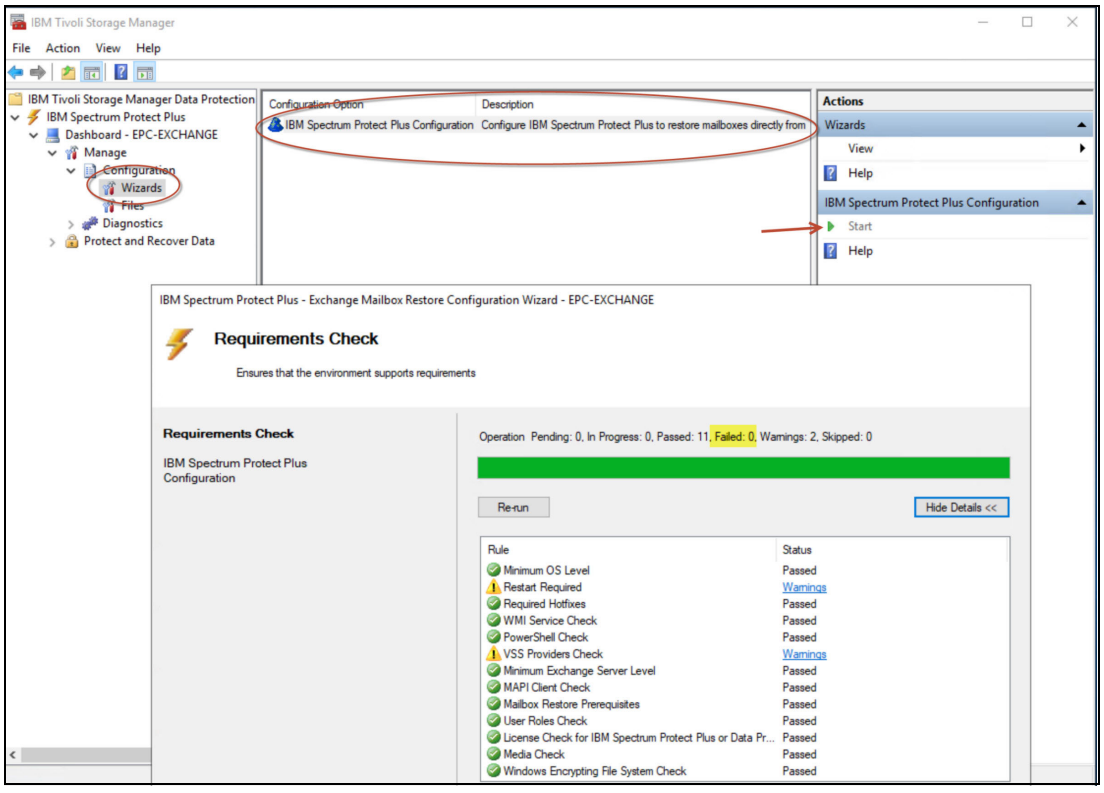

Figure 12-43 IBM Spectrum Protect Configuration wizard

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 12-44. This error is shown because it is not recommended to perform the recovery with the exchange server.

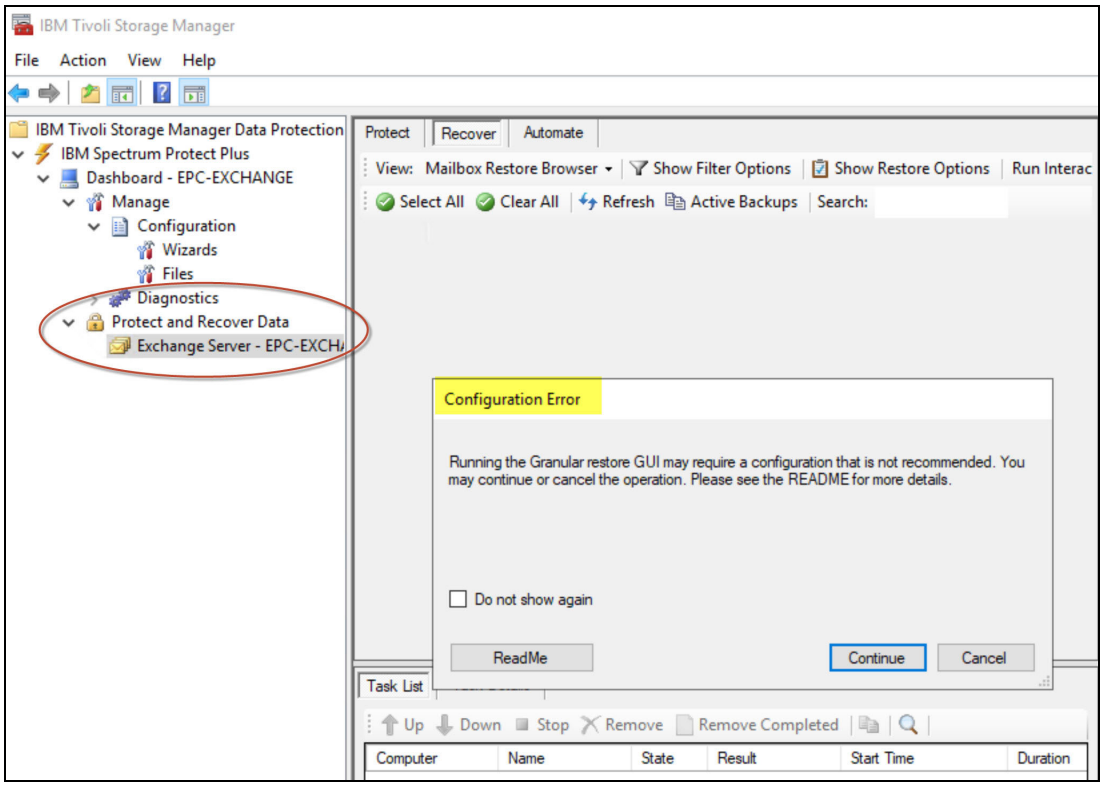

Figure 12-44 Protect and Recover Data

8. Click **ReadMe** to see information that is similar to the information that was described in 12.2.1, “Granular restore remote package installation” on page 332 (see Figure 12-45).

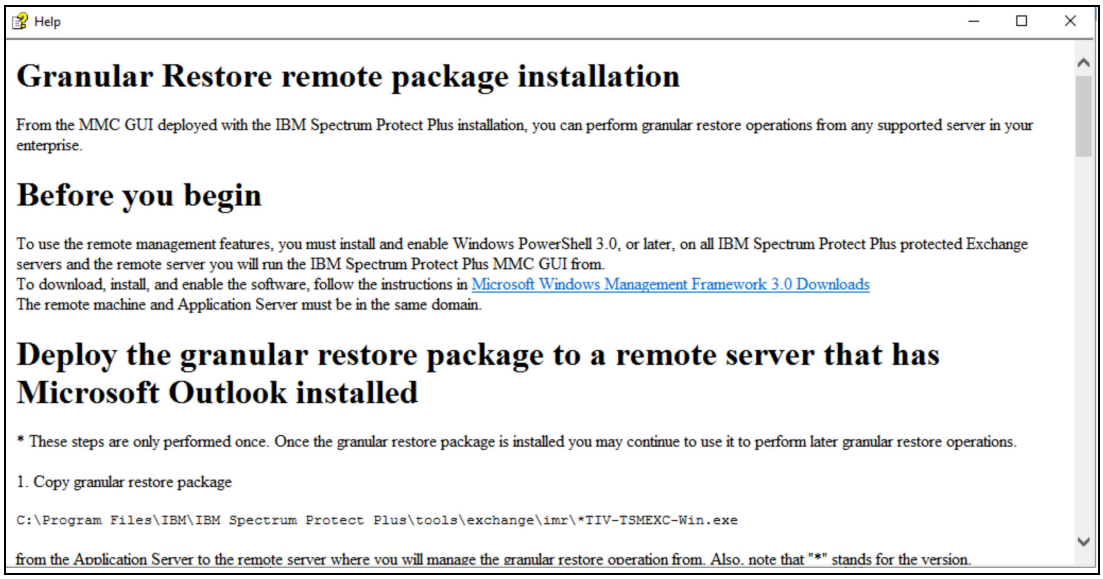

Figure 12-45 Installing ReadMe Granular Restore remote package

9. 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 12-46.

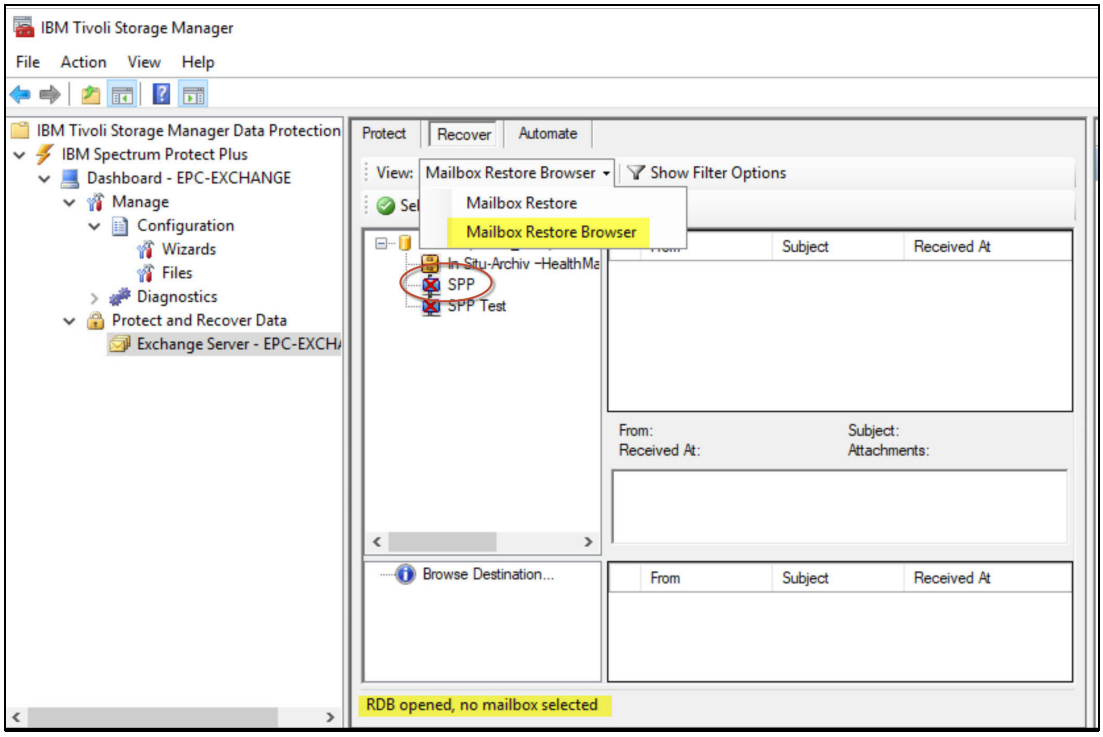

Figure 12-46 Mailbox Restore Browser

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 12-47). 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.

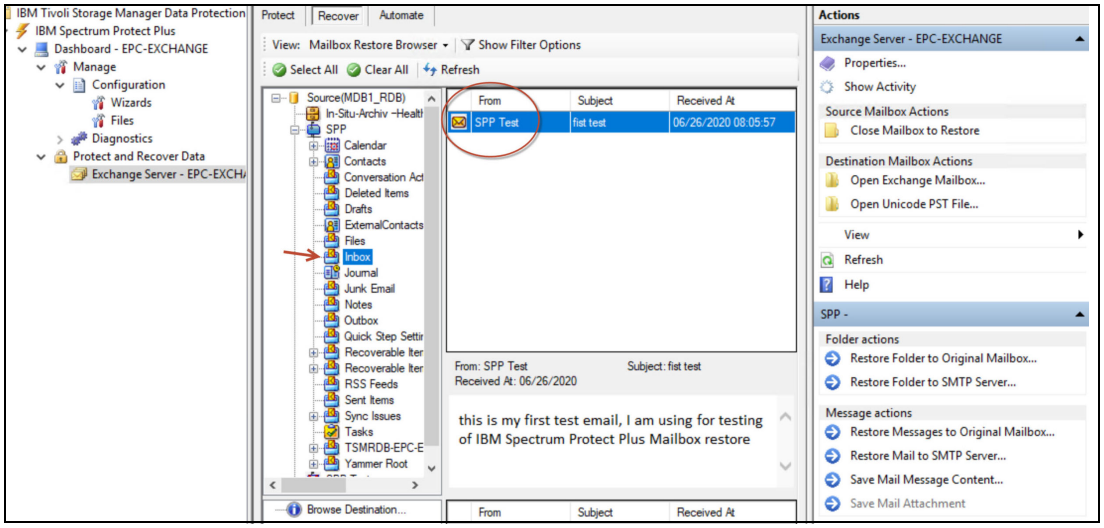

Figure 12-47 Item select

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 12-48.

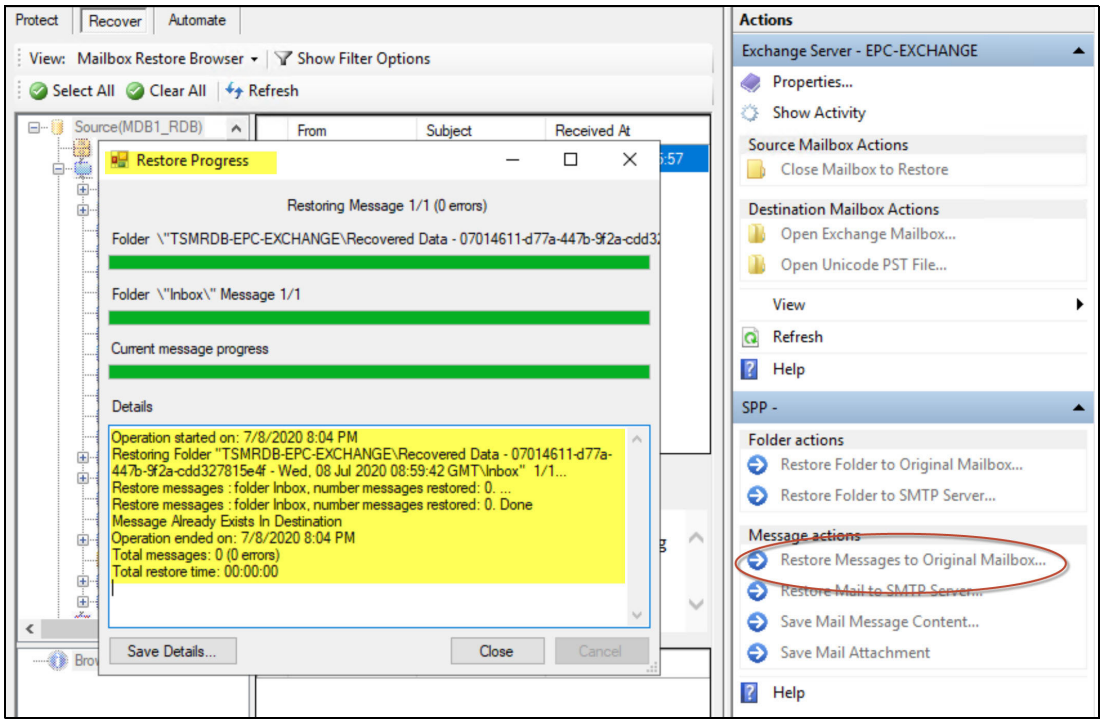

Figure 12-48 Restore progress: Restoring message

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 cancelling the running job, as shown in Figure 12-49.

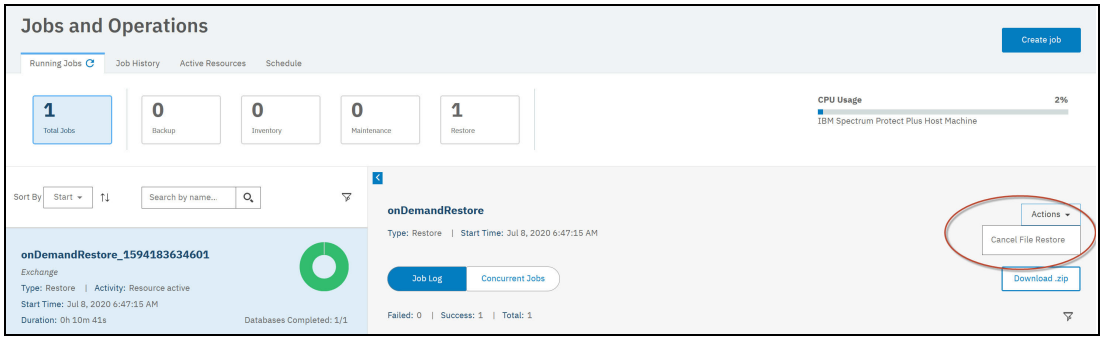

Figure 12-49 Cleanup in Jobs and Operations: Running Jobs

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 12-50.

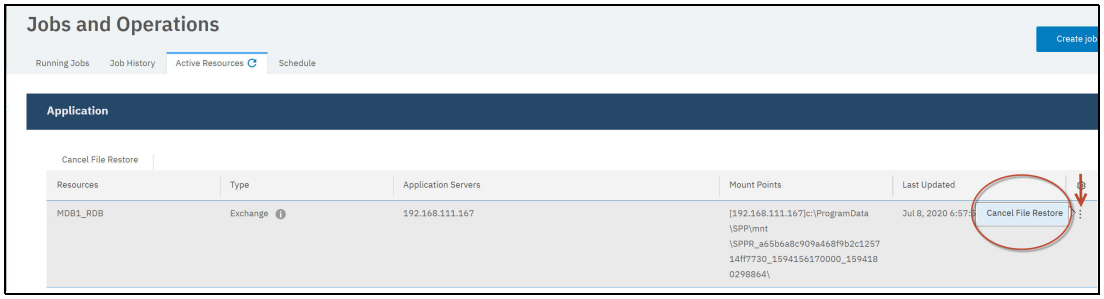

Figure 12-50 Clean up in Jobs and Operations: Active Resources

The Job History Job Logs includes the detailed log of the cleanup procedure and is confirmed with a success message, as shown in Figure 12-51.

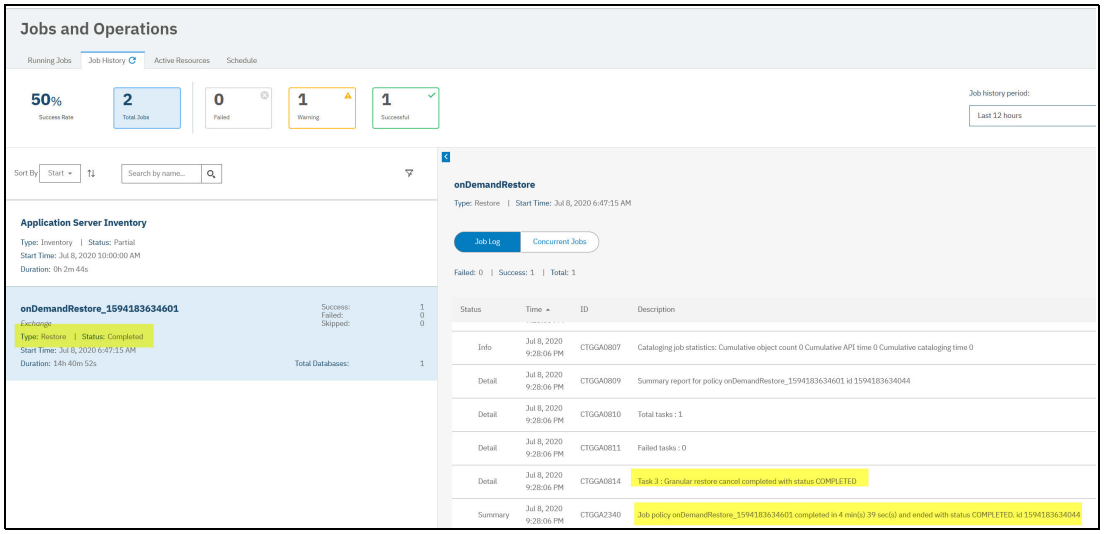

Figure 12-51 Job History cleanup

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 12-52.

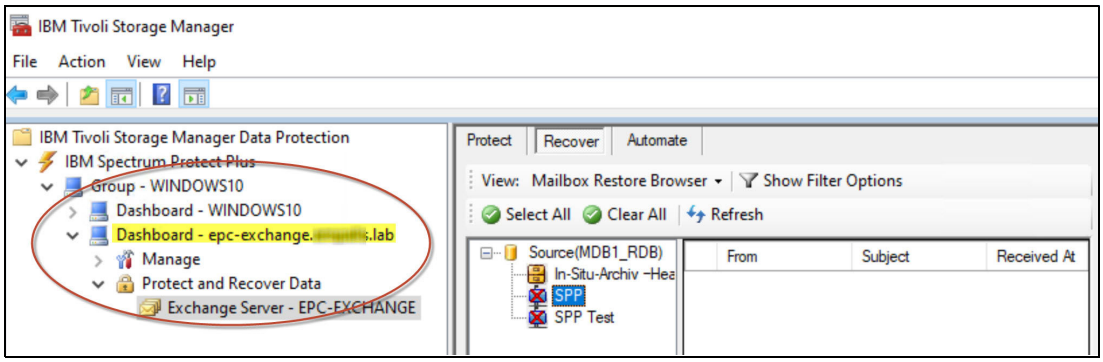

Figure 12-52 Group Dashboard view

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 357.

13

Backing up and restoring Microsoft 365 data

Although Microsoft 365 is hosted on Microsoft Azure Cloud infrastructure, it is the responsibility of the customer to ensure data security and integrity and protect Microsoft 365 accounts from the following possible threats:

- Erroneous or deliberate data corruption
- External or internal attacks
- Encryption or deletion due to malicious software
- Deletion of retired users data

This chapter explains how to effectively protect Microsoft 365 data by using IBM Spectrum Protect Plus.

Product name update: Microsoft Corporation announced new product names, effective 21 April 2020, for its Office 365 offerings for small and medium businesses. With this announcement, all small and medium business plans transitioned to the new Microsoft 365 brand.

In IBM Spectrum Protect Plus V10.1.6, the user interface and documentation use the original product name, Office 365. For more information, see this Microsoft 365 blog entry.

This chapter includes the following topics:

13.1 Solution overview

Microsoft 365 is a cloud based-solution. All of your organization and user account data is stored in the Microsoft Azure infrastructure. To effectively protect the data, all account data must be retrieved from Microsoft 365 and stored in another location. This other location can be on-premises in a data center or running on a supported public cloud. Figure 13-1 shows how data is transferred for backup and restore by using IBM Spectrum Protect Plus.

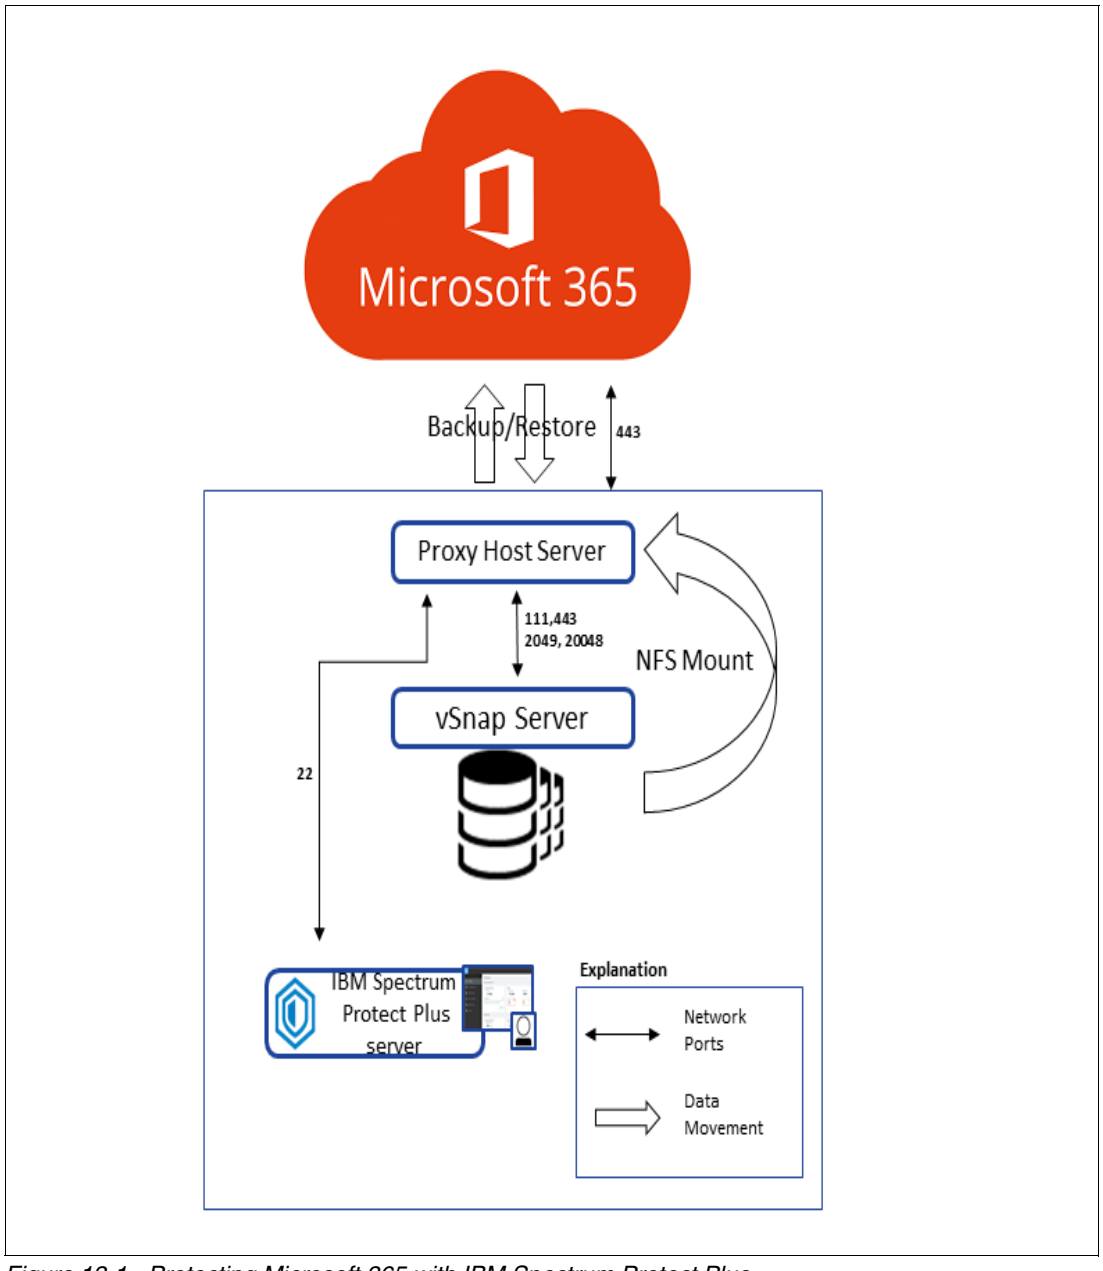

Figure 13-1 Protecting Microsoft 365 with IBM Spectrum Protect Plus

For more information about the IBM Spectrum Protect Plus components that are shown in Figure 13-1, see Chapter 1, “IBM Spectrum Protect Plus product architecture and components” on page 1.

IBM Spectrum Protect Plus use the proxy host server to connect to the Microsoft 365 cloud. The protected data is stored on the vSnap server. The vSnap server mounts its snapshots to the Proxy Host Server as NFS shares.

13.2 Prerequisites

To protect Microsoft 365 application with IBM Spectrum Protect Plus, you must register the application with Azure Active Directory and grant appropriate access permissions. An active Microsoft 365 subscription and a Microsoft 365 administrative userID are also required.

Figure 13-2 shows the supported subscriptions in the Global Microsoft 365 Regions.

IBM Spectrum Protect Plus	**Microsoft 365 Business** Basic, Business Standard, Business Premium editions	**Office 365 for Enterprise** E1, E3, and E5 editions	**Office365 for Education** A1, A3, and A5 editions	**Office 365 for Firstline Workers** F3 edition	**Microsoft 365 for Enterprise** E3 and E5 editions
	Former product name: Office 365 Business: Business, Essentials, and Business Premium edition		**Former product name**: Office 365 Education edition	**Former product name:** Microsoft 365 F1	
V10.1.5	✓	✓	✓	✓	✓
V10.1.6	✓	✓	✓	✓	

Figure 13-2 Coverage matrix for application levels supported by IBM Spectrum Protect Plus

The registration process is described in 13.2.2, **Microsoft 365 application registration and API permissions** on page 368.

The management of cloud data also requires a proxy host server that must be installed separately. For more information, see 13.2.1, Proxy host server on page 367.

For more information about prerequisites, see IBM Knowledge Center.

13.2.1 Proxy host server

The proxy host server is a new component for IBM Spectrum Protect Plus. It is designed to work as a proxy to perform data protection operations for Microsoft 365 workloads. In the IBM Spectrum Protect Plus GUI, it is referred to as *Application Server*.

As shown in Figure 13-1 on page 366, the proxy host server must run on a dedicated virtual or physical server and a Linux based operating system.

Supported operating systems are Red Hat or CentOS 7.X based Linux systems. Installation can be performed on a physical or virtual machine.

The same packages must be installed on the proxy host server, including the following examples:

- Java 8
- Libicu
- NFS client

Depending on the Linux distribution and edition, these packages are preinstalled or can easily be added by using the corresponding package manager. For more information, see the operating system's manual.

The SSH connection port 22 must be opened and a user with the required *sudo* permissions must be created. The login with SSH does not prompt for a password and is verified during IBM Spectrum Protect Plus Microsoft 365 configuration.

The following minimum hardware requirements must be met:

- 4 GB RAM
- 4 CPUs or vCPUs
- 5 GB temporary disk space during backup workloads in addition to the operating system requirements

For more information about all of the requirements and open network ports, see IBM Knowledge Center.

13.2.2 Microsoft 365 application registration and API permissions

For IBM Spectrum Protect Plus to access Microsoft 365 accounts, it must be registered as an application and with suitable access rights. Complete the following steps:

1. Log in to your Azure Active Directory as an administrative user. If App registrations are not bookmarked, select **All services** → **App registrations**, as shown in Figure 13-3.

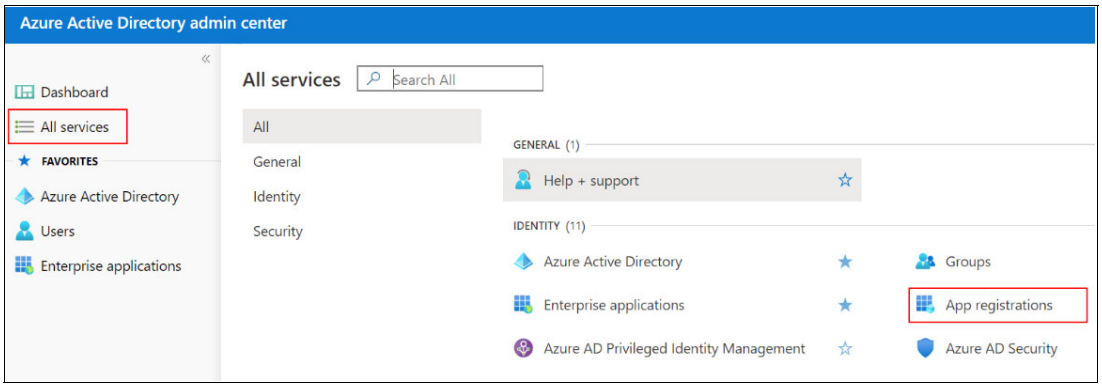

Figure 13-3 Azure AD App registration

2. Select **New registration**, as shown in Figure 13-4.

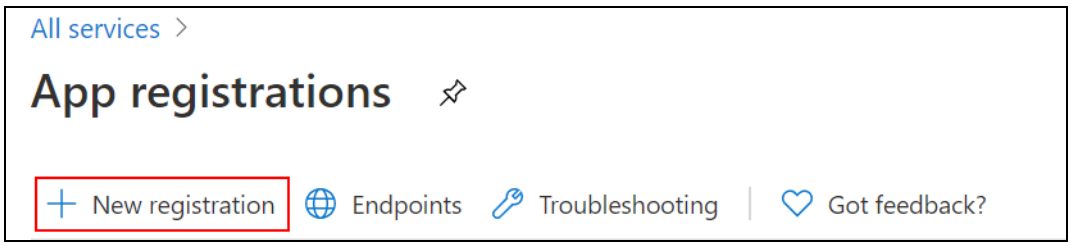

Figure 13-4 New registration

3. Enter a descriptive name for the new App registration and ensure that the supported account type is set to the wanted organization directory, as shown in Figure 13-5 on page 369.

Note: Real names, IDs, and IP addresses are blanked out in the figures for confidentiality reasons.

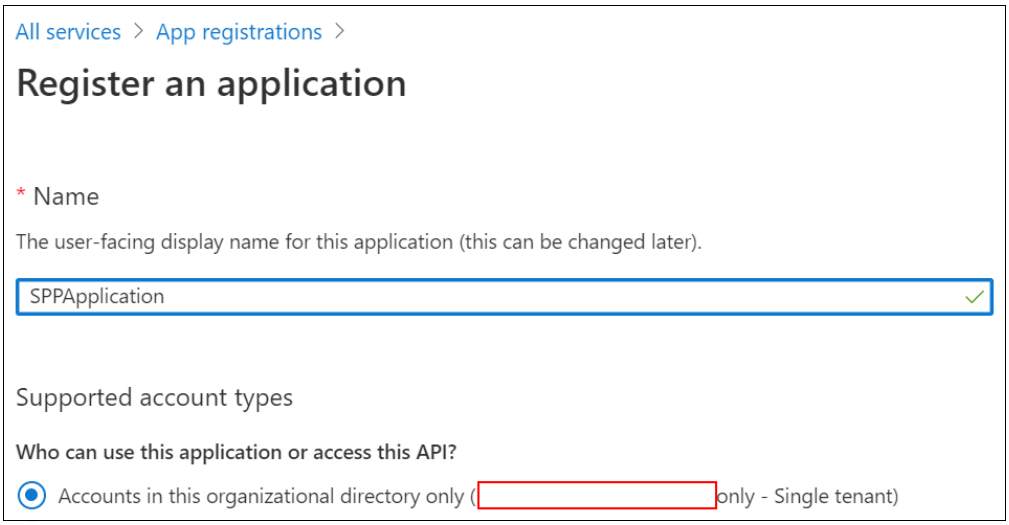

Figure 13-5 Register an application

After completing this page, the Application (client) ID, Directory (tenant) ID, and Object ID are shown. These IDs are used by IBM Spectrum Protect Plus to connect and authenticate to Microsoft 365 services (see Figure 13-6).

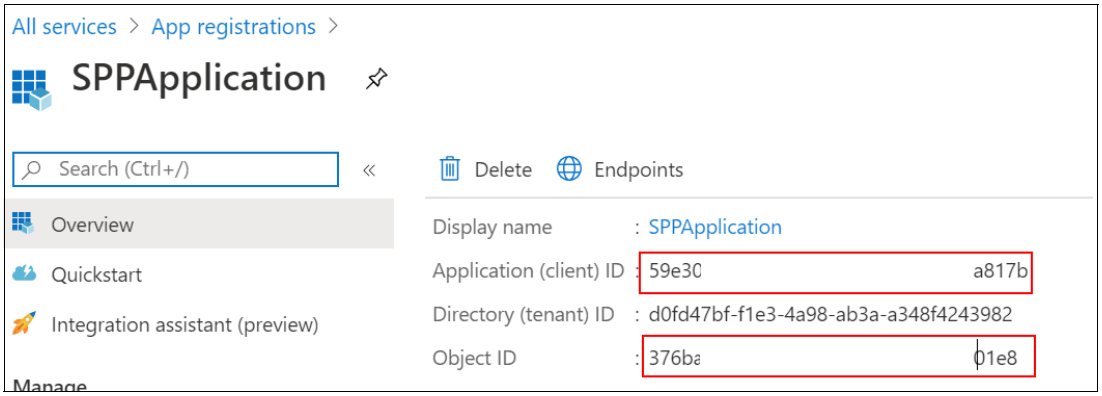

Figure 13-6 Application IDs

4. After creating the App registration, a Client secret must be generated. Select **Manage** → **Certificates & Secrets?** → **New Client Secret**, as shown in Figure 13-7.

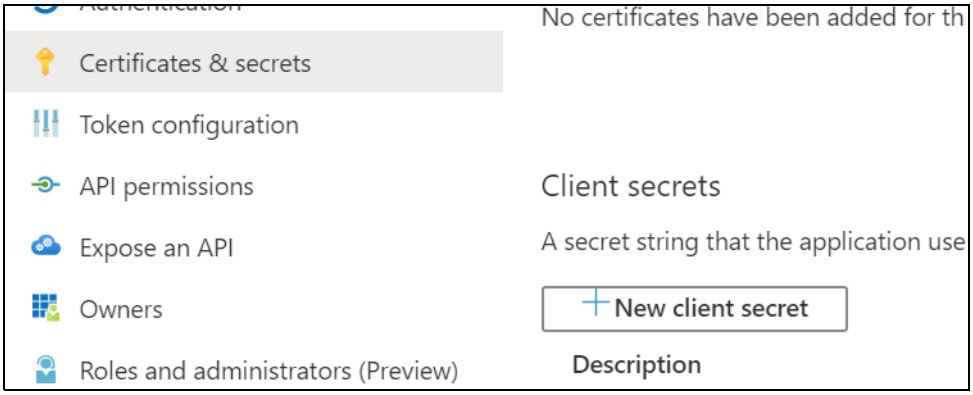

Figure 13-7 Certificates and secrets

5. Enter a descriptive name and an expiration period. Then, click **Add** to generate the client secret, as shown in Figure 13-8. When the secret is set to expire, it must be renewed in advance to avoid loss of access to the Microsoft online resources.

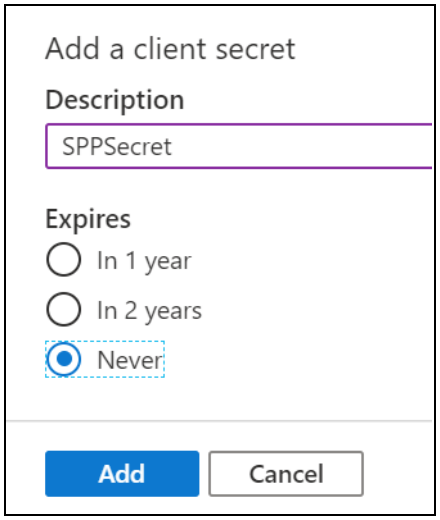

Figure 13-8 Add client secret

6. After completing the App registration, grant the API permissions as listed in Table 13-1. S

Table 13-1 API Permissions

API	Permission	Delegated or Application
Azure Active Directory Graph	User.Read.All	Delegated
Azure Active Directory Graph	Directory.Read.All	Application
Exchange	full_access_as_app	Application
Microsoft Graph	Calendars.ReadWrite	Application
Microsoft Graph	Contacts.ReadWrite	Application
Microsoft Graph	Files.ReadWrite.All	Application
Microsoft Graph	Mail.ReadWrite	Application
Microsoft Graph	Sites.ReadWrite	Application
Microsoft Graph	User.Read	Application
Microsoft Graph	User.Read.all	Delegated

7. To grant API permissions, select **API Permissions** → **Add permissions**.

   An example of the assigned permissions is shown in Figure 13-9 on page 371. Admin consent *must* be granted for configured permissions to use them.

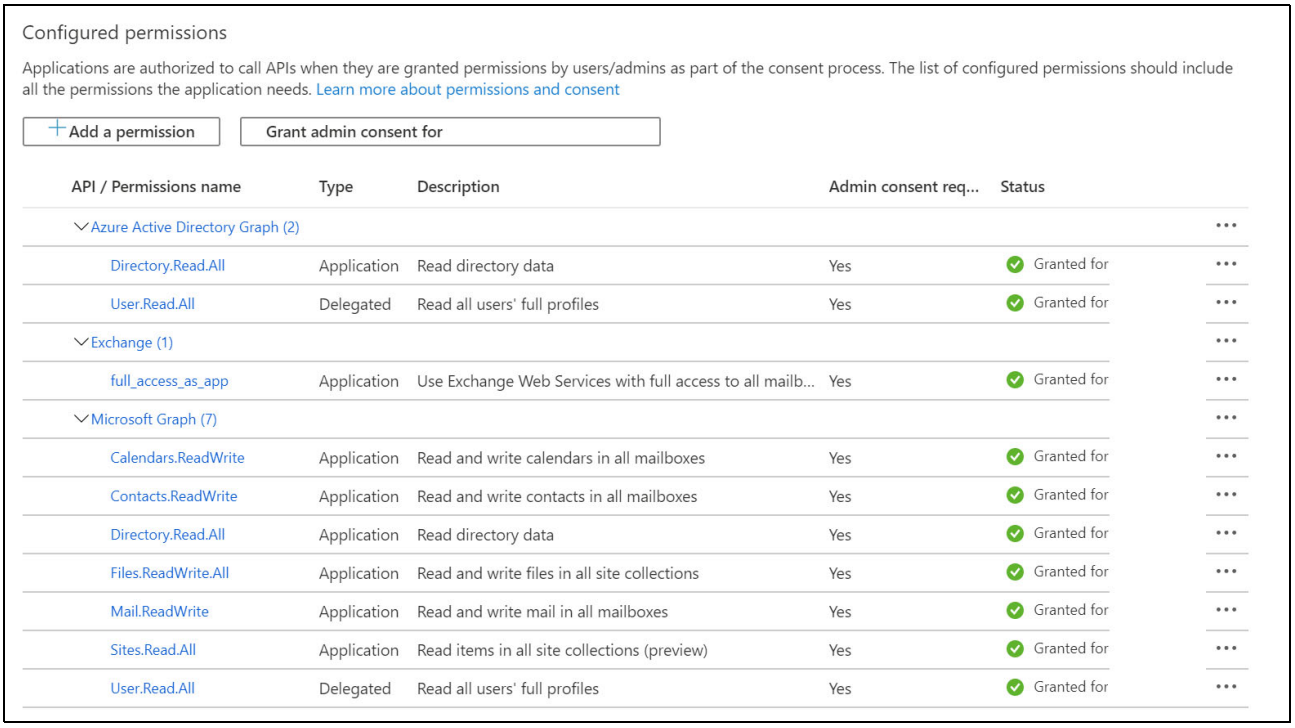

Figure 13-9 Configured API permissions

13.3 IBM Spectrum Protect Plus configuration for Microsoft 365

Open the IBM Spectrum Protect Plus Server UI and log in to the Dashboard. Select **Manage Protection** → **Cloud Management** → **Office 365.** Expand **Manage Application Server** and click **Add Application Server**.

Enter the credentials that were obtained from Azure Active Directory page, as shown in Figure 13-10.

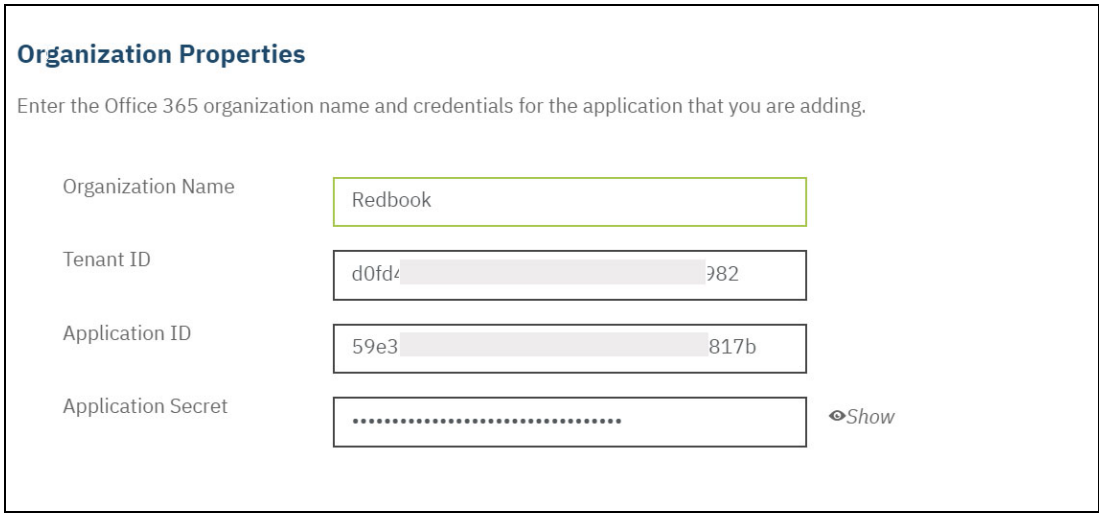

Figure 13-10 Organization properties

For more information about how to define these properties in Azure Active Directory, see 13.2.2, **Microsoft 365 application registration and API permissions** on page 368.

For the Proxy properties (see Figure 13-11), enter the host address or DNS name and user ID of the predefined Application Server, as described in 13.2.1, “Proxy host server”.

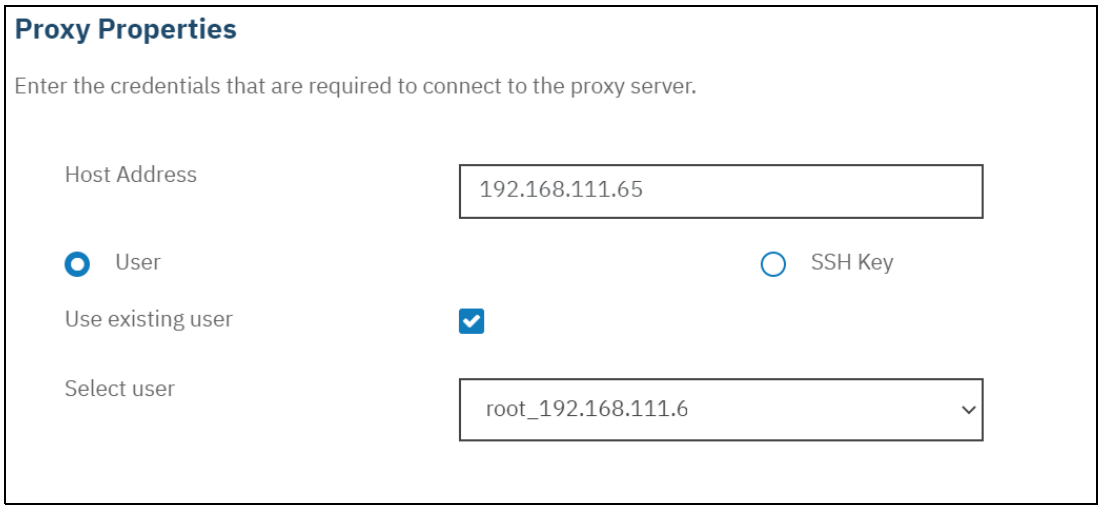

Figure 13-11 Proxy properties

Also, the maximum number of concurrent sessions (Microsoft 365 accounts) can be set. The default is 10, as shown in Figure 13-12.

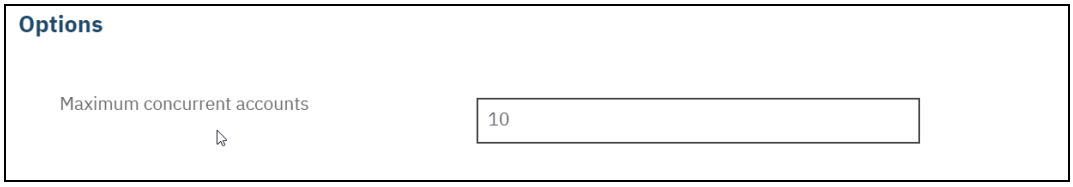

Figure 13-12 Options

After clicking **Save**, IBM Spectrum Protect Plus verifies the settings and registers the application server.

If the verification was successful, the message that is shown in Figure 13-13 is displayed.

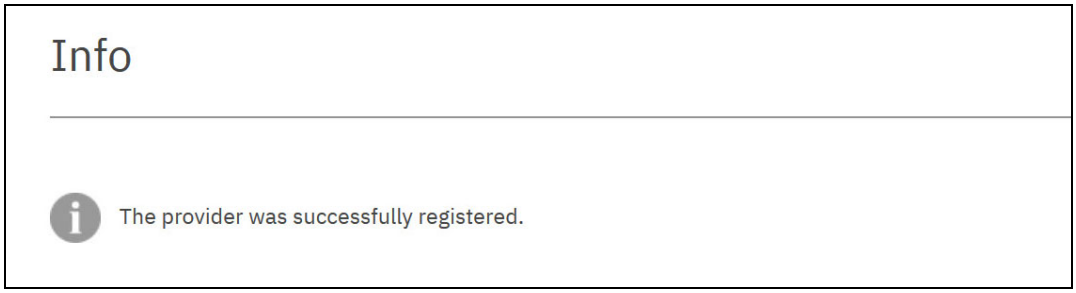

Figure 13-13 Success message

A test can be run to verify the correct setup, as shown in Figure 13-14.

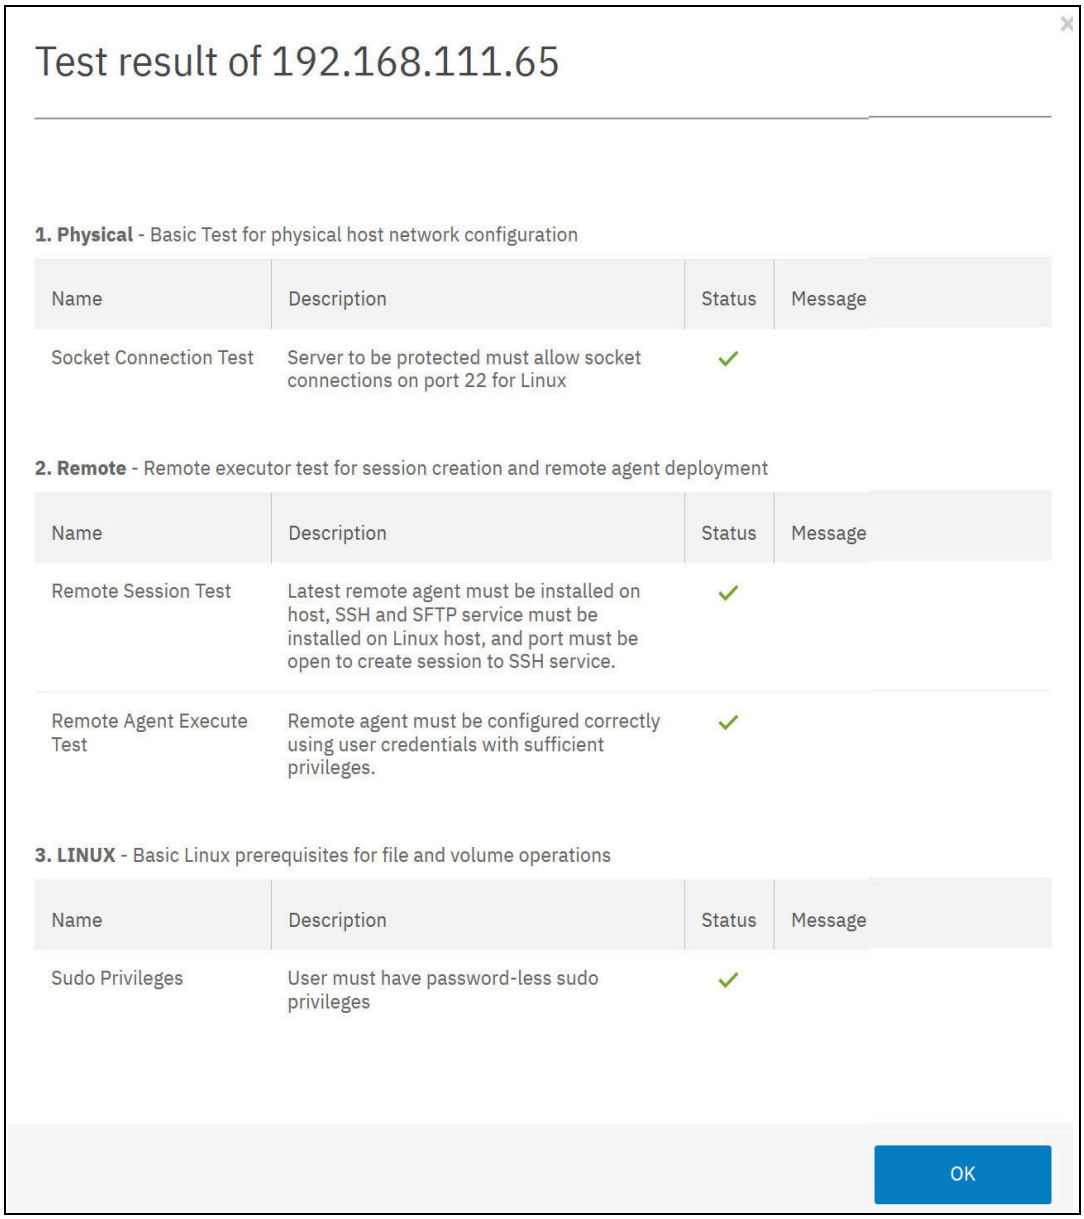

Figure 13-14 Test results

13.4 Protecting Microsoft 365 accounts

Backing up Microsoft 365 data is done on an incremental forever basis. After the initial full backup of data, only the changed items are backed up for any subsequent period as defined in the SLA.

An SLA can be assigned at different organization levels of the Microsoft 365 account:

- Calendars
- Contacts
- Mailbox
- OneDrive

13.4.1 Planning considerations

Microsoft 365 is hosted on a Microsoft Azure Infrastructure that offers high availability and redundancy for all data that is stored in Microsoft 365. Microsoft manages the updates of the infrastructure, but does not configure or administer the products.

Other than the possibility to run simple restore requests within Microsoft 365 by relying on certain recycle bins, no protection is available for other scenarios, such as the following example:

- Corruption or encryption of user data by malicious software
- Preserving deleted user data over default deletion period
- Erroneous or deliberate manipulation of data

These risks must be addressed by the customer. The use of IBM Spectrum Protect Plus provides protection against any type of potential data loss by allowing Microsoft 365 data to be backed up in a different and independent location and for long-term retention.

Integrating IBM Spectrum Protect or Object Storage into your Microsoft 365 protection strategy helps to provide more restore points on lower-cost storage and with lower total cost of ownership (TCO).

13.4.2 Configuring Microsoft 365 protection

Open the IBM Spectrum Protect Plus Server GUI and log in to the dashboard. Select **Manage Protection** → **Cloud Management** → **Office 365.** You can then select the organization to protect from the Organizations view, as shown in Figure 13-15.

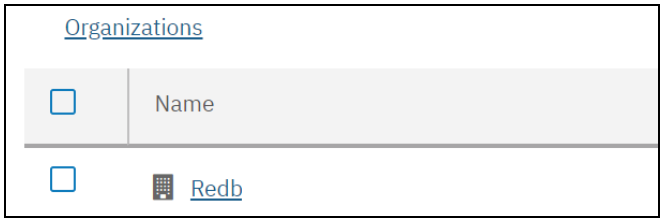

Figure 13-15 Organization Redb

By clicking the Organization, the included Microsoft 365 accounts can be browsed, as shown in Figure 13-16.

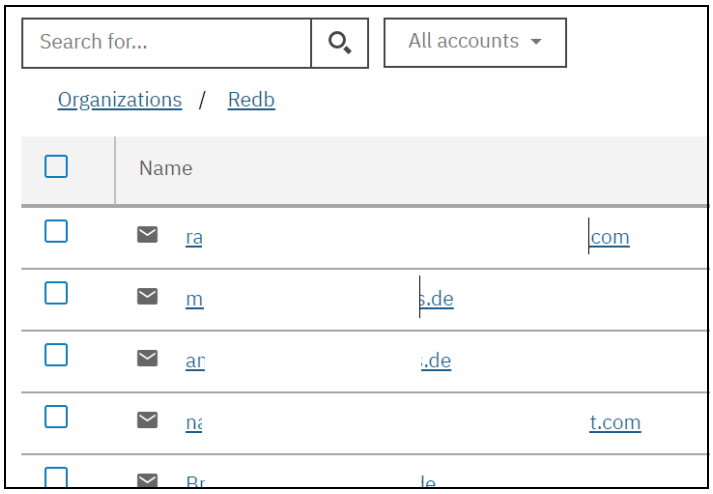

Figure 13-16 Sample Microsoft 365 accounts (real names are redacted)

When selecting a single mail address, the drill-down scope can be changed to Calendars, Contacts, Mailbox, and OneDrive (see Figure 13-17).

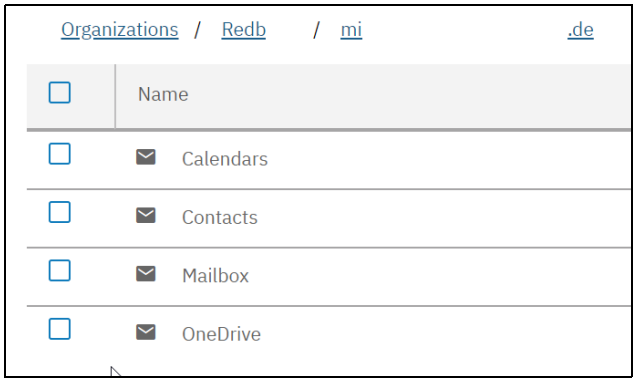

Figure 13-17 Microsoft 365 details

A best practice is to assign an SLA that contains the general requirements of the organization at the Organization level. This SLA is inherited by all underlying accounts and items. For any accounts that feature special SLA requirements, an exception can be made at the wanted level.

Search feature

Instead of assigning SLAs manually, the Search box can be used to filter the Microsoft 365 accounts. You can search for any string that is part of the Microsoft 365 account name.

Also, in the toggle box that is shown in Figure 13-18, accounts can be filtered to show only accounts that are not in an SLA or in any defined SLA.

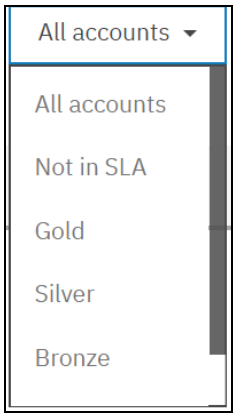

Figure 13-18 Search toggle

13.4.3 Restoring Microsoft 365 data

With IBM Spectrum Protect Plus, complete Microsoft 365 accounts or the following items can be restored:

- Calendars
- Contacts
- Mailbox
- OneDrive

The restore can be performed to either:

- Original account
- Other Microsoft 365 account

Data can be restored into target accounts to:

- Original location
- Alternative restore location (path)

Performing restore operations

To perform the restore operation, complete the following steps:

1. Open the IBM Spectrum Protect Plus Server GUI and log in to the Dashboard. Select **Manage Protection** → **Cloud Management** → **Office 365,** and click **Create Job,** as shown in Figure 13-19.

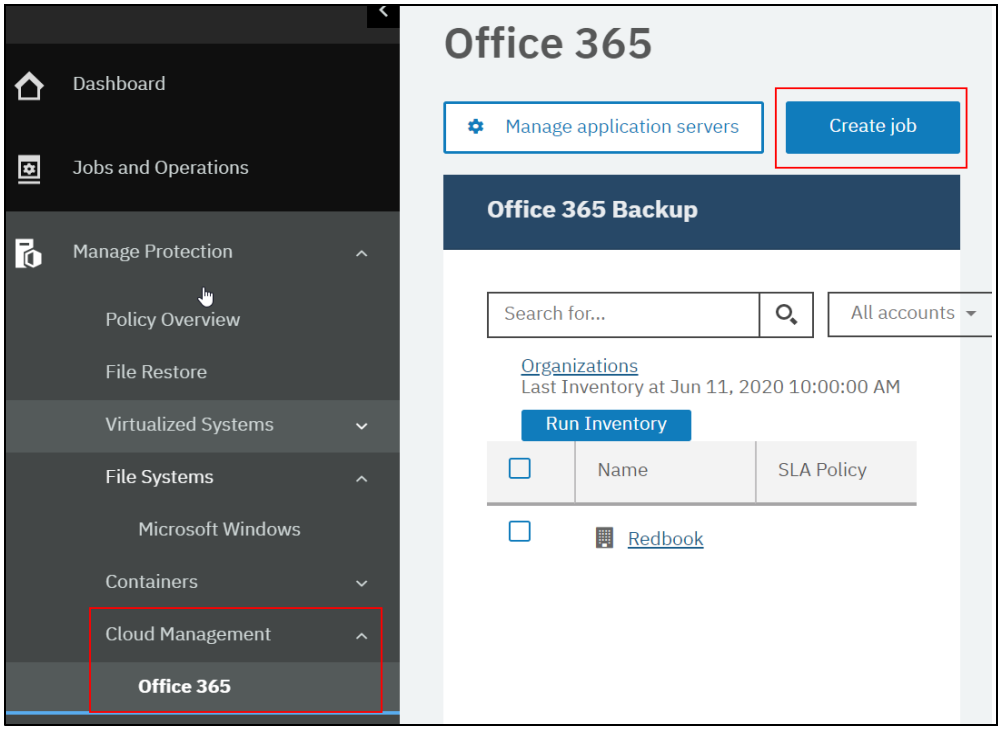

Figure 13-19 Create Job selection

2. In the next window, click **Select** (see Figure 13-20).

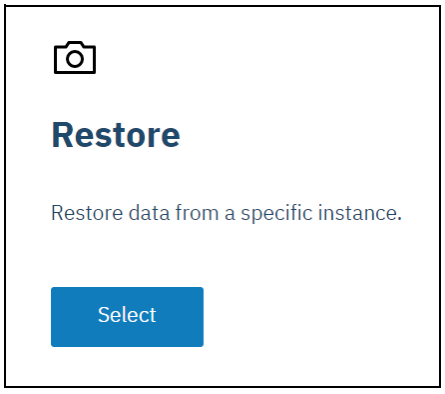

Figure 13-20 Select Restore

3. In the Select Source page, the Organizations can be browsed and specific Microsoft 365 accounts can be selected, as shown in Figure 13-21.

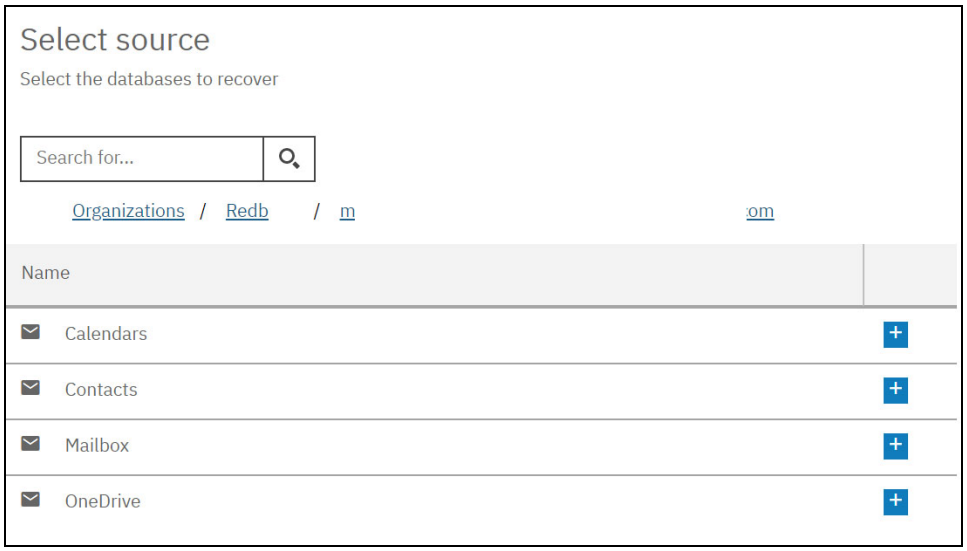

Figure 13-21 Select Source page

Inside the Accounts, specific Items can be added to Restore by clicking the plus sign (**+**) button, as shown in Figure 13-22.

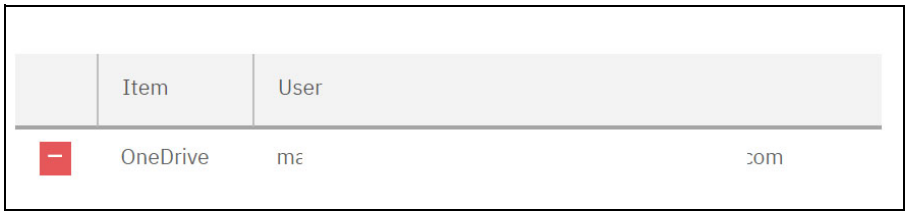

Figure 13-22 Select item for Restore

Instead of selecting items manually, the Search box can be used to find specific accounts or items.

4. Selecting the wanted items and click **Next**. A Snapshot type and destination must be configured. The Type of Restore is On-demand Snapshot and the selection of Location Type and Location depends on the vSnap server configuration and Sites, as shown in Figure 13-23.

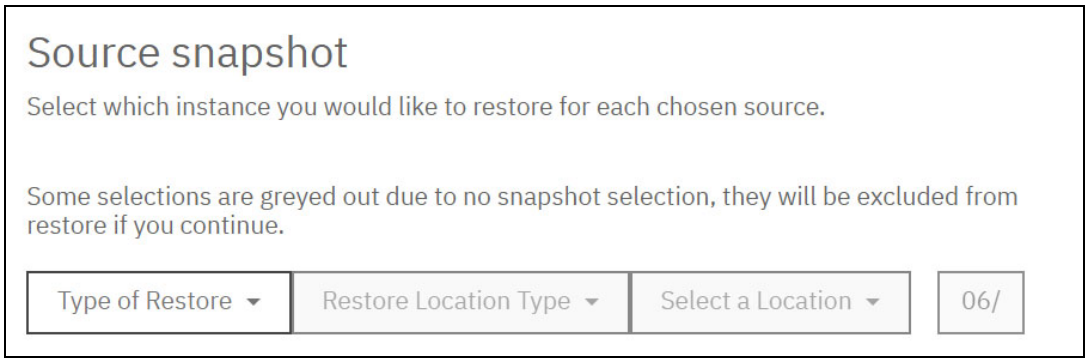

Figure 13-23 Snapshot type selection

5. After the Snapshot Type is selected, click **Backup** to choose a restore point (see Figure 13-24 in which the restore point is highlighted). Click **Next** to open the Restore dialog.

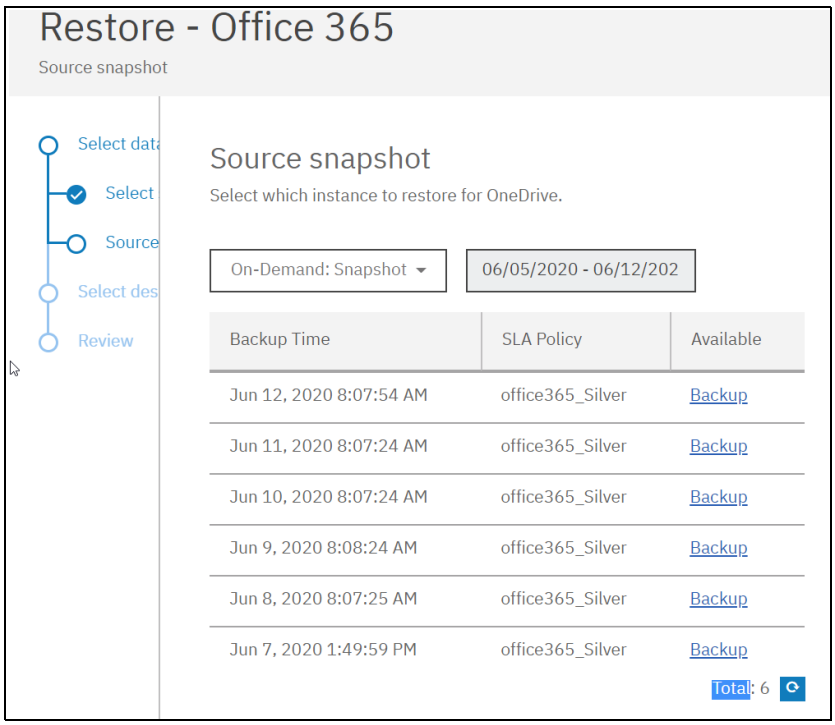

Figure 13-24 Restore point selection

6. Highlight the snapshot to restore by clicking **Backup** in the Available column. To proceed to the Restore destination window, click **Next**.

   Here, you also can define whether the Restore is performed into the original account or in any other account in the same organization.

   The Restore Path is optional and can be any subfolder (if the subfolder does not exist, it is created) inside the defined Microsoft 365 account. A slash is used to separate between folder and subfolders, as shown in example Figure 11-23.

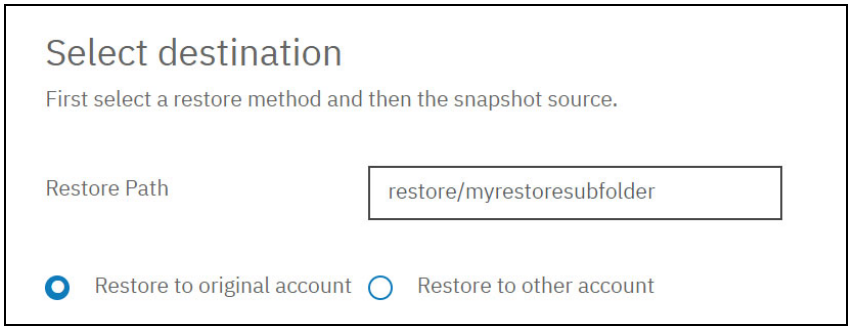

Figure 13-25 Restore Destination

7. The next window shows a summary of the Microsoft 365 restore job (see Figure 13-26). Click **Submit**.

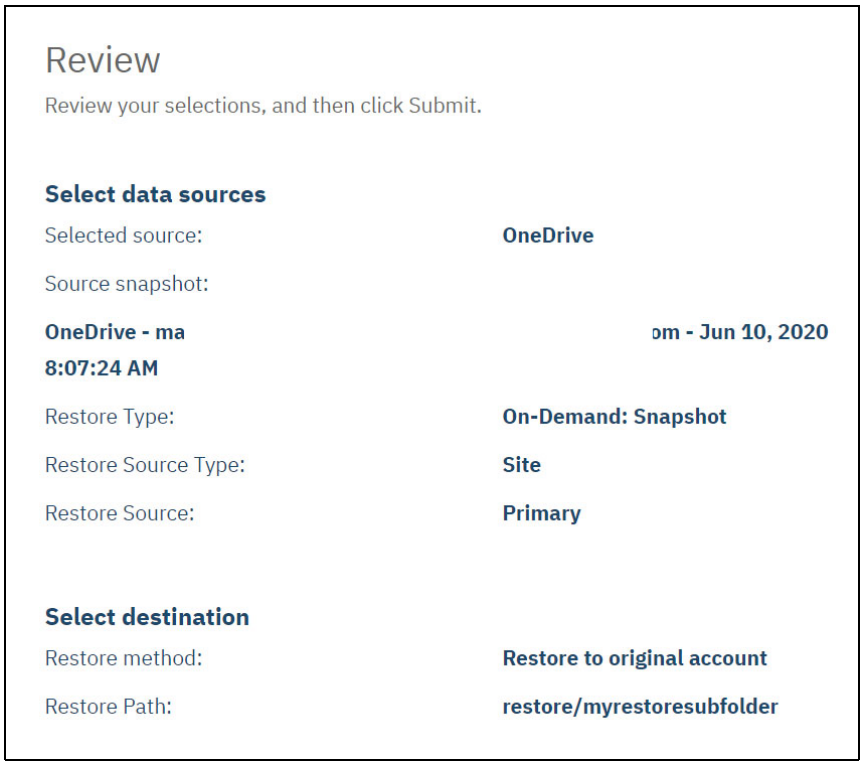

Figure 13-26 Restore summary

13.5 Exchange Hybrid Environments

An Exchange Hybrid Environments combines on-premises Exchange hosted mailboxes and Microsoft 365 accounts in the same Exchange domain. IBM Spectrum Protect Plus provides data protection for Microsoft Exchange Server and this function can be used to protect the on-premises mailboxes.

The Microsoft 365 function uses data protection for mailboxes that are hosted in Microsoft Azure. Therefore, a combination of two solutions can be used to protect the complete Exchange Hybrid Environment.

However, mailboxes that are moving from on-premises Exchange to Microsoft Azure and vice versa must be tracked.

14

Backing up and restoring containers

This chapter presents the conceptual foundations of containers and the open source container orchestration engine Kubernetes. It also discusses the integration with IBM Spectrum Protect Plus Container Backup Support that manages backup and restore operations by using Kubernetes custom resources. Finally, it introduces the Red Hat Enterprise Kubernetes product that is called Red Hat OpenShift.

This chapter includes the following topics:

14.1 Containers and orchestration with Kubernetes

Containers are a new way of packaging and running applications that are quickly changing the IT landscape. Containers allow multiple separate applications to share the underlying hardware; however, unlike traditional virtual, there is no need for a different operating system instance for each running container. Instead, containers rely on distinctive features to isolate applications. This approach allows containers to be smaller and lighter weight than traditional virtual machines (VMs). It also enables changes in application design, making it more practical to have many small related pieces of functionality that are known as *microservices*.

Although it is possible to run containers on an individual machine, the use of containers at scale is only practical with a container management platform. Although several platforms currently exist, Kubernetes is the de-facto standard container orchestration platform. It forms the basis of most cloud-based container services, such as IBM Kubernetes Service and an enterprise-class container management platform as OpenShift. A complete exploration of Kubernetes and OpenShift is beyond the scope of this document; however, the following core concepts are worth highlighting:

- State

  Containers can be stateless or stateful. With stateless containers, the code that is needed for the container to operate is part of the container image, and any configuration information or data it requires to function is outside the container.

  No data must be persisted inside the container across reboots. They do not require a traditional restore; just redeploy them if needed. Front-end web servers are a typical example. However, stateful containers must store data and have that data available. A typical example is the database back end of a web application.

- Registries

  Container images are stored in registries. Several large public registries are used in which images for “off the shelf” applications, such as Nginx, are published. Most enterprises that write in-house applications also have private registries that are hosted by a cloud provider or run internally.

- Persistent volumes

  When containers must store data that lives through restarts, they use persistent volumes. As the name indicates, these storage volumes are available until deleted. Kubernetes manages access to them through a construct known as a persistent volume claim (PVC).

- Persistent Volume Claim

  A request to a storage system for a place to store data; that is, PVC. With dynamic provisioning, a developer or DevOps engineer can deploy an application with a persistent volume claim, and a volume automatically creates and assigns.

- Configuration files

  Kubernetes uses YAML files to describe the wanted state of containers and other constructs, such as volumes. These files define the various properties of a specific object and can be updated to instruct Kubernetes to update the configuration. Although a Kubernetes web interface is available, this feature is the primary mechanism for telling Kubernetes what to do.

- Etcd

  Within Kubernetes, a central configuration management database that is known as Etcd is available. It runs on a particular type of node that is known as a *master node*. In a typical deployment, it runs on at least three master nodes, which makes it highly available.

- Networking

  Kubernetes allows for the use of several different software-defined networking technologies. Although significant differences exist between them, running containers can talk only to other containers within the same cluster in most cases unless the container's definition includes special instructions.

- Container Storage Interface (CSI)

  Kubernetes uses the CSI as an abstraction layer for underlying storage technologies. Storage vendors can write a CSI driver, and then Kubernetes dynamically provisions persistent volumes as needed. CSI also supports other storage functions, such as snapshots, though not all CSI enabled storage includes support for the snapshot function. The CSI standard replaces a legacy approach that is known as in-tree storage plug-ins that Kubernetes is deprecating.

- Namespaces and labels

  Kubernetes provides several constructs to make it easier to manage larger-scale deployments where multiple teams share a Kubernetes instance. Namespaces provide a way to create separate islands; for example, to keep two teams from seeing each other's containers. Labels allow a user to tag containers so that operations, such as performing a backup, can occur without itemizing everything that requires backing up.

- Config Maps and secrets

  Config maps provide a way to pass configuration data to a running container without hardcoding that information in the container image that is stored in the registry or the YAML file definition. Secrets provide similar functions for sensitive information that requires extra privacy and security.

14.1.1 Kubernetes and virtualization analogies

When learning about a new construct, such as Kubernetes, it can be helpful to think of it in the context of a well-known entity, such as virtualization. Although this learning tool can useful initially, all such analogies are inherently inexact.

Kubernetes (commonly stylized as k8s) is an open source container orchestration system for automating application deployment, scaling, and management. If you compare Kubernetes with the vCenter Server, it unifies resources from individual hosts that are sharing them among the entire cluster. It unifies resources by managing the assignment of VMs to the hosts and the allocation of resources to the VMs within a specific host that is based on the policies that the system administrator sets.

The critical distinction between containers and VMs is that VMs share the hardware, but each VM has an independent but isolated operating system instance. Containers share the core operating system (kernel), but individual applications running on it are separated and unaware of each other.

Figure 14-1 on page 384 shows how the three deployment models characteristics differ. Traditional deployment workloads examples include physical and bare metal deployments, files, databases, and applications. The virtualized deployment consists of multiple VMs, each with a separate operating system on a single physical server's CPU cores. The container-based deployment, allows multiple applications to be within the same operating system but with their own independent sets of libraries and isolated storage and networking.

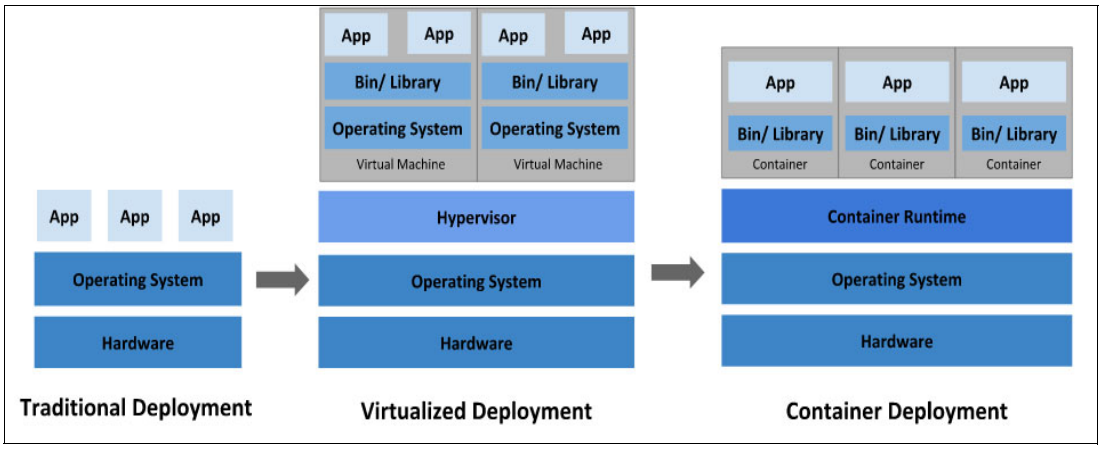

Figure 14-1 Application deployment models

For more information, see the following resources:

- What is Kubernetes?
- Kubernetes Documentation: Concepts
- OpenShift Documentation web page

Within Kubernetes, nodes are the underlying machines that run the containers. Nodes can have one or more roles, and the two most common roles are worker nodes and master nodes. *Master nodes* run the command and control components of Kubernetes that provide the orchestration and intelligence. *Worker nodes* are responsible for running the containers. Comparing Kubernetes to VMware, the master nodes provide functions that are similar to what is found in the vCenter or vCloud director; worker nodes are analogous to ESXi servers.

The most straightforward construct for the containers is a pod. A *pod* contains one or more containers. Real-world workloads often use higher-level constructs, such as daemon sets, stateful sets, and deployments. These constructs define aspects of the environment, such as:

- How many copies of a POD are running.
- Their placement relative to each other.
- Exposing the pods in the network.
- How they use storage.

The distinctions between the nodes are beyond the scope of this document. Although a like-for-like construct is not available within traditional VMware, thinking of these as a vApp is a starting point.

Pods that are running in Kubernetes are stateless by default. Being stateless means that when restarted, data that is generated during the previous run is not accessible. No direct equivalent is available to this behavior in VMware. Persistent data volumes in Kubernetes solve this problem and can originate from several underlying storage technologies, such as the following examples:

- AWS Elastic Block Storage
- Azure Volume
- vSphere Volume
- NFS
- iSCSI
- CephRBD
- Portworx

The coupling between the underlying storage and Kubernetes evolved significantly. The current direction is to integrate through an abstraction layer that is known as CSI. The CSI functions in a roughly equivalent manner to vVols within VMware, which allows for dynamic provisioning of storage volumes within Kubernetes as new pods deploy. The mapping of the volumes to the consuming pods occurs through a construct that is known as a *persistent volume claim*. Within VMware, this mapping is similar to indicating which VMs can access which VMDK file.

14.1.2 Applications and containers

Unlike traditional applications, applications are not installed at all when running as containers. Instead, developers and DevOps engineers deploy from a container image registry. After they are deployed, one or more copies of the container image run until they complete or a Kubernetes user stops them, at which point the lifecycle of the container ends.

Many container images are stateless, meaning they do not need to retain any data between runs. The image includes the code that is necessary to operate along with any libraries that are needed or other components. Kubernetes passes along configuration items through constructs, such as config maps and secrets. The Kubernetes metadata database stores and secures those items and expects them to be reasonably sized.

Some applications, however, require data to be retained across multiple runs or restarts of the container. A classic example is a relational database. Redeploying an empty database is of no use if the data is gone. To handle this issue, Kubernetes introduces the idea of persistent data volumes that can be attached to stateful containers. These attachments are known as *claims*.

Kubernetes inherently offers various capabilities to make applications resilient, whether they are stateless or stateful. It monitors running groups of containers (pods) and ensures that the requested number of copies of are running always. It can also move them around between the physical worker nodes that are running the containers. If a worker node fails, this behavior protects against routine hardware failures. The underlying infrastructure design drives the availability of the container image registry and persistent data.

Although this behavior helps protect against infrastructure failure by providing redundancy, it does not provide resiliency. It is not possible to return to a known good state to deal with human error, ransomware, or other higher-level modes of failure. Production-ready resiliency that is equivalent to that resiliency that is found in systems requires an external data protection approach.

14.1.3 What to protect within Kubernetes

Protecting Kubernetes containers, includes:

- Cluster metadata

  Kubernetes stores most of the cluster's operational state in the ETCD cluster on the master nodes. A backup of ETCD is always considered a best practice for production clusters. However, ETCD is not the only consideration. It is also essential to have backups of some of the static assets that are created during deployment or updated later; for example, PKI assets or certificates that are used by the Kubernetes API server and any secrets or config maps.

- Application Data:

  In addition to all of the metadata that is necessary to reconstitute Kubernetes, recovering stateful pods is useless unless the persistent data that is associated with those pods also is recovered. This process involves protecting the stateful data on persistent volumes. Containerized applications also need crash or application-consistent backups, depending on the requirements of the application.

14.2 IBM Spectrum Protect Plus Integration with Kubernetes

In this section, we describe two use cases that show how IBM Spectrum Protect Plus integrates with Kubernetes.

14.2.1 Use cases and personas

The following use cases are described:

- Protect persistent data of containerized workloads or applications
- Data reuse

Personas that are involved in the uses cases:

- Backup (architect/admin) Jose (operations):
  - Jose installs data protection for containers that provided predefined SLA options to be available for Jane. Jose likes self-service capabilities so he can allow Jane to assign SLA to PVCs without his involvement.
  - Jose is concerned about resiliency and fault domain and security (offload considerations and reliance on Kubernetes security model for access controls).
- Dev/Test/Restore (application developer) Jane (developer):
  - She needs to protect her application and likes the idea of reusing data snapshots for testing.
  - She uses the native Kubernetes interface (kubectl) to assign backup/restore policies to perform backup/restores of PVC.
  - Capture copy of production data.
  - Restore snapshot to an alternative container in the same cluster.

Jane is the Kubernetes administrator who must ensure that her workloads are up and running. Jose is the backup administrator, and he is ultimately responsible for ensuring protecting the data across the company. Because the organization uses many data types and platforms, Jose needs a single centralized way of viewing and interacting with the data protection capabilities rather than dealing with a half dozen different point solutions.

Jane focuses on all aspects of container management, including protection, and her method of administration is by way of the kubectl CLI. Fortunately, IBM Spectrum Protect Plus accommodates both with a self-service approach in which Jane manages and performs data protection operations through native kubectl commands that operate on YAML files or by using the IBM Spectrum Protect Plus RBAC enabled web interface.

14.2.2 Solution architecture, planning, and design

This section provides a broad understanding of how to design and implement an IBM Spectrum Protect Plus solution in a Kubernetes environment with a focus on concepts and decisions rather than individual steps or commands. For procedural-oriented guidance, see IBM Knowledge Center.

IBM Spectrum Protect Plus focuses on protecting persistent data within Kubernetes environments. Unlike stateless containers or images, persistent data cannot be redeployed if it is lost because it is lost forever. As with other data types, it is essential to ensure that the data is not only captured, but is usable. To accomplish this goal, IBM Spectrum Protect Plus integrates with the CSI API layer of Kubernetes.

By using this layer, IBM Spectrum Protect Plus makes API calls to create a point-in-time consistent snapshot of the CSI volume. It then can mount the snapshot on an alternative container (known as a *Data Mover*) that is part of IBM Spectrum Protect Plus and transfer the data to the vSnap server for storage along with the other backup types.

IBM Spectrum Protect Plus uses the CSI to start and persist snapshots for Kubernetes PVC by using Ceph block storage in version 10.1.5 and 10.1.6. An application developer uses the native Kubernetes command line (kubectl) to perform this function in a self-service manner and can create Kubernetes-aware backup scheduling automation.

The backup Service Level Agreements (SLAs), also known as *backup policies* in IBM Spectrum Protect Plus, provide control for scheduling, snapshot retention, and copy to vSnap server. If the original volume is damaged or lost, the snapshot or copy backups on the vSnap servers can facilitate recovery. Figure 14-2 shows how Kubernetes' Backup Support is deployed in the Kubernetes environment and interacts with IBM Spectrum Protect Plus.

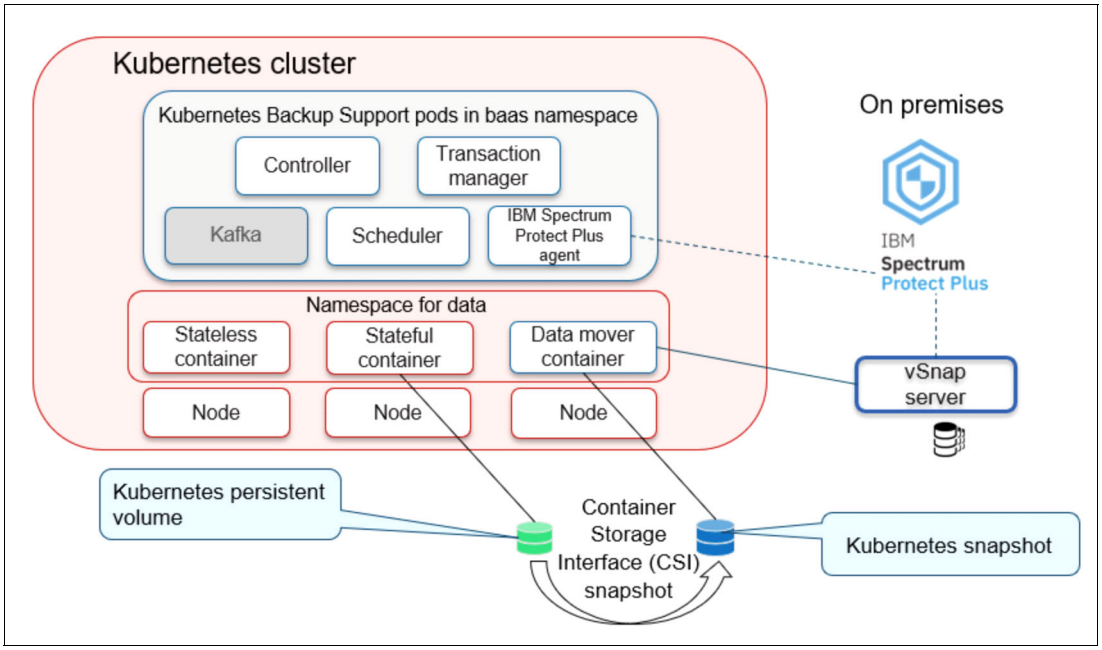

Figure 14-2 Kubernetes' Backup Support

This type of integration offers the following advantages over other strategies, such as backing up underlying storage directly or putting backup clients inside individual PODs:

- It avoids the need to inject special software or components into the pod or deployment definitions of the workloads IBM Spectrum Protect Plus protects (zero touch).
- It reduces the reliance on specific hardware that is underneath Kubernetes or access to the hardware at all. It requires only that the storage is CSI-based and support the snapshot function of the CSI API. This aspect is especially useful when cloud-based Kubernetes services are used where the cloud provider controls the Kubernetes nodes and infrastructure.
- It allows IBM Spectrum Protect Plus to use Kubernetes native capabilities for grouping workloads when defining jobs and policies, such as namespaces
- It provides a holistic view of volume recovery, taking advantage of snapshots on storage and the data that is stored in the vSnap server repository, which enables greater flexibility.
- It enables native integration for self-service by using typical Kubernetes constructs, such as kubectl and YAML files, while also using the interfaces and reporting structure within IBM Spectrum Protect Plus.

  This deep integration makes Kubernetes backups a first-class citizen with the same management and governance capabilities as other data types.

An operator governs the IBM Spectrum Protect Plus components that are installed within Kubernetes. It, in turn, manages several deployments in a microservices architecture (see Figure 14-3).

Component	Replica	Request CPU	Limit CPU	Request Memory	Limit Memory
baas-spp-agent	1	2	3	800Mi	1000Mi
baas-cert-monitor	1	250m	1	50Mi	250Mi
baas-datamover	1	100m	500m	500Mi	1000Mi
baas-kafka	1	500m	2	600Mi	2Gi
baas-scheduler	1	100m	750m	150Mi	500Mi
baas-controller	1	250m	1	50Mi	250Mi
baas-transaction-manager	3	200m	1	100Mi	500Mi
baas-transaction-manager-worker	3	200m	2	250Mi	500Mi
baas-transaction-manager-redis	3	50m	200m	50Mi	250Mi

Figure 14-3 IBM Spectrum Protect Plus components in Kubernetes

14.3 Installing the IBM Spectrum Protect Plus service in Kubernetes

In this section, we describe the various aspects of installing the IBM Spectrum Protect Plus service in Kubernetes.

14.3.1 Installation prerequisites

Before installing the Kubernetes data protection components, always ensure that the IBM Spectrum Protect Plus server and vSnap server versions are up to date to avoid experiencing negative impacts from existing APARs.

For more information about prerequisites, including a list of supported operating systems, Kubernetes versions, and CSI drivers, see IBM Knowledge Center.

Online or offline installation

Installing IBM Spectrum Protect Plus Kubernetes Backup Support is done by using one of the following methods:

- Downloading and installing the Helm package from IBM Helm Charts Repository and IBM Entitled Registry.
- The use of a download from IBM Passport Advantage® Online to facilitate offline installations in cases where outbound internet access is constrained.

The use of the Helm-based approach requires creating a pull secret. Pull secrets are used by registries, such as the IBM registry, to ensure that the request to pull-down a container image is coming from an entitled user.

The offline method is also registry-based. Rather than pulling the images down from the IBM public registry, it instead pushes them first from the downloaded package into the private registry that is used by Kubernetes. The installer must know about your private Docker registry so it can push the container images for the IBM Spectrum Protect Plus product to it. Then, the Helm charts are used but point to the private registry instead.

Note: Separate installers are used for each approach. Attempting to use the offline installer with the IBM registry results in the installer attempting to push to it rather than pull from it and fail.

Although the installer can run from any Linux machine, it expects the following components to be in place:

- Kubectl: Kubectl is the Kubernetes command-line management tool. It must be installed and configured to communicate with the cluster IBM Spectrum Protect Plus is configured to protect.
- Helm: Helm is a Kubernetes tool for automation software installation. It requires initialization to work with a specific cluster.
- Docker: Docker must be installed and the Docker service must be running to facilitate offline installations, even when a containerization engine other than Docker is used.

Installing the Helm client

In many cases, Helm and tiller are installed in the environment because they are the recommended way of deploying complex applications into Kubernetes environments. The IBM Spectrum Protect Plus installation package includes a Helm setup script (`helm_install_k8s.sh`) that is designed for use when the Helm client is not present.

This script deploys a specific version of the Helm client from the v2 family, which can be older than the version of Helm that is used in the environment. Helm v3 represents a substantial change; the v3 family no longer uses tiller on the cluster.

Although the Helm setup script that is provided by IBM Spectrum Protect Plus attempts to configure RBAC for Helm correctly by creating a service account for the tiller deployment and a `clusterrolebinding`, it does not force the redeployment of tiller if it is running on the cluster.

After running the script, use a basic command, such as **`helm is`** to confirm that the Helm client is functioning correctly. If errors occur, you might need to debug the RBAC configuration for the running tiller instance or use Helm reset to remove the configuration, followed by rerunning the SPP provided script. The script expects the tiller deployment to be running as a service account. You can confirm this configuration by checking the deployment YAML by running the following command:

```
kubectl --namespace kube-system get deploy tiller-deploy -o yaml
...
serviceAccount: tiller
serviceAccountName: tiller
....
```

That service account should have a clusterrolebinding to the cluster-admin role:

```
kubectl describe clusterrolebinding tiller
Name:         tiller-clusterrolebinding
Labels:       <none>
Annotations:  <none>
Role:
  Kind:  ClusterRole
  Name:  cluster-admin
Subjects:
  Kind            Name    Namespace
  ----            ----    ---------
  ServiceAccount  tiller  kube-system
```

In early versions of 10.1.6, the `helm_install_k8s.sh` script leaves behind the Helm and tiller binaries in a subdirectory that is named `linux-amd64` in whatever directory from where you ran the script. If you ran the script from a directory that is part of the installer structure, it must be removed before proceeding to avoid an error during the installation. Failure to remove it or run it from a directory outside the installation path results in installations failing with the following message:

```
Error: grpc: received message larger than max (81905241 vs. 20971520)
```

Also, the installer includes a connectivity test capability. It attempts to connect from wherever it is running to the IBM Spectrum Protect Plus server, even if it is running outside the cluster.

Enabling the VolumeSnapshotDataSource feature

IBM Spectrum Protect Plus Kubernetes Backup Support creates snapshots of persistent volumes to preserve a point-in-time copy of data. For restore or “copy to vSnap” operations, these snapshots must be cloned, which requires a special feature to be enabled in the Kubernetes environment.

In Kubernetes 1.16, this “VolumeSnapshotDataSource” feature is in Alpha state. If you want to use IBM Spectrum Protect Plus Kubernetes Backup Support in a Kubernetes 1.16 cluster, a so-called *feature gate* must be established to enable the required alpha feature.

For more information about how to enable an alpha feature in Kubernetes 1.16 see, IBM Knowledge Center.

Verifying whether the Metrics Server is running

The Metrics Server is used by the Kubernetes Backup Support scheduler component to determine the resources that are used by concurrent data mover instances.

If the Metrics Server does not return data, the number of data movers that are used for backup operations is limited, which might negatively affect performance; therefore, it is recommended to have the Metrics Server correctly configured in the Kubernetes cluster.

For more information about how to validate and enable the Metrics Server see, IBM Knowledge Center.

Creating an image-pull secret (for external registry only)

For more information about how to create an image-pull secret for a secure Docker image registry, see IBM Knowledge Center.

14.3.2 Preparing the installer configuration file

The Kubernetes Backup Support installer uses a configuration file (`baas_config.cfg`) to read various parameters during the installation process. In either case, you must gather the following configuration information from your cluster before the installation runs:

- IBM Spectrum Protect Plus user credentials (BAAS_ADMIN and BAAS_PASSWORD)

  These credentials are within the IBM Spectrum Protect Plus application when the Kubernetes plug-in registers. Because these credentials sit in a configuration file that can be run by the Kubernetes administrator rather than by the IBM Spectrum Protect Plus administrator, a separate user within the IBM Spectrum Protect Plus application must be defined. For increased security, specify an empty string ("") for BAAS_PASSWORD. You are prompted for the password when you run the deployment script. If you must specify a password in the configuration file for automated test deployments, ensure that the file is stored in a secure location.

- IBM Spectrum Protect Plus server IP address (SPP_IP_ADDRESS)

  This IP address is the IP of the IBM Spectrum Protect Plus “master server” rather than the vSnap server.

- Cluster CIDR address (CLUSTER_CIDR)

  This address is obtained by examining the output of the **`kubectl cluster-info dump`** command. The formatting of the output can differ between Kubernetes platforms that are based on the CNI plug-in that is used. This network is *not* the network on which the worker nodes or pre-existing pods sit.

- Cluster API server (CLUSTER_API_SERVER_IP_ADDRESS and CLUSTER_API_SERVER_PORT)

  Kubernetes is an API-driven environment that facilitates interactions with Kubernetes. It uses an API server to represent itself rather than requiring direct interactions with all of the component services. This information is available from several sources and in some cases, might be obtained from a cloud provider, such as IBM Kubernetes Service by way of the cloud providers interface in addition to getting it by using Kubernetes native commands. To obtain this information from Kubernetes directly, run the following command:

  ```
  kubectl get endpoints -n default -o yaml
  ```

  If multiple IP addresses are listed, use any one that is reachable from the SPP installation.

- Cluster name

  This name is used in IBM Spectrum Protect Plus. Although it does not need to match a name on the Kubernetes side, it should be meaningful to both the Kubernetes and IBM Spectrum Protect Plus administrators.

- Product registry details (PRODUCT_IMAGE_REGISTRY)

  This detail tells the installer where to get the IBM Spectrum Protect Plus Kubernetes plugin-container images. It might refer to the IBM public registry or a private registry. If an external registry is used, create the Docker pull secret for the registry in use and indicate the name of the secret in the PRODUCT_IMAGE_REGISTRY_SECRET_NAME field. If an internal registry is used, enter an empty string ("").

Example configuration file

Example 14-1 shows a `baas_config.cfg` that uses an internal (private) registry.

Example 14-1 Example baas_config.cfg using a private registry

```
# =========================
# BaaS Config File (v0.1.2)
# =========================

# ------------------------------------------
# GLOBAL BAAS CREDENTIALS
# ------------------------------------------

BAAS_ADMIN="isppadmin"
# PRODUCT ACCOUNT PASSWORD (leave empty "" to be prompted for it)
BAAS_PASSWORD=""

# ------------------------------------------
# PRODUCT LICENSE CHECK
# ------------------------------------------

LICENSE="ACCEPTED"

# ------------------------------------------
# IBM SPECTRUM PROTECT PLUS (SPP) CONFIGURATION
# ------------------------------------------

SPP_IP_ADDRESSES="10.0.240.222"
SPP_PORT="443"

# SPP AGENT NODEPORT
```

```
# Leave empty "" to have a nodePort value automatically assigned
SPP_AGENT_SERVICE_NODEPORT=""

# -------------------------------------------
# NETWORKPOLICY CONFIGURATION OPTIONS
# -------------------------------------------

CLUSTER_CIDR="192.168.0.0/16"

CLUSTER_API_SERVER_IP_ADDRESS="10.0.240.150"
CLUSTER_API_SERVER_PORT="6443"

CLUSTER_NAME="Prague-K8s-Cluster"

# -------------------------------------------
# PRODUCT DEPLOYMENT AND CONFIGURATION OPTIONS
# -------------------------------------------

PRODUCT_NAMESPACE="baas"

OPERATOR_NAMESPACE="default"

PRODUCT_TARGET_PLATFORM="K8S"

PRODUCT_LOCALIZATION="en_US"

PRODUCT_LOGLEVEL="INFO"

# ----------------------
# PRODUCT IMAGE REGISTRY
# ----------------------

PRODUCT_IMAGE_REGISTRY="10.0.240.150:5000"

PRODUCT_IMAGE_REGISTRY_NAMESPACE="baas"

PRODUCT_IMAGE_REGISTRY_SECRET_NAME=""
```

14.3.3 Running the installer

With the prerequisites in place, running the installer creates output that is similar to the truncated view that is shown in Example 14-2.

Example 14-2 Example Installer output

```
./ baas_install.sh -i
Script baas_install.sh started at Wed Jul 29 12:13:00 CEST 2020
### Starting prerequisites check at Wed Jul 29 12:13:00 CEST 2020 ###
Sourcing BaaS configuration file baas_config.cfg...
WARNING: No secret name for imagePullSecret for image registry provided in
baas_config.cfg. May not be required for local cluster registry.
Using product localization en_US as defined in baas_config.cfg for the deployment.
Targeting generic Kubernetes (K8S) as container orchestration platform as defined
in baas_config.cfg for the deployment.
Found BaaS version 10.1.6 of Helm chart baas-k8s/baas
```

```
10.0.240.222 is already an ipaddress
10.0.240.150 is already an ipaddress
Validating port values in baas-k8s/baas/values.yaml
Checking for kubectl command line tool on local system...
Checking for active connection to target Kubernetes cluster...
Kubectl is connected to the Kubernetes target cluster
>>kubernetes-admin@kubernetes<<.
Checking for helm client availability on local system...
Checking for proper initialization of helm client/server on local system...
Checking for Docker running on the local system...
Requesting password for BAAS ADMIN ACCOUNT (isppadmin)...

-----------------------[BAAS ADMIN PASSWORD REQUIRED]-----------------------
Please enter the isppadmin ACCOUNT PASSWORD:********
Please repeat the isppadmin ACCOUNT PASSWORD:********
------------------------------------------------------------------------

-----------------------[CONFIRMATION REQUIRED]------------------------
Would you like to verify the connectivity to IBM Spectrum Protect Plus server?
Please enter 'yes' to continue: yes
------------------------------------------------------------------------
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  2983    0  2983    0     0   4598      0 --:--:-- --:--:-- --:--:--  4596
Successfully verified connectivity to IBM Spectrum Protect Plus server
10.0.240.222
Created namespace >>baas<< for product secrets.
secret/baas-secret created
### Starting installation of product release baas version 10.1.6 at Wed Jul 29
12:13:31 CEST 2020 ###
-----------------------[CONFIRMATION REQUIRED]------------------------
Please confirm to continue with installation to target cluster
>>kubernetes-admin@kubernetes<<:
Please enter 'yes' to continue: yes
------------------------------------------------------------------------
Using existing namespace >>baas<< for product deployment.

Loading Docker images from ./images/baas-10.1.6.tar.gz into local Docker image
repository...
Loaded image: baas-transaction-manager-redis:10.1.6
Loaded image: baas-cert-monitor:10.1.6
Loaded image: baas-controller:10.1.6
Loaded image: baas-scheduler:10.1.6
Loaded image: baas-transaction-manager-worker:10.1.6
Loaded image: baas-transaction-manager:10.1.6
Loaded image: baas-datamover:10.1.6
Loaded image: baas-spp-agent:10.1.6

Pushing Docker images to image registry at 10.0.240.150:5000...
Finished pushing docker images to image registry at 10.0.240.150:5000.

Deploying product release >>baas<< version 10.1.6 to target cluster
>>kubernetes-admin@kubernetes<<
NAME:   baas
LAST DEPLOYED: Wed Jul 29 12:14:25 2020
```

```
NAMESPACE: baas
STATUS: DEPLOYED

NOTES:

------------------------------------------------------------
Product: IBM Spectrum Protect Plus Kubernetes Backup Support
------------------------------------------------------------

Release baas was deployed in namespace baas.

Waiting for components to start up...
------------------------------------------------------------------------
waiting... 5:00 minutes
Pod baas-controller-7ccd77b756-tpl4r started.
Pod baas-kafka-5d48874f7b-vqmzc started.
Pod baas-scheduler-6fcfdc9656-nwnmk started.
Pod baas-spp-agent-96444fc75-xslq4 started.
Pod baas-transaction-manager-6594674598-gzpnt started.
------------------------------------------------------------------------
kubectl get pods -l app.kubernetes.io/name=baas -n baas
NAME                                        READY   STATUS    RESTARTS   AGE
baas-controller-7ccd77b756-tpl4r            0/1     Running   0          6s
baas-kafka-5d48874f7b-vqmzc                 0/1     Running   0          6s
baas-scheduler-6fcfdc9656-nwnmk             0/1     Running   1          6s
baas-spp-agent-96444fc75-xslq4              0/1     Running   0          6s
baas-transaction-manager-6594674598-gzpnt   0/3     Running   1          6s
------------------------------------------------------------------------
All pods are running.
All resources are installed successfully.
Installation is completed.
Product release >>baas<< version 10.1.6 has been installed in namespace >>baas<<
at Wed Jul 29 12:15:02 CEST 2020.
Script baas_install.sh finished at Wed Jul 29 12:15:02 CEST 2020.
```

During the installation process, the Kubernetes Backup Support automatically registers to the IBM Spectrum Protect master server. You can validate this registration by accessing the GUI and selecting **Manage Protection** → **Containers** → **Kubernetes**.

The cluster appears with the name that you specified in the `baas_config.cfg` file, as shown in Figure 14-4.

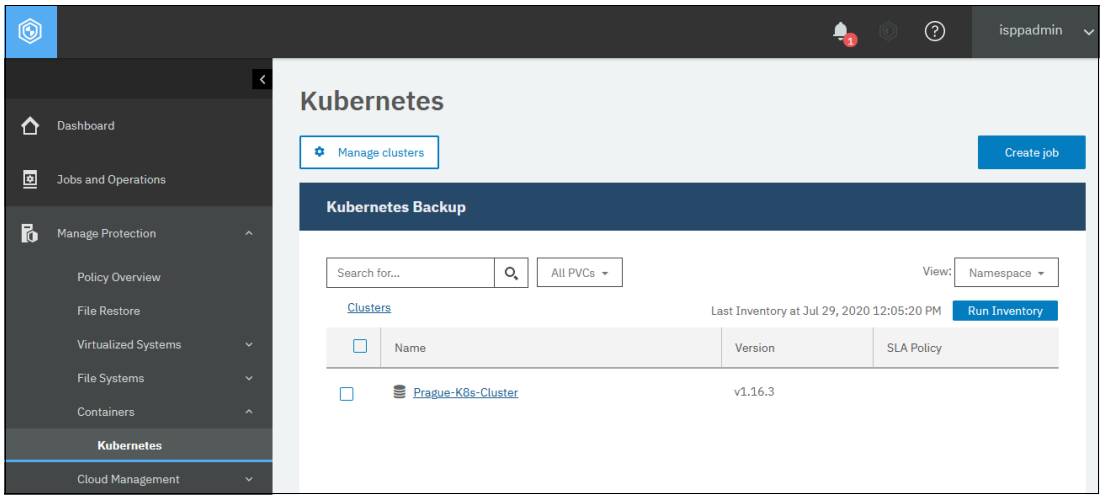

Figure 14-4 List Kubernetes cluster in GUI

If an error occurs within the automatic registration, the Kubernetes cluster can be registered manually by using the IBM Spectrum Protect Plus application. This process is performed from within the IBM Spectrum Protect GUI by selecting **Manage Protection** → **Containers** → **Kubernetes** → **Manage Clusters**. It requires the user to enter the following information:

- The cluster name: This name must match the value in the `baas_config.cfg` file that is used for installing the Kubernetes side IBM Spectrum Protect Plus components. Although it can be an IP, the use of a meaningful name is recommended.
- Host address: This address is the address of the Kubernetes master service. In production environments with multiple masters, this address is represented by a service or load balancer.
- Port number: This port is used to connect to the IBM Spectrum Protect Plus agent within Kubernetes. It can be obtained by running the **`kubectl get service -n baas | grep baas-spp-agent`** command.
- User ID: This ID is not a traditional user name; rather, it is generated by the IBM Spectrum Protect Plus Kubernetes installer and stored as a secret within Kubernetes. It can be obtained by running the following command:

  ```
  echo "`kubectl get secret baas-secret -n baas -o yaml | /bin/grep datamoveruser
  | cut -d: -f2 | tr -d ' ' | base64 -d`"
  ```

  Because this ID appears cryptic, an identity must be created for it by using a meaningful name. This process can be done in the GUI by selecting **Accounts** → **Identity**.
- Password: As with the user ID, this information is generated by the installer within Kubernetes. It can be obtained by running the following command:

  ```
  echo "`kubectl get secret baas-secret -n baas -o yaml | /bin/grep
  datamoverpassword | cut -d: -f2 | tr -d ' ' | base64 -d`"
  ```

After defining the cluster, use the test function (as shown in Figure 14-5) followed by the inventory function to ensure functionality.

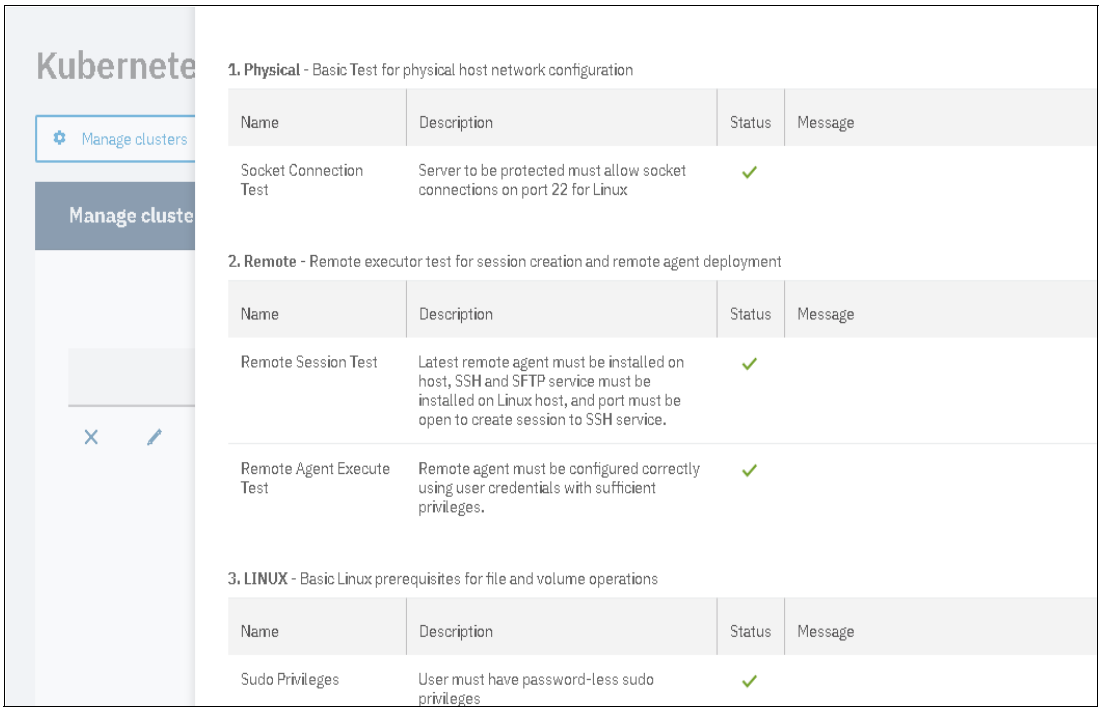

Figure 14-5 Testing the connection to the Kubernetes cluster

You can now browse through the namespaces and see individual persistent volumes, as shown in Figure 14-6.

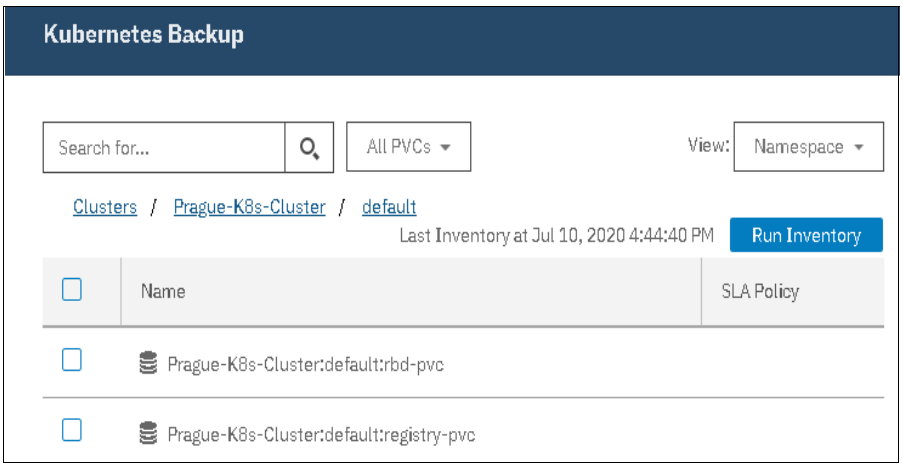

Figure 14-6 Listing persistent volumes

14.4 Protecting data

With the IBM Spectrum Protect Plus components for Kubernetes Backup Support successfully installed and configured, the data protection operations are set up next.

14.4.1 Defining SLAs for Kubernetes Backup Support

As for all other workloads, IBM Spectrum Protect Plus uses SLAs as the foundation for managing data protection for containers. With other workload types, IBM Spectrum Protect Plus provides traditional backup style operations where data moves from the production environment to the vSnap server. With Kubernetes, IBM Spectrum Protect Plus offers two levels of protection. Because the product uses CSI snapshots to create a point-in-time copy of data in a volume, it can also retain those snapshots on the production storage and revert to them for quick recovery.

In addition, store backups within IBM Spectrum Protect Plus vSnap server repositories for longer-term and lower-cost protection. Having an extra copy in a vSnap server also protects against primary storage issues that affect production volumes and snapshots.

Figure 14-7 shows an example SLA for Kubernetes Backup Support. It creates CSI snapshots every 6 hours and retains them for one day. It mounts the latest snapshot to a data mover once per day and copies it to a vSnap in the “Primary” site. The retention of data in the vSnap server is three days in this example.

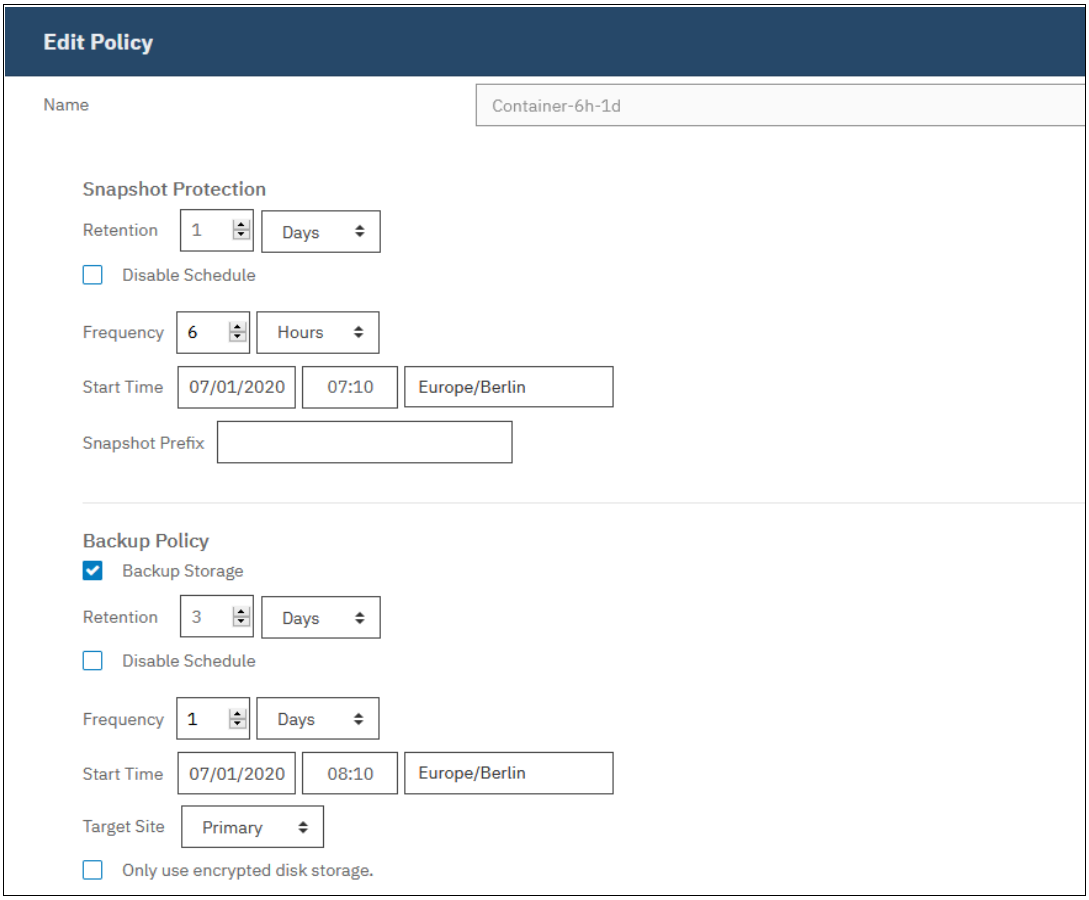

Figure 14-7 Example SLA for protection of persistent volumes

In the SLA Policy overview that is shown in Figure 14-8, the frequency for snapshot creation and for vSnap server copy are summarized.

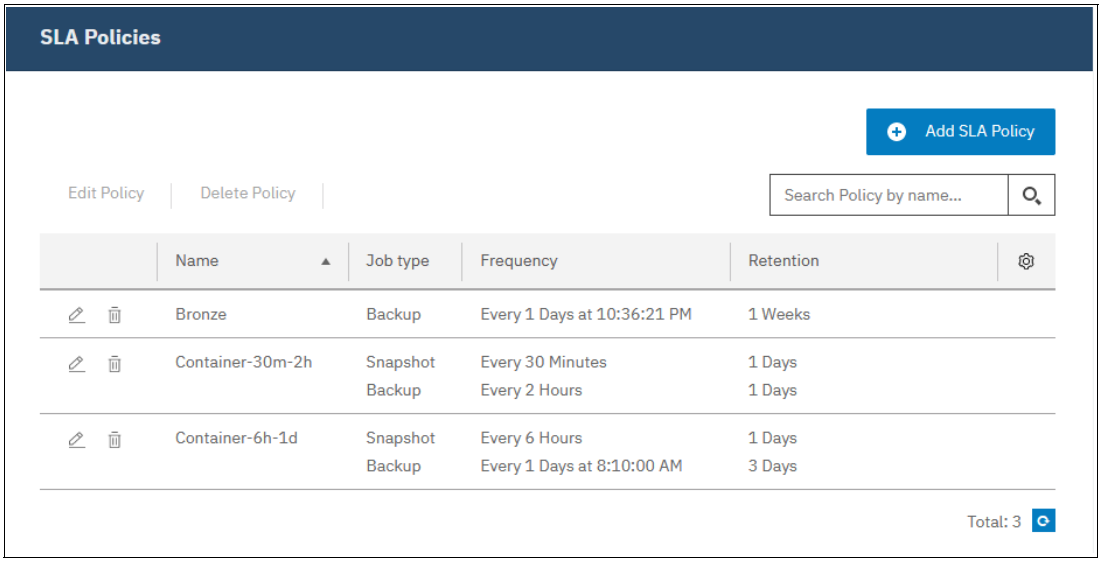

Figure 14-8 SLA Policy overview

14.4.2 Assigning SLAs to protect persistent volumes

IBM Spectrum Protect Plus always included a self-service focus with strong RBAC capabilities, allows workload administrators to leverage the easy to use web UI to perform data protection operations. With Kubernetes, this capability extends to the native tools as well. Native tools, such as kubectl and YAML files, manage constructs, such as SLAs and backup schedules, by using custom resource definitions.

Defining an SLA is only part of the picture. Within IBM Spectrum Protect Plus, an administrator must discover data and bind it to an SLA for data protection operations, such as snapshots and backup copies to occur. In the case of Kubernetes, the following methods are available for declaring which persistent volumes to protect:

- Everything in the cluster: As the name suggests, this option protects all volumes. However, use it with caution in cases where multiple storage classes are in use and some of them cannot perform CSI snapshots.
- Namespace level backups: This option backs up all persistent volumes within a specific namespace by using the indicated SLA.
- Label-based backups: This method backs up all persistent volume claims with a specified label that is attached to them. The label is on the claim rather than the pod or deployment.
- PVC names: This method allows the user to specify a static list of PVCs to protect.

Namespace boundaries are a critical consideration when configuring data for backup. Figure 14-9 shows a data mover container that exists within the individual namespaces.

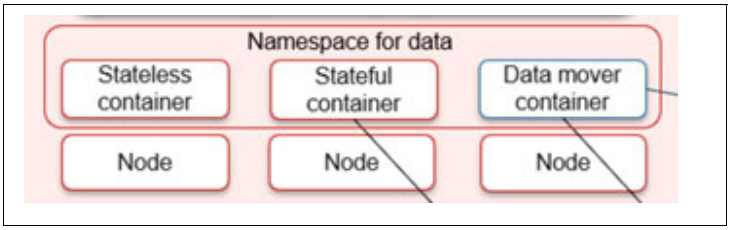

Figure 14-9 IBM Spectrum Protect and containers interaction

Note: During the initial installation instructions, the installation instructions specify to register a pull secret in each namespace that requires protection so that Kubernetes can pull the image from the registry. When developers or DevOps engineers create namespaces, they must complete this procedure for the new namespaces as well or backups fail.

The following methods are available to assign SLAs to persistent data:

- By using the GUI
- By using the native interface

These methods are described next.

Assigning SLAs by using the GUI

After a Kubernetes cluster is registered in IBM Spectrum Protect Plus and the Application Inventory job finishes, you can use the GUI by selecting **Manage Protection** → **Containers** → **Kubernetes** and browse through the discovered cluster resources, which can be viewed by Namespace or by Label.

Select one or more volumes and assign an SLA in the same way as you assign an SLA to any other workload, such as a VM or application.

Figure 14-10 shows a PVC `demo-pvc` in namespace `demo`, which was assigned to SLA `Container-6h-1d`.

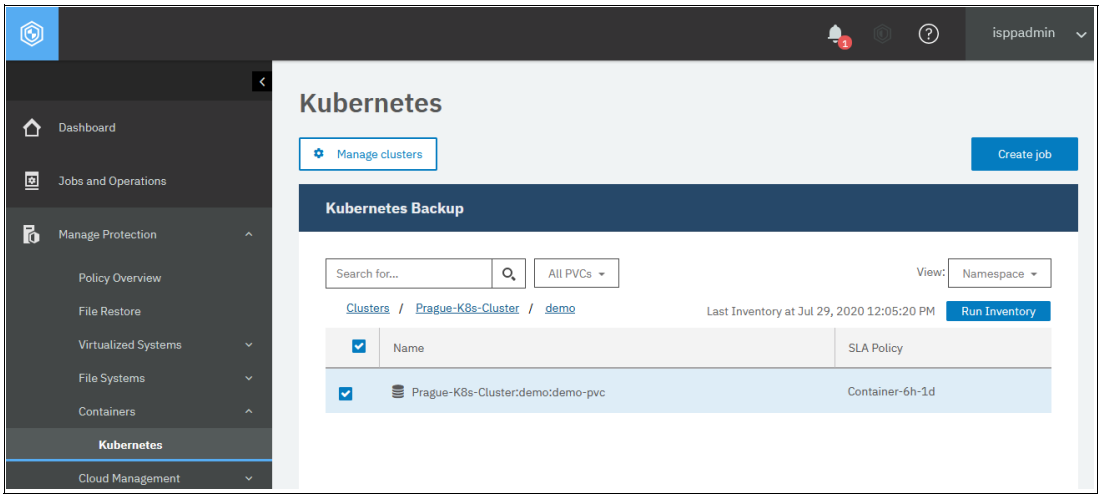

Figure 14-10 Volume assigned to a SLA

Assigning SLAs by using the native interface

IBM Spectrum Protect Plus always featured a self-service focus with strong RBAC capabilities that allowed workload administrators to leverage the easy to use web UI to perform data protection operations. Instead of the use of the UI, the native Kubernetes command line tool kubectl and YAML files can be used to achieve the same goal.

Example 14-3 shows how a Kubernetes administrator can assign SLA `Container-6h-1d` to PVC `demo-pvc` by using kubectl and a YAML file.

Example 14-3 Assign a SLA to a PVC using kubectl

```
[root@prague-k8-master spp]# more demo-pvc-backup.yaml

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq
metadata:
  name: demo-pvc
  namespace: demo
spec:
  requesttype: backup
  sla: Container-6h-1d
  encryption: no
  volumesnapshotclass: csi-rbdplugin-snapclass
[root@prague-k8-master spp]# kubectl apply -f demo-pvc-backup.yaml
```

Assigning SLAs to containerized applications

As described in 14.1.2, “Applications and containers” on page 385, modern software development uses the approach of building large systems or applications by realizing different functions through a set of multiple, interconnected containers instead of writing a single, monolithic piece of code.

Because many of these containers might rely on persistent data that is stored in different persistent volumes, it makes sense to assign all of these volumes to the same SLA to ensure that data is protected at the same point-in-time and that it is retained for the same amount of days.

Labels: In addition to assigning SLAs to single PVCs, IBM Spectrum Protect Plus can assign a SLA to a set of PVCs based on a label.

Labels can be assigned to volumes by the Kubernetes administrator or application developer. This approach ensures that, for example, all volumes that belong to a specific application are protected by the same SLA, without the need to determine and manually assign the SLA to all individual volumes.

In the GUI, labels can be selected by clicking **Manage Protection** → **Containers** → **Kubernetes** → **View** → **Labels**.

In the native command line, backup requests are created by applying a YAML file. In this example, we must change the request type from Backup to BackupLabel or alternatively to BackupNamespace.

After the SLA is assigned, the backups are automatically scheduled or can be triggered by manually starting the corresponding SLA.

Example 14-4 shows how to determine which backups exist for a specific PVC.

Example 14-4 Show existing backups for a PVC

```
[root@prague-k8-master ~]# kubectl describe baasreq demo-pvc -n demo
Name:          demo-pvc
Namespace:     demo
Labels:        <none>
API Version:   baas.io/v1alpha1
Backupstatus:  Ready
Kind:          BaaSReq
Metadata:
  Creation Timestamp:  2020-07-01T14:53:02Z
  Generation:          191
  Resource Version:    96604301
  Self Link:           /apis/baas.io/v1alpha1/namespaces/demo/baasreqs/demo-pvc
  UID:                 7aae5504-569d-4984-9a31-b449fa187e4c
Spec:
  Inprogress:   None
  Instanceid:   ed642c72eba28ecb46be77532aba3536
  Origreqtype:  backup
  Requesttype:  backup
  Size:         1073741824
  Sla:
    Container-6h-1d
  Spppvcname:           Prague-K8s-Cluster:demo:demo-pvc
  Volumesnapshotclass:  csi-rbdplugin-snapclass
Status:
  Snapshotname:  2101.snapshot.26
  Timestamp:     2020-07-29 06:23:41
  Type:          COPY
  Snapshotname:  spp-1004-2106-17398fb76c5
  Timestamp:     2020-07-29 05:23:34
  Type:          FAST
  Snapshotname:  spp-1004-2106-17397b1e16d
  Timestamp:     2020-07-28 23:23:35
  Type:          FAST
  Snapshotname:  spp-1004-2106-17396684aa9
  Timestamp:     2020-07-28 17:23:35
  Type:          FAST
  Snapshotname:  spp-1004-2106-173951eb445
  Timestamp:     2020-07-28 11:23:36
  Type:          FAST
  Snapshotname:  2101.snapshot.25
  Timestamp:     2020-07-28 06:23:43
  Type:          COPY
  Snapshotname:  2101.snapshot.24
  Timestamp:     2020-07-27 06:23:45
  Type:          COPY
  Snapshotname:  2101.snapshot.23
  Timestamp:     2020-07-26 06:23:46
  Type:          COPY
Volumename:      demo-pvc
Events:          <none>
```

```
[root@prague-k8-master ~]# kubectl get volumesnapshot -n demo
NAME                          AGE
spp-1004-2106-173951eb445     22h
spp-1004-2106-17396684aa9     16h
spp-1004-2106-17397b1e16d     10h
spp-1004-2106-17398fb76c5     4h13m
```

14.5 Restoring data

Like backups, restores can be started by using the UI or the native command line.

Restoring data by using the GUI

Restores can be started through the GUI by selecting **Create Job** → **Restore** → **Kubernetes**, and selecting a snapshot or copy backup image by using a specific timestamp and finally by specifying a new name for the restored volume. After this process is complete, you can review the summary window and start the restore (see Figure 14-11).

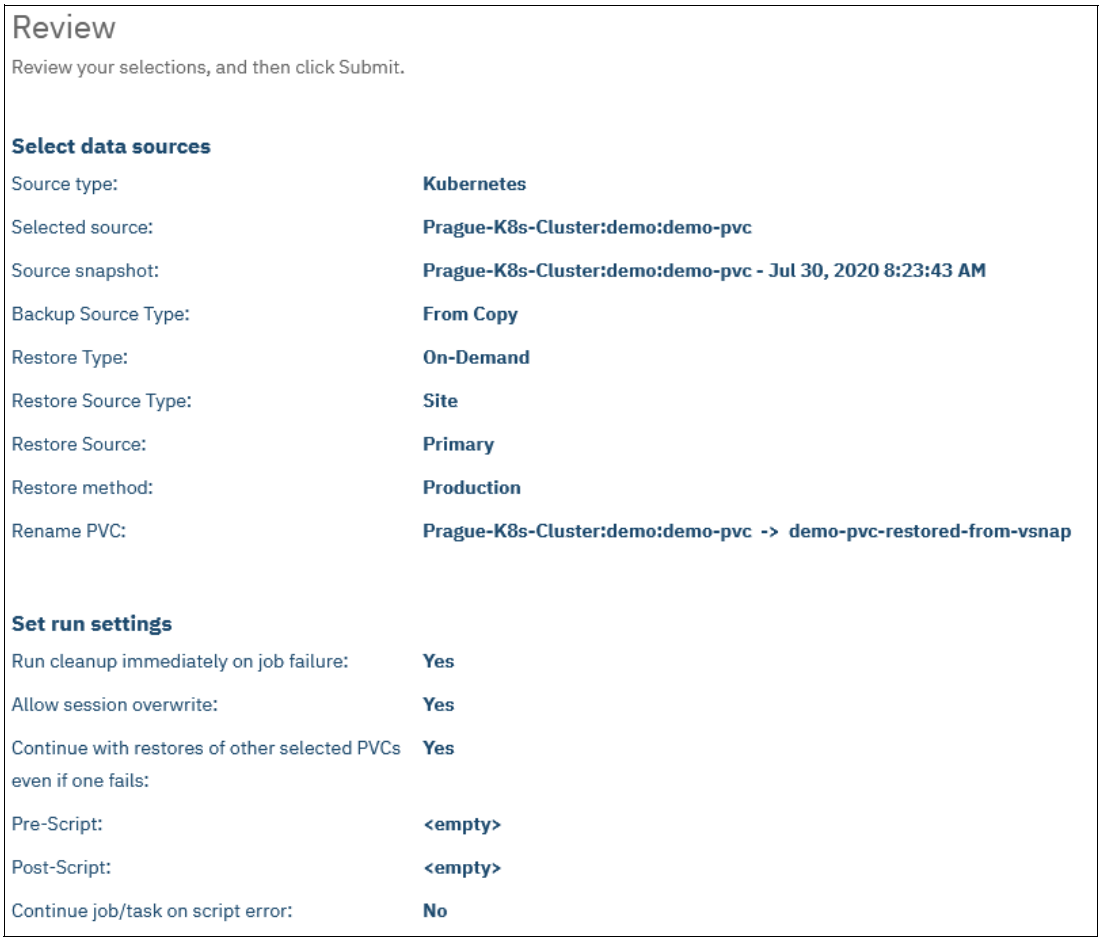

Figure 14-11 Restore PVC backup from vSnap server copy to new volume

In the example that is shown in Figure 14-11 on page 403, a vSnap server copy of volume `demo-pvc` is restored to a new volume, which is named `demo-pvc-restored-from-vsna`. After the restore is started, a new volume is provisioned, a temporary data mover is started in the corresponding namespace `demo`, and the contents of the backup image are returned to the new PVC.

Example 14-5 shows the new target volume and the temporary data mover.

Example 14-5 Show restore target volume and temporary data mover

```
[root@prague-k8-master ~]# kubectl get pvc -n demo
NAME                          STATUS   VOLUME                                     CAPACITY   ACCESS MODES
demo-pvc                      Bound    pvc-1f122c61-6396-4b51-9cc7-862cf5770522   1Gi        RWO
demo-pvc-restored-from-vsnap  Bound    pvc-86eaf822-a1df-4206-948f-6f8e1257039a   1Gi        RWO

[root@prague-k8-master ~]# kubectl get pods -n demo
NAME                                                              READY   STATUS    RESTARTS   AGE
demo-pod                                                          1/1     Running   1          166d
restore-demo-pvc-ondemandrestore-1596101549124-15961015491nczxg   1/1     Running   0          44s
```

When the restore is completed, the Kubernetes administrator can decide how to continue; for example, mount the new volume to a pod or container.

For more information about how to restore data by using the UI, see IBM Knowledge Center.

Restoring data by using the native interface

As for doing backups by using the kubectl command line interface, restores are also performed by using YAML files. Example 14-6 shows the various parameters that must be specified for the restore.

Example 14-6 YAML file for restoring a PVC

```
#------------------------
# Filename: filename.yaml
#------------------------

apiVersion: "baas.io/v1alpha1"
kind: BaaSReq

metadata:
  name: name_of_restore_request
  namespace: namespace
spec:
  requesttype: restore
  pvcname: pvc_name
  targetvolume: target_volume_for_restore
  storageclass: storage_class_of_target_volume
  restorepoint: timestamp_of_backup
  restoretype: fast | copy
```

In addition to the original and target name of the PVC, you also must specify the timestamp of the backup image that you want to restore. A list of available backup images can be determined by using the kubectl describe **baasreq** command, as shown in Example 14-4 on page 402.

For more information about how to restore data by using the native command line, see IBM Knowledge Center.

15

Replication and additional copies

This chapter discusses different approaches to increase the level of protection of backup data by replicating data or creating additional copies.

Such additional copies can be established in the following ways:

- Replication of backup-storage data from one vSnap server to another vSnap server
- Additional Copies to secondary backup storage:
  - To standard Object Storage:
    - Incremental copy to Object Storage
    - Incremental copy to an IBM Spectrum Protect cloud-container storage pool
  - Archive Object Storage or tape:
    - Full copy to a Long-Term Object Storage (for example, IBM Cloud Object Storage Archive or Amazon Glacier)
    - Full copy to a Long-Term IBM Spectrum Protect directory container for tapes storage pool
- Backing up data to multiple vSnap servers by using multiple SLAs

Important: None of the techniques that are discussed in this chapter are moving or migrating data from one storage tier to another. All methods that are described are establishing extra copies. However, different lifecycles can be applied to each copy, which keeps “hot data” on a fast storage while placing “cold data” on a slower, low-cost storage.

This chapter includes the following topics:

15.1 Reasons to create more copies of backup data

In IBM Spectrum Protect Plus, vSnap servers are the primary target for backup data. A vSnap is a pool of storage that receives data from production systems for the purposes of data protection and reuse. The vSnap server consists of one or more disks and can be scaled up (by adding disks to increase capacity) or scaled out (by distributing the workload to multiple vSnap servers to increase overall performance).

Redundant Array of Independent Disks (RAID) technology can be used to protect vSnap servers against data loss that is caused by a hardware failure of a single disk - either by using software RAID inside the vSnap or by using hardware or software features of the underlying storage system that provides the capacity to the vSnap pool.

However, there might be reasons to create additional copies of backup data, for instance, to be able to recover from a disaster or to store long-term data on lower-cost storage.

15.2 Extra copies: Overview and options comparison

This section presents and compares various options for creating additional copies.

15.2.1 Replication of backup data

Backup data can be replicated between vSnap servers. This replication includes backups of VMs, applications, and IBM Spectrum Protect Plus server catalog backups.

Replication of data between vSnap servers can be done for various reasons, including the following examples:

- Protecting the backup data against the complete loss of a vSnap system.
- Protecting the backup data against the loss of a whole data center (assuming that the replication target is in a remote data center).
- Storing a copy of data in a remote site with a longer retention when not enough space exists in the primary site; for example, replicating data from small branch offices to a large central data center.
- Migration of a vSnap system.

Note: As an option, additional copies can be created from the replication target vSnap. An example includes replicating backup data daily from a vSnap in a small branch office to a remote vSnap in a central data center and from there, creating a full copy to an Object Storage or to a IBM Spectrum Protect server tape pool monthly.

15.2.2 Additional copies to Object Storage

Backup data that is in a vSnap server can be copied to an Object Storage provider. At the time of this writing, the following Object Storage providers are supported:

- Amazon S3
- IBM Cloud Object Storage
- Microsoft Azure Blob Storage
- S3 Compatible Object Storage

Data can be copied to an Object Storage provider in two ways:

- Incremental copy to a “hot” Object Storage tier. This approach stores backup data with longer retentions at a moderate cost and supports instant access recovery.
- Full copy/archive to a “cold” Object Storage tier. This approach stores long-term data at the lowest cost and is intended to be used for monthly, quarterly, or annual full backups. The backup data is bundled in a tar file and uploaded as a single object. With this two-step backup approach, the restore operation also is a two-step process, which does not allow for instant access recovery and is slower than the incremental copy.

15.2.3 Additional copies to a repository server

Backup data that is in a vSnap server can be copied to a repository server. At the time of this writing, the only supported repository server type is IBM Spectrum Protect.

Data can be copied to an IBM Spectrum Protect server in two ways:

- Incremental copy to a “hot” storage tier. This approach stores backup data with longer retentions to an IBM Spectrum Protect container storage pool (directory container or cloud container) and supports instant access recovery.
- Full copy/archive to a “cold” storage tier. This approach stores long-term data at lower cost into an IBM Spectrum Protect Tape storage pool (directory container for Tapes). The backup data is bundled in a tar file and uploaded as a single object. In this case, the restore operation also is a two-step process, which does not allow for instant access recovery and is slower than the incremental copy.

For more information about how to configure IBM Spectrum Protect Plus for the creation of additional full and incremental copies to an IBM Spectrum Protect server, see 15.6, “Creating incremental and full copies of backup data to an IBM Spectrum Protect server” on page 448.

15.2.4 Dual-site backup using multiple SLAs

All methods that were described in this chapter have one thing in common: Backup objects that are ingested into a primary vSnap server are *copied* over to another location (for example, a secondary vSnap, an Object Storage, or an IBM Spectrum Protect server). The additional copy belongs to the same point-in-time as the primary backup object.

By assigning a virtual machine (VM) or an application to more than one SLA backups can be created of the same entity in two different locations (sites) and at different points in time. However, in this case, it is no longer a copy because additional vSnaps are used for more frequent backups. The key difference is the fact that the primary data is being backed up twice, which results in two independent “primary” backup versions of the same VM or application.

Important: Increasing the backup frequency adds load to the hypervisor (for example, VMware snapshots must be created more often) or to the application (for example, because it is set to back up mode more often). Because of network latency, additional backups to a distant site might experience slower performance and cannot benefit from compression because it is being used when replicating backup data between two vSnap servers.

15.2.5 Comparing the options

Creating additional copies of backup data can be established in various ways. The choice of a particular method and type of storage, is driven by answers to the following questions or requirements:

- How fast can the copy be accessed in case of a restore request?
- Is an instant access recovery possible?
- Is the target storage appropriate for long-term archiving?
- Is the target storage physically disconnected from the production environment and protected against malicious attacks, for example, Computer viruses (“Air-Gap”)?
- How much cost is assigned with the implementation of the secondary storage?

The following table summarizes the different options and their characteristics:

Table 15-1 Comparison of different options to create additional copies of backup data

Copy Option	Access to data	Long-term Readiness	Air Gap	Cost	Comment
vSnap Replication	Fast	No	No	High	Incremental copy to fast storage, compressed (opt.)
incremental copy to Object Storage	Medium	Medium	Medium	Moderate	Incremental copy to Object Storage
incremental copy to IBM Spectrum Protect	Medium	Medium	Medium	Moderate to high	Only with IBM Spectrum Protect container storage pools
full copy to Object Storage (Archive)	Slow	Yes	Medium	Low	Object Storage archive storage, no instant access recovery
full copy to IBM Spectrum Protect (Archive)	Slow	Yes	Yes	Moderate	Tape, no instant access recovery
Dual-Site Backup	Fast	No	No	High	Two initial full backups required

15.3 Replicating backup data

The backup data that is stored by a vSnap server can be replicated to a second vSnap server as an operational protection. This replication applies to backups of VMs, databases, applications, and IBM Spectrum Protect Plus server catalog backups.

The replication target (receiving) vSnap server must belong to a different site than the source (sending) vSnap server. For more information about the relationship and options of sites and vSnap servers for replication, see the chapter *Replication considerations* in the IBM Spectrum Protect Plus Blueprints.

15.3.1 Configuring vSnap replication

To replicate data from a primary vSnap server to a secondary vSnap server, both vSnap servers must be configured on the IBM Spectrum Protect Plus server and belong to different sites, as described in 15.2.4, “Dual-site backup using multiple SLAs” on page 408.

In addition, the following configuration tasks must be completed to enable replication between vSnap servers:

- A replication partnership is configured between all vSnap servers.
- The replication policy is configured as part of the SLA.

Important: If your environment includes a mixture of encrypted and decrypted vSnap servers, select the **Only use encrypted disk storage** option to replicate data to encrypted vSnap servers. If this option is selected and only decrypted vSnap servers are available, the associated replication job fails.

vSnap replication partnership

To configure a replication partnership between two vSnap servers from the IBM Spectrum Protect Plus GUI, complete the following steps:

1. Select **System Configuration** → **Backup Storage** → **Disk** and then, click the **Edit** icon vSnap server to enter properties page.

   Figure 15-1 shows the properties page of a vSnap server in the IBM Spectrum Protect Plus GUI. The Configure Storage Partners section does not list any partner yet.

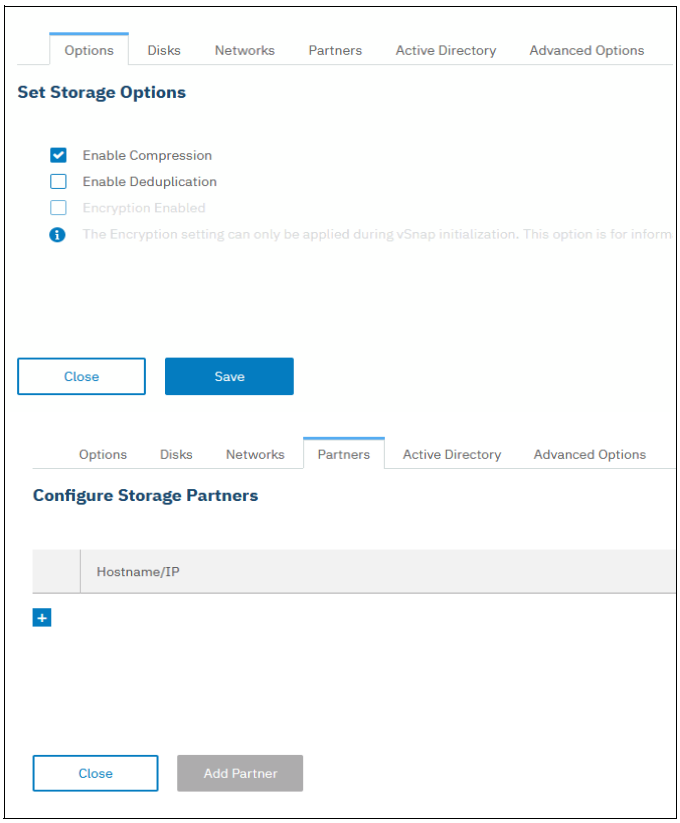

Figure 15-1 vSnap server properties page

Note: The vSnap CLI command **`vsnap partner show`** also does not list any partner now.

2. Click the plus sign that is below the Configure Storage Partners section and choose the partner server from the pull-down menu. Now, click **Add Partner** to establish the partnership, as shown in Figure 15-2.

Figure 15-2 Add vSnap server as storage partner in the vSnap server properties page

After the partner is added, it is listed as a storage partner with information about the site that contains the partner vSnap server, as shown in Figure 15-3.

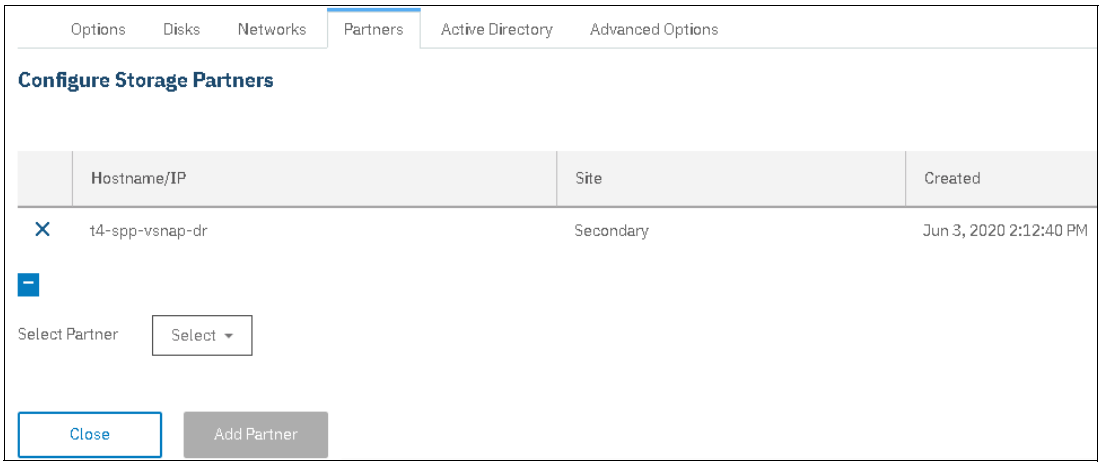

Figure 15-3 vSnap Partner server added

In Example 15-1, we run the vSnap CLI command **`vsnap partner show`** to list the vSnap partner server.

Example 15-1 List vSnap partner server on the Primary site vSnap

```
[serveradmin@t4-spp-vsnap ~]$ vsnap partner show
ID                               | PARTNER TYPE | MGMT ADDRESS    | API PORT | SSH PORT
-------------------------------------------------------------------------------------
4b34983b741447a6a41a96a62525cf87 | vsnap        | t4-spp-vsnap-dr | 8900     | 22
```

3. Run the same command that is shown in Example 15-2 on the second vSnap server for the Secondary site, which lists also the partner vSnap server.

Example 15-2 List vSnap partner server on the Secondary site vSnap

```
[serveradmin@t4-spp-vsnap-dr home]$ vsnap partner show
```

```
ID                                | PARTNER TYPE | MGMT ADDRESS | API PORT | SSH PORT
----------------------------------------------------------------------------------------
9b44fb0c808b434688b7e8e70a6b391a | vsnap        | t4-spp-vsnap | 8900     | 22
```

Replication partnership between both vSnap servers is now configured and can be used by an SLA in both directions.

SLA replication policy

As a next step, choose the SLA that must be configured to use replication as Operational Protection. Complete the following steps:

1. From the IBM Spectrum Protect Plus GUI, select **Manage Protection** → **Policy Overview**. Click the **Edit** icon to enter the properties page of the wanted SLA where replication should be added. In our example, we use the Bronze SLA, as shown in Figure 15-4.

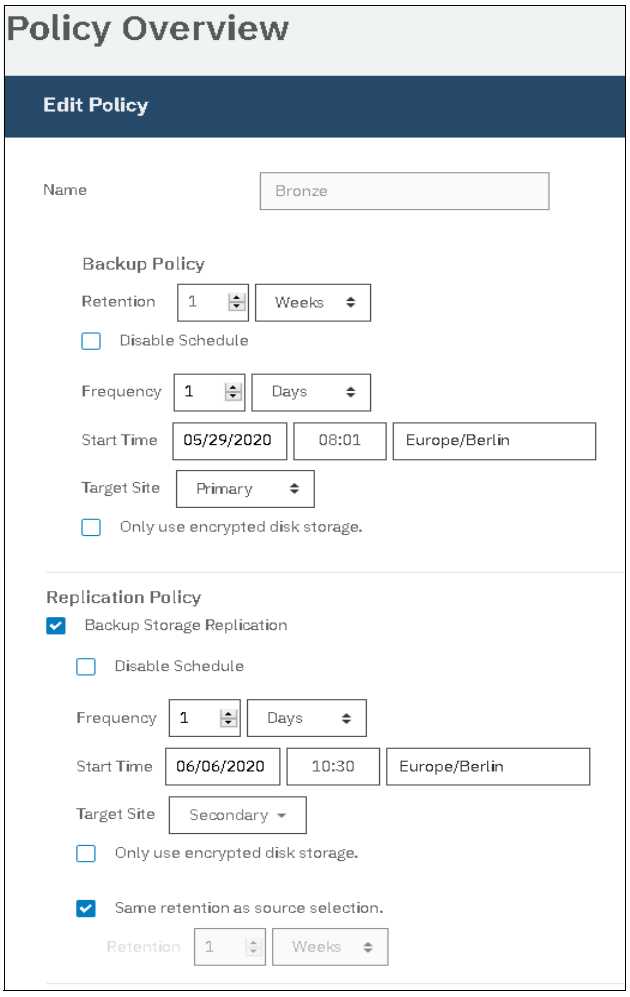

Figure 15-4 Bronze SLA properties page

2. The Replication Policy section is not enabled by default. Select the **Backup Storage Replication** option, adjust the **Start Time** as required and choose **Secondary** as the Target Site, as shown in Figure 15-4. Click **Save** to save your changes.

We now configured replication from the Primary site to the Secondary site within the Bronze SLA. From a system perspective, data is replicated from the vSnap server `t4-spp-vsnap`, which belongs to the Primary site to vSnap server `t4-spp-vsnap-dr`, which belongs to the Secondary site. The relationship from vSnap server to site also is shown by selecting **System Configuration** → **Backup Storage** → **Disk** in the GUI, as shown in Figure 15-5.

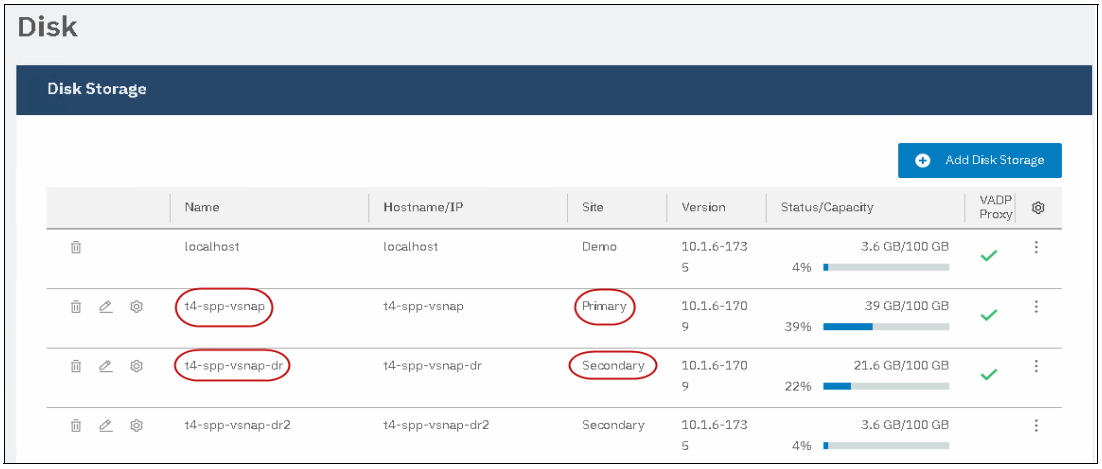

Figure 15-5 vSnap server overview

15.3.2 Running the vSnap replication

The replication process runs as part of the backup job that was created for a corresponding SLA. Because we added replication to the Bronze SLA that we use to backup VMware VMs, the backup job `vmware_spp_red_sql` is used. You can wait until replication starts at the configured time or start it manually by clicking **Jobs and Operations** → **Schedule** → **Actions Menu** icon ⋮ drop-down menu behind the Bronze Job. Then, select **Start** and then, **Replicate**, as shown in Figure 15-6 and Figure 15-7 on page 414.

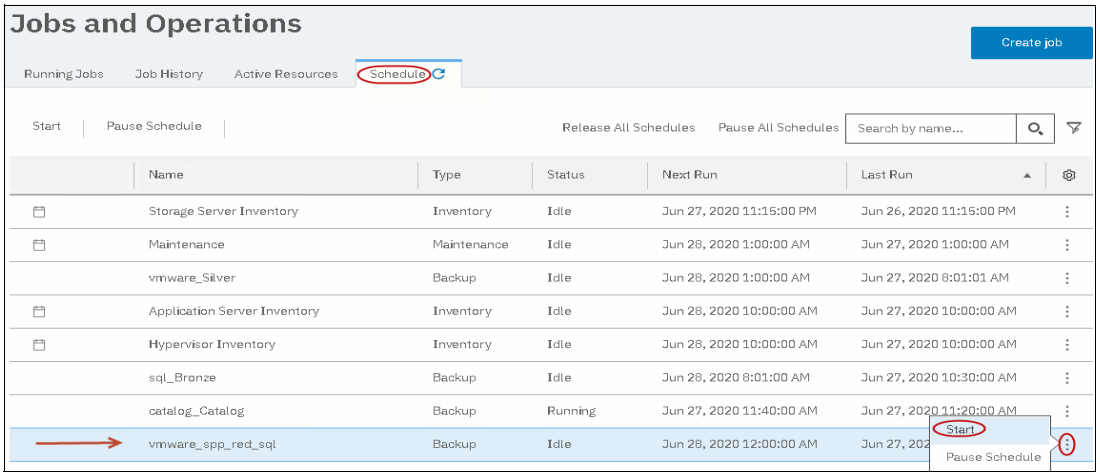

Figure 15-6 Select vmware job and Start one operation

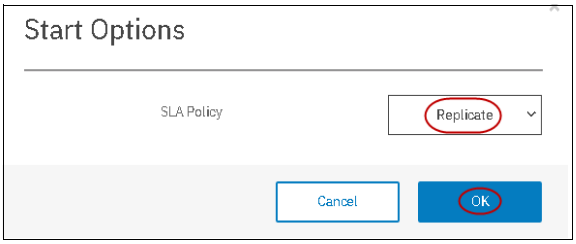

Figure 15-7 Start Option dialog

Note: If more than one vSnap server was selected as a target for replication, the replication process chooses the server with the most free space.

15.3.3 Determining space that is replicated on the Target vSnap

To determine how much space a replicated vSnap is using, click **System Configuration** → **Backup Storage** → **Disk** in the GUI, as shown in Figure 15-5 on page 413. The vSnap capacity is shown in the column Status/Capacity.

For more information about vSnap storage utilization, select **Reports and Logs** → **Reports** → **vSnap Storage Utilization Report** → ▶ as shown in Figure 15-8.

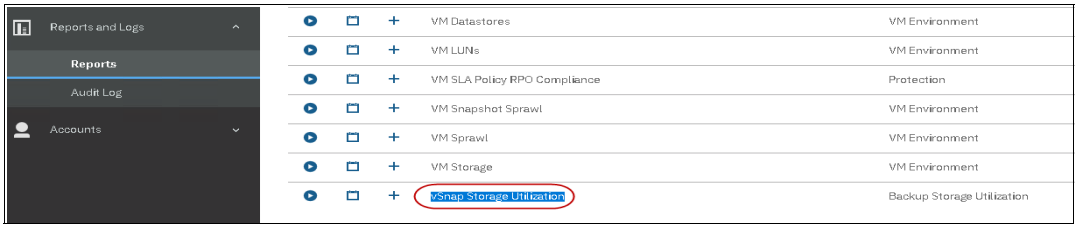

Figure 15-8 vSnap Storage Utilization Report generated

For example, you can generate one report to obtain vSnap Storage Utilization with Replica Destination Volumes with the **including** option, as shown in Figure 15-9 on page 415, and another report with Replica Destination Volumes with the **excluding** option, as shown in Figure 15-10 on page 415. Then, the results can be compared.

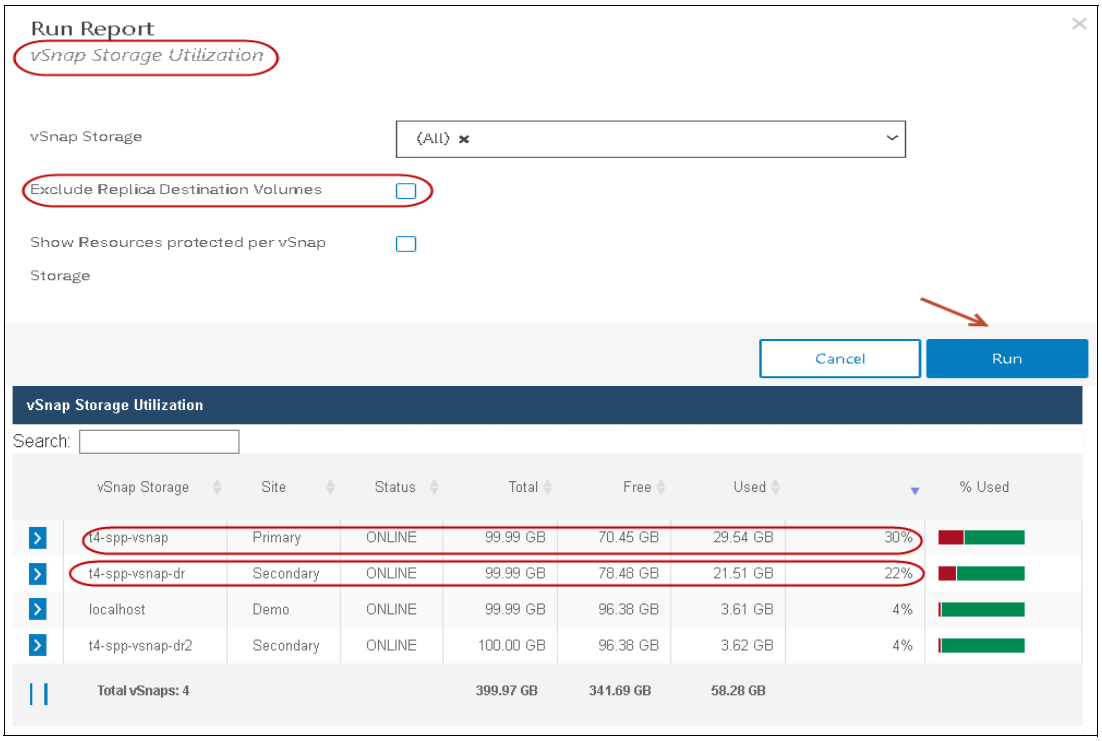

Figure 15-9 vSnap Storage Utilization with Replica Destionation Volumes including, vSnap-dr at 22%

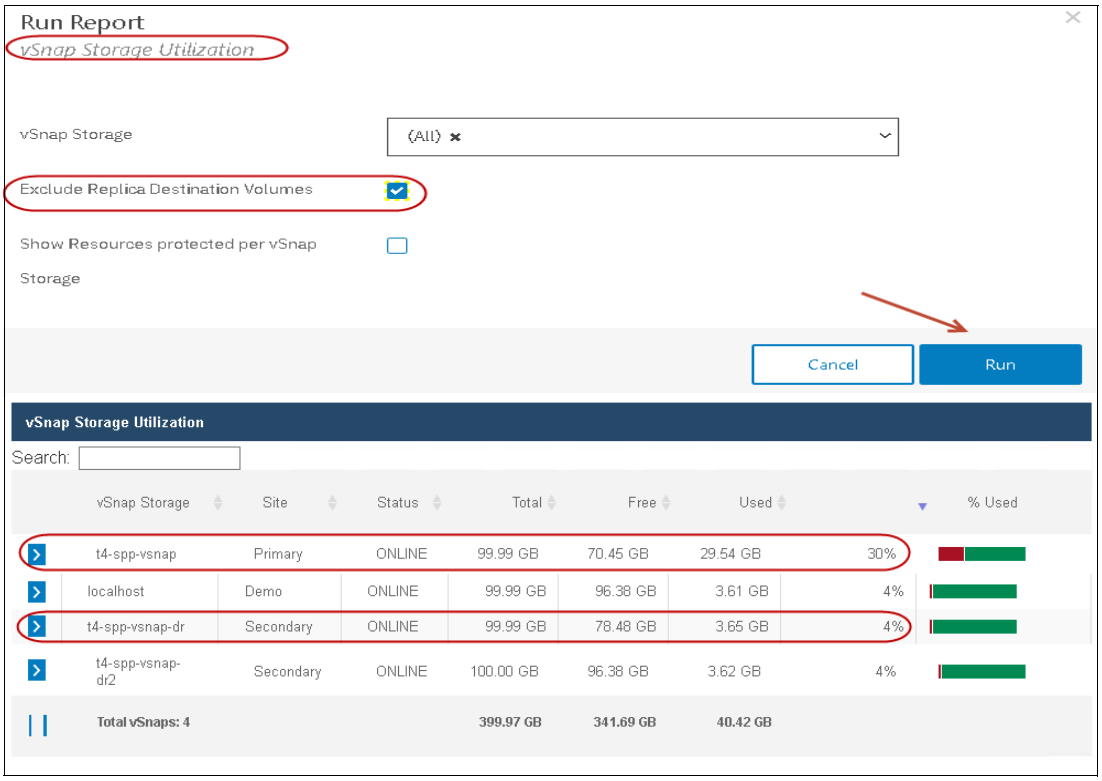

Figure 15-10 vSnap Storage Utilization with Replica Destination Volumes excluding, vSnap-dr at 4%

In this case, the result of the comparison is that space that is replicated on the Target vSnap is 18% of the total vSnap disk space.

15.3.4 vSnap commands for data replication

As shown in Example 15-1, a vSnap partnership is created between `t4-spp-vsnap` and `t4-spp-vsnap-dr`. When starting a replication task of an SLA policy, the following actions are completed on the source and target vSnap servers:

1. A new replication target volume is created on the target vSnap server.
2. A relationship for the source and target volume is created on the source and target vSnap servers.
3. An empty snapshot on the target volume is created.
4. A replication session is created on the source vSnap server.
5. The target snapshot is updated on the target vSnap server after the replication session is completed.

Example 15-3 shows these steps in detail. The replication job creates a target volume (ID 3) on the target vSnap server `t4-spp-vsnap-dr`. The vSnap CLI command **`vsnap volume show`** is listing it.

Example 15-3 Replication target volume is created on the target vSnap server

```
[serveradmin@t4-spp-vsnap-dr home]$ vsnap volume show
ID | TYPE       | POOL | IS CLONE | TOTAL   | FREE    | USED     | NAME                                          | TAGS
---------------------------------------------------------------------------------------------------------------------
1  | filesystem | 1    | No       | 99.99GB | 96.34GB | 24.00KB  | vsnap_metadata_cloud                          | N/A
2  | lun        | 1    | No       | 2.00TB  | 2.00TB  | 14.30MB  | spp_1728a608d05_spp_1004_2002_1727A158CDC     | N/A
3  | filesystem | 1    | No       | 99.99GB | 96.34GB | 635.00KB | spp_1728a60a7e4_spp_1004_2002_1727A13274D     | N/A
```

In the next step, the replication relationship between both volumes is configured on both vSnap servers. The vSnap CLI command **`vsnap relationship show`** displays these relationships, as shown in Example 15-4, for the source vSnap server `t4-spp-vsnap` and in Example 15-5 for the target vSnap server `t4-spp-vsnap-dr`.

Example 15-4 vSnap relationship on the source server created

```
[vsnapadmin@t4-spp-vsnap T4-VM-SQL_ESCC]$ vsnap relationship show
ID                               | PARTNER ADDR    | PARTNER TYPE | LOCAL ROLE | LAST SYNC | LOCAL VOL                | REMOTE VOL
-------------------------------------------------------------------------------------------------------------------------------------------------------------------
--
df31519488e093914295b620d6cf3254 | t4-spp-vsnap-dr | vsnap        | primary    | COMPLETED | spp_1004_2002_1727A158CDC |
spp_1728a608d05_spp_1004_2002_1727A158CDC
8786a0e1a7b0af626eab9a76d714705b | t4-spp-vsnap-dr | vsnap        | primary    | COMPLETED | spp_1004_2002_1727A13274D |
spp_1728a60a7e4_spp_1004_2002_1727A13274D

[vsnapadmin@t4-spp-vsnap T4-VM-SQL_ESCC]$ vsnap relationship show --id df31519488e093914295b620d6cf3254

ID: df31519488e093914295b620d6cf3254
PARTNER ID: 4b34983b741447a6a41a96a62525cf87
PARTNER TYPE: vsnap
PARTNER ADDR: t4-spp-vsnap-dr
LOCAL ROLE: primary
LOCAL POOL ID: 1
LOCAL VOL ID: 3
LOCAL VOL NAME: spp_1004_2002_1727A158CDC
REMOTE POOL ID: 1
REMOTE VOL ID: 2
REMOTE VOL NAME: spp_1728a608d05_spp_1004_2002_1727A158CDC
LAST SYNC STATUS: COMPLETED
LAST SYCNED SNAP ID: 57
LAST ATTEMPT SNAP ID: 57
CREATED: 2020-06-06 16:04:45 UTC
UPDATED: 2020-06-06 16:04:45 UTC
```

Example 15-5 vSnap relationship on the target server created

```
[serveradmin@t4-spp-vsnap-dr home]$ vsnap relationship show
ID                               | PARTNER ADDR | PARTNER TYPE | LOCAL ROLE | LAST SYNC | LOCAL VOL                                | REMOTE VOL
-----------------------------------------------------------------------------------------------------------------------------------------------------------
df31519488e093914295b620d6cf3254 | t4-spp-vsnap | vsnap        | replica    | COMPLETED | spp_1728a608d05_spp_1004_2002_1727A158CDC | spp_1004_2002_1727A158CDC
8786a0e1a7b0af626eab9a76d714705b | t4-spp-vsnap | vsnap        | replica    | COMPLETED | spp_1728a60a7e4_spp_1004_2002_1727A13274D | spp_1004_2002_1727A13274D

[serveradmin@t4-spp-vsnap-dr home]$ vsnap relationship show --id df31519488e093914295b620d6cf3254
```

```
ID: df31519488e093914295b620d6cf3254
PARTNER ID: 9b44fb0c808b434688b7e8e70a6b391a
PARTNER TYPE: vsnap
PARTNER ADDR: t4-spp-vsnap
LOCAL ROLE: replica
LOCAL POOL ID: 1
LOCAL VOL ID: 2
LOCAL VOL NAME: spp_1728a608d05_spp_1004_2002_1727A158CDC
REMOTE POOL ID: 1
REMOTE VOL ID: 3
REMOTE VOL NAME: spp_1004_2002_1727A158CDC
LAST SYNC STATUS: COMPLETED
LAST SYCNED SNAP ID: 56
LAST ATTEMPT SNAP ID: N/A
CREATED: 2020-06-06 16:04:46 UTC
UPDATED: 2020-06-06 16:04:46 UTC
```

The parameters Partner Address, Local Volume ID, Local Volume Name, Remote Volume ID and Remote Volume Name are used to identify which vSnap servers and volumes are used for this relationship.

A snapshot is created by running the **`vsnap snapshot show`** command, as shown in Example 15-6.

Example 15-6 New snapshot created

```
[vsnapadmin@t4-spp-vsnap fs2]$ vsnap snapshot show
ID | PARENT ID | CREATED                | NAME                         | VERSION ID
-------------------------------------------------------------------------------------
64 | 6         | 2020-06-20 09:17:48 UTC | spp_1010_2001_6_172d10600cd | 11635605720081780423
```

When the replication is running, a session is created for the replication on the source vSnap server. In Example 15-7 the command **`vsnap session show`** lists the replication session on the source server `t4-spp-vsnap`. The status field indicates whether it is an `Active` session or a `Completed` session.

Example 15-7 List replication session using the vSnap CLI

```
[vsnapadmin@t4-spp-vsnap fs2]$ vsnap session show
ID  | RELATIONSHIP                     | PARTNER TYPE | LOCAL SNAP | REMOTE SNAP | STATUS    | SENT     | STARTED                | ENDED
-------------------------------------------------------------------------------------------------------------------------------------------
106 | b3b0d887744e23aec938528f00d2a546 | vsnap        | 44         | 59          | ACTIVE    | 2.93GB   | 2020-06-20 10:59:22 UTC | N/A

[vsnapadmin@t4-spp-vsnap fs2]$ vsnap session show --id 106

ID: 106
RELATIONSHIP ID: b3b0d887744e23aec938528f00d2a546
PARTNER ID: N/A
PARTNER TYPE: vsnap
REPL ADDRESS: 10.0.250.43
LOCAL SNAP ID: 44
LOCAL SNAP NAME: spp_1010_2001_6_172bbbaee99
REMOTE SNAP ID: 59
REMOTE SNAP NAME: spp_1010_2102_172d162ecbd_Silver
PRIORITY: 50
STATUS: ACTIVE
CANCELLED: No
SENT: 2.93GB
QUEUED: 2020-06-20 10:59:20 UTC
STARTED: 2020-06-20 10:59:22 UTC
ENDED: N/A
MESSAGE: Transferred 2.93GB of 17.34GB; 17% complete; Average throughput 100.06MB/s

[vsnapadmin@t4-spp-vsnap fs2]$ vsnap session show --id 106

ID: 106
RELATIONSHIP ID: b3b0d887744e23aec938528f00d2a546
PARTNER ID: N/A
PARTNER TYPE: vsnap
REPL ADDRESS: 10.0.250.43
LOCAL SNAP ID: 44
LOCAL SNAP NAME: spp_1010_2001_6_172bbbaee99
REMOTE SNAP ID: 59
REMOTE SNAP NAME: spp_1010_2102_172d162ecbd_Silver
PRIORITY: 50
STATUS: COMPLETED
CANCELLED: No
SENT: 17.37GB
QUEUED: 2020-06-20 10:59:20 UTC
STARTED: 2020-06-20 10:59:22 UTC
ENDED: 2020-06-20 11:02:32 UTC
MESSAGE: Completed
```

After the replication session completes, the snapshot on the target volume is updated and details of the snapshot are available (see Example 15-8) by using the vSnap CLI command **vsnap snapshot show**.

Example 15-8 Snapshot on the replication target is created

```
[serveradmin@t4-spp-vsnap-dr home]$ vsnap snapshot show
ID | PARENT ID | CREATED                | NAME                              | VERSION ID
-------------------------------------------------------------------------------------------------
65 | 4         | 2020-06-20 11:05:41 UTC | spp_1010_2102_172d1685d94_Silver | 7656150810181148115

[serveradmin@t4-spp-vsnap-dr home]$ vsnap snapshot show --id 65

ID: 65
NAME: spp_1010_2102_172d1685d94_Silver
PARENT ID: 4
PARENT NAME: spp_172d162ed86_spp_1010_2001_172bbb8773f__group7e7ae94b6398496
POOL ID: 1
POOL NAME: primary
HAS CLONES: No
USED SPACE: 0.00KB
VERSION ID: 7656150810181148115
CREATED: 2020-06-20 11:05:41 UTC
UPDATED: 2020-06-20 11:05:41 UTC
```

The relationship of the volumes on the source and the target server is updated (see Example 15-9) for the source vSnap server and in Example 15-10 for the target vSnap server.

Example 15-9 vSnap relationship on the source server completed

```
[vsnapadmin@t4-spp-vsnap fs2]$ vsnap relationship show --id b3b0d887744e23aec938528f00d2a546

ID: b3b0d887744e23aec938528f00d2a546
PARTNER ID: 4b34983b741447a6a41a96a62525cf87
PARTNER TYPE: vsnap
PARTNER ADDR: t4-spp-vsnap-dr
LOCAL ROLE: primary
LOCAL POOL ID: 1
LOCAL VOL ID: 6
LOCAL VOL NAME: spp_1010_2001_172bbb8773f__group7e7ae94b6398496985ca5f389ab892a
REMOTE POOL ID: 1
REMOTE VOL ID: 4
REMOTE VOL NAME: spp_172d162ed86_spp_1010_2001_172bbb8773f__group7e7ae94b6398496
LAST SYNC STATUS: COMPLETED
LAST SYCNED SNAP ID: 65
LAST ATTEMPT SNAP ID: 65
CREATED: 2020-06-20 10:59:20 UTC
UPDATED: 2020-06-20 10:59:20 UTC
```

Example 15-10 vSnap relationship on the target server completed

```
[serveradmin@t4-spp-vsnap-dr home]$ vsnap relationship show --id b3b0d887744e23aec938528f00d2a546

ID: b3b0d887744e23aec938528f00d2a546
PARTNER ID: 9b44fb0c808b434688b7e8e70a6b391a
PARTNER TYPE: vsnap
PARTNER ADDR: t4-spp-vsnap
LOCAL ROLE: replica
LOCAL POOL ID: 1
LOCAL VOL ID: 4
LOCAL VOL NAME: spp_172d162ed86_spp_1010_2001_172bbb8773f__group7e7ae94b6398496
REMOTE POOL ID: 1
REMOTE VOL ID: 6
```

```
REMOTE VOL NAME: spp_1010_2001_172bbb8773f__group7e7ae94b6398496985ca5f389ab892a
LAST SYNC STATUS: COMPLETED
LAST SYCNED SNAP ID: 65
LAST ATTEMPT SNAP ID: 65
CREATED: 2020-06-20 10:59:22 UTC
UPDATED: 2020-06-20 10:59:22 UTC
```

15.4 Additional copies to Object Storage, tape, or archival storage

As an alternative to replicating vSnap data to another vSnap, the following approaches are also possible:

- **Additional Copies to standard Object Storage** to create additional incremental copies of vSnap backup data into a hot storage tier.

  Supported Object Storage providers are Amazon S3, IBM Object Storage, Microsoft Azure, S3 Compatible Storage (Dell EMC, RStor, and OpenIO) and IBM Spectrum Protect with container storage pools.

  When vSnap data is copied to Object Storage, a full copy is created during the first copy operation. Subsequent copies are incremental and capture cumulative changes since the last copy. Data on Object Storage is compressed, but not deduplicated.

  Copping snapshots to Object Storage is useful if you want relatively fast backup and recovery times and do not require the longer-term protection, lower cost, and security benefits that are provided by tape or Object Storage archive.

  Note: Copy to standard Object Storage operates at a block level (incremental copies) and must copy only the changed block data since the last copy. IBM Spectrum Protect Plus can reconstruct any of the recovery points from that point. Archive Object Storage or tape are different because they are always a full copy.

- **Additional Copies to archive Object Storage or tape** to create additional full copies of backup data into a “cold” storage tier.

  Archive Object Storage archive is a long-term storage method that copies data to one of the following storage services: Amazon Glacier, IBM Cloud Object Storage Archive Tier, or Microsoft Azure Archive.

  Tape storage means that data is stored on physical tape media or in a virtual tape library that is connected to an IBM Spectrum Protect server. By storing tape volumes at a secure, offsite location that is not connected to the internet, you can help to protect your data from online threats, such as malware and hackers.

  Copying snapshots to tape or Object Storage archive provides extra cost and security benefits. However, because copping to these storage types requires a full data copy, the time that is required to copy data increases. In addition, the recovery time can be unpredictable and the data might take longer to process before it is usable.

15.4.1 Prerequisites

To create additional copies of vSnap data to Object Storage, the following requirements must be met:

- A disk cache area is present on the vSnap server.

- Virtual cloud devices are excluded from the multipath configuration.
- Certificates from private certificate authority (CA) are trusted if a private CA is used.
- Network communication is allowed between vSnap and Object Storage provider.
- An Object Storage provider is configured.

To create additional copies of vSnap data to tape or archival storage, the requirements are similar to the Object Storage requirements; however, the vSnap disk cache area is not required for full archival operations.

Note: Object Storage copy SLA only sends the latest backup, not the intermediate backups.

Preparing the disk cache area

The vSnap disk cache area is used as a temporary staging area for objects that are pending upload to the Object Storage provider endpoint. Starting at version 10.1.6, this option is considered optional for vSnap server deployments.

During restore operations, the disk cache area is used to cache downloaded objects and store any temporary data that might be written into the restore volume.

The cache area must be configured in the form of an XFS file system that is mounted at `/opt/vsnap-data` on the vSnap server. If this mount point is not configured, copy or restore jobs fail with the following error:

```
Cloud functionality disabled: Data disk /opt/vsnap-data is not configured.
```

Note: For new virtual vSnap appliance deployments starting at version 10.1.3, a pre-configured disk cache area of 128 GB size is mounted to `/opt/vsnap-data`.

For all custom vSnap server installations and virtual vSnap appliances that were upgraded from version 10.1.1, the cache area must be configured manually.

For more information about sizing, configuring, or increasing the cache area, see the latest version of the IBM Spectrum Protect Plus Blueprints.

Multipath configuration

During copy operations to Object Storage, IBM Spectrum Protect Plus attaches and detaches virtual cloud devices on vSnap servers. If the multipath configuration is enabled on the vSnap server by using `dm-multipath`, the configuration can interfere with the copy operation.

Tip: To avoid this interference, the virtual cloud devices must be excluded from the multipath configuration.

Add the following lines under the blacklist section of the multipath configuration file `/etc/multipath.conf`:

```
blacklist {
    device {
        vendor "LIO-ORG"
        product ".*"
    }
}
```

After making this change, reload the multipath configuration by using the following command:

```
$ sudo systemctl reload multipath
```

Trusting private CA certificates

The following CA certificates can be used:

- Certificates signed by private CA

  In addition to the certificate, the root/intermediate certificate of the private CA must be added to the system certificate store in each vSnap server by completing the following steps:

  a. Log in to the vSnap server console as the server admin user and upload any private CA certificates (in PEM format) to a temporary location.

  b. Copy each certificate file to the system certificate store directory (`/etc/pki/ca trust/source/anchors/`) by running the following command:

  ```
  sudo cp /tmp/private-ca-cert.pem /etc/pki/ca-trust/source/anchors/
  ```

  c. To incorporate the newly added custom certificate and update the system certificate bundle, run the following command:

  ```
  $ sudo update-ca-trust
  ```

- Certificates signed by public CA

  If the Object Storage endpoint uses a public CA-signed certificate, no special action is required. The vSnap server validates the certificate by using the default system certificate store.

- Self-signed certificates

  If the Object Storage endpoint or repository server uses a self-signed certificate, you must specify the certificate in Privacy Enhanced Mail (PEM) format when you register the Object Storage or repository server in the IBM Spectrum Protect Plus user interface and then follow the same procedure to add it.

Note: If the Object Storage endpoint or repository server uses a self-signed certificate, the certificate must be specified while registering the Object Storage or repository server in the IBM Spectrum Protect Plus user interface. Because this process occurs during the registration of the Object Storage provider in the IBM Spectrum Protect Plus UI, it does not need to be prepared in advance.

Allowing network communication between vSnap and object providers

The ports that are used for communication between vSnap servers and Object Storage or repository server endpoints and need to be allowed through any existing firewalls.

Table 15-2 Outgoing vSnap server firewall connections

Port	Initiator	Target	Description
443	T vSnap Cserver P	Object Storage Server enpoints	Allows the vSnap server to communicate with Amazon Simple Storage Service (S3), Microsoft Azure, or IBM Cloud Object Storage endpoints.
9000	T vSnap Cserver P	Repository Server endpoints	Allows the vSnap server to communicate with IBM Spectrum Protect (repository server) endpoints.

Any firewalls or network proxies that perform SSL interception or Deep Packet inspection for traffic between vSnap servers and Object Storage endpoints might interfere with SSL certificate validation on vSnap servers. This interference can also cause Object Storage copy job failures. To prevent this interference, the vSnap servers must be exempted from SSL interception and inspection in the firewall or proxy configuration.

15.4.2 Preparing Object Storage providers

This section describes how to prepare Object Storage providers for use with IBM Spectrum Protect Plus.

Important: Native lifecycle management performed by the Object Storage provider is not supported. IBM Spectrum Protect Plus manages the lifecycle of uploaded objects automatically by using an incremental-forever approach in which older objects can still be used by newer snapshots. Automatic or manual expiration of objects that are outside of IBM Spectrum Protect Plus lead to data corruption.

Amazon S3 copy requirements (“hot” incremental copies)

When the Object Storage provider is registered in IBM Spectrum Protect Plus, a bucket in one of the supported storage tiers must be specified:

- S3 Standard
- S3 Intelligent-Tiering
- S3 Standard-Infrequent Access
- S3 One Zone-Infrequent Access

For performance reasons, the selected Region should be close to the geographical region where the vSnap server system is located.

Note: Amazon S3 supports various storage classes for storing objects into a bucket. IBM Spectrum Protect Plus “hot” incremental copies store objects to the default storage class, which is `S3 STANDARD`.

Complete the following steps to add one Amazon S3 Storage bucket:

1. Log in to the AWS console at: `https://console.aws.amazon.com/s3/`.
2. Click **Create bucket**.
3. In the **Bucket name** field, enter a DNS-compliant name for your bucket.
4. In **the Region** field, choose the AWS Region where you want the bucket to be stored, as shown in Figure 15-11.

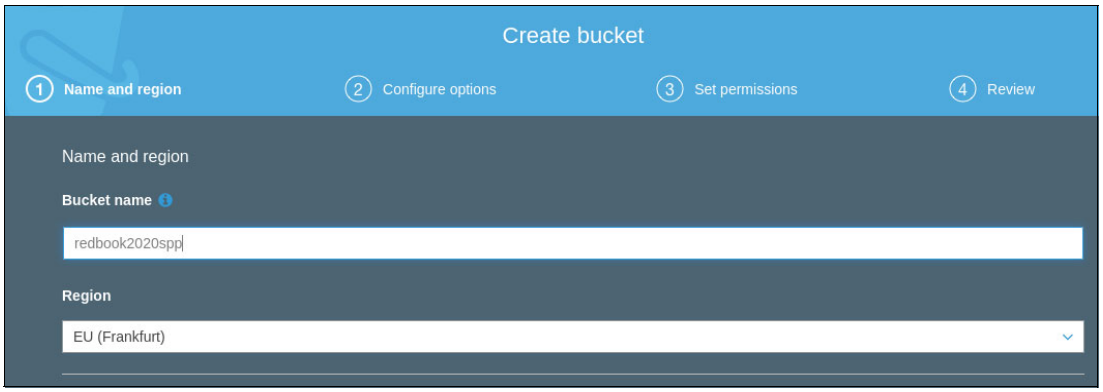

Figure 15-11 Create Amazon S3 Bucket

5. In the Bucket settings for Block Public Access, keep the values set to the defaults, as shown in Figure 15-12.

Figure 15-12 Amazon S3 Block Public Access values

6. Select **Next** to review your settings and then, click **Create Bucket.** The Amazon S3 Bucket is created (see Figure 15-13).

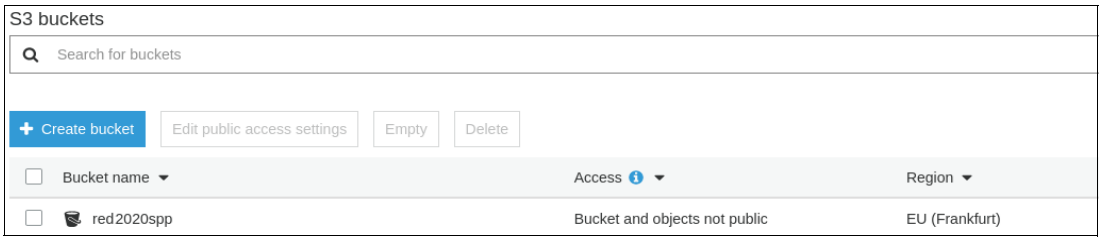

Figure 15-13 Amazon S3 Bucket is created

Complete the following steps to create a User Account for Amazon S3 bucket access:

1. Log in to the AWS console at: `https://console.aws.amazon.com`.
2. Enter `IAM` into the Find Services search bar and select **Manage Access to AWS resources**.
3. Select **Users** in the left pane and then, click **Add User**, as shown in Figure 15-14 and Figure 15-15.

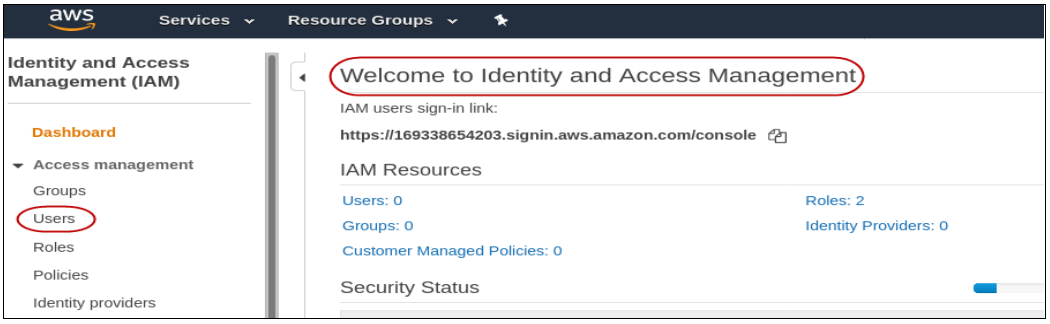

Figure 15-14 Selecting Users

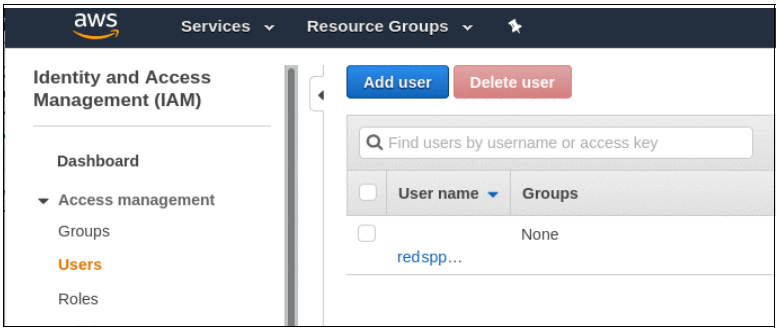

Figure 15-15 Adding Amazon S3 bucket user

4. Create a username and select only **Programmatic access**. Then, click **Next**.

> **Note:** This user account is restricted to S3 only.

5. Select **Attach existing policies**. In the Filter policies text field, enter S3 and select only **AmazonS3FullAccess**. Then, click **Next** (see Figure 15-16). Optionally, add tags to the user account.

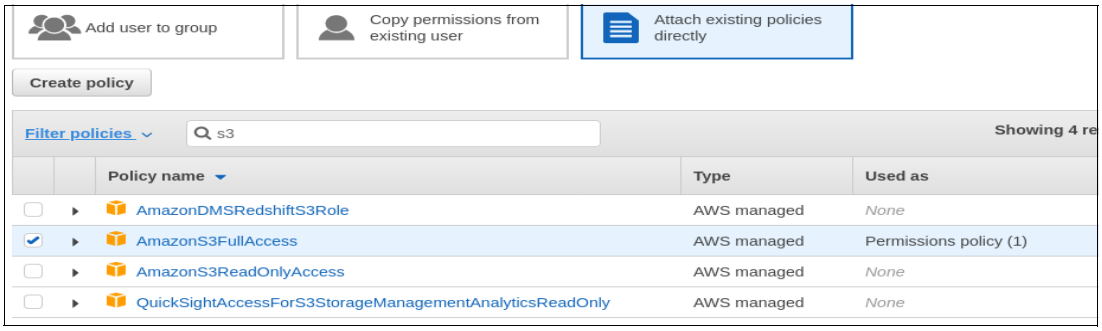

Figure 15-16 Amazon S3 User Access policies

6. Click **Create User**. Click **Show** ion the Secret Access key column and keep this key and the Access key ID, as shown in Figure 15-17.

.

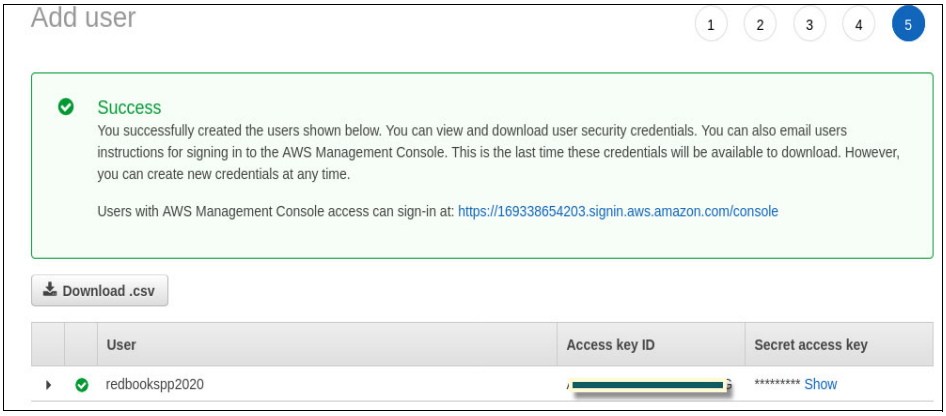

Figure 15-17 Amazon S3 bucket user is created

Amazon S3 archive Object Storage requirements (“cold” full copies)

When the Object Storage provider is registered in IBM Spectrum Protect Plus, a bucket in one of the supported storage tiers must be specified:

- S3 Standard
- S3 Intelligent-Tiering
- S3 Standard-Infrequent Access
- S3 One Zone-Infrequent Access

IBM Spectrum Protect Plus directly uploads data files to the Glacier tier. Some small metadata files are stored in the default tier for the bucket. A copy of these metadata files is also placed into the Glacier tier for disaster recovery purposes.

The vSnap server must communicate with IBM Spectrum Protect (repository server) endpoints.

Complete the following steps to add an Amazon S3 Archive Object Storage bucket:

1. Log in to the AWS console at: `https://console.aws.amazon.com/s3/`.
2. Create one Amazon S3 Bucket and User as described in “Amazon S3 copy requirements (“hot” incremental copies)” on page 422.

3. Create one Amazon S3 Archive LifeCycle Policy (for more information, see this web page), as shown in Figure 15-18.

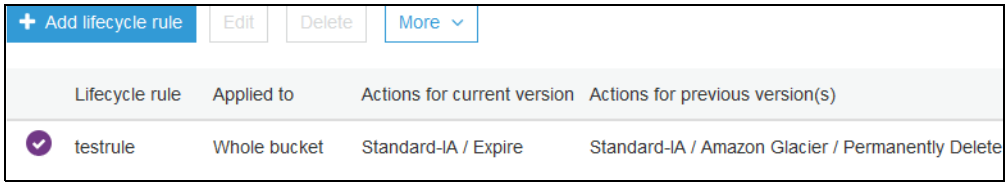

Figure 15-18 Creation of Amazon S3 Archive Bucket

IBM Cloud copy requirements (“hot” incremental copies)

When the Object Storage provider is registered in IBM Spectrum Protect Plus, a bucket must be specified. If the specified bucket includes a Write Once Read Many (WORM) policy that locks objects for a specific time, IBM Spectrum Protect Plus automatically detects the configuration and deletes snapshots after the WORM policy removes the lock. The bucket must have the Name Index setting enabled.

For performance reasons, the selected Region should be close to the geographical region where the vSnap server system is located.

Complete the following steps to add a IBM Cloud Object Storage bucket:

1. Log in to the AWS console: `https://cloud.ibm.com`.
2. Click **Create Source**.
3. Select **Object Storage** from the provider list.
4. Select **Standard Pricing Plan** → **Add Service Name** → **Select a Resource Group** (if a group was configured) → **Create**, as shown in Figure 15-19.

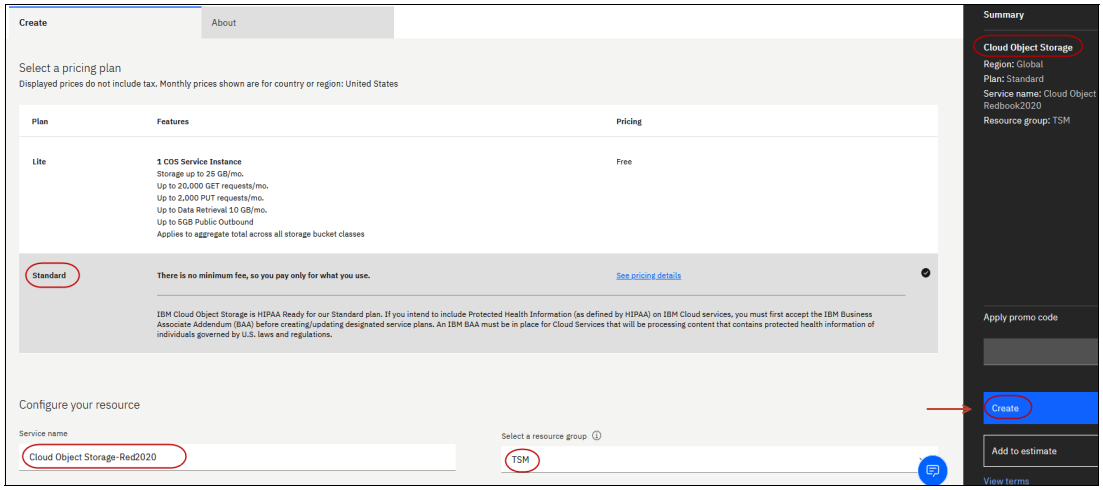

Figure 15-19 Creating IBM Cloud Object Storage

5. Select **IBM Cloud** → **Resource List**, as shown in Figure 15-20.

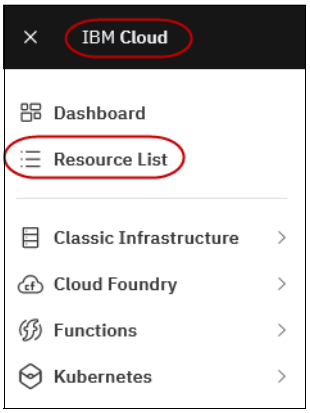

Figure 15-20 Selecting Resource List

6. Select **Storage** → **Choose Object Storage has been created**, as shown in Figure 15-21.

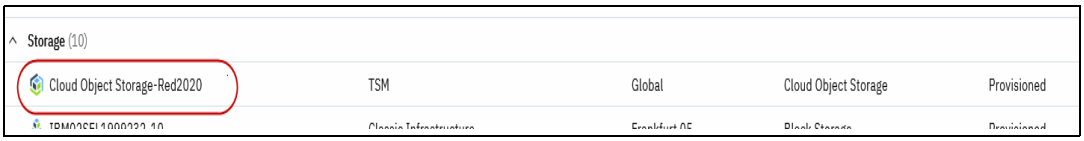

Figure 15-21 IBM Cloud Object Storage is created

7. Create the bucket, depending on the customer needs, as shown in Figure 15-22. Complete the following fields:
   - Unique Bucket name
   - Residency
   - Location
   - Storage Class

   Click **Create Bucket**.

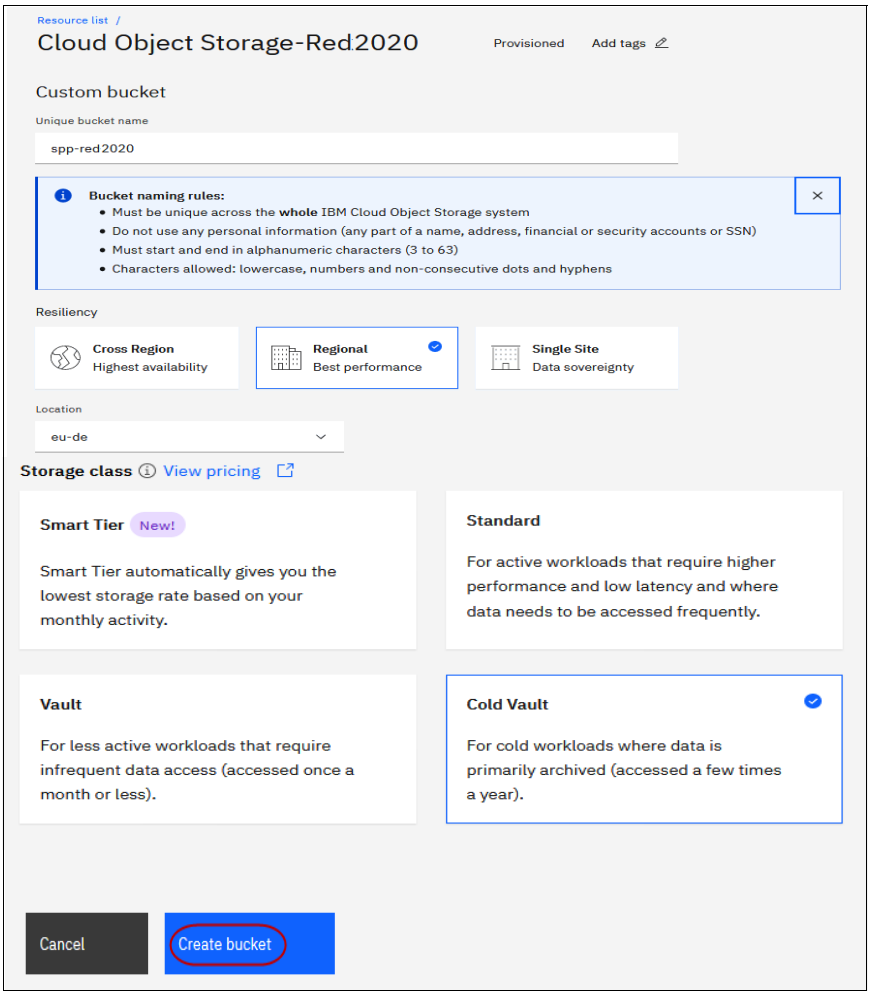

Figure 15-22 Creating IBM Cloud Bucket

8. Create a Service Credential Bucket user by clicking **Service Credentials** → **New Credential** → **Select Credential name** → **Role writer** → **Click Include HMAC Credential ON** → **Add** (see Figure 15-23 and Figure 15-24 on page 429).

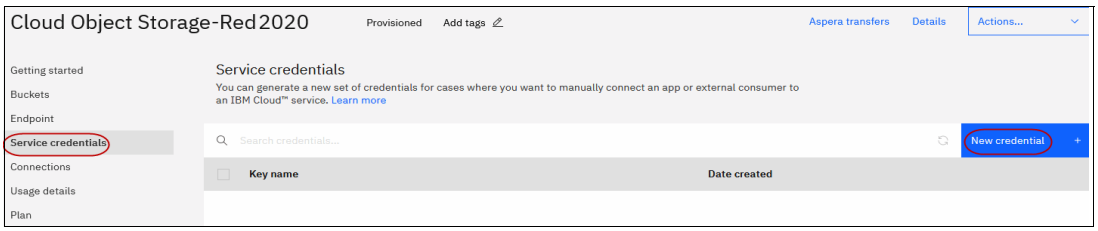

Figure 15-23 Creating IBM Cloud Service Credentials

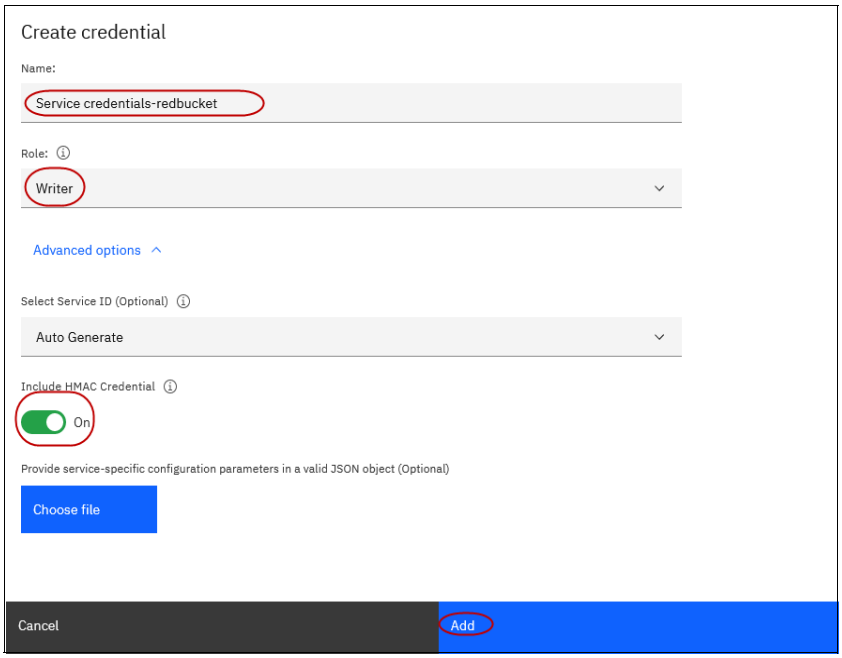

Figure 15-24 Selecting IBM Cloud Service Credentials Options

Displaying Credential User Resume, obtained Access Key ID, Secret Access Key, and Endpoints values as shown in Figure 15-25.

Figure 15-25 IBM Cloud Credential Resume

IBM Cloud Archive Object Storage Requirements (“cold” full copies)

When the Object Storage provider is registered in IBM Spectrum Protect Plus, a bucket must be specified. If the specified bucket includes a WORM policy that locks objects for a specific time, IBM Spectrum Protect Plus automatically detects the configuration and deletes snapshots after the WORM policy removes the lock.

IBM Spectrum Protect Plus creates a single lifecycle management rule on the bucket to migrate data files to the archive tier. The bucket must have the Name Index setting enabled.

For performance reasons, the selected Region should be close to the geographical region where the vSnap server system is located.

Complete the following steps to add an IBM Cloud Object Storage Archive bucket:

1. Create one normal bucket and user as described in “Amazon S3 copy requirements (“hot” incremental copies)” on page 422.
2. Before creating the bucket, in the Bucket Configuration window**,** select **Archive Rule Create** option, if its status is disabled for this bucket, as shown in Figure 15-26. Then, click **Save**.

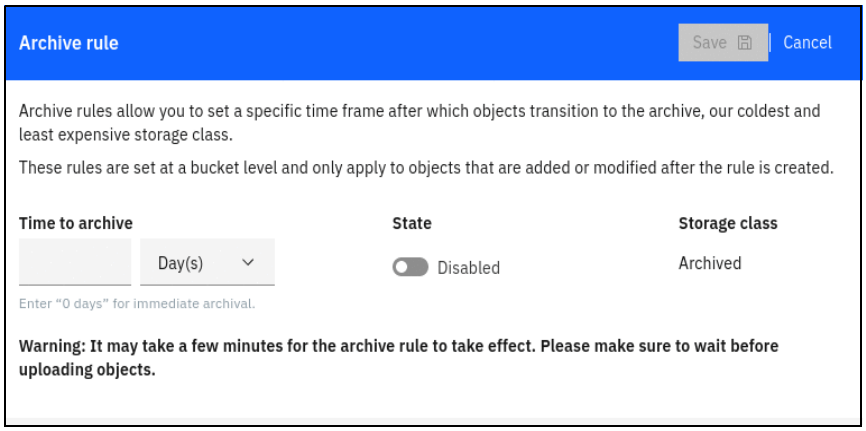

Figure 15-26 Creating IBM Cloud Archive Bucket

Note: Consider the following points:

- Object Storage in IBM Cloud supports an archive tier for long-term data. IBM Spectrum Protect Plus creates a single lifecycle management rule on the bucket to migrate data files to the archive tier.
- Because on-premises IBM Cloud Object Storage systems do not support the archive tier, archiving (“cold” full copies) are supported on IBM Cloud only.

Microsoft Azure copy requirements (“hot” incremental copies)

When the Object Storage Microsoft Azure is registered in IBM Spectrum Protect Plus, a container in a hot or cool storage must be specified.

You can use a cool storage tier for cost-effective, long-term storage. However, it is more costly to restore data from a cool storage tier than from a hot storage tier.

Complete the following steps to add a Microsoft Azure hot storage container:

1. Log in to the Azure portal and click **Storage accounts** (see Figure 15-27).

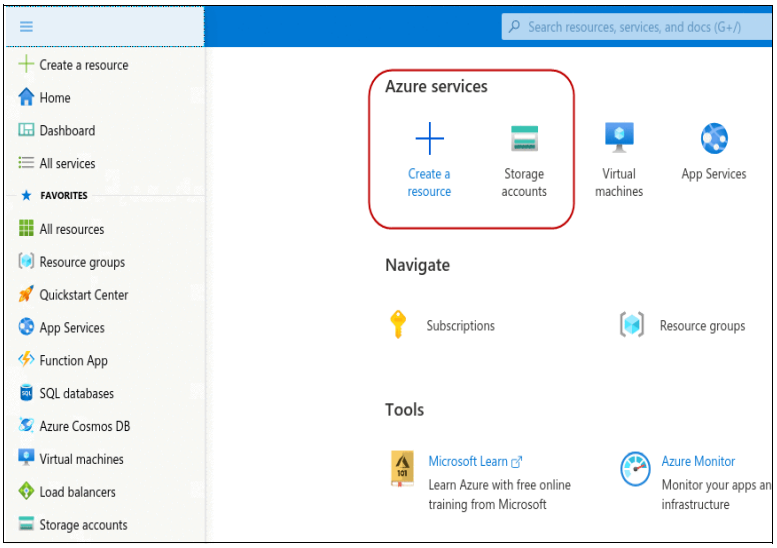

Figure 15-27 Creating Azure Services Storage accounts

2. Click **Create Account** to create the storage account. Select a resource group or create a resource group for the new storage account.
3. Enter the name for the storage account and choose the location. Select the Access tier (in our example, **Hot** storage) and click **Next** (see Figure 15-28).

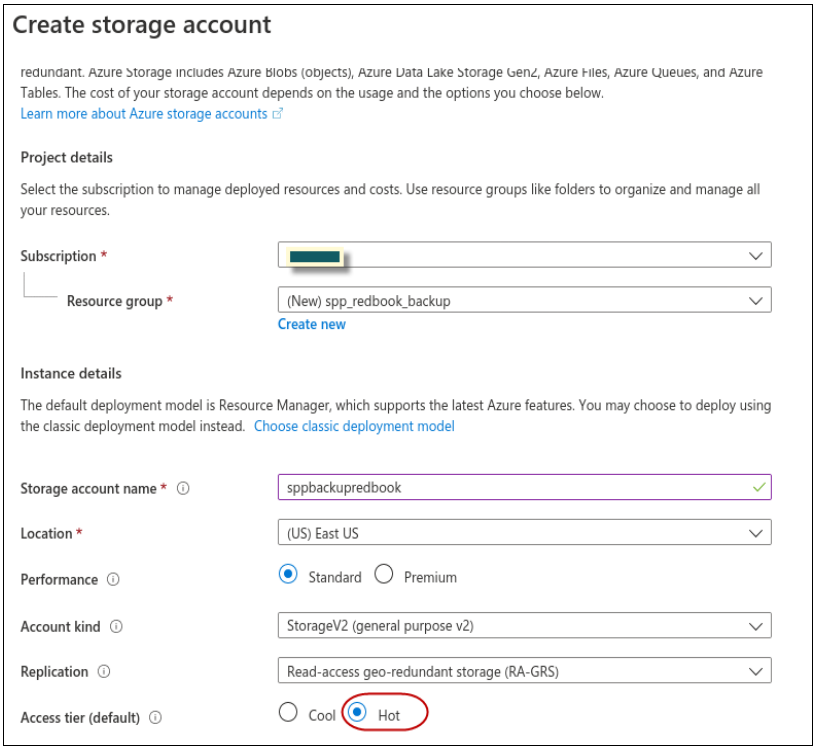

Figure 15-28 Defining Azure Storage account Hot tier and region

4. Keep the options at their default settings, including the settings in the Advanced tab, as shown in Figure 15-29. Click **Next**.

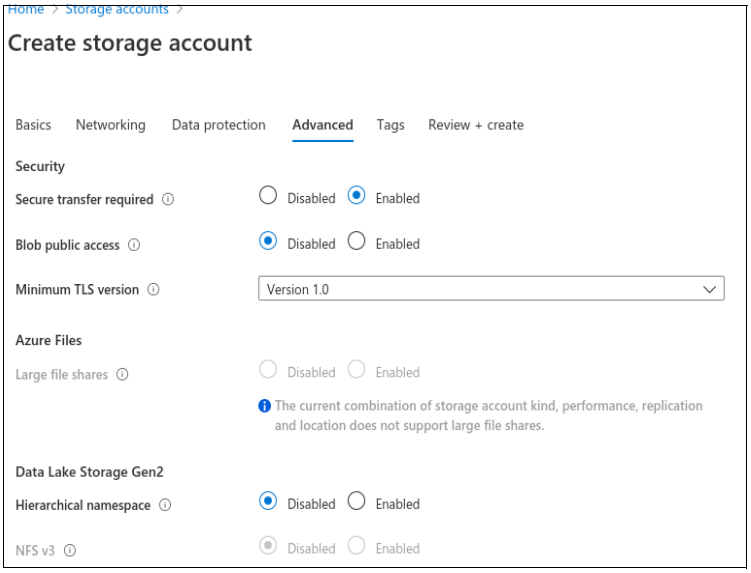

Figure 15-29 Microsoft Azure Container Advanced Options window

5. Optionally, create a tag for the new storage account and click **Review+Create** to review the settings. Verify all settings and click **Create**, as shown in Figure 15-30.

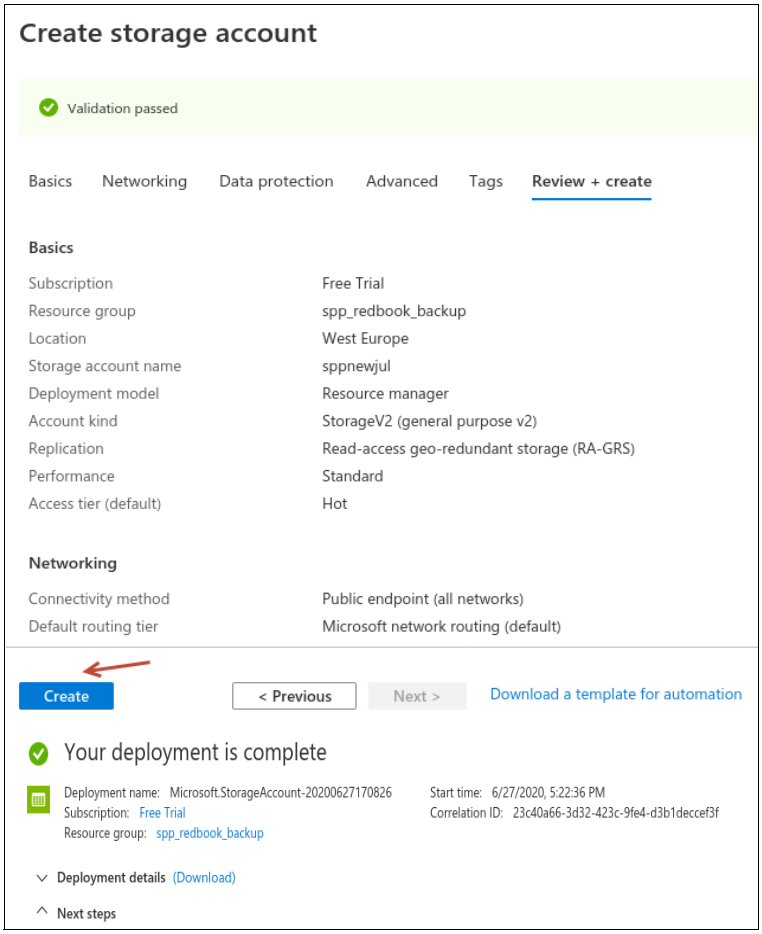

Figure 15-30 Microsoft Azure Storage Accounts created

6. To find Access keys for the storage account that was created, click **Storage account** and select **Access Keys**, as shown in Figure 15-30

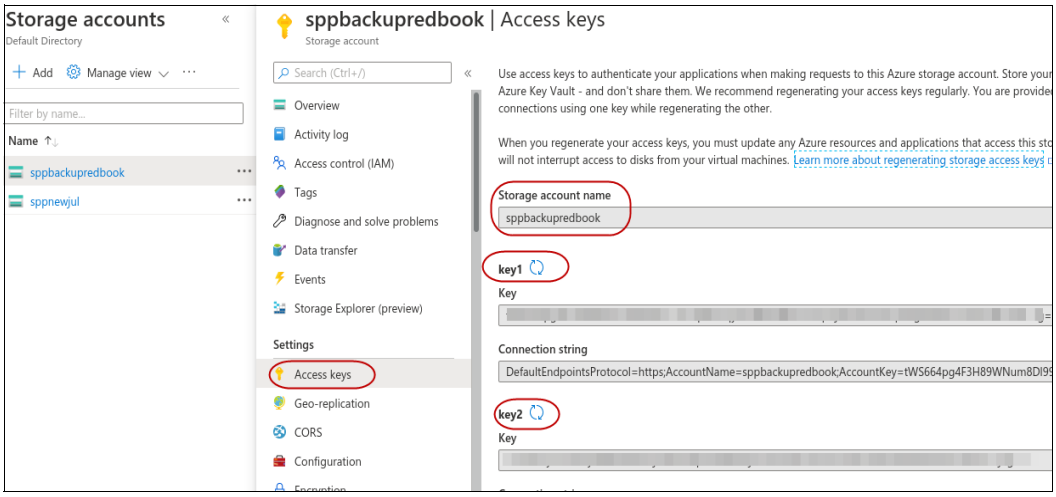

Figure 15-31 Microsoft Azure Access keys

7. Create a Blob container by selecting **Containers** under Blob service in the navigation pane, as shown in Figure 15-32.

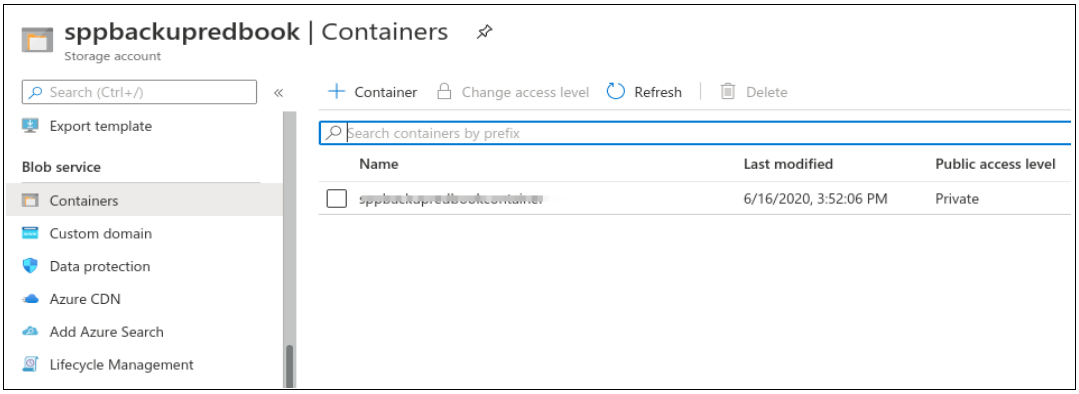

Figure 15-32 Microsoft Azure Blob container creation

8. Enter a name for the storage container and click **OK**. A blob container for copy use is created, as shown in Figure 15-33.

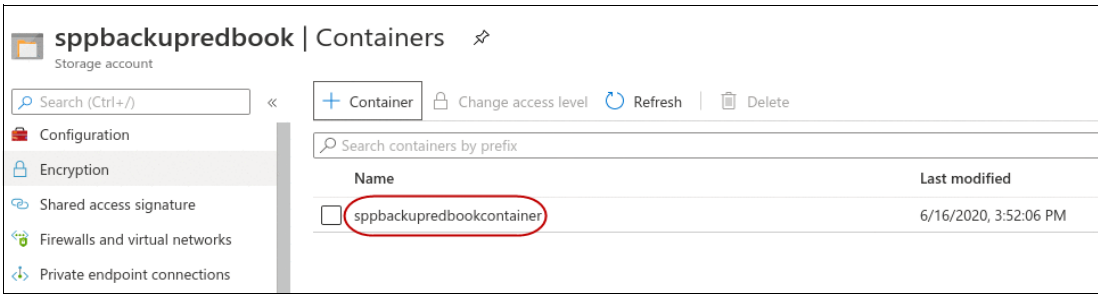

Figure 15-33 Azure Storage Container is created

Microsoft Azure Archive Object Storage requirements (“cold” full copies)

When the Object Storage provider is registered in IBM Spectrum Protect Plus, a container in a cool storage account must be specified. IBM Spectrum Protect Plus moves files between tiers on demand. Data files are immediately moved to the archive tier and temporarily returned to the hot tier during restore operations only.

Some small metadata files are stored in the default tier for the container. A copy of these metadata files is also placed in the archive tier for disaster recovery purposes.

Complete the following steps to add a Microsoft Azure cold storage container:

1. Log in to the Azure portal and click **Storage accounts**.
2. Complete the steps that are described in “Amazon S3 copy requirements (“hot” incremental copies)” on page 422, but choose **Cold** as the default access, as shown in Figure 15-34.

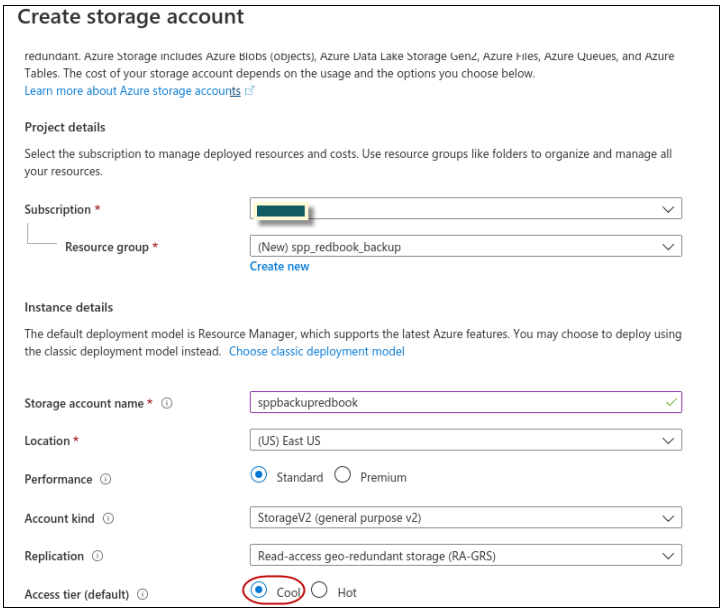

Figure 15-34 Defining Microsoft Azure Storage account Cold tier and region

S3 compatible Object Storage

In addition to backing up data to Amazon Storage (S3) and IBM Cloud Object Storage, you can back up data to other S3-compatible Object Storage providers.

Before you back up data in a production environment to any other S3 compatible Object Storage, ensure that the Object Storage is validated for use with IBM Spectrum Protect Plus.

The different compatible providers are listed in Table 15-3, along with some useful links to the documentation of those providers. For more information about the latest version of this table, see this web page.

Table 15-3 S3 Compatible Object Storage

Object Storage	IBM Spectrum Protect Plus version	Supported IBM Spectrum Protect Plus buckets (standard or archive)	For more information
Red Hat Ceph Object Storage	V10.1.6	Standard Object Storage buckets	About configuring and administering Ceph S3 Object Storage, see this web page.
RStor Object Storage	V10.1.6	Standard Object Storage buckets	See the RStor website.
Dell EMC Elastic Cloud Storage	V10.1.6	Standard Object Storage buckets	See the Quick Start guide.
OpenIO Object Storage	V10.1.6	Standard Object Storage buckets	See the Quick Start guide.

15.4.3 Preparing repository server storage

Note: For IBM Spectrum Protect Plus, the repository server must be an IBM Spectrum Protect server Version 8.1.8 or later.

IBM Spectrum Protect server requirements

When IBM Spectrum Protect is used as the Object Storage provider for IBM Spectrum Protect Plus, you cannot prepare a bucket in advance. IBM Spectrum Protect Plus creates a uniquely named bucket for its own use, on demand.

The corresponding ObjectClient in IBM Spectrum Protect must be assigned to a policy domain that uses a default management class. This class points to a directory- or cloud-container storage pool.

IBM Spectrum Protect Plus offers the Archive option, which stores additional full copies of backup data into an IBM Spectrum Protect Tape storage pool. To support the selection of different target storage types (“storage classes”) by using the S3 protocol, IBM Spectrum Protect server version 8.1.8 was enhanced with a new storage pool type (“ColdDataCache”) and a new policy domain type (“ObjectDomain”). Consequently, all IBM Spectrum Protect ObjectClient nodes that store data for IBM Spectrum Protect Plus must be assigned to an ObjectDomain in IBM Spectrum Protect.

If you are planning to copy or archive IBM Spectrum Protect Plus data to an IBM Spectrum Protect server, three configurations are available (as listed in Table 15-4). Choosing the one to configure depends on which scenario applies to your data protection needs.

Table 15-4 Possible configurations of IBM Spectrum Protect Repository Server

User scenario	Purpose	Steps
Copying to standard Object Storage when you are running daily or less frequent copies to standard Object Storage.	Copy data to standard Object Storage (Hot).	To copy data to standard Object Storage to the IBM Spectrum Protect server, must create a cloud-container or directory-container storage pool, and set up the object agent component of IBM Spectrum Protect. For more information, see steps 2 - 4 in “Setting up and configuring data transfer communication”.

User scenario	Purpose	Steps
Copying to tape when you are creating a weekly or less frequent full-copy of your data to tape storage.	When you copy data to tape, a full copy of the data is created at the time of the copy process (Cold).	To copy data to tape, you must create a cloud-container or directory-container storage pool for tape, and a cold-data-cache storage pool on the IBM Spectrum Protect server. For more information, see steps 1-4 in “Setting up and configuring data transfer communication”.
Mixture of standard Object Storage and long-term copying to tape.	Secure your data in incremental backups on the IBM Spectrum Protect server, and retaining data on tape for longer term security.	This scenario is a combination of the previous cases; that is, data is stored to tape and on standard Object Storage at the IBM Spectrum Protect server. The required data storage pools also are set up for both scenarios.

Note: Copying to an IBM Spectrum Protect server with cloud containers instead of copying directly from the IBM Spectrum Protect Plus server to S3 Object Storage can be done for several reasons, including the following examples:

- Deduplication and compression
- Additional Object Storage targets can be added
- Consolidated management can be done by using IBM Spectrum Protect server.

Setting up and configuring data transfer communication

Complete the following steps to set up and configure the data transfer communication between IBM Spectrum Protect Plus and the IBM Spectrum Protect server:

1. Create a tape storage pool and a cold-data-cache storage pool for copying data to tape. For more information, see this IBM Knowledge Center web page.
2. Configure an object policy domain that defines the rules that control the backup services for IBM Spectrum Protect Plus. For more information, see this IBM Knowledge Center web page.
3. Set up standard Object Storage if you are copying data to a standard storage pool or tape. For more information, see this IBM Knowledge Center web page.
4. Add an object agent for copying data that provides a gateway between the IBM Spectrum Protect Plus server and the IBM Spectrum Protect server. For more information, see this IBM Knowledge Center web page.
5. Add and configure an object client for copying data. For more information, see this IBM Knowledge Center web page.

Note: For more information about how to create a ColdDataCache storage pool in IBM Spectrum Protect 8.1.10, see IBM Knowledge Center.

For more information about how to configure an ObjectDomain in IBM Spectrum Protect, see IBM Knowledge Center.

For more information about how to use IBM Spectrum Protect to create additional copies of IBM Spectrum Protect Plus vSnap data (“hot” or “cold”) see 15.6, “Creating incremental and full copies of backup data to an IBM Spectrum Protect server” on page 448.

15.4.4 Configuring an Object Storage provider in IBM Spectrum Protect Plus

After the Object Storage providers are prepared, they must be registered in the IBM Spectrum Protect Plus user interface.

Because all object providers (Amazon S3, IBM Cloud Object Storage, Microsoft Azure Blob, and IBM Spectrum Protect) use the Amazon S3 protocol, the process for registering them to IBM Spectrum Protect Plus is similar as well.

Adding Amazon S3 Object Storage as a backup storage provider

Complete the following steps to add Amazon S3 as a backup storage provider to IBM Spectrum Protect Plus:

1. Select **System Configuration** → **Backup Storage** → **Object Storage**.
2. Click **Add Object Storage**.
3. Select **Amazon S3** from the provider list.
4. Complete the following fields in the **Object Storage Registration** pane:
   - Enter a meaningful **name** to identify the Object Storage.
   - Select the Amazon **AWS region** endpoint of the Object Storage.
   - Use an existing key or specify the **Key name**, **Access key**, and **Secret key**.
5. Click **Get Buckets** to obtain a list of prepared buckets, as shown in Figure 15-35.

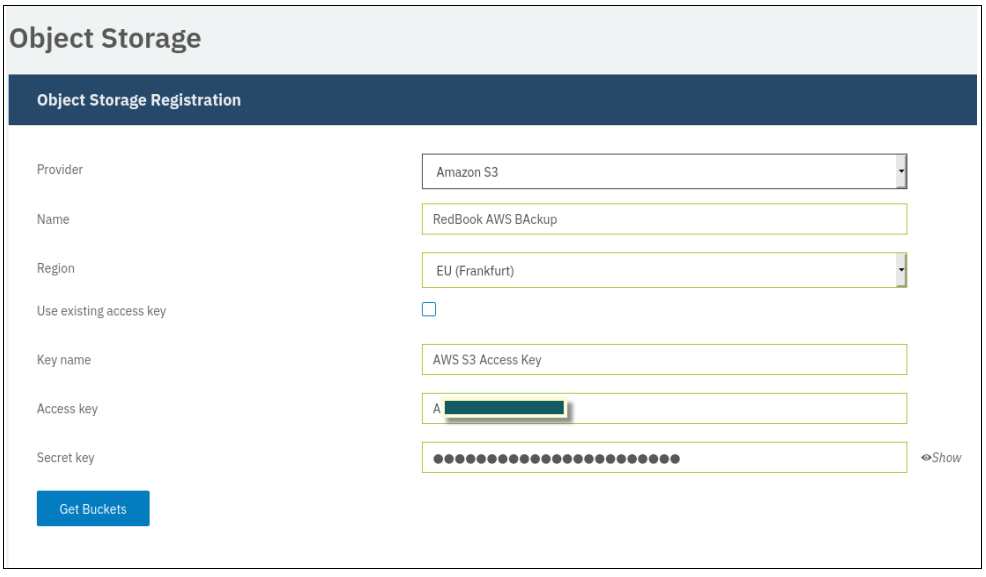

Figure 15-35 Amazon S3 Object Storage Registration window

6. Select the bucket to be used for copy operations (incremental copy to a “hot” tier).
7. (Optional) Select the bucket to be used for archive operations (to “cold” tier).
8. Click **Register** to add the provider to the Object Storage servers table.

Adding IBM Cloud Object Storage as a backup storage provider

Complete the following steps to add IBM Cloud Object Storage as a backup storage provider to IBM Spectrum Protect Plus:

1. Select **System Configuration** → **Backup Storage** → **Object Storage**.
2. Click **Add Object Storage**.
3. Select **IBM Cloud Object Storage** from the provider list.
4. Complete all fields in the **Object Storage Registration** pane:
   - Enter a meaningful **name** to identify the Object Storage.
   - Select the **endpoint** of the Object Storage.
   - Use an existing key or specify the **Key name**, **Access key**, and **Secret key**.
   - Use an existing **Certificate**, or alternatively, upload or copy and paste a new certificate (No certificate is required for public IBM Cloud Object Storage.)
5. Click **Get Buckets** to obtain a list of prepared buckets, as shown in Figure 15-36.
6. Select the bucket to be used for copy operations (incremental copy to a “hot” tier).
7. (Optional) Select the bucket to be used for archive operations (to “cold” tier).
8. Click **Register** to add the provider to the Object Storage servers table.

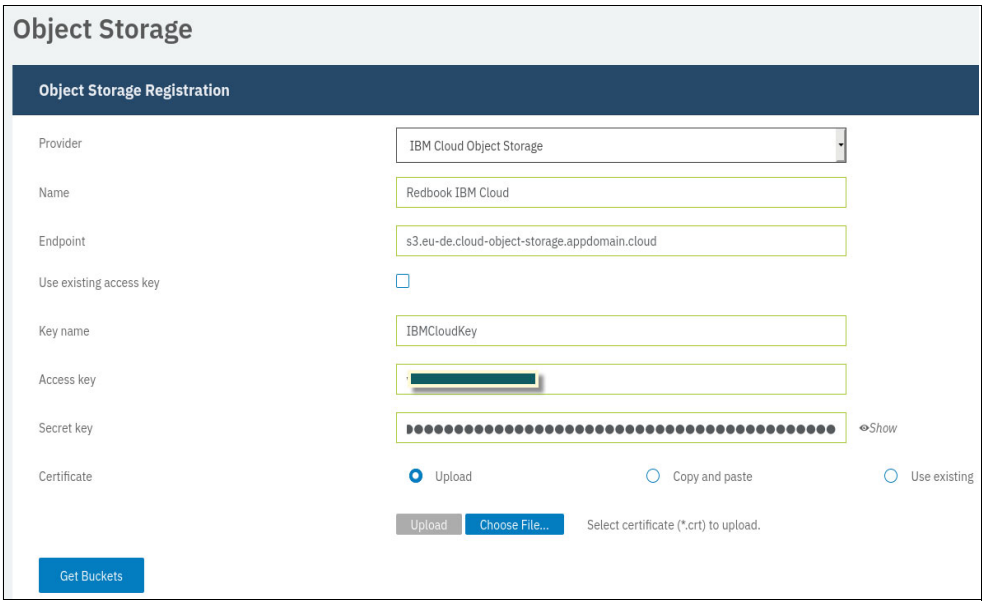

Figure 15-36 Register IBM Cloud Object Storage

Adding Microsoft Azure Object Storage as a backup storage provider

The following is a summary of tasks that need to be performed to add Microsoft Azure Cloud Storage as a backup storage provider to IBM Spectrum Protect Plus:

1. Select **System Configuration** → **Backup Storage** → **Object Storage**.
2. Click **Add Object Storage**.
3. Select **Microsoft Azure Blob Storage** from the provider list.
4. Complete all fields in the **Object Storage Registration** pane:
   - Enter a meaningful **name** for the Object Storage.
   - Select the **endpoint** of the Object Storage.

- Use an existing key or specify the **Key name**, **Storage Account Name**, and **Storage Account Shared Key**.

5. Click **Get Buckets** to obtain a list of prepared buckets.
6. Select the bucket to be used for copy operations (incremental copy to a “hot” tier).
7. Optional: select the bucket to be used for archive operations (to “cold” tier).
8. Click **Register** to add the provider to the Object Storage servers table, as shown in Figure 15-37.

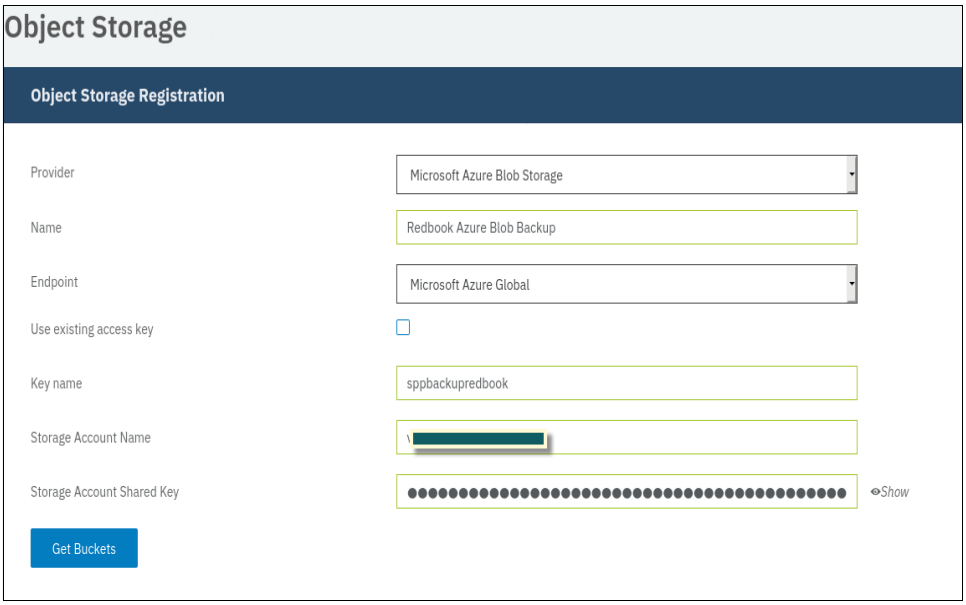

Figure 15-37 Register Microsoft Azure Blob Storage

Adding S3 compatible Object Storage backup provider

Complete the following steps to add S3 Compatible Storage as a backup storage provider to IBM Spectrum Protect Plus:

1. Select **System Configuration** → **Backup Storage** → **Object Storage**.
2. Click **Add Object Storage**.
3. Select **S3 Compatible Object Storage** from the provider list.
4. Complete all fields in the **Object Storage Registration** pane:
   - Enter a meaningful **name** to identify the Object Storage.
   - Select the **endpoint** of the Object Storage.
   - Use either an existing key or specify the **Key name**, **Access Key**, and **Secret Key**.
5. Click **Get Buckets** to obtain a list of prepared buckets.
6. Select the bucket to be used for copy operations (incremental copy to a “hot” tier).
7. (Optional) Select the bucket to be used for archive operations (to “cold” tier).
8. Click **Register** to add the provider to the Object Storage servers table.

Adding an IBM Spectrum Protect server as a repository server provider:

After you add IBM Spectrum Protect as a repository server provider to IBM Spectrum Protect Plus., the wizard provides you with the endpoint for communicating with the object agent on the server, and the access ID, secret key, and certificate for connecting securely.

Certificates can be obtained from the IBM Spectrum Protect server Operations Center by selecting **Server** → **Object Agent** → **Agent Certificate**. Alternatively, the certificate can be obtained from the IBM Spectrum Protect Plus appliance by running the following command:

```
$ openssl s_client -showcerts -connect <ip-address>:9000 </dev/null 2>/dev/null |
openssl x509
```

Complete the following steps to add and register an IBM Spectrum Protect server as a backup storage provider:

- Select **System Configuration** → **Backup Storage** → **Repository Server**.
- Click **Add Repository Server**.
- Complete all fields in the **Register Repository Server** pane:
  - Enter a meaningful **Name** to identify the repository server.
  - Enter the high-level-address (HLA) of the repository server object agent in the **Hostname** field.
  - Enter the TCP port that is used by the object agent (default: 9000).
  - Use an existing key or specify the **Key name**, **Access key**, and **Secret key**.
  - Use a **Certificate**, or alternatively upload or copy and paste a certificate.
- Click **Register** to add the provider to the Object Storage servers table, as shown in Figure 15-38.

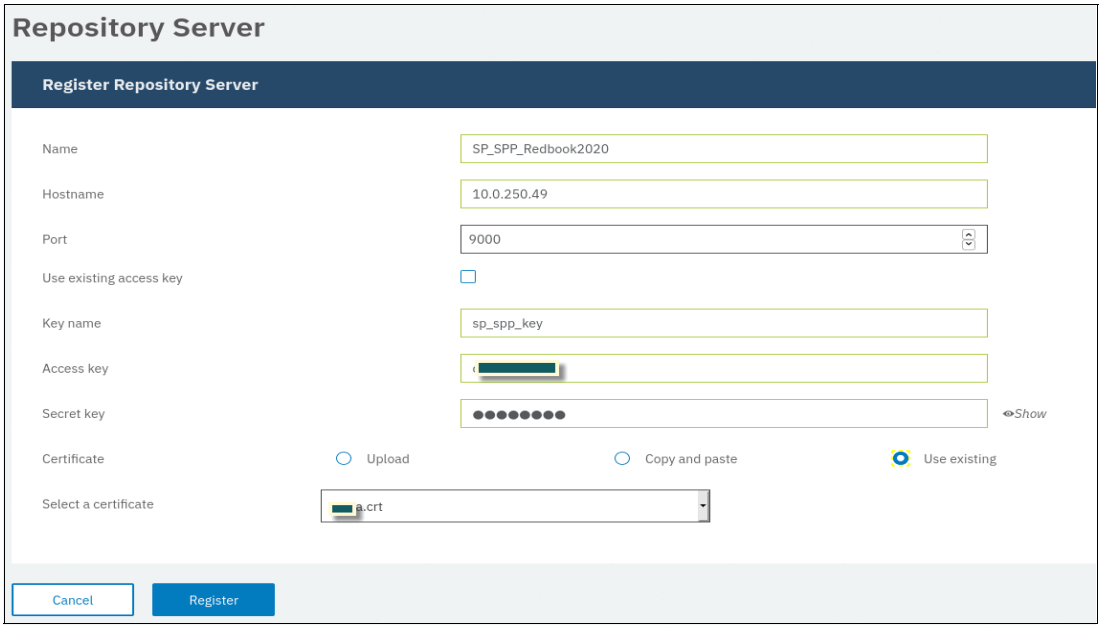

Figure 15-38 Register Repository Server

Note: When registering Object Storage providers, different buckets can be selected for incremental copy or full archival operations.

When registering IBM Spectrum Protect as a repository server, the selection of different buckets for copy and archive operations is *not* required. The configuration of the ObjectDomain in the IBM Spectrum Protect server automatically ensures that “hot” data is stored in a container storage pool and “cold” data is stored to the ColdDataCache storage pool (which is migrated to tape later on).

15.4.5 Configuring additional copies to Object storage in the SLA

IBM Spectrum Protect Plus Service Level Agreement (SLA) policies are used to define parameters for backup jobs. To better understand how these parameters affect the creation of additional copies to Object Storage, the basic structure of a SLA is described next.

A SLA is divided into two main sections:

- Operational Protection: This section includes all parameters that belong to backup data, which is stored into vSnap servers:
  a. Backup Policy: Controls Retention, Schedule, and Target Site for the primary VM, application, or a IBM Spectrum Protect Plus catalog backup.
  b. Replication Policy: Controls Retention, Schedule, and Target Site for vSnap replication operations.

Note: The Recovery Point Objective (RPO) of a backup object is always controlled by the parameters of the Main Policy in the Operational Protection section.

- Additional Copies: This section includes all parameters that belong to the creation of additional copies of data from a vSnap server into object, archive, or tape storage:
  a. Standard Object Storage: Controls Retention, Schedule, Source, and Target for incremental copies to a “hot” storage tier (Object Storage or IBM Spectrum Protect container storage pool).
  b. Archive Object Storage: Controls Retention, Schedule, Source, and Target for full copies to a “cold” storage tier (Object Storage archive storage or IBM Spectrum Protect tape storage pool).

In the parameters of the Additional Protection section, you allow can control when and where to create additional copies and how long these copies are retained.

Important: Implement different schedules for the main policy, replication policy, and the creation of additional copies. The additional copies should be created after the operational backups and replication complete.

Additional copies to a “hot” storage tier often are kept for a short term. Configure a short retention or select the **Same retention as source selection** option.

Additional copies to a “cold” storage tier (Archives) are usually kept for a long time. Select a retention of months or years. Some Object Storage providers charge an extra cost if data is deleted from their archive tier before the minimum lifetime is reached.

The following SLA parameters control the source of a copy operation:

- Backup Policy Destination: The source for the copy or archive operation is the target site that is defined in the Backup Policy section.
- Replication Policy Destination: The source for the copy or archive operation is the target site that is defined in the Replication Policy section. This option is available only when Backup Storage Replication is enabled.

The following SLA parameters control the Destination of a copy operation:

- Object Storage Servers: The target for the copy or archive operation is an Object Storage provider.
- Repository Servers: The target for the copy or archive operation is a IBM Spectrum Protect server.
- Target: The Object Storage system or repository server to which you want to copy or archive data. For archive operations, only Object Storage targets that include a defined archive bucket are shown in this list.

The example in Figure 15-40 on page 443 shows how to configure a SLA Policy Backup using the IBM Spectrum Protect Plus GUI. Select **Manage Protection** → **Policy Overview** → **Add SLA Policy**, as shown in Figure 15-39.

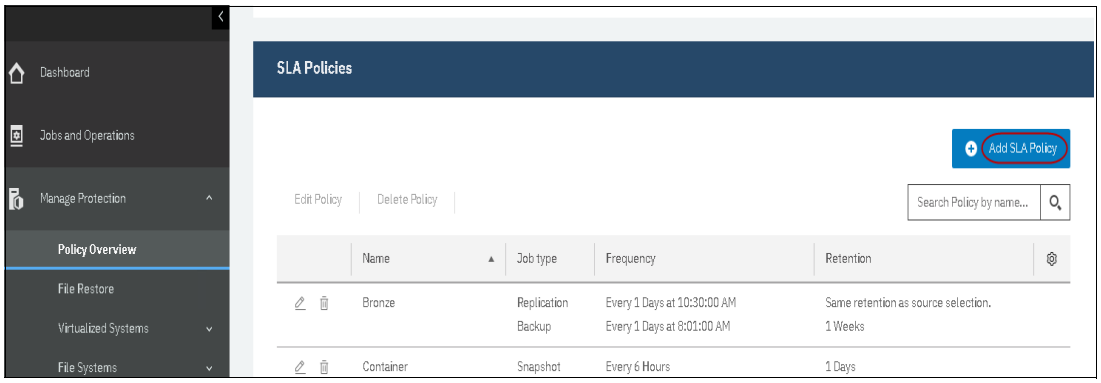

Figure 15-39 Adding an SLA Policy Backup

The example in Figure 15-40 on page 443 shows creating a daily incremental copy of data from the primary vSnap to a IBM Spectrum Protect server container storage pool (by using the same retention that is used for the primary backup) and a monthly full copy to a tape pool at the same IBM Spectrum Protect server (using a retention of 3 months).

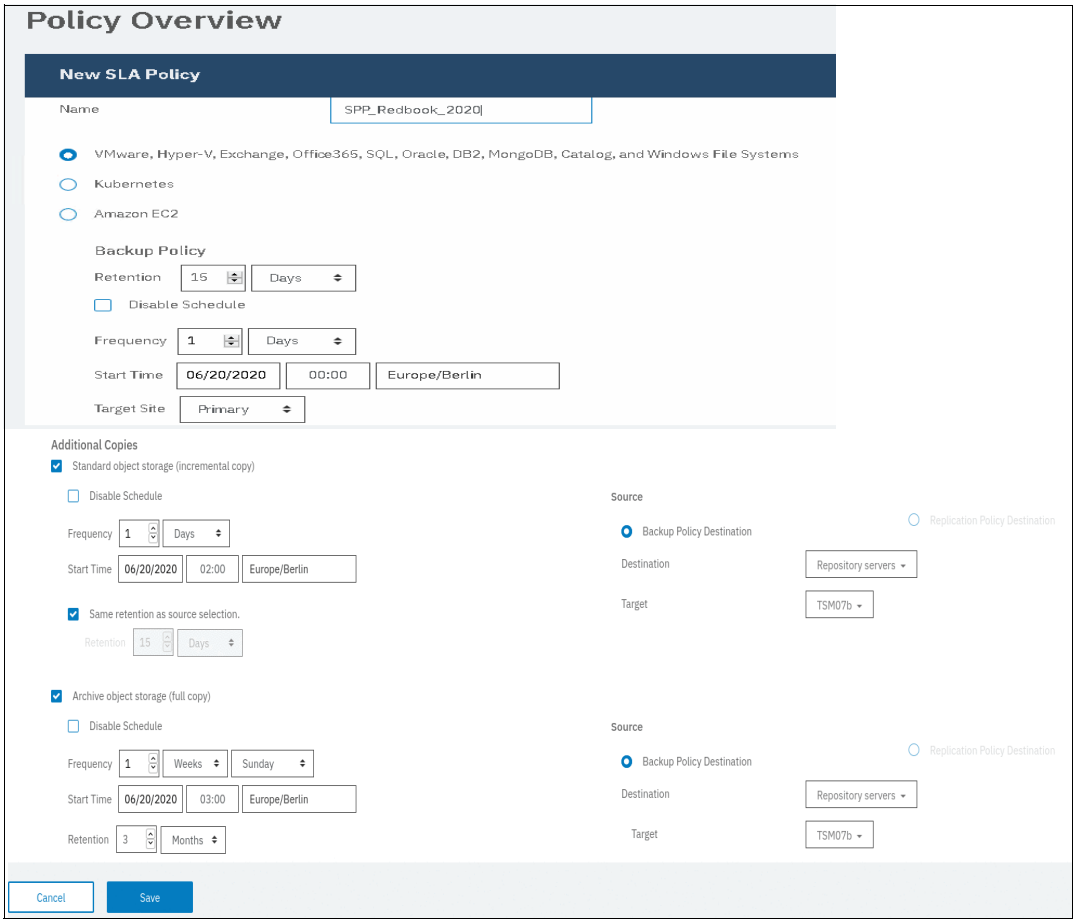

Figure 15-40 Configuring an SLA for additional copies to a repository server example

For more information about how to create an SLA policy, see IBM Knowledge Center.

15.4.6 vSnap commands for Object Storage data

After the Object Storage is added (see Figure 15-36 on page 438) in the IBM Spectrum Protect Plus GUI, the Object Storage server is listed with details about its standard Object Storage bucket, as shown in Figure 15-41 and Figure 15-42 on page 444.

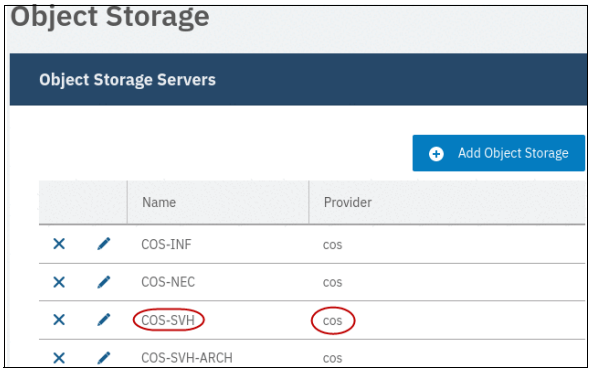

Figure 15-41 Object Storage Servers created

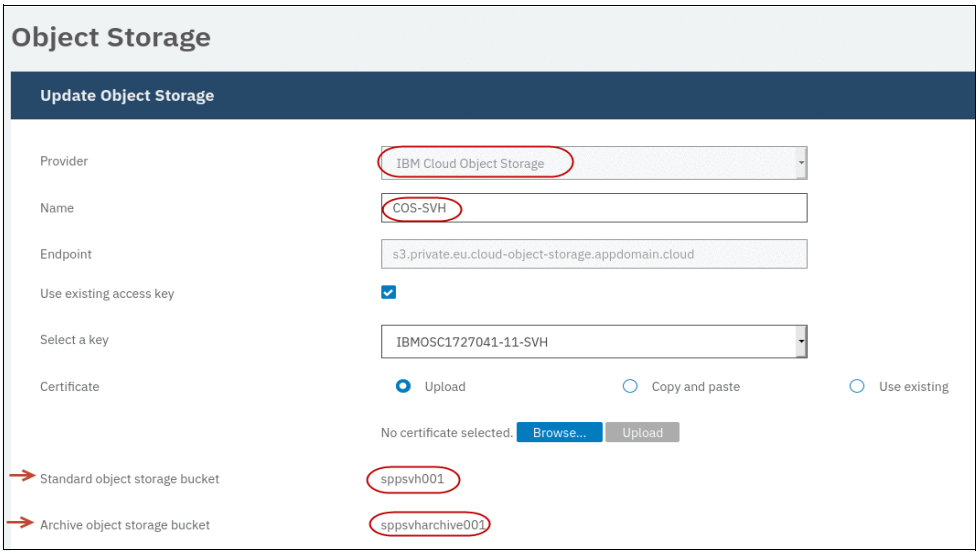

Provider: IBM Cloud Object Storage
Name: COS-SVH
Endpoint: s3.private.eu.cloud-object-storage.appdomain.cloud
Use existing access key: ☑
Select a key: IBMOSC1727041-11-SVH
Certificate: ● Upload ○ Copy and paste ○ Use existing
No certificate selected.

Standard object storage bucket: sppsvh001
Archive object storage bucket: sppsvharchive001

Figure 15-42 Object Storage Servers details

Example 15-11 shows the use of the vSnap CLI command **`vsnap cloud partner show`** to list the Object Storage partner server and its bucket `sppsvh001`.

Example 15-11 List Object Storage partner server on the vSnap

```
[serveradmin@t4-spp-vsnap ~]$ vsnap cloud partner show
ID                               | PARTNER TYPE | MGMT ADDRESS                                       | API PORT | PROVIDER | BUCKET
----------------------------------------------------------------------------------------------------------------------------------
-
de1a9c3882094c71986c06e7bad18bd7 | cloud        | s3.private.eu.cloud-object-storage.appdomain.cloud | 80       | cos      | sppsvh001
```

As shown in Example 15-11, an Object Storage partnership is created between `t4-spp-vsnap` and the Object Storage bucket `sppsvh001`. When a Copy Object Storage task of an SLA policy is started, the following actions occur on the source and target Object Storage server:

1. A new Object Storage target volume is created on the vSnap server.
2. A relationship for the source and target volume is created on the source vSnap and target Object Storage.
3. An empty snapshot is created on the source volume.
4. A copy Object Storage session is created on the source vSnap server.
5. The target snapshot is updated on the source vSnap server after the copy Object Storage session is completed.

As shown in Example 15-12, the copy Object Storage job creates a target volume (ID `9d01a010593441669d6c055554ec420c`) on the target vSnap server `t4-spp-vsnap`. The vSnap CLI command **`vsnap cloud volume show`** lists it.

Example 15-12 Copy Object Storage target volume is created on the target vSnap server

```
[serveradmin@t4-spp-vsnap ~]$ vsnap cloud volume show
ID                               | PARTNER ID                       | TYPE       | IS CLONE | NAME                                        | TAGS
-------------------------------------------------------------------------------------------------------------------------------------------
9d01a010593441669d6c055554ec420c | de1a9c3882094c71986c06e7bad18bd7 | filesystem | No       | spp_1007_2112_16efaed8584__group0_95_       | N/A
8358edb231c74a74b37fdf1a928e33ad | de1a9c3882094c71986c06e7bad18bd7 | filesystem | No       | spp_1008_2114_16efc8e5b4c__group0_95_       | N/A
a76874dea062484fadcda57202e40d46 | de1a9c3882094c71986c06e7bad18bd7 | lun        | No       | spp_1014_2107_16FCA57CF92
```

Next, the Object Storage relationship between both volumes is configured on the vSnap server. The vSnap CLI command **`vsnap cloud relationship show`** lists these relationships as shown in Example 15-13 for the source vSnap server `t4-spp-vsnap` and the Remote Copy Object Storage volume target.

Example 15-13 vSnap relationship on the source server created

```
[serveradmin@t4-spp-vsnap ~]$ vsnap cloud relationship show
ID                               | PARTNER ID                       | PARTNER TYPE | LAST SYNC | LOCAL VOL                                               | REMOTE
VOL
------------------------------------------------------------------------------------------------------------------------------------------------------------------
---------------------------------------------
e09e5526309f47acaa2d3a8585cbb57b | de1a9c3882094c71986c06e7bad18bd7 | cloud        | COMPLETED | spp_1007_2112_16efaed8584__group0_95_                   |
spp_1007_2112_16efaed8584__group0_95_

[serveradmin@t4-spp-vsnap ~]$ vsnap cloud relationship show --id e09e5526309f47acaa2d3a8585cbb57b

ID: e09e5526309f47acaa2d3a8585cbb57b
PARTNER ID: de1a9c3882094c71986c06e7bad18bd7
PARTNER TYPE: cloud
LOCAL POOL ID: 1
LOCAL VOL ID: 5
REMOTE VOL ID: 0b7a2b6c04aa4bf2b3e5bc10b8327663
LOCAL VOL NAME: spp_1007_2112_16efaed8584__group0_95_
REMOTE VOL NAME: spp_1007_2112_16efaed8584__group0_95_
LAST SYNC STATUS: COMPLETED
LAST SYCNED SNAP ID: 4720
LAST ATTEMPT SNAP ID: 4778
LAST SYCNED DB CHECKSUM: 43dc837b24356f0908c516f10be646b3
LAST SYNCED BIN CHECKSUM: f2737219118a3695eee632c8a58c0945
CREATED: 2020-06-21 08:29:16 UTC
UPDATED: 2020-06-21 08:29:16 UTC
```

A empty Object Storage Snapshot is created by running the **`vsnap cloud snapshot show`** command, as shown in Example 15-14. The details of the snapshot are available by running the **`vsnap cloud snapshot show --id 7196997217654727064 --partner_id de1a9c3882094c71986c06e7bad18bd7`** command.

Example 15-14 Empty Object Storage snapshot created

```
[serveradmin@t4-spp-vsnap ~]$ vsnap cloud snapshot show
ID                   | PARTNER ID                       | PARENT ID                        | CREATED                 | NAME
| RESTORABLE
------------------------------------------------------------------------------------------------------------------------------------------------------------------
----------------------
7196997217654727064  | de1a9c3882094c71986c06e7bad18bd7 | d0332652c53045599ce2f98f74a86fc8 | 2020-06-21 00:08:07 UTC |
spp_1008_2435_172d5c91288_SVH-VM-BACKUP-LNX         | Yes

[serveradmin@t4-spp-vsnap ~]# vsnap cloud snapshot show --id 7196997217654727064 --partner_id de1a9c3882094c71986c06e7bad18bd7


ID: 7196997217654727064
VERSION: 7196997217654727064
VOLUME: d0332652c53045599ce2f98f74a86fc8
NAME: spp_1008_2435_172d5c91288_SVH-VM-BACKUP-LNX
PARTNER ID: de1a9c3882094c71986c06e7bad18bd7
PARTNER TYPE: cloud
CLOUD OBJECT COUNT: No
USED SPACE: 0.00KB
CREATED: 2020-06-21 00:08:07 UTC
UPDATED: 2020-06-21 00:08:07 UTC
RESTORABLE: Yes
FILESYSTEM TYPE: N/A
PATHS: N/A
SOURCE SIZE: N/A
```

When the copy Object Storage is running, a session is created for the Object Storage on the source vSnap server. As shown in Example 15-15, the **vsnap cloud session show** command is used to list the copy Object Storage session on the source server `t4-spp-vsnap`. The status field indicates if it is an `Active` session or a `Completed` session.

Example 15-15 Copy Object Storage session using the vSnap CLI

```
[serveradmin@t4-spp-vsnap ~]# vsnap cloud session show
ID                                | ACTION | RELATIONSHIP                     | LOCAL SNAP | VERSION             | STATUS    | SENT      | STARTED
| ENDED
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------
----------------------------
1b8d1f80f9124e318b85192b74942e38 | upload | 1d90b09e27c14bc8b7d0756ff051c0d9 | 4787       | 7196997217654727064 | ACTIVE    | 55.13MB   | 2020-06-21 07:27:42
UTC | N/A


[serveradmin@t4-spp-vsnap ~]# vsnap cloud session show --id 1b8d1f80f9124e318b85192b74942e38

ID: 1b8d1f80f9124e318b85192b74942e38
PARTNER ID: de1a9c3882094c71986c06e7bad18bd7
PARTNER TYPE: cloud
RELATIONSHIP ID: 1d90b09e27c14bc8b7d0756ff051c0d9
LOCAL SNAP ID: 4787
LOCAL SNAP NAME: spp_1008_2174_193_172d436fe9b
REMOTE SNAP NAME: spp_1008_2435_172d5c91288_SVH-VM-BACKUP-LNX
PRIORITY: 50
STATUS: ACTIVE
SENT: 55.13MB
QUEUED: 2020-06-21 07:27:19 UTC
STARTED: 2020-06-21 07:27:42 UTC
ENDED: N/A
MESSAGE: Transferred 55.13MB of 105.62MB; 52% complete; Average throughput 15.3 MB/s
VERSION: 7196997217654727064
REMOTE VOL ID: d0332652c53045599ce2f98f74a86fc8
MAX OBJECT SIZE: 16777216
NUM OBJECTS WRITTEN: No
BUCKET ID: sppsvh001
ACTION: upload
REBASE: No
IS CANCELLED: No
LOCAL POOL ID: Yes
LOCAL VOL ID: 193
CLONE VOL ID: N/A
CLONE VOL NAME: N/A

[serveradmin@t4-spp-vsnap ~]# vsnap cloud session show
ID                                | ACTION | RELATIONSHIP                     | LOCAL SNAP | VERSION             | STATUS    | SENT      | STARTED
| ENDED
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------
----------------------------
1b8d1f80f9124e318b85192b74942e38 | upload | 1d90b09e27c14bc8b7d0756ff051c0d9 | 4787       | 7196997217654727064 | COMPLETED | 105.62MB  | 2020-06-21 07:27:42
UTC | 2020-06-21 07:39:08 UTC

[serveradmin@t4-spp-vsnap ~]# vsnap cloud session show --id 1b8d1f80f9124e318b85192b74942e38

ID: 1b8d1f80f9124e318b85192b74942e38
PARTNER ID: de1a9c3882094c71986c06e7bad18bd7
PARTNER TYPE: cloud
RELATIONSHIP ID: 1d90b09e27c14bc8b7d0756ff051c0d9
LOCAL SNAP ID: 4787
LOCAL SNAP NAME: spp_1008_2174_193_172d436fe9b
REMOTE SNAP NAME: spp_1008_2435_172d5c91288_SVH-VM-BACKUP-LNX
PRIORITY: 50
STATUS: COMPLETED
SENT: 105.62MB
QUEUED: 2020-06-21 07:27:19 UTC
STARTED: 2020-06-21 07:27:42 UTC
ENDED: 2020-06-21 07:39:08 UTC
MESSAGE: Completed
VERSION: 7196997217654727064
REMOTE VOL ID: d0332652c53045599ce2f98f74a86fc8
MAX OBJECT SIZE: 16777216
NUM OBJECTS WRITTEN: No
BUCKET ID: sppsvh001
ACTION: upload
REBASE: No
IS CANCELLED: No
LOCAL POOL ID: Yes
LOCAL VOL ID: 193
CLONE VOL ID: N/A
CLONE VOL NAME: N/A
```

15.5 Configuring a multi-site backup

All methods to create additional copies of data that were discussed in this chapter use a single SLA that controls the creation of the primary backup and all subsequent replication or copy operations to secondary storage. Therefore, all of these additional *copies* belong to the same point-in-time as the primary backup object. The *Recovery Point Objective* (RPO) of a backup object is always controlled by the parameters of the Main Policy in the Operational Protection section of a SLA.

You can improve (reduce) the RPO by assigning a VM, an application, or an IBM Spectrum Protect Plus catalog backup to more than one SLA in parallel. Similar to replication, this creates backups of the same entity (VM, application, or IBM Spectrum Protect Plus catalog backup) in two different locations (sites). The key difference is the fact that the primary data is being backed up twice, which results in two independent “primary” backup versions of the same entity.

Example 1: SLA with Replication

A VM is assigned to a single SLA with a main policy that runs a daily backup at 8 PM and a replication policy replicating data to another vSnap server at 8 AM.

The result is to have two copies of backup data with an RPO of 24 hours.

Example 2: Multi-Site/Multi-SLA Backup

A VM is assigned to two SLAs at the same time. One SLA runs a daily backup at 8 PM and the second SLA runs a daily backup at 8 AM. Both use the same retention, but different target sites and no replication.

The result is to have two copies of backup data with an RPO of 12 hours.

Note: A similar result can be achieved with a single, replication-enabled SLA by changing the schedule from daily to twice daily.

In addition to the fact that the RPO can be improved and independent copies can be created with Multi-SLA backup, it is important to understand that the method has some limitations and drawbacks compared to a replication approach:

- Increasing the backup frequency adds more load to the hypervisor (for example, VMware snapshots need to be created more often) or to the application (for example, because it is set to back up mode more often).
- Additional backups to a distant site might negatively impact performance because they cannot benefit from compression as it is being used when replicating backup data between two vSnap servers.
- Microsoft SQL server backups and IBM Db2 backups are limited to a single SLA when log backup is enabled. Having such applications assigned to two different SLAs at the same time with log backup enabled, this might lead to unpredictable results and require manual intervention.
- An Object storage Copy retains compression during transfer and storage, but not deduplication.
- Databases Applications, as the secondary copy, are not copying logs to any S3 storage (Object Storage or Repository Server).

- File indexing on Object Storage copies is not allowed. This decision was made to keep the catalog size manageable. The index is limited to 2 billion files.

Note: Object Storage repository contains a full version and not partial data. It means that it is possible to restore those copies from any vSnap, regardless of where they are located.

15.6 Creating incremental and full copies of backup data to an IBM Spectrum Protect server

In this example, we show you how to create additional copies of backup data from IBM Spectrum Protect Plus into an IBM Spectrum Protect server.

Note: In this example, we assume that we have a running version 10.1.6 IBM Spectrum Protect Plus server with a dedicated vSnap server and a running version 8.1.8 IBM Spectrum Protect server on a Linux operating system with a tape library attached. A directory container pool and a tape storage pool also exist.

Both types of additional copies (“hot” incremental copy to a directory container pool and a “cold” full copy/Archive to a tape storage pool) are configured.

The following tasks must be performed:

- Preparation of the Spectrum Protect server:
  - Define a ColdDataCache storage pool
  - Define a new Policy Domain of type ObjectDomain
  - Start an ObjectAgent and register an ObjectClient
- Register the IBM Spectrum Protect server as a Repository Server in the IBM Spectrum Protect Plus user interface
- Create an SLA that performs backups to a vSnap and creates regular additional copies to Spectrum Protect
- Run the SLA and observe the job results

15.6.1 Preparing the IBM Spectrum Protect server

In this section, we describe how to define a ColdDataCache storage pool, an Object Domain, and an Object Client.

Define a ColdDataCache storage pool

With IBM Spectrum Protect server version 8.1.8, a ColdDataCache storage pool type is used as an intermediate cache to buffer data that shall be copied from an ObjectClient (in this case, an IBM Spectrum Protect Plus vSnap server) to a IBM Spectrum Protect tape storage pool.

In this example, we use the IBM Spectrum Protect server administrative command line to define a ColdDataCache storage pool named SPP-TAPECACHE, which uses four directories in a local file system and a tape storage pool (TAPEPOOL) as the next storage pool in the hierarchy:

```
define stgpool SPP-TAPECACHE stgtype=colddatacache nextstgpool=TAPEPOOL \
directory=/tsm/stg/colddatacache1,/tsm/stg/colddatacache2,/tsm/stg/colddatacache3,
/tsm/stg/colddatacache4
```

```
ANR2200I Storage pool SPP-TAPECACHE defined (device class SPP-TAPECACHEDEVCLASS)
```

Note: When creating a storage pool of type ColdDataCache, a device-class is automatically created. The name of this device-class is <STGP_NAME>DEVCLASS.

Defining an ObjectDomain

After the ColdDataCache storage pool is created, an ObjectDomain must be defined. As part of the command that defines an ObjectDomain, the following storage pools must be specified:

- A StandardPool, which must be a directory- or cloud-container storage pool that is used for “hot” incremental copies of vSnap backup data.
- A ColdPool, which must be a ColdDataCache storage pool that is used for “cold” full copies/archives of vSnap backup data.

The following command creates an ObjectDomain named SPP and uses a directory-container storage pool that is named DIRECTORYPOOL for hot data and formerly defined ColdDataCache storage pool that is named SPP-TAPECACHE for cold data:

```
define OBJECTDomain SPP STANDARDPool=DIRECTORYPOOL COLDPool=SPP-TAPECACHE

ANR1500I Policy domain SPP defined.
ANR1510I Policy set STANDARD defined in policy domain SPP.
ANR1520I Management class STANDARD defined in policy domain SPP, set STANDARD.
ANR1520I Management class COLD defined in policy domain SPP, set STANDARD.
ANR1530I Backup copy group STANDARD defined in policy domain SPP, set STANDARD,
management class STANDARD.
ANR1530I Backup copy group STANDARD defined in policy domain SPP, set STANDARD,
management class COLD.
ANR1538I Default management class set to STANDARD for policy domain SPP, set
STANDARD.
ANR1514I Policy set STANDARD activated in policy domain SPP.
```

Note: When defining a policy domain of type ObjectDomain, a policy set and two management classes (STANDARD and COLD) are automatically defined and activated.

The management classes are configured in a way that the application (in this case SPP) controls the data lifecycle.

Example 15-16 shows the copy group parameters that have been automatically configured.

Example 15-16 Management classes and copy group parameters in an ObjectDomain

```
Protect: ESCC-SP-SRV>query copy group SPP active

Policy        Policy       Mgmt         Copy         Version-     Version-      Retain      Retain
Domain        Set Name     Class        Group         s Data       s Data       Extra        Only
Name                       Name         Name           Exists      Deleted     Versions     Version
---------     ---------    ---------    ---------    --------     --------    --------     -------
SPP           ACTIVE       COLD         STANDARD            1            1           0           0
SPP           ACTIVE       STANDARD     STANDARD            1            1           0           0
```

Defining an ObjectClient

The easiest way to register an ObjectClient to an IBM Spectrum Protect server is using the IBM Spectrum Protect Operations Center (OC).

Note: When you register an ObjectClient to a IBM Spectrum Protect server for the first time, the Operation Center wizard prompts you to create an ObjectAgent, which is enabling the IBM Spectrum Protect server to import data through the Amazon S3 protocol.

After the IBM Spectrum Protect Operations Center wizard prepares the configuration, the agent must be started from the IBM Spectrum Protect server operating system command line.

If the ObjectAgent is running (for example, because other ObjectClients use the IBM Spectrum Protect server), this task can be skipped.

Complete the following steps:

1. Log in to the IBM Spectrum Protect Operations Center that manages the IBM Spectrum Protect server and start the ObjectClient registration wizard (see Figure 15-43 on page 450) by selecting **Clients** → **Add Client** → **Object Client**. Then, click **Next**.

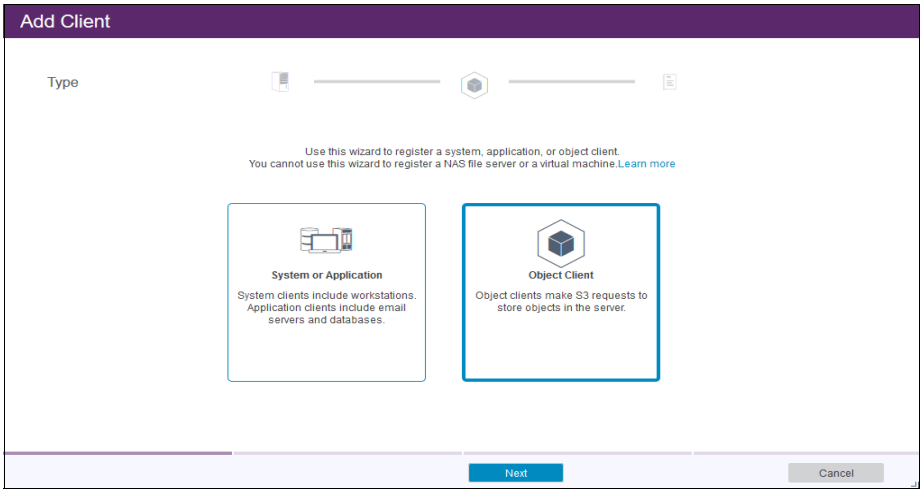

Figure 15-43 Add ObjectClient wizard in the IBM Spectrum Protect Operations Center

2. In the next dialog, choose the IBM Spectrum Protect server instance from a drop-down menu. We select the server **ESCC-SP-SRV**, which does not yet have an ObjectAgent defined. The wizard automatically prompts you to create an ObjectAgent and for a name and port number for the listener. In this example (as shown in Figure 15-44), we use S3OBJECTAGENT as the name and use the default tcp port 9000.

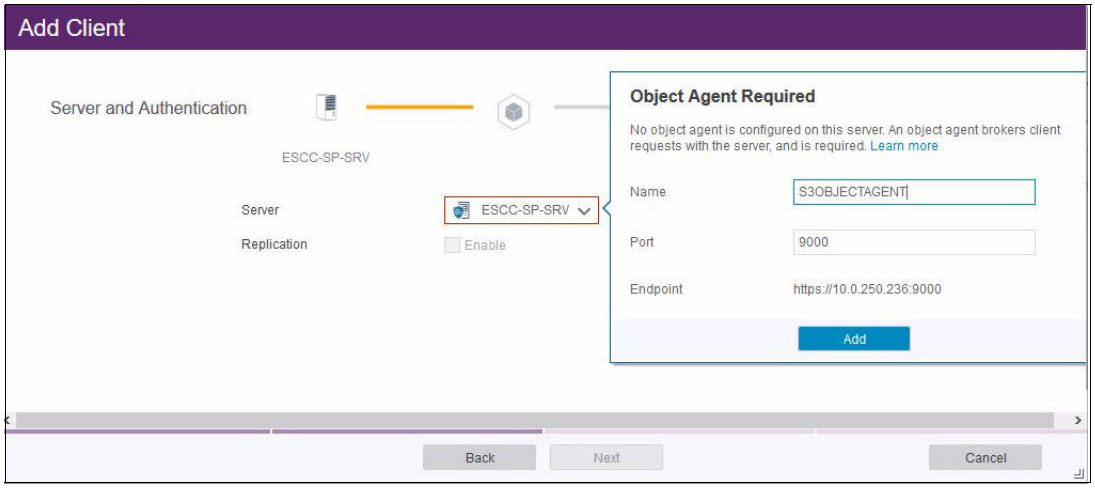

Figure 15-44 Specifying the ObjectAgent details

3. When the ObjectAgent configuration is complete, it must be started through the operating system command line. To obtain the necessary commands, click **View steps**, as shown in Figure 15-45.

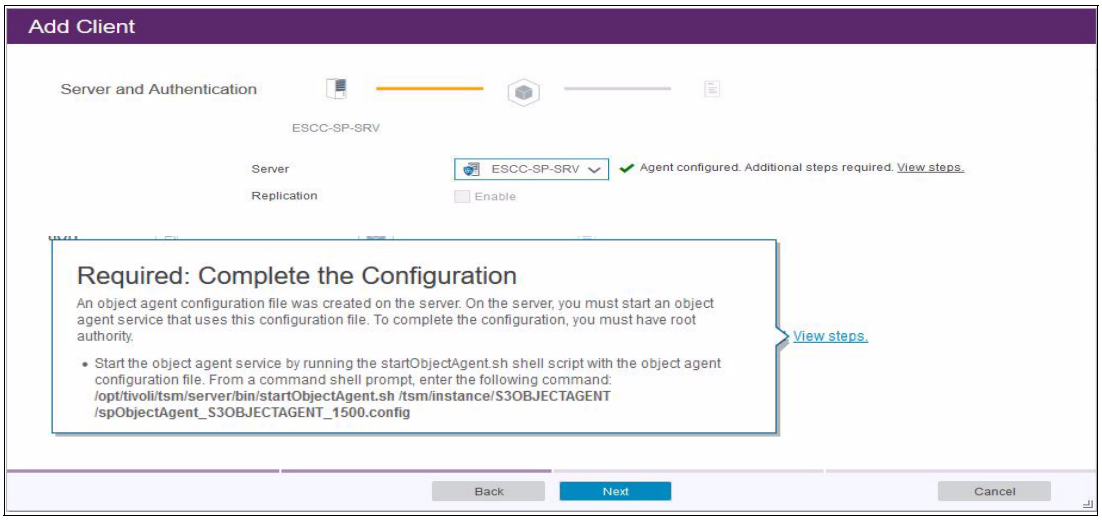

Figure 15-45 Obtaining the command to start the ObjectAgent by using the OS command line

4. Log in to the IBM Spectrum Protect server operating system and review the ObjectAgent configuration file that was created by the IBM Spectrum Protect Operations Center wizard, as shown in Example 15-17.

Example 15-17 ObjectAgent configuration file

```
[root@escc-sp-srv ~]# cat
/tsm/instance/S3OBJECTAGENT/spObjectAgent_S3OBJECTAGENT_1500.config
objagentexe=\"/opt/tivoli/tsm/server/bin/spObjectAgent\"
svcname="IBM Spectrum Protect object agent S3OBJECTAGENT:1500"
```

```
agentname=S3OBJECTAGENT
agentport=9000
serverhla=10.0.250.236
serverlla=1500
keystore="/tsm/instance/S3OBJECTAGENT/agentcert.p12"
pwdfile="/tsm/instance/S3OBJECTAGENT/agentcert.pwd"
serverkeypub="/tsm/instance/cert256.arm"
agentconfig=\"/tsm/instance/S3OBJECTAGENT/spObjectAgent_S3OBJECTAGENT_1500.config\"
servercertname="TSM Self-Signed Certificate"
```

5. Start the ObjectAgent by using the **startObjectAgent.sh** script, as shown in Example 15-18.

Example 15-18 ObjectAgent started

```
[root@escc-sp-srv ~]# /opt/tivoli/tsm/server/bin/startObjectAgent.sh
/tsm/instance/S3OBJECTAGENT/spObjectAgent_S3OBJECTAGENT_1500.config
Creating service spObjectAgent_S3OBJECTAGENT_1500 ...
Starting spObjectAgent ... Succeeded
```

6. After the ObjectAgent is started, the ObjectAgent registration continues in the IBM Spectrum Protect Operations Center. As shown in Figure 15-46, specify an ObjectClient name and contact details and then, click **Next**.

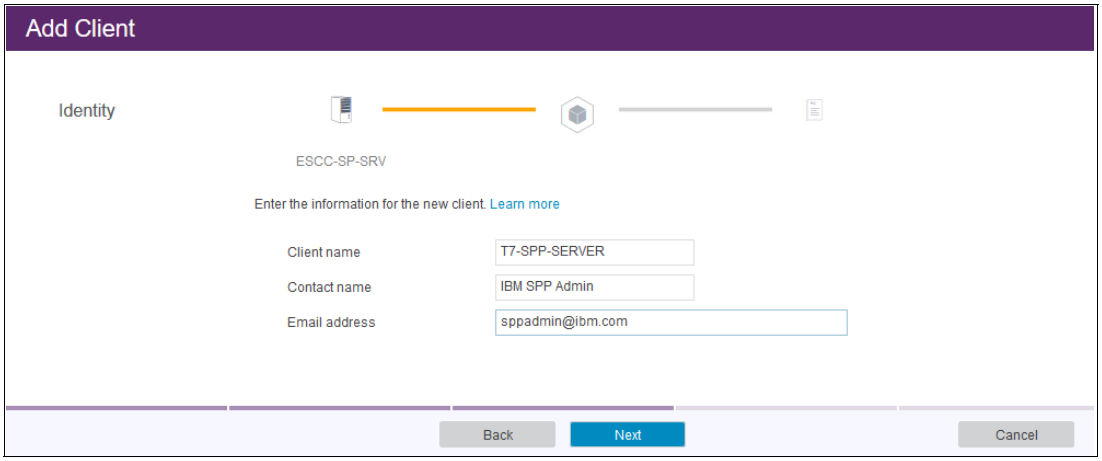

Figure 15-46 Specifying the ObjectClient name

7. In the next window, the policy domain must be specified. In this example (see Figure 15-47), we select the ObjectDomain **IBM Spectrum Protect Plus**, which was prepared previously.

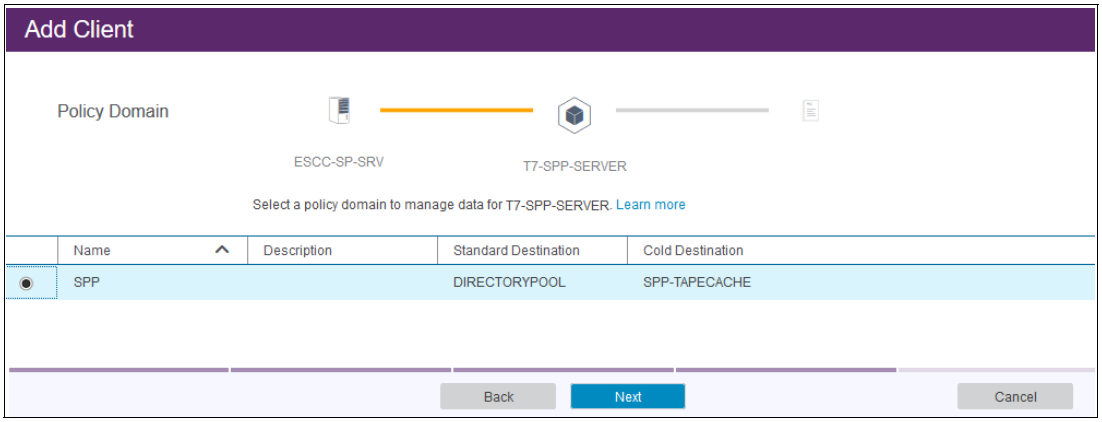

Figure 15-47 Select the ObjectDomain

> **Note:** The wizard should show the Standard- and Cold copy destinations for the ObjectDomain.

8. Click **Next** to configure the client and open the final window of the Add ObjectClient wizard, as shown in Figure 15-48 on page 454.

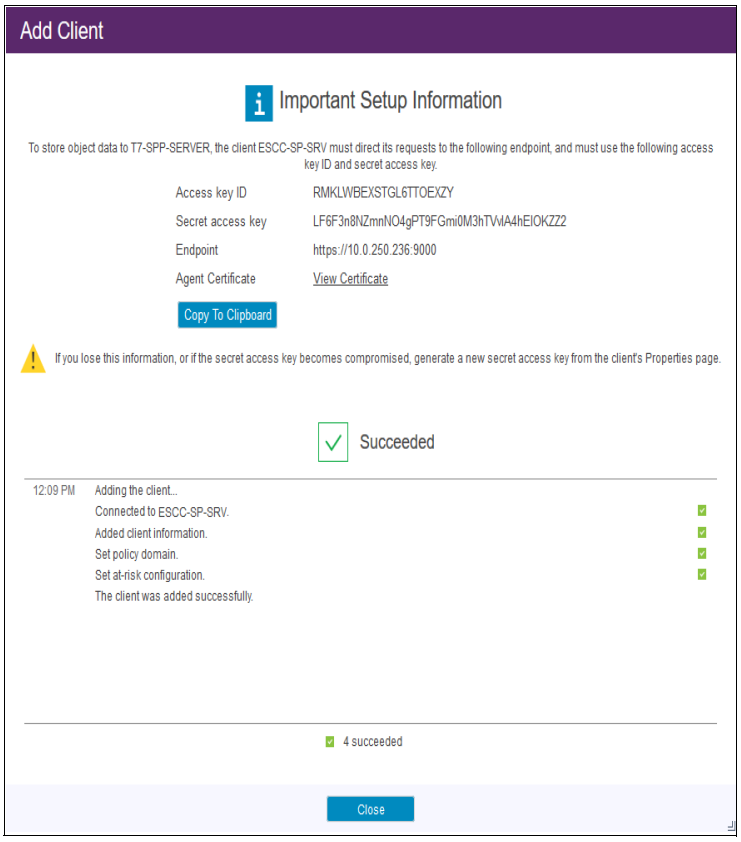

Figure 15-48 Add ObjectClient wizard summary information

Important: The summary window of the Add ObjectClient wizard contains all of the information that is needed later in the IBM Spectrum Protect Plus user interface to register a Repository Server:

- S3 access key id
- S3 secret access key
- ObjectAgent URL
- ObjectAgent certificate

Write down, copy, and store this information in a safe place because it is needed to complete the configuration procedure.

The preparation of the IBM Spectrum Protect server is now complete. As a next step, the IBM Spectrum Protect server must be registered in IBM Spectrum Protect Plus as a Repository server. This process is described next.

15.6.2 Registering an IBM Spectrum Protect server as Repository server in IBM Spectrum Protect Plus

Complete the following steps to register a Repository server in IBM Spectrum Protect Plus:

1. Log in to the IBM Spectrum Protect Plus user interface and select **System Configuration** → **Backup Storage** → **Repository Server**.
2. Specify a meaningful name, the FQDN, port number (default is 9000) and all the access credentials that were obtained during the final step of the IBM Spectrum Protect Operations Center Add ObjectClient configuration wizard, as shown in Figure 15-49.

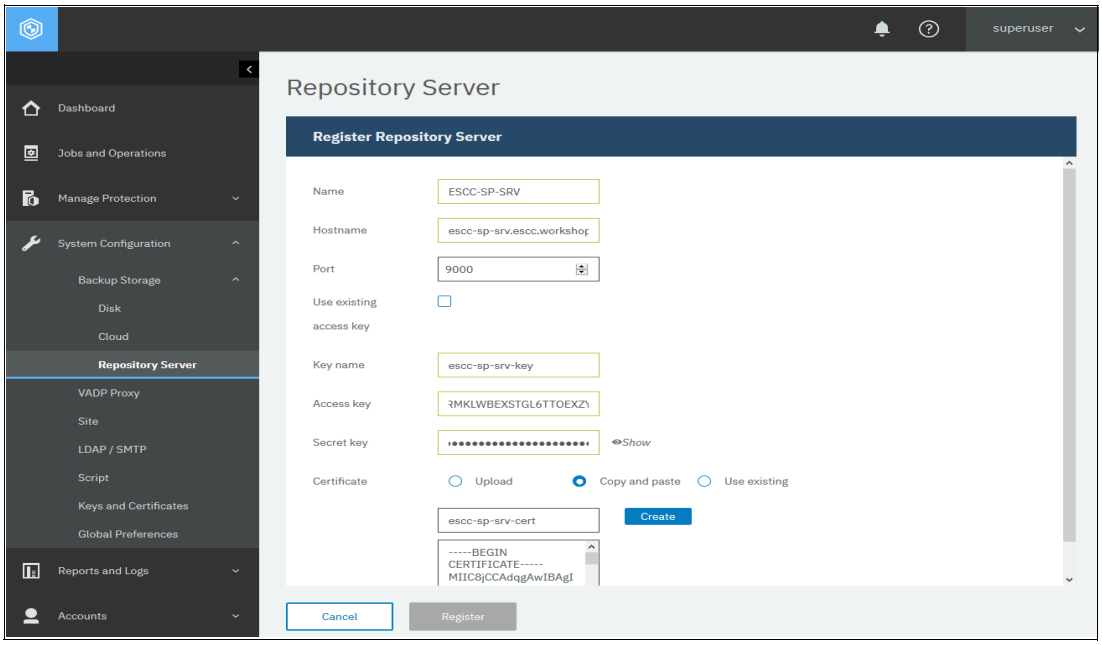

Figure 15-49 Adding a Repository server to IBM Spectrum Protect Plus

Tip: In the example that is shown in Figure 15-49, we paste the S3 access key and the content of the ObjectAgent certificate file directly into the corresponding dialog boxes.

As an alternative, these credentials can be registered to IBM Spectrum Protect Plus in advance by selecting the **System Configuration** → **Keys and Certificates** menu and then, selecting from a drop-down menu in the Add Repository Server dialog.

3. Click **Register** to start the registration process. When this process is complete, you receive a confirmation message, as shown in Figure 15-50.

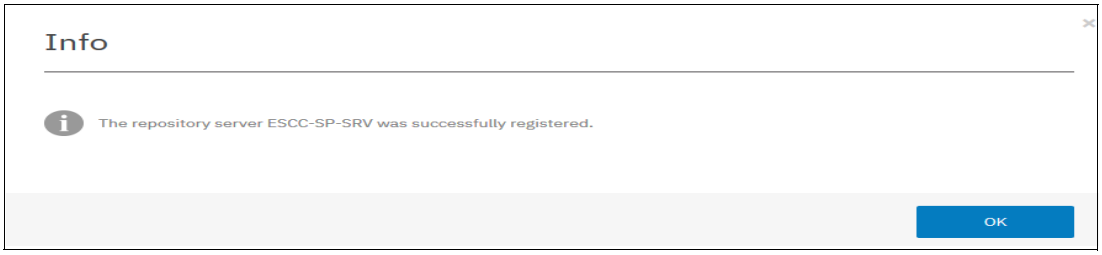

Figure 15-50 Repository Server configuration completed successfully

4. Click **OK** to complete the Add Repository server wizard and continue to define a SLA that uses this Repository server to create additional copies of backup data.

15.6.3 Creating an SLA that creates regular additional copies to IBM Spectrum Protect

In this example, we define an SLA with the following specifications:

- Name: Primary-with-Tape-Archive
- Type: VMware, Hyper-V, Exchange, Office365, SQL, Oracle, DB2, MongoDB, Catalog, and Windows File Systems
- Backup Policy:
  - Retention: 7 days
  - Frequency: daily
  - Target Site: Primary
- Replication Policy: No replication between vSnaps is used in this example
- Additional Copies - Standard Object Storage:
  - Retention: Same as source (7 days)
  - Frequency: daily
  - Source: Backup Policy Destination (means: data from original vSnap - not from a replica)
  - Destination: Repository Server
  - Target: ESCC-SP-SRV
- Additional Copies - Archive Object Storage:
  - Retention: 3 months
  - Frequency: Weekly, on Sundays
  - Source: Backup Policy Destination (means: data from original vSnap - not from a replica)
  - Destination: Repository Server
  - Target: ESCC-SP-SRV

In the IBM Spectrum Protect Plus user interface, click **Manage Protection** → **Policy Overview** → **Add SLA Policy** to define the new SLA. Figure 15-51 shows the configuration of the Main Policy (backup to vSnap for operational protection).

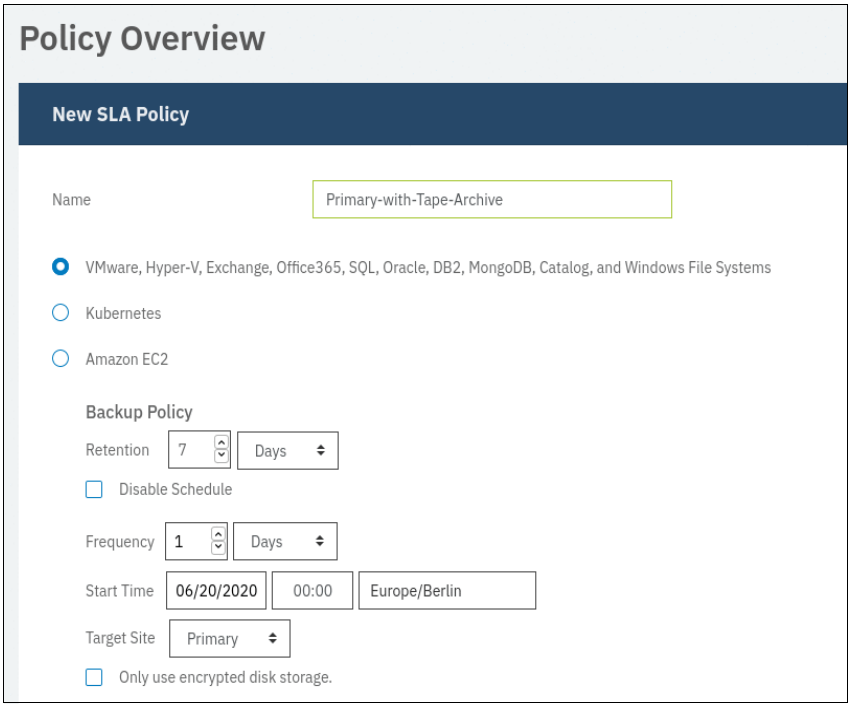

Figure 15-51 New SLA Backup Policy

Figure 15-52 shows the configuration for the additional copies to IBM Spectrum Protect.

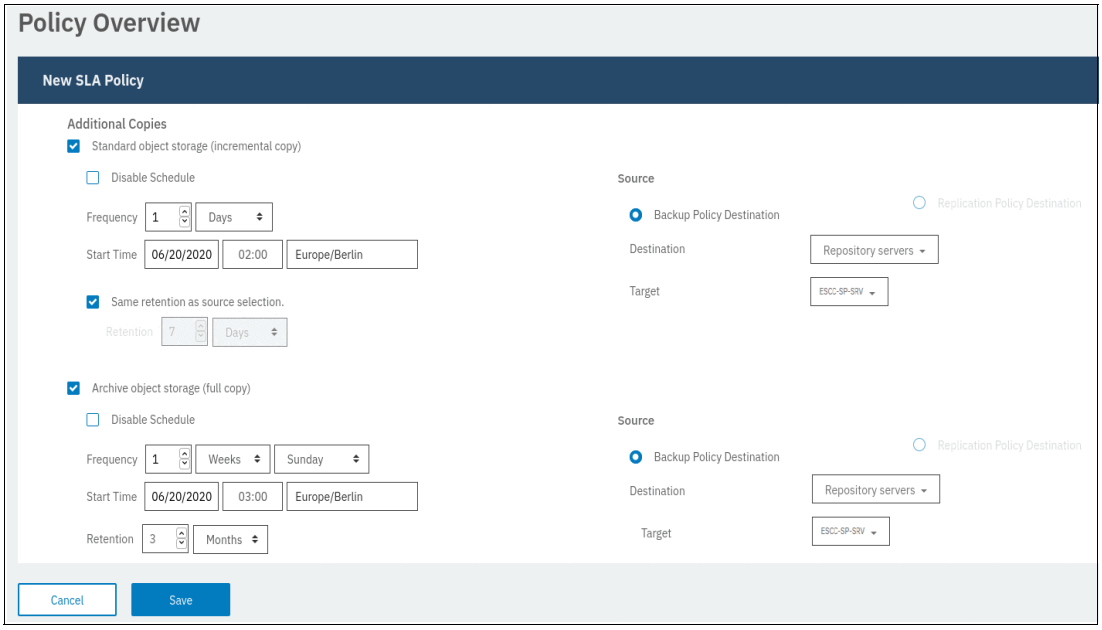

Figure 15-52 New SLA Additional Copies to a Repository Server

Tip: In this example, we backup data to a vSnap server in the primary site and create additional copies to a IBM Spectrum Protect server from the same vSnap.

Alternatively, it is possible to deploy an additional vSnap server in a different site (for example, “Secondary”), replicate data from the primary site’s vSnap to the secondary site and to create the additional copies to Spectrum Protect from the secondary vSnap.

To configure this replication in the SLA, the SLA parameter “Source” must be changed from “Main Policy Destination” to “Replication Policy Destination”.

Click **Save** to finish the SLA definition.

15.6.4 Running the SLA and observe the job results

To validate that the configuration was successful, you do not need to wait until the configured start time. We assign the “Primary-with-Tape-Archive” SLA to a VM and then run it manually.

Because the SLA consists of three sections (“Backup to vSnap”, “Copy”, and “Archive”) we run it three times in a sequential manner. Complete the following steps:

1. Create a backup of the VM into a vSnap, as shown in Figure 15-53.

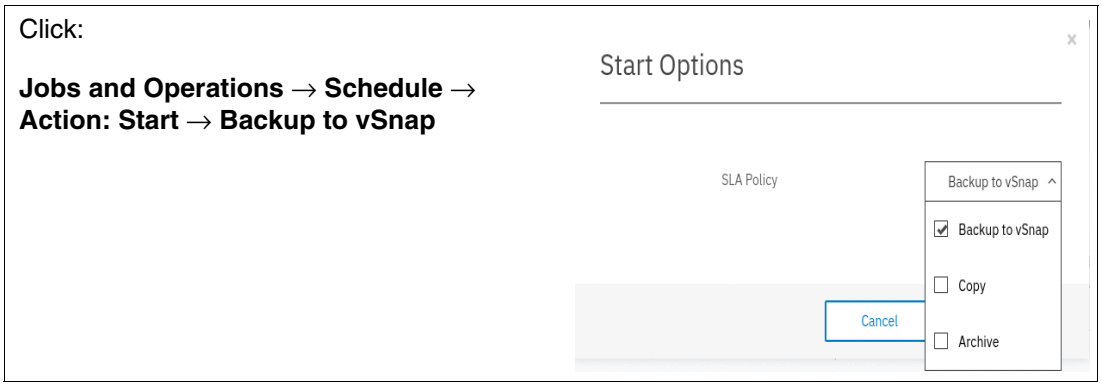

Figure 15-53 Backup of a VM to vSnap

2. Wait until the backup job is complete. As shown in Figure 15-54, the data is copied from the vSnap to the directory container storage pool on the IBM Spectrum Protect server. This full copy is used for the first run, but is incremental for all subsequent job runs.

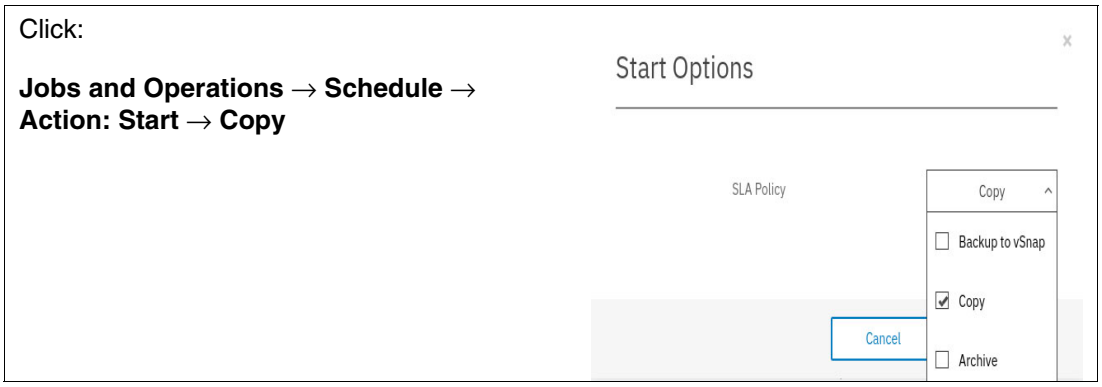

Figure 15-54 Copying data from vSnap to “Hot” Object Storage

3. Wait until the copy backup job is complete. As a next step, as shown in Figure 15-55, the data is copied from the vSnap to the tape storage pool on the IBM Spectrum Protect server. This copy is always a full copy.

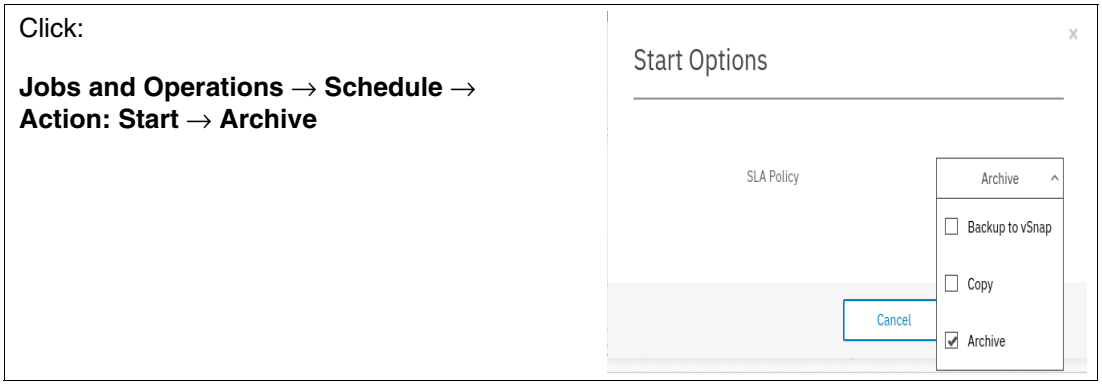

Figure 15-55 Copying data from vSnap to "Cold" Tape storage

4. Confirm that the backup and copy jobs completed successfully. Figure 15-56 shows the different job logs.

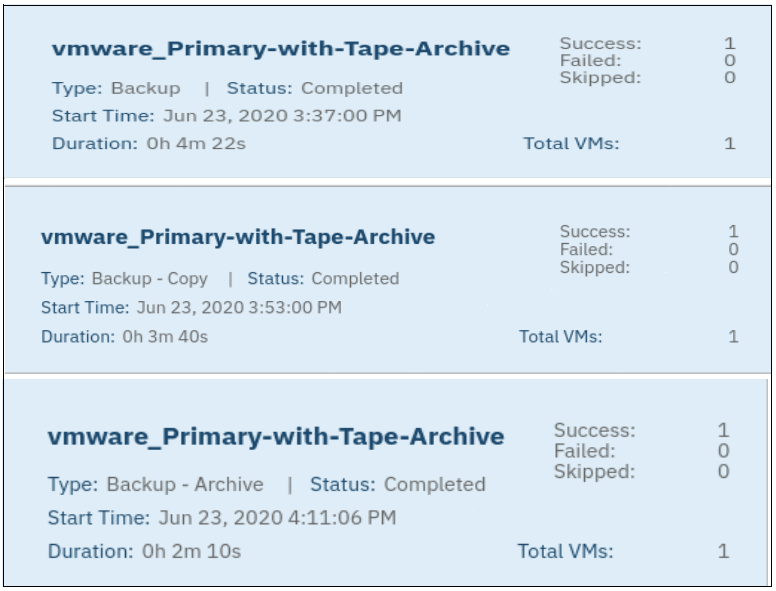

Figure 15-56 Primary-with-Tape-Archive SLA job logs

Note: IBM Spectrum Protect Plus Copy Object Storage is set to fail if nothing is available to copy. The SLA job status can be one of three types: Completed, Partial, or Failed. Completed status is when a recovery point exists for this SLA job that runs. If no recovery points exist for that job, it is marked with a Partial or Failed status.

16

REST API

This chapter provides an introduction into the REST APIs and how IBM Spectrum Protect Plus leverages this technology.

We demonstrate the REST API documentation and how users can discover and use the API for monitoring, configuring, and administration tasks. We also show how to discover API calls and their parameters to build a customized Python script for user-started backups of a single virtual machine (VM) from command line by using REST API services.

Finally, the `sppclient` utility is introduced. It is a Python module that encapsulates REST API complexity and reuses code for reoccurring tasks. The `sppclient` utility features a set of more than 20 Python scripts that demonstrate how to realize workflows with REST instead of the use of the GUI.

This chapter includes the following topics:

16.1 REST API overview

A REST application programming interface (API) enables client applications to access and manage resources on a server by sending requests and receiving responses using the Hypertext Transfer Protocol (HTTP) or HTTPS protocol. The IBM Spectrum Protect Plus Representational State Transfer (REST) API provides access to the product resources by way of so-called endpoints and enables other applications to interact with the product by sending requests and receiving responses by using the HTTP or HTTPS protocols. You can extend the capabilities of the IBM Spectrum Protect Plus application by developing scripts that use the API.

In general, REST (also known as *RESTful*) APIs share the following characteristics:

- RESTful systems use stateless protocols. That means the server does not hold any record of previous interactions with clients. Every client request and interaction need to be handled based entirely on information that comes from the client.
- Stateless components can be easily redeployed after failures occurred and can be started.
- RESTful APIs are popular in web deployments, cloud computing, and micro services. RESTful APIs scale out well as they can be started, stopped, and restarted easily when needed because they do not need to preserve any client data or session states.
- RESTful clients send requests to a resource's URI and expects a response; for example, in HTML, XML, and JSON. While there are several response formats, JSON is the most widely adopted format. IBM Spectrum Protect Plus uses the JSON format for requests and responses over HTTP.
- The following REST operations are available for HTTP:
  - GET
  - HEAD
  - POST
  - PUT
  - PATCH
  - DELETE
  - CONNECT
  - OPTIONS
  - TRACE

IBM Spectrum Protect uses the GET, POST, PUT, and DELETE operations.

An example how a REST application programming interface (API) enables client applications is shown in Figure 16-1. How to access and manage resources on a server by sending requests and receiving responses by using the HTTP or HTTPS protocols is shown.

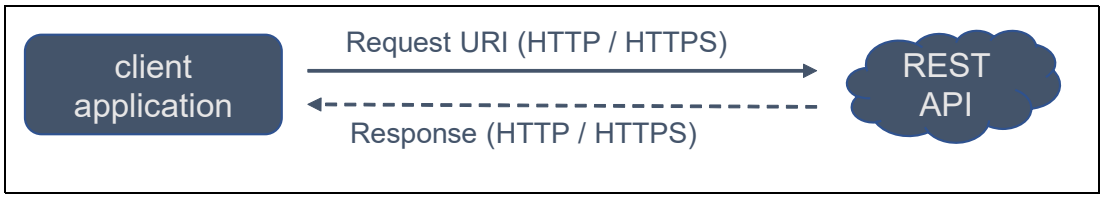

Figure 16-1 REST API client/server communication

16.2 IBM Spectrum Protect Plus REST API

IBM Spectrum Protect Plus uses a RESTful API to handle communications between the front-end (GUI) and the application server. This API can be also used by other client applications to realize tasks, such as automation, integration, configuration, and data collection.

By using the REST API, IBM Spectrum Protect Plus can integrate with automation tools, such as Jenkins, Puppet, and Ansible. Furthermore, any programming language that can make REST calls can be used (for example, Python, JavaScript, and PowerShell). The necessity to have a deeper look at the REST API mostly becomes relevant when the IBM Spectrum Protect Plus Environment is scaling out.

The GUI might not be always the best fit for every client for the following reasons:

- Increasing number of Hypervisors and VMs to configure and protect.
- Amount and variety of applications.
- Quantity of vSnap servers and VADP proxy servers
- The need to check in an automated way the SLA's status (for example, if it failed, or ended partial, or completed successfully).
- Monitoring the SPP environment in terms of resource utilization (CPU, RAM, vSnap pool capacity), job logs, and warnings or error messages.

16.2.1 REST API documentation

IBM Spectrum Protect Plus 10.1.6 introduces a REST API guide that provides documentation, useful examples, and scenarios to help you with the API. The guide is available at the following resources:

- In the product web GUI and in IBM Knowledge Center as a PDF.
- On the IBM Spectrum Protect and IBM Spectrum Protect Plus YouTube channel.
- An introduction to the REST API documentation is available on YouTube.

The REST API PDF describes the essentials of REST in detail. It provides guidance to session handling and four basic HTTP methods for applications to interact with: GET, POST, PUT, and DELETE. Additionally, it contains plenty of RESTful API examples that demonstrate interaction, basic functions such as retrieving a session ID, assigning a VM, and starting a job, JSON filtering, pagination, and a large subset of available REST API endpoints.

Also, the REST API documentation is integrated into the GUI for interaction with the API. It can be start from the drop-down menu in the upper right of the GUI, as shown in Figure 16-2.

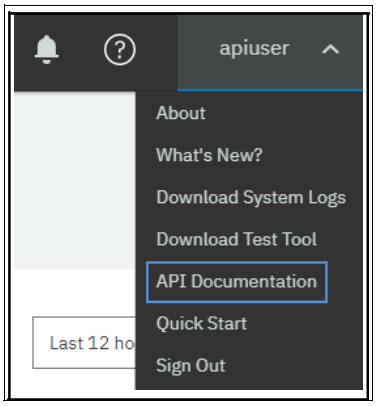

Figure 16-2 API Documentation

The reference describes the general information about the API design along with guidance and examples to script and interact with IBM Spectrum Protect Plus. In the reference, the topics are grouped and correspond to the product GUI menu structure. For example, the user finds the API functions related to a “vSnap” under **System Configuration** → **Backup Storage: Disk,** as shown in Figure 16-3.

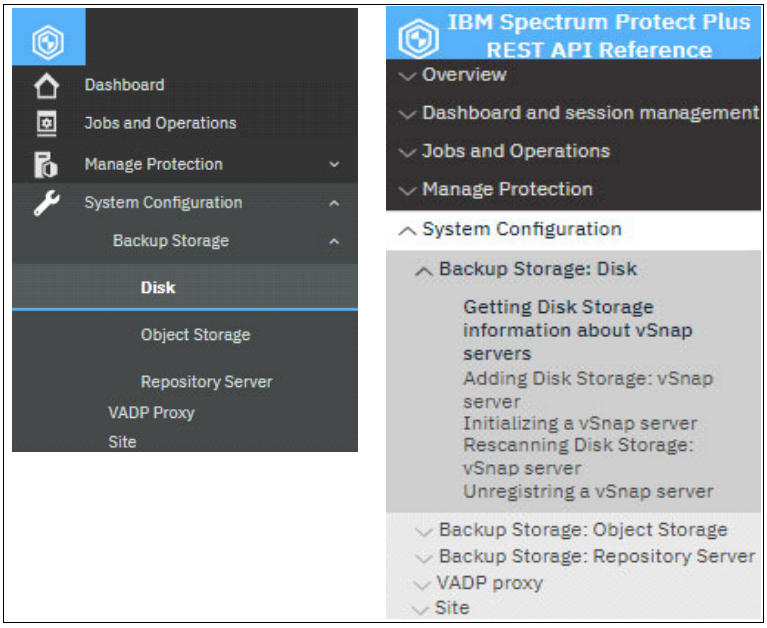

Figure 16-3 Using the REST API Reference

An example for available operations for the API endpoint site is shown in Figure 16-4.

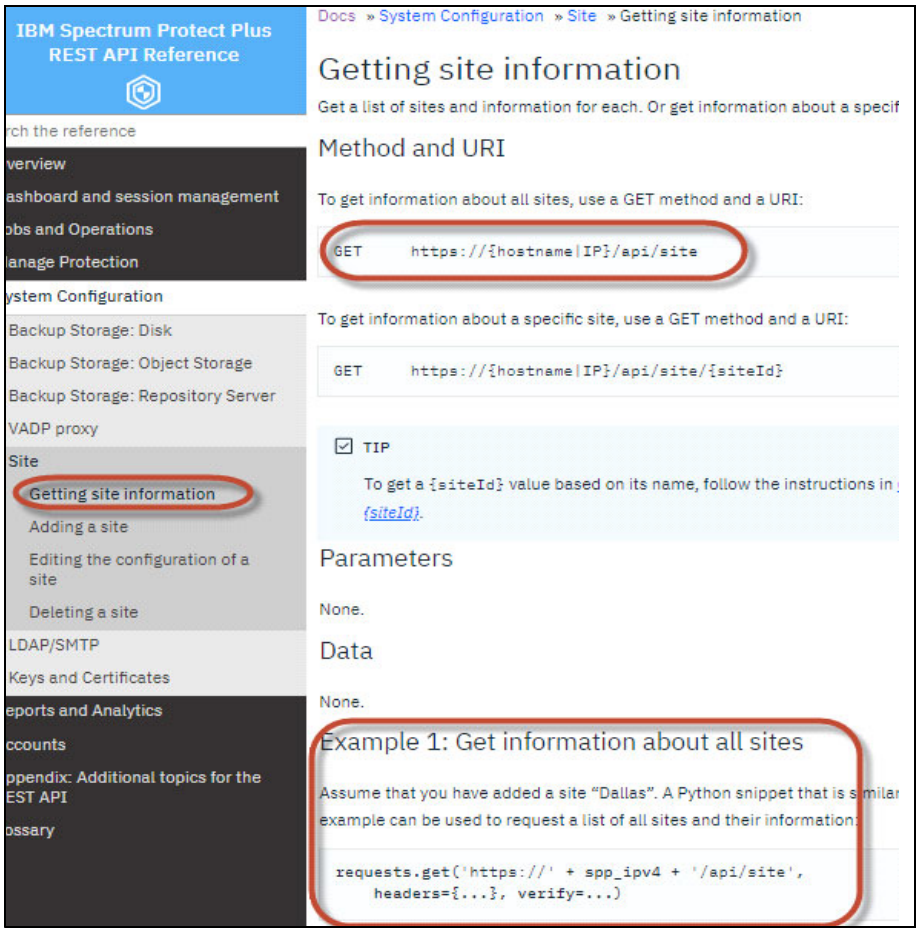

Figure 16-4 API endpoint site operations

This publication describes the process to identify more endpoints in 16.4.2, "Trace GUI REST operations by using Firefox" on page 474.

For most API operations, the API documentation provides the following resources:

- Screen captures of the equivalent GUI function with detailed explanations.
- Sample code in Python, which can be adopted by the user for integration into customized scripts.
- Information about how to:
  - Design the API request (GET, POST, UPDATE, and DELETE)
  - How to define parameters and filters, pagination, and so on

16.3 Discovering the REST API

This section demonstrates various methods to query the REST API. As a prerequisite for the different methods, we suggest using the following or similar tools:

- Firefox browser with another browser extension called *RESTclient*. For more information, see this web page.
- The `curl` tool and the tool `jq`. The `jq` tool is a lightweight and flexible command-line JSON processor that provides formatting, filtering, mapping, and transformation functions. For more information about the jq tool, see this web page.
- Python 3.x for the operating system of your choice. The sppclient modules are written in Python language and require extra modules and the Python interpreter in version 3 or higher.

During exploration of the REST API endpoint that is called *site*, we demonstrate the use of several different methods. However, all methods are following the sequence that is shown in Figure 16-5:

1. The client utility (for example, curl, Python libraries) connects to the API and transmits username and password in an HTTP header. In the case of valid credentials, the API responds with a unique authentication token. The IBM Spectrum Protect Plus API expects an authentication token as the `sessionid` attribute in all other API requests.
2. The client performs a POST operation.
3. The client performs a GET operation on the API endpoint `site`. Although a GET operation retrieves information only from the API with no changes, a POST or PUT operation creates or updates information by using the endpoint with data that must be specified in the request body. In step 2, 3, and 4, the `sessionid` is used as authentication token.
4. The session ends and the authentication token is invalidated.

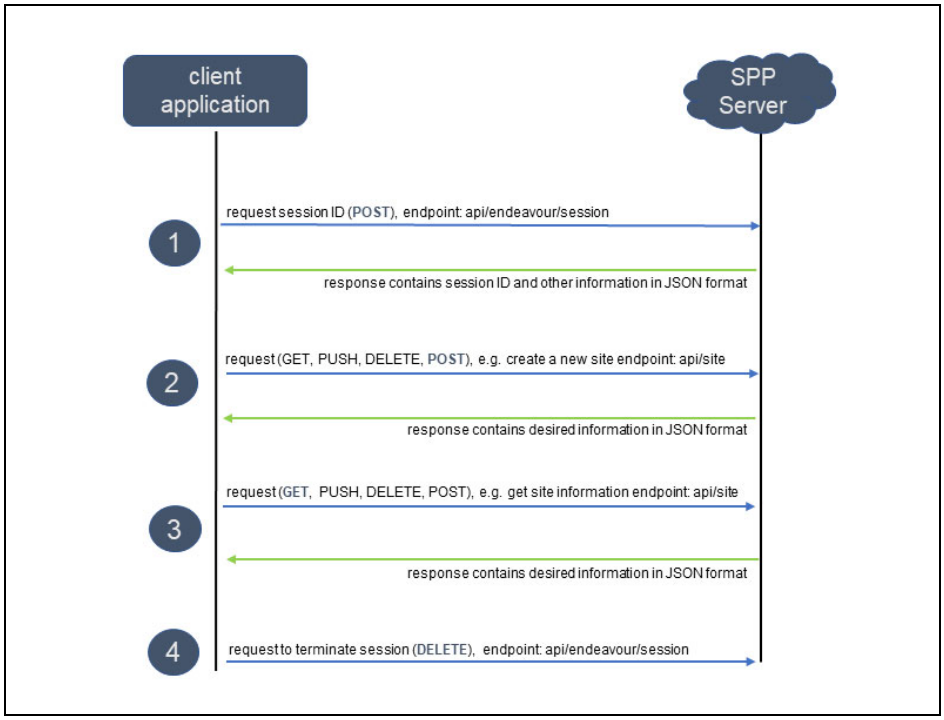

Figure 16-5 REST API call to endpoint site, including request of authentication token

Next, we describe how to realize the sequence in Figure 16-5 on page 466 by using the following methods:

- CURL
- Firefox RESTclient

16.3.1 CURL

We demonstrate the steps that are required to perform the steps that are shown in Figure 16-5 on page 466 by using the command-line utilities `curl` and `jq`. This simple example demonstrates the use of the REST operations POST, GET, and DELETE.

Log in and retrieve a session ID

The initial step is to establish a connection to the API and pass a username and password in the HTTP header. This POST operation is sent to the API endpoint `/api/endeavour/session`.

After successful authentication, the HTTP status code equals 200 and the response is in JSON format. An excerpt of the JSON response message is shown in Example 16-1. The most important information from this response is the value of the key `sessionid`.

Example 16-1 Retrieve login information with sessionid from the API with curl (truncated)

```
command: curl -sS -k -X POST -H 'Accept:application/json' -H 'Content-type:application/json' --user
"restapiuser:pass4AP!" https://spphost/api/endeavour/session

response: {"sessionid":"49531dd8345c417aa7d47b4179ecc03a","user":{"links":{"self":{"rel":"self"," href":
"https://spphost/api/security/user/1001"},"up":{"rel":"up","href":"https://spphost/api/security/user"},"changeMetadat
a":{"rel":"action","href":"https://spphost/api/endeavour/session/metadata/NATIVE_USER:1000:restapiuser?action=changeM
etadata"},"changePassword":{"rel":"action","href":"https://spphost/api/security/user/1001?action=changePassword"}},"n
ame":"restapiuser","type":"NATIVE_USER","typeDisplayName":"Native
User","tenantId":1000,"loginCount":18,"lastLogin":1560864770707,"failedLogin":0,"lastFailedLogin":null,"lastPasswordU
pdate":1560778818736,"passwordAge":0,"passwordExpiresAt":0,"passwordExpired":false,"accountDisabled":...
```

This authentication token is used for more communication with the API until the token expires or the user invalidated the token by sending a DELETE command to the endpoint. The response message contains information that is related to the user account and so-called Hypermedia as the Engine of Application State (HATEOAS*)* components. HATEOAS is a component of a REST architecture that enables the client to interact with the server and discover information dynamically though hypermedia or response links.

The client does not require up front knowledge of methods of resource interaction. Example 16-1 shows such a response links to change a user's password (see the `changePassword` response link in Example 16-2). The response from the API endpoint contains more links or API endpoints that are related to session or user management.

Example 16-2 Retrieve a sessionid from the API with curl and Python pretty-print formatting

```
command: curl -sS -k -X POST -H 'Accept:application/json' -H 'Content-type:application/json' --user
"restapiuser:pass4AP!" https://spphost/api/endeavour/session | python -m json.tool

response:
{
    "sessionid": "49531dd8345c417aa7d47b4179ecc03a",
    "user": {
        "links": {
            "self": {
                "rel": "self",
                "href": "https://spphost/api/security/user/1001"
            },
            "up": {
                "rel": "up",
                "href": "https://spphost/api/security/user"
            },
            "changeMetadata": {
                "rel": "action",
```

```
                "href":
"https://spphost/api/endeavour/session/metadata/NATIVE_USER:1:restapiuser?action=changeMetadata"
            },
            "changePassword": {
                "rel": "action",
                "href": "https://spphost/api/security/user/1001?action=changePassword"
            }
        },
        "name": "restapiuser",
        "type": "NATIVE_USER",
        "typeDisplayName": "Native User",
        "tenantId": 1,
        "loginCount": 74,
        "lastLogin": 1561473242064,
        "failedLogin": 0,
...
}
```

In addition to the previous example, we use the Python module `json.tool` to format the JSON response and transform it from a single line string into a formatted output with indents and key-value pairs per line.

For advanced JSON parsing, filtering, and transformations of the JSON responses, we recommend the tool `jq`. For more information about this tool, see this web page.

The `jq` tool is a flexible command-line JSON processor and provides powerful options to extract values from the JSON responses. If you pipe the JSON output to the `jq` tool (see Example 16-3), the output is formatted similar to the pretty-print function of the `json.tool` Python module.

Example 16-3 Retrieve a sessionid from the API with curl and JQ for pretty-print formatting

```
command: curl -sS -k -X POST -H 'Accept:application/json' -H 'Content-type:application/json' --user
"restapiuser:pass4AP!" https://spphost/api/endeavour/session | jq .
result: the formatted result looks almost identical to the output of the Python formatting
```

The `jq` tool provides functions beyond the pretty-print formatting. It provides extra functions, such as customized formatting, key value transformations, calculations, and filtering. A simple filter on the key `sessionid` is used in Example 16-4 to display only the value of the wanted attribute that is reused for more authentication.

Example 16-4 Retrieve a sessionid from the API with curl and JQ for pretty-print formatting and filtering

```
command: curl -sS -k -X POST -H 'Accept:application/json' -H 'Content-type:application/json' --user
"restapiuser:pass4AP!" https://spphost/api/endeavour/session | jq '.sessionid'

response: "49531dd8345c417aa7d47b4179ecc03a
```

Listing and creating a site

As a first example, we create another site by using the REST API. Therefore, we specify the `sessionid` within the HTTP header and use the POST operation. Also, we must specify the site information with the **`data`** parameter (`-d` or `--data`). In Example 16-5 on page 469, we specify values for the properties `name`, `description`, and `defaultSite`. In addition, we declare the parameters `-iL` to include the HTTP response status code in the `curl` output.

Note: When specifying the parameters `-iL`, piping `curl` output to the `jq` tool fails. The extra parameters are modifying the output in a way that it is only a JSON formatted string and thus, the `jq` tool no longer can parse and format the output. Therefore, we specify the parameters `-iL` for demonstration purposes during dedicated examples in which we must check the HTTP response code. For example, response status code 201 indicates that a new resource was successfully created.

Example 16-5 Creating a site by running the REST command

```
command: curl -iL -sS -k -X POST -H 'Accept:application/json' -H 'Content-type:application/json' -H
'x-endeavour-sessionid:49531dd8345c417aa7d47b4179eccO3a' https://spphost/api/site -d '{"name" : "REST_site",
"description":"site created via REST", "defaultSite":"false"}'


response:
HTTP/1.1 201
X-Application-Context: zuul:443
Date: Fri, 28 Jun 2019 11:33:18 GMT
Content-Type: application/json;charset=UTF-8
Transfer-Encoding: chunked

{"links":{"self":{"rel":"self","href":"https://spphost/api/site/3106","hreflang":null,"media":null,"title":null,"type
":null,"deprecation":null},"up":{"rel":"up","href":"https://spphost/api/site","hreflang":null,"media":null,"title":nu
ll,"type":null,"deprecation":null},"edit":{"rel":"update","href":"https://spphost/api/site/3106","hreflang":null,"med
ia":null,"title":null,"type":null,"deprecation":null},"delete":{"rel":"delete","href":"https://spphost/api/site/3106"
,"hreflang":null,"media":null,"title":null,"type":null,"deprecation":null},"usedby":{"rel":"related","href":"http://l
ocalhost:8082/api/endeavour/association/resource/site/3106?action=listUsingResources","hreflang":null,"media":null,"t
itle":null,"type":null,"deprecation":null},"netapp":{"rel":"related","href":"https://spphost/api/netapp?siteId=3106",
"hreflang":null,"media":null,"title":null,"type":null,"deprecation":null},"vsphere":{"rel":"related","href":"https://
spphost/api/vsphere?siteId=3106","hreflang":null,"media":null,"title":null,"type":null,"deprecation":null},"resourcep
rovider":{"rel":"related","href":"https://spphost/api/site/3106/resourceprovider","hreflang":null,"media":null,"title
":null,"type":null,"deprecation":null}},"id":"3106","name":"REST_site3","description":"site created via
REST","defaultSite":false,"throttles":null,"demo":false,"rbacPath":"root:0/site:0/site:3106"}
```

To confirm that the site creation was successful, run a GET operation on the site endpoint, as shown in Example 16-6.

Example 16-6 Query sites by running REST command

```
command: curl -sS -k -X GET -H 'Accept:application/json' -H 'Content-type:application/json' -H
'x-endeavour-sessionid:49531dd8345c417aa7d47b4179eccO3a' https://spphost/api/site | jq -c '.sites[] | { name, id,
description, defaultSite}'

output:
{"name":"Secondary","id":"2000","description":"secondary site","defaultSite":false}
{"name":"Primary","id":"1000","description":"Primary Site","defaultSite":true}
{"name":"Demo","id":"3000","description":"Demo site","defaultSite":false}
{"name":"REST_site","id":"3106","description":"site created by REST command","defaultSite":false}
```

A site can be deleted by running a REST command without any confirmation, as shown in Example 16-7. The response status code 204 indicates that the resource was deleted successfully. The API method DELETE is run on endpoint **`/api/site/siteid`**, whereas the `siteid` must be replaced by the integer value of the corresponding site; for example, `3106` for the site that was created by using REST API.

Example 16-7 Deleting a site via REST command

```
command: curl -IL -sS -k -X DELETE -H 'Accept:application/json' -H 'Content-type:application/json' -H
'x-endeavour-sessionid:49531dd8345c417aa7d47b4179eccO3a' https://spphost/api/site/3106

response: HTTP/1.1 204
X-Application-Context: zuul:443
Date: Fri, 28 Jun 2019 11:07:48 GMT
```

Deleting the session

When the user completes the wanted tasks, it is recommended to invalidate the session by sending a DELETE request to the session endpoint, as shown in Example 16-8. This DELETE request is comparable to gracefully logging out from the GUI. This requirement is not a hard requirement and a login with the same user credentials returns a new `sessionid`, whether the user gracefully logged out previously.

Example 16-8 Delete the authentication token

```
command: curl -iL -sS -k -X DELETE -H 'Accept:application/json' -H 'Content-type:application/json' -H
'x-endeavour-sessionid:49531dd8345c417aa7d47b4179eccO3a' https://spphost/api/endeavour/session
```

```
response: HTTP/1.1 204
X-Application-Context: zuul:443
Date: Fri, 28 Jun 2019 11:39:19 GMT
```

16.3.2 Firefox RESTclient

We now demonstrate how to perform the steps that are shown in Figure 16-5 on page 466 to retrieve a `sessionid` and create a site with the Firefox browser add-on that is named RESTclient.

The advantage over the use of `curl` is that you can reuse frequently used queries and headers. The RESTclient generates the HTTP request that is based on several building blocks, such as Authentication header and Content type header. If you installed the RESTclient plug-in, you can start the tool within the browser by clicking the  icon.

Preparing the RESTclient

When you start the plug-in for the first time, you must specify the IBM Spectrum Protect Plus user credentials to form a base64 encoded authentication. Complete the following steps:

1. Select **Authentication** → **Basic Authentication**, as shown in Figure 16-6.

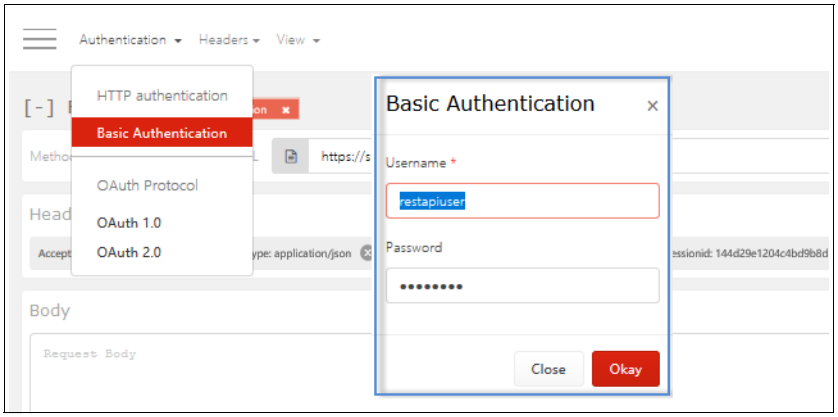

Figure 16-6 Create basic authentication

2. Define the headers (see Figure 16-7) to specify the type of accepted responses from the server and the message format for request from the client to the server (Content-Type).

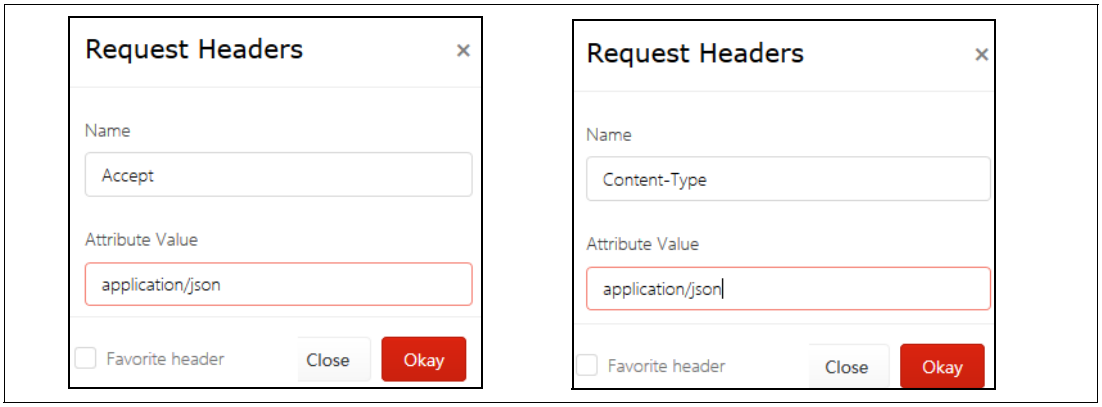

Figure 16-7 Request headers

Retrieve a sessionid

Complete the following steps to retrieve a `sessionid` by using RESTclient:

1. Select the POST operation.
2. Enter the URI of the session endpoint (for example, `https://spphostname/api/endeavour/session`) and click **Send** to submit the request to the IBM Spectrum Protect Plus Server.

   The RESTclient displays the response status code of 200, which indicates that the POST operation was successful. The corresponding **curl** command is generated at the bottom of the browser window (see Figure 16-8).

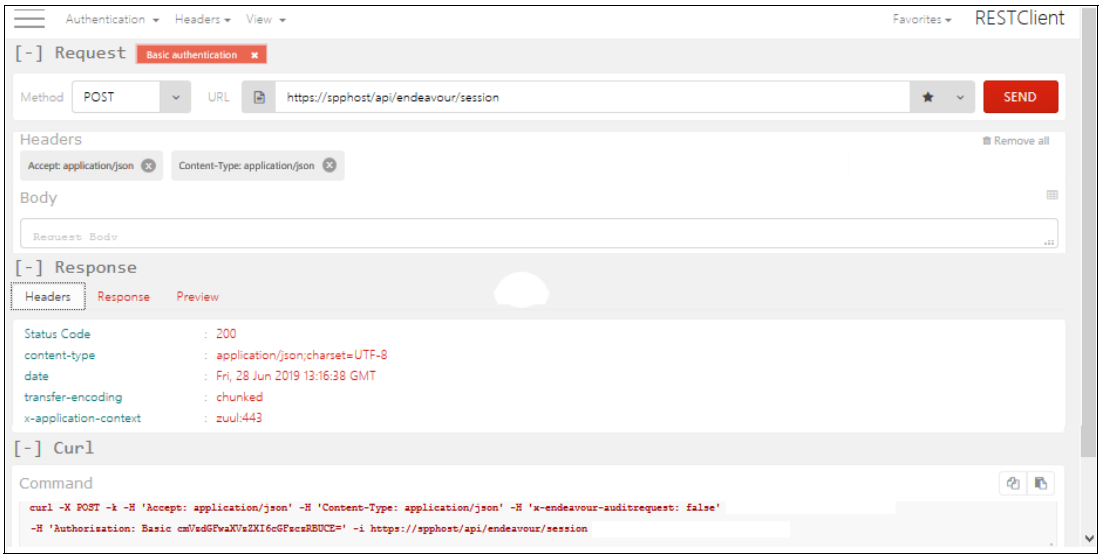

Figure 16-8 RESTclient with response header, status code and curl command

3. Select the **Preview** tab in the Response section to view the formatted JSON response from the API. The response contains session and user-related information from which we are interested only in the `sessionid` key-value pair, which is highlighted in Figure 16-9.

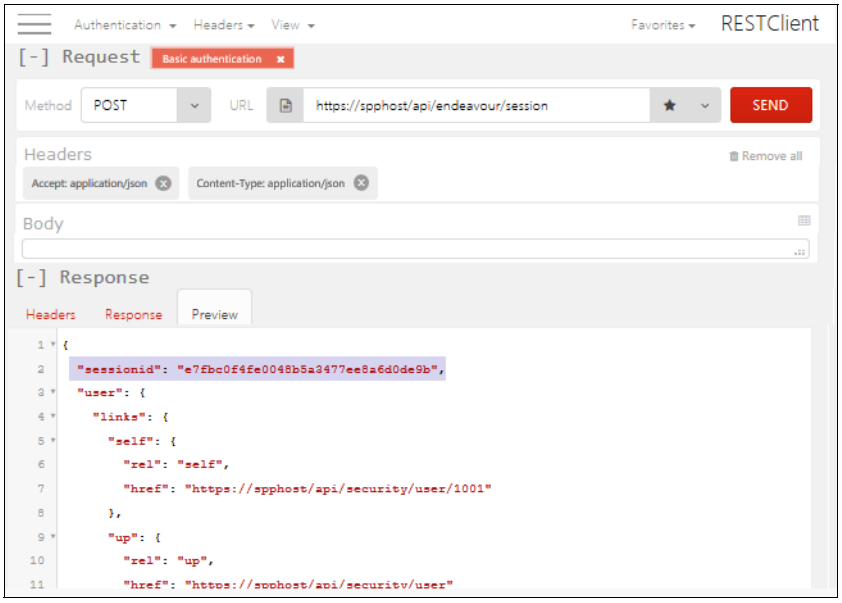

Figure 16-9 RESTclient with formatted output of the response

4. Copy the value of `sessionid` attribute and create another header, which is called `x-endeavour-sessionid` and paste the value of the previously retrieved `sessionid`, as shown in Figure 16-10.

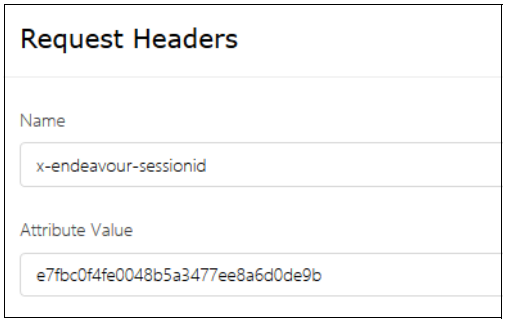

Figure 16-10 RESTclient: create request header x-endeavour-sessionid

Creating a site with RESTclient

As shown in Example 16-5 on page 469, we create a site by using the RESTclient. In the body field, we enter the JSON formatted text that is shown in Example 16-9.

Example 16-9 JSON formatted string for HTTP body

```
{ "name":"RESTclient", "defaultSite":"false", "description":"site created with
RESTclient" }
```

The RESTclient displays the response status code `201` (created successfully) and the **`curl`** command, as shown at the bottom of Figure 16-11, which accomplishes the same result.

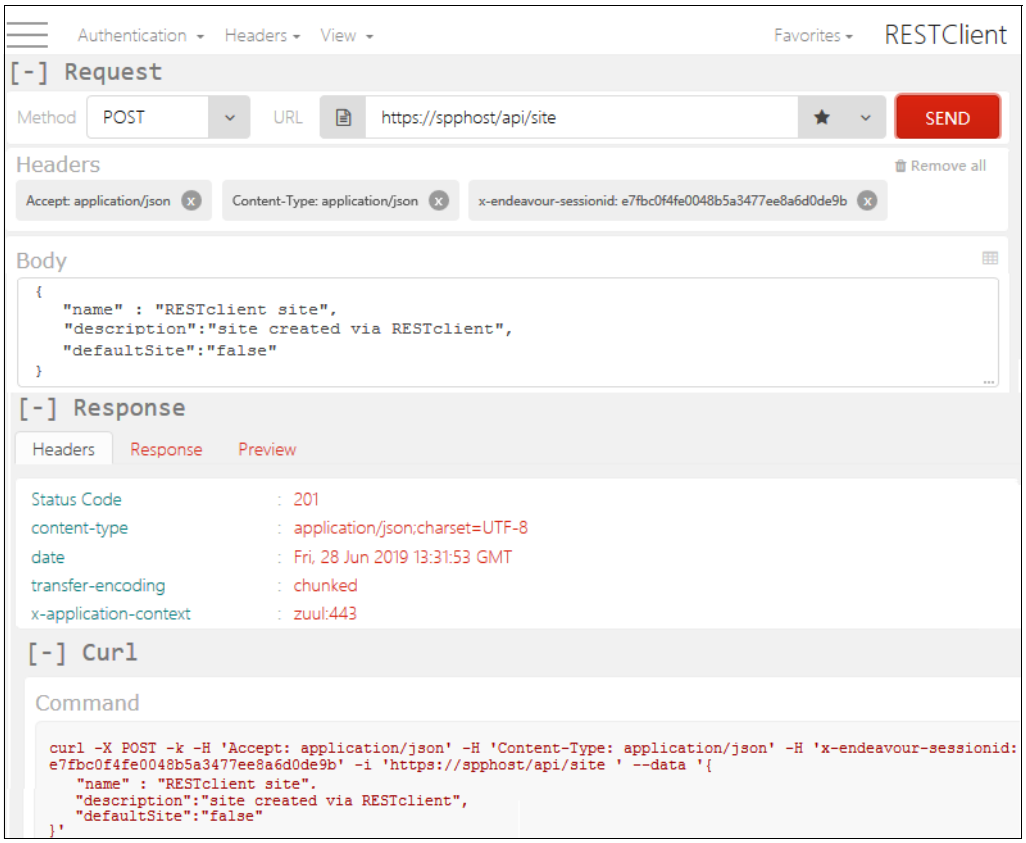

Figure 16-11 Create a new site via RESTclient

16.4 Use Case: Starting a VM backup of VMs by using a REST API with Python

This section describes the use case of starting a VM backup by using a REST API with Python.

16.4.1 Overview

You might want to start a backup of a single VM or a subset of available VMs, but not all VMs that are assigned to a specific SLA. This backup can be achieved by using the GUI and selecting the wanted VMs or by running a suitable REST operation. The following use cases show when a REST API-driven backup might be useful:

- Use of the GUI is not an option because of large amounts of hypervisors and VMs, which results in time-consuming GUI operations.
- An immediate required backup is a frequently occurring task.
- A backup of a VM is part of a customized and automated deployment workflow; for example, a VM is created by using the hypervisors REST API or other deployment techniques, and then it is added to a SLA at the IBM Spectrum Protect Plus Server by using scripts/REST operations. Users might want to create a backup and potentially a test restore or clone immediately as part of their deployment process.

Based on the client's requirements, we demonstrate how to identify the required steps to trigger an immediate VM backup by using REST commands. As the flow chart in Figure 16-12 on page 474 outlines, two main options are available to identify the REST endpoints or URIs, the suitable methods, and input and return values:

- The API endpoints are described in the API documentation.
- The IBM Spectrum Protect Plus GUI is communicating to the back-end server by way of REST calls (client side-started) and we can monitor the network traffic within the browser. A browser's development function is useful here because most modern browsers provide such functions:
  - Firefox calls them *web developer functions* (CTRL + SHIFT + E).
  - Chrome provides the *inspect* (CTRL + SHIFT + I) function.
  - In Microsoft Edge, select **More Tools** → **Developer tools (F12)**.

  As we demonstrate next, the browser inspect method is helpful because it helps to identify not only the wanted API endpoint, but also the real information about action, input parameters, and their values, and the API response status codes and API request responses in JSON format.

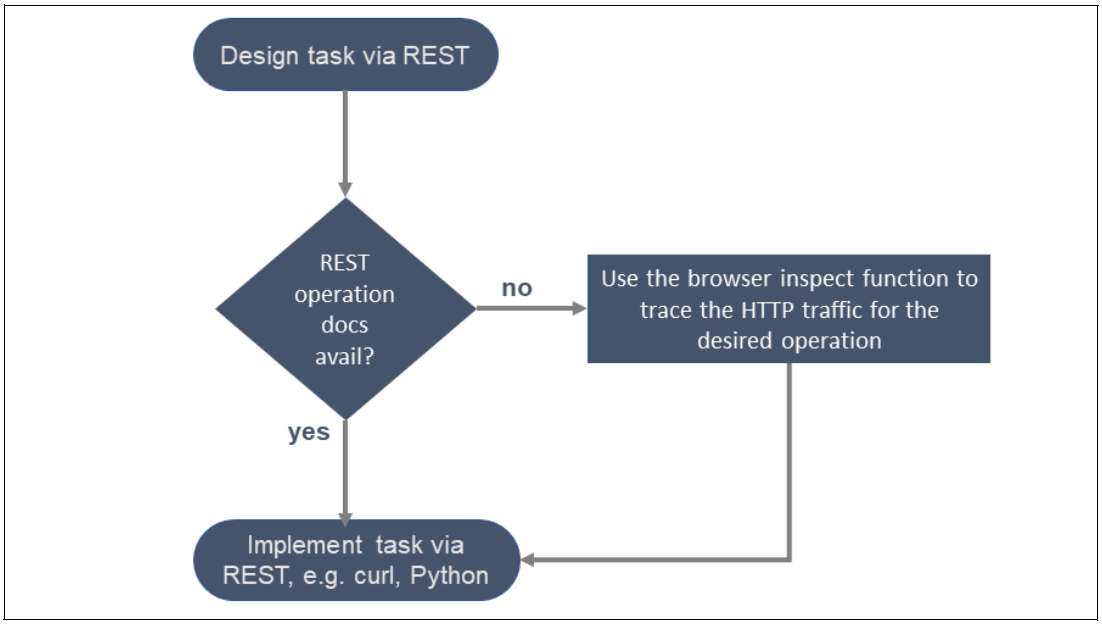

Figure 16-12 Flowchart: identify API calls

16.4.2 Trace GUI REST operations by using Firefox

In this use case preparation, we use the Firefox built-in Web Development tool (see Figure 16-13) to trace the network traffic that is generated by the IBM Spectrum Protect Plus Web GUI. The IBM Spectrum Protect Plus GUI runs REST API operations to the IBM Spectrum Protect Plus server and is populating the user interface with the required information based on the response. The web developer function can be started by selecting **Tools** → **Web Developer** → **Network** (keyboard shortcut CTRL + SHIFT + E).

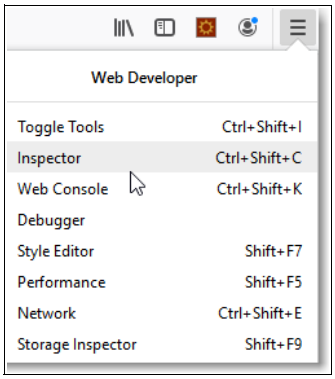

Figure 16-13 Firefox built-in Web Developer tool

After the network monitoring tool is started, return to the IBM Spectrum Protect Plus GUI and start a backup job for one or two VMs without running the entire SLA, which backs up all assigned VMs and not only a subset.

Select **Manage Protection** → **Hypervisors** → **VMware** and choose the wanted vCenter and the VMs. Before clicking **Run**, switch to the Web Developer window and clear the content.

For ease of debugging, it is recommended to switch from the GUI back to the Network monitoring window immediately after the backup run is started to pause the network monitoring.

As shown in Figure 16-14, the three functions are highlighted in red: clear, pause network recording, and the Network tab. After the backup is started from the GUI, the Networking monitoring window shows one or more entries. We are interested in the line with the API operation to the request the URL `/ngp/hypervisor?action=adhoc`.

When the specific line is selected, more information can be obtained from the Header tab (see Example 16-10) and the Params tab (see Example 16-11).

Example 16-10 Header tab shows the API operation POST and the API request URL

```
Request URL:https://spphost/ngp/hypervisor?action=adhoc
Request method:POST
```

From Example 16-10, we identified the required action (`POST`) and the request URL. From Example 16-11, we determined the parameter names and their values that must be passed with an API request within the request body, which is also referred to as *data* (in `curl`) or *payload* (in Firefox). Although `slaPolicyName` (for example, Gold, Silver, and Bronze) and `subtype` (`vmware` or `hyperv`) are self-explaining, the `resource` parameter is an array of one or more custom-built URLs, one for each VM to be backed up.

Example 16-11 Params tab shows the request payload

```
{ "slaPolicyName":"Silver", "subtype":"vmware",
  "resource":["https://spphost/api/hypervisor/1001/vm/1bebb22a3857d152c364c67579be380c?from=hlo"] }
```

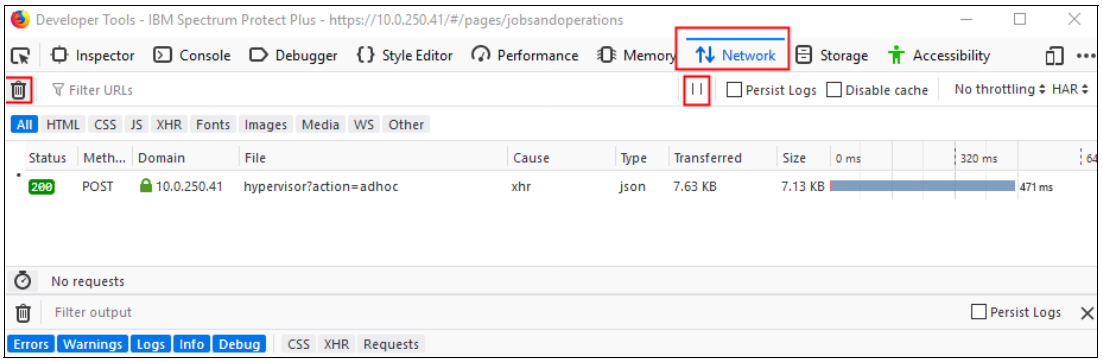

Figure 16-14 web developer function to inspect network traffic

In addition to the API URL and the method, we identified that we must pass more information with the API operation POST as data. This information can be captured during tracing of the network traffic with the web developer (see Figure 16-14).

The more complex resource parameters are built upon the API endpoint, which includes the `hypervisor id` and the `unique id` of the VM, as shown in Example 16-12.

Example 16-12 API endpoint URI

```
https://spphost/api/hypervisor/hypervisorId/vm/vmID?from=hlo
```

In 16.4.3, “Python code overview” on page 476, we describe the basic steps to implement a Python script that triggers a VM backup. The example does not reuse functions of the `sppclient` Python package; instead, it uses the native *request library*. The code does not implement command-line parsing for input parameters, such as `vmnames` to be backed up.

Exception handling, logging, and parameterization of such a script is highly recommended and is part of a real implementation in most cases. However, it is not part of the examples that follow. For the sake of ease, we use basic required steps and split it into pieces and define functions:

- To log in to the API with username and credentials and retrieve a `sessionid`.
- Log out and invalidate the `sessionid`.
- To retrieve (GET) information from the API by specifying an endpoint.
- Create a list of all registered hypervisors.
- Create a list of all VMs for a specific hypervisor.
- For starting a backup of a list of VMs

Outlook

In a productive implementation of such a Python script, the user might want to specify a list of VMs by their names instead of using their unique ID or pass a comma-separated value list of VMs to the script. The programmer must retrieve a list of hypervisors, then retrieve all VMs of each hypervisor and iterate through the list of objects (list of VMs) to find the wanted VM by its name. From hereon, it is possible to use the objects VM ID to build the resource string.

These extra steps make such a script much more comfortable and dynamically usable. However, because we are intending to introduce the basic steps to achieve a task over the introduction of Python programming techniques, we are keeping the functions as simple as possible.

16.4.3 Python code overview

For the sake of clarity, we split the script into smaller pieces and comment parts of the code briefly. In our Python script, we require several external modules that must be imported. The most important module is the `request` module, which is required to send HTTP requests to the IBM Spectrum Protect Plus API. After importing modules, we create an authentication token of type `HTTPBasicAuth` from the username and password, which is then sent to the API to retrieve a `sessionid`.

Although the log in function in Example 16-13 is responsible to retrieve the `sessionid` by way of a PUT request operation, the `logout` function in Example 16-14 on page 477 invalidates the `sessionid` by sending a DELETE operation to the API URL `/api/endeavour/session`.

The `sessionid` and its value are reused as parameters in HTTP headers throughout this script until the session is stopped at the end of the script. Each request returns a request status code that can be evaluated to determine whether the operation was carried out as expected.

Example 16-13 Login function - POST operation

```
import requests
import getpass
from requests.auth import HTTPBasicAuth
import urllib3
import sys

urllib3.disable_warnings(urllib3.exceptions.InsecureRequestWarning)

username  = 'restapiuser'
password  = getpass.getpass()

host      = "https://spphost:443/"
session   = requests.Session()
```

```
def login():
   authToken = HTTPBasicAuth(username, password)
   endpoint  = "/api/endeavour/session"
   url       = host + endpoint

   requestHeaders = {'Accept' : 'application/json', 'Content-type' : 'application/json'}

   response      = session.post(url, headers=requestHeaders, auth=authToken, verify=False)
   responseJson  = response.json()
   responseCode  = response.status_code

   sessionid     = responseJson['sessionid']

   if (responseCode != 200):
      print("request not successful. responseCode: " + responseCode)
      sys.exit(1)
   else:
      return sessionid
```

Although the log-in function that is shown in Example 16-13 on page 476 expects a return value of 200 for a successful authentication, the log-out function that is shown in Example 16-14 is expecting a response return code of 204 for successful DELETE operation. For a list of REST API response codes, see 16.6, “API response code” on page 485.

Example 16-14 Logout function - DELETE operation

```
def logout(sessionid):

    url        = host + "api/endeavour/session"
    requestHeaders = {'Accept': 'application/json', 'Content-type': 'application/json'}
    requestHeaders.update({'X-Endeavour-Sessionid': sessionid})
    response   = session.delete(url, headers=requestHeaders, verify=False)
    responseCode    = response.status_code

    if (responseCode != 204):
        print("request not successfull. responseCode: " + str(responseCode))
        sys.exit(1)
    else:
        print("\nLogout successful")
```

To query information from the API, the `query_endpoint` function is updating the header with the new `sessionid` so that the specification of username and password is no longer required. Also, the function is sending a request **`GET`** command to the specified URL. The `listHypervisors` function that is show in Example 16-15 is querying the API for a list of hypervisors and evaluating the response. The response contains the response status code and the response to the request in JSON format.

If the `response.status_code` returns 200 for a successful operation, the JSON response is stored in the `responseJson` variable. With Python, we can easily access the objects and variables by referencing them by their name. From a query of the endpoint/api/hypervisor with `curl`, we identified the structure of the returned JSON starting and determined that the response contains an array that is called `hypervisors`, which can be directly accessed by using **`hypervisorList=responseJson['hypervisors']`**. Inside the loop for each hypervisor, one line is printed with only selected attributes (here, the `id` and `name`).

Example 16-15 Code: function to query hypervisors

```
def query_endpoint(sessionid, endpoint):

    url = host + endpoint
    print("endpoint to query: " + url)
```

```
    requestHeaders = {'Accept': 'application/json', 'Content-type': 'application/json'}
    requestHeaders.update({'X-Endeavour-Sessionid': sessionid})

    return session.get(url, headers=requestHeaders, verify=False)

def listHypervisors(sessionid):
    endpoint = "/api/hypervisor"
    response = query_endpoint(sessionid, endpoint)
    if response.status_code == 200:
        print()
        hypervisor_fmt = " {:<5.5s} | {:<30.30s}"
        print(hypervisor_fmt.format("ID", "Hypervisor Name"))
        print("-" * 40)
        responseJson = response.json()
        hypervisorList = responseJson['hypervisors']

        for hypervisor in hypervisorList:
            print(hypervisor_fmt.format(hypervisor['id'], hypervisor['name']))
    print()
```

After the first functions are defined, we call the functions from the main function (see Example 16-16). We modify the main function while the script capabilities expand.

Example 16-16 Add main function and sample code to retrieve all registered Hypervisors

```
def main():
   sessionid = login()
   listHypervisors(sessionid)
   logout(sessionid)

if __name__ == "__main__":
    main()
```

When the script is started, the user is prompted to provide the user password. The username `restapiuser` and the server hostname `spphost` are hardcoded in this script. The function from Example 16-15 on page 477 generates the list that is shown in Example 16-17. In a later step, we reuse a hypervisor's ID to retrieve a list of all VMs that are registered to the hypervisor.

Example 16-17 Output of configured hypervisors

```
py listHypervisors.py
Password: *********


 ID    | Hypervisor Name
----------------------------------------
 1001  | vcenter-b.escc.workshop
 1002  | vcenter-a.escc.workshop
 1003  | vcenter-c.escc.workshop

Logout successful
```

To get a list of VMs for a specified hypervisor (by ID), a new endpoint URL is created. As shown in Example 16-18, we query hypervisor `vcenter-b.escc.workshop`, which results in the API URL: `/api/hypervisor/1001/vm`.

Example 16-18 Python code: query VMs by hypervisor id

```
def listVMs(sessionid, hypervisorID):
    endpoint = "/api/hypervisor/" + str(hypervisorID) + "/vm"
    response = query_endpoint(sessionid, endpoint)
    responseCode = response.status_code
```

```
    responseJson = response.json()
    print()

    vmlist_fmt = " {:<33.33s} | {:<20.20s} | {:<5s}"
    print(vmlist_fmt.format("VM ID", "VM Name", "Hypervisor ID"))
    print("-" * 60)

    for vm in responseJson['vms']:
        print(vmlist_fmt.format(vm['id'], vm['name'], str(vm['hypervisorManagementServerID'])))
```

With the function parameters `sessionid` and `hypervisorID=1001`, the list that is shown in Example 16-19 is generated from the function **`listVMs`**.

Example 16-19 List of VMs by Hypervisor (ID=1001)

```
 VM ID                                | VM Name              | Hypervisor ID
------------------------------------------------------------------------------
 d332c4589e51c5be1290538f25a73b3a | t4-vm-sql-template   | 1001
 d42842a87ba64c15b13f535986a240a5 | t4-vm-lx-template    | 1001
 1bebb22a3857d152c364c67579be380c | t4-vm-win-sklm1      | 1001
 a25b9e25c02b3db9dd0cfe4f1f514ffd | t4-mgmt-win          | 1001
```

As shown in Example 16-20, we define the function **`startBackup`** where we specify the endpoint to trigger the backup **`/ngp/hypervisor?action=adhoc`** and create the other parameters that must be specified as data of the HTTP POST operation. The most complex part is the definition of the resource array that contains one or more API URLs, one for each VM to be backed up.

Example 16-20 Python code to initiate a vm backup

```
def startBackup(sessionid, subtype, slaPolicyName, vmsToBackup):

    endpoint = "/ngp/hypervisor?action=adhoc"
    url = host + endpoint
    resource=[]
    requestHeaders = {'Accept': 'application/json', 'Content-type': 'application/json'}
    requestHeaders.update({'X-Endeavour-Sessionid': sessionid})

    for vm in vmsToBackup:
        vmresource = 'https://10.0.250.41/api/hypervisor/' + vm['hypervisorId'] + "/vm/" + vm['vmid'] + "?from=hlo"
        resource.append(vmresource)

    reqData = json.dumps({'subtype': subtype, "slaPolicyName": slaPolicyName, "resource": resource})
    print(reqData)

    response = session.post(url, headers=requestHeaders, data=reqData, verify=False)
    responseCode = response.status_code
    responseJson = response.json()
    print("response status code:" + str(responseCode))
    if(responseCode != 200):
        print("response JSON:" + json.dumps(responseJson))
```

Each resource element is built form the API endpoint for hypervisors plus the hypervisor ID and the VM ID, followed by the HTTP parameter **`?from=hlo`**. We can search through two primary catalogs in IBM Spectrum Protect Plus: `recovery` and `hlo`. `Recovery` shows resources that were been backed up and are available for restore. The `hlo` catalog lists resources that were inventoried from various registered providers. Example 16-21 shows a valid resource that is based on Example 16-20.

Example 16-21 Example URI

```
https://spphost/api/hypervisor/1001/vm/1bebb22a3857d152c364c67579be380c?from=hlo
```

Finally, we update the main function, declare the list **vmsToBackup**, and add two JSON dictionaries to the list with the properties `vmid`, `hypervisorid`, and `subtype`, as shown in Example 16-22. This list is passed to the function **startBackup** and the resource list with the API URLs is created.

Example 16-22 Python code: Updated main function with list of VMs to back up

```
def main():
    sessionid = login()

    #listHypervisors(sessionid)
    #listVMs(sessionid, 1001)

    vmsToBackup=[]
    vmsToBackup.append({'vmid':'1bebb22a3857d152c364c67579be380c','hypervisorId':'1001', 'subtype':'vmware'} )
    vmsToBackup.append({'vmid':'a25b9e25c02b3db9dd0cfe4f1f514ffd','hypervisorId':'1001', 'subtype':'hyperv'} )

    startBackup(sessionid=sessionid, subtype="vmware", slaPolicyName="Silver", vmsToBackup=vmsToBackup)

    logout(sessionid)
```

After the code that is shown in Example 16-22, the output looks similar to the output that is shown in Example 16-23. The script ran twice. Although the first run was successful, the second run failed with `status_code 404` because an SLA Policy of the same type (Silver) is still running. Only one job per SLA can be started at the same time.

Example 16-23 Output of backup REST command

```
py backupVM.py
{"subtype": "vmware", "slaPolicyName": "Silver", "resource":
["https://10.0.250.41/api/hypervisor/1001/vm/1bebb22a3857d152c364c67579be380c?from=hlo",
"https://10.0.250.41/api/hypervisor/1001/vm/d42842a87ba64c15b13f535986a240a5?from=hlo"]}

response status code:200 ==> OK

Logout successful

py backupVM.py
{"subtype": "vmware", "slaPolicyName": "Silver", "resource":
["https://10.0.250.41/api/hypervisor/1001/vm/1bebb22a3857d152c364c67579be380c?from=hlo",
"https://10.0.250.41/api/hypervisor/1001/vm/d42842a87ba64c15b13f535986a240a5?from=hlo"]}

response JSON:{"statusCode": 404, "response": {"rc": "Failed to run an already running SLA Policy \"Silver\""}}

Logout successful
```

16.5 The sppclient: a Python library for REST operations

The `sppclient` is a Python project that is provided by the IBM Spectrum Protect Plus development team. It is available from GitHub at this web page.

The client is provided AS-IS without guarantee to run and with no support. The project is a collection of scripts as part of the module `sppclient` with more than 20 examples for SPP REST API scripts written in Python.

The integration of the `sppclient` project eases interaction with SPP's REST API and is intended for testing, monitoring, and automation. The module is registered in the Python Package Index (PyPI) and can be installed with all other module requirements by running the **`pip install sppclient`** command.

For various aspects in the day-to-day business of using IBM Spectrum Protect Plus, the REST API can be helpful to simplify and automate repeating tasks and monitoring activities. The `sppclient` project is a convenient and easy way forward to implement script-based solutions for IBM Spectrum Protect Plus.

Although we demonstrated a simple Python implementation to run VM backups with native Python commands and no logging or exception handling, script is available (`vmwareadhocbackup.py`) under the sample section of the Git repository. This script realizes the same operations as the use case implementation that is described in 16.4.3, “Python code overview” on page 476. However, it adds capabilities, such as wrapper functions of the `sppclient` module, logging, and command-line parameter parsing.

Note: The `sppclient` with it’ sample scripts is provided as-is without support or guarantee to run error free. The `sppclient` is provided for demonstration purposes of REST API capabilities with Python and is not an official IBM tool nor the development of more scripts part of product plans. The `sppclient` and the sample scripts are built by users, testers, and developers to fulfill their needs for testing and monitoring.

Although can contribute to the project if you developed your own solutions that are based on the `sppclient` module, do not expect active issue tracking, development, or bug fixing.

16.5.1 The sppclient scripts: General usage information

All scripts are providing a list of valid parameters (**`--help`**) and at minimum expecting the parameters **`--host`**, **`--user`**, and **`--pass`** whereas the host is defined with the IP address or the fully qualified hostname, which is also resolvable from the IBM Spectrum Protect Plus server. Specify the format of the hostname with HTTPS protocol and the port number; for example, `https://t4-spp-server.workshop.escc`.

16.5.2 The sppclient script overview and selected examples

The number of script examples is growing and subject to change. Therefore, this section is not intended to be a comprehensive reference for all samples with input and output data.

We attach a list of script names while we introduce a subset of the list (highlighted in **bold** in the following list) with some other information. Most of the scripts are self-explaining by their script name or include a small help section that describes the purpose of the script.

The following sample scripts were implemented and made available in the GitHub repository:

- **appassigntosla.py**
- vmwaretestrestore.py
- createslapolicy.py
- deletesite.py
- filerestore.py
- **get_alerts.py**
- get_sessions.py
- **getJSON.py**
- joblist.py
- registerhypervisor.py
- registervsnap.py
- runjob.py
- spplogcollect.py
- sppvmbackupinfo.py
- sqladhocbackup.py
- sqlcopies.py
- sqlrestore.py
- storageList.py
- **systemInfo.py**

- validatevsnapstatus.py
- vmware_chargeback.py
- **vmwareadhocbackup.py**
- vmwareassigntosla.py
- vmwaretestrestore.py

appassigntosla.py

Use this script to assign one or more databases to an SLA policy. A comma-separated list of database names along with the wanted SLA name must be specified. The names are case-sensitive.

Syntax

```
Usage: appassigntosla.py [options]

Options:
  -h, --help       show this help message and exit
  --user=USERNAME  SPP Username
  --pass=PASSWORD  SPP Password
  --host=HOST      SPP Host, (ex. https://172.20.49.49)
  --type=TYPE      Application type: sql, oracle or db2
  --dbs=DBS        Database name(s) (comma seperated)
  --sla=SLA        SLA Policy Name
```

Example

```
python appassigntosla.py --user=restapiuser --pass=pass4AP! --host=https://10.0.250.41:443 --type=sql
--sla=Bronze --dbs="ESCC",”SQL_ITSO”

INFO:logger:Adding db ESCC to SLA Bronze
INFO:logger:Adding db SQL_ITSO to SLA Bronze
INFO:logger:dbs are now assigned
```

get_alerts.py

The intention of this script is to provide a command-line equivalent to the alert message display in the GUI. Various optional parameters can be specified to filter by timeframe, acknowledged state, and alert type, along with an optional full text search. The script shows only alerts and not job-related messages (check `get_messages.py`).

Syntax

```
Usage: get_alerts.py [options]

Options:
  -h, --help            show this help message and exit
  --user=USERNAME       SPP Username
  --pass=PASSWORD       SPP Password
  --host=HOST           SPP Host, (ex. https://172.20.49.49)
  --type=TYPE           type of alert: ERROR or WARN (optional)
  --ack=ACK             acknowledged: True or False (optional)
  --sort=SORT           sort order: DESC or ASC (optional)
  --timeframe=TIMEFRAME specify how many hours to look backwards: [int] (optional)
  --search=SEARCH       search within the alert message text (optional)
```

Example

```
python get_alerts.py --host="https://spphost" --user=restapiuser --pass=pass4AP!  --type=warn --search=low

   last occurrence   | Type | ackn | description

-------------------------------------------------------------------------------------------------------------------
2019-07-04 16:43:59 | WARN | True | 10.0.250.48 (Vsnap): free disk space 19,370 MB (18.92% free) lower than
threshold 20%.


python get_alerts.py --host="https://spphost" --user=restapiuser --pass=pass4AP!

last occurrence   | Type  | description
-------------------------------------------------------------------------------------------------
2019-06-19 15:45:42 | ERROR | Job vmware_Silver (id=1004, session=1,560,951,941,215) failed.
2019-06-19 17:32:22 | WARN  | Job vmware_Silver (id=1004, session=1,560,958,337,185) partially succe
2019-07-02 22:03:48 | WARN  | Job sklm_clone_twice (id=1007, session=1,562,097,600,075) partially su
```

```
2019-07-03 08:16:49 | ERROR | Job sklm_clone_twice (id=1007, session=1,562,134,593,162) failed.
2019-07-03 09:41:03 | WARN  | Job vmware_Silver (id=1004, session=1,562,139,576,243) partially succe
2019-07-03 10:43:32 | ERROR | Job vmware_Gold (id=1008, session=1,562,143,338,077) failed.
2019-07-04 16:43:59 | WARN  | 10.0.250.48 (Vsnap): free disk space 19,370 MB (18.92% free) lower tha
2019-07-08 13:36:39 | WARN  | Job Application Server Inventory (id=1010, session=1,562,585,757,121)
```

getJSON.py

The `getJSON.py` script is a simple Python script to query the IBM Spectrum Protect Plus REST API endpoints with GET operations. It is similar to what a user can do by using **`curl`** commands. However, the script provides some advantages over `curl`.

The `getJSON.py` handle the session (log in and log out) for the user; by using `curl`, a `sessionid` must be retrieved first and passed along in all GET operations in the headers. Also, the `getJSON.py` script generates a pretty-formatted output or an unformatted output (raw). It implements a verbose option that then also includes HATEOAS information, such as discoverable links.

The script is helpful if a user is evaluating the REST API responses step by step and might want to apply more commands, such as **`grep`** or **`jq`** for filtering of key-value-pairs or to identify the structure of the JSON formatted response for more programming.

The script provides another parameter that is intended for API discovery during script programming. When the user specifies the **`-a`** parameter, the **`--host`**, **`--user`**, and **`--pass`** parameters are obsolete because this information is stored in a file that is called `auth.txt`. The structure and content of the file are shown in Example 16-24.

Example 16-24 Structure of auth.txt file

```
host=https://10.0.250.41
username=restapiuser
password=pass4AP!
```

Syntax

```
Usage: getJSON.py [options]

Options:
  -h, --help            show this help message and exit
  --user=USERNAME       SPP Username
  --pass=PASSWORD       SPP Password
  --host=HOST           SPP Host, (e.g. https://172.20.49.49)
  --endpoint=URL, --url=URL
                        API endpoint, e.g. --endpoint="api/site/{siteID}"
  --filter=FILTER       optional, filter as JSON,  e.g.:
                        [{"property": "type", "op": "=", "value": "WARN"}]
  --sort=SORT           optional, sort as JSON,  e.g.:  [{"property": "name", "direction": "DESC|ASC"}]
  --pagesize=PAGESIZE   optional, number of max results
  -v                    verbose information, incl. links objects
  -a                    use file with host address and user credentials, use exclusivly to user, pass, and host
  -r                    do not format JSON, display RAW message
```

Example

The following example demonstrates the use of the **`-a`** flag. Because the **`-v`** (verbose) parameter is not specified, the HATEOAS information is not included in the output.

In addition, we limit the amount of returned object to a single object by specifying **`--pagesize=1`** and running the query on endpoint `api/site`. The filtering is occurring on the server and not within the script code.

The server returns only one site object in the response while more sites are defined. The total number of objects in the sites list `"total": 5`.

The simple example demonstrates how to query the API and gives the user or programmer an idea of the JSON response and its structure. This information is useful if the objects and its members are accessed; for example, in Python code.

```
py getJSON.py -a --endpoint="api/site" --pagesize=1

Endpoint:       https://10.0.250.41/api/site
restURL:        https://10.0.250.41/api/site?pageSize=1
{
    "page": 1,
    "sites": [
        {
            "defaultSite": true,
            "description": "Primary Site",
            "id": "1000",
            "name": "Primary",
            "rbacPath": "root:0/site:0/site:1000",
            "throttles": null
        }
    ],
    "total": 5
}
```

systemInfo.py

This script is inspired by the IBM Spectrum Protect command **query system** and intends to provide an at-a-glance overview of system and storage utilization, health, and database backups.

Example

```
py systemInfo.py -a

================================================
   SPP Server information:
================================================

   cpuUtil                 :         20.99
   memory_size             :         47.01
   memory_util             :          0.62
   compressionRatio        :          1.94
   deduplicationRatio      :          1.36
   sizeFreeAllStorage      :        841.15
   sizeTotalAllStorage     :        962.55
   sizeUsedAllStorage      :        121.40
   unavailable             :             1

================================================
   filesystems: 4
================================================

catalog name  | status |  GB Total |  GB used | GB free | % used |  type
-------------------------------------------------------------------------
Configuration | NORMAL |     48.11 |     2.78 |   45.32 |   5.78 |  None
File          | NORMAL |    145.55 |     0.46 |  145.09 |   0.32 |  None
System        | NORMAL |     25.35 |     4.64 |   20.71 |  18.32 |  None
Recovery      | NORMAL |     48.11 |     0.55 |   47.56 |   1.14 |  None

================================================
   catalog backups (backup, replication & offload)
================================================

   jobName                 :   catalog_SPP-backup
   status                  :            COMPLETED
   duration                :              0:04:10
   results                 :            COMPLETED
   start                   :  2019-04-04 07:00:00
   end                     :  2019-04-04 07:04:10
   nextFireTime            :  2019-04-04 12:00:00

jobName            | state    |backup time         |type    | policy| expires on
-------------------------------------------------------------------------------------
catalog_SPP-backup | COMPLETED|2019-07-01 07:00:00 |catalog | BACKUP | 2019-04-04 07:00:00
catalog_SPP-backup | COMPLETED|2019-07-01 19:00:00 |catalog | BACKUP | 2019-04-04 19:00:00
catalog_SPP-backup | COMPLETED|2019-07-01 12:00:00 |catalog | REPLIC | 2019-04-04 12:00:00
```

vmwareadhocbackup.py

This script can be used to trigger an ad hoc backup of a VM by specifying the VM's name. The **sla** parameter is optional and required only if the VM is assigned to more than one SLA.

Only one job is available per SLA running. Therefore, if the user intends to run a second VM backup with the same SLA, the user must wait for the first run to complete. The current implementation of the script allows only one VM name to be passed as argument.

Syntax

```
Usage: vmwareadhocbackup.py [options]

Options:
  -h, --help       show this help message and exit
  --user=USERNAME  SPP Username
  --pass=PASSWORD  SPP Password
  --host=HOST      SPP Host, (ex. https://172.20.49.49)
  --filter=FILTER  Filter for unique datacenter, cluster or folder name in case VM name is not unique (optional)
  --vm=VM          VM Name
  --sla=SLA        SLA policy to run if VM is assigned to multiple
```

Example

```
python vmwareadhocbackup.py --user=restapiuser --pass=pass4AP! --host=https://10.0.250.41:443
--vm="t4-vm-win-sklm1" --sla="Silver"

INFO:logger:Running backup job for vm t4-vm-win-sklm1
```

16.6 API response code

Table 16-1 lists the HTTP status codes. For additional information, see this web page.

Table 16-1 Status codes

Status code	Description
200 OK	The request completed successfully.
201 Created	A new resource was created successfully. The resource's URI is available from the response's Location header.
204 No Content	An update to a resource was applied successfully
400 Bad Request	The request was malformed. The response body includes an error that provides more information.
401 Unauthorized	Log in attempt with invalid credentials.
403 Forbidden	Generally related to permissions through Role Base Access Control.
404 Not Found	The requested resource did not exist.
405 Method Not Allowed	URL is unsupported.
500 Unrecoverable Error	Diagnosed in Virgo log.
503 Service Unavailable	Too many requests are going to the same controller.

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this paper.

IBM Redbooks

The following IBM Redbooks publication *Protecting the VMware Environment with IBM Spectrum Protect,* REDP-5252, provides more information about the topic in this document. This publication might be available in softcopy only.

You can search for, view, download, or order this document 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 Knowledge Center:

  https://www.ibm.com/support/knowledgecenter/en/SSNQFQ/landing/welcome_ssnqfq.html
- IBM Spectrum Protect Plus Support:

  https://www.ibm.com/mysupport/s/topic/0TO50000000IQWtGAO/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
- IBM Spectrum Protect Plus V10.1.x Update History:

  https://www.ibm.com/support/pages/update-history-ibm-spectrum-protect-plus-v101x
- Download Spectrum Protect Plus 10.1.6.x:

  https://www.ibm.com/support/pages/download-information-version-1016x-ibm-spectrum-protect-plus
- Download Spectrum Protect Plus 10.1.6x interim fixes:

  https://www.ibm.com/support/pages/node/6254732

- Known Issues and Limitations with IBM Spectrum Protect Plus V10.1.6.x:

  https://www.ibm.com/support/pages/known-issues-and-limitations-ibm-spectrum-protect-plus-v1016x

- vSnap Installation and User's Guide:

  https://www.ibm.com/support/knowledgecenter/SSNQFQ_10.1.6/spp/b_ispp_vsnap_guide_6.pdf?view=kc

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

REDP-5532-01

ISBN 0738459194

Printed in U.S.A.